



NSCT Prep

Free MCQ Practice for NSCT Test Preparation

All Subjects

13,410 Multiple Choice Questions · 10 Subjects

Subjects Included:

#	Subject	MCQs
1	AI / Machine Learning & Data Analytics	1,530
2	Computer Networks & Cloud Computing	900
3	Cyber Security	1,440
4	Databases	1,530
5	Data Structures & Algorithms	1,080
6	Operating Systems	1,080
7	Problem Solving & Analytical Skills	1,440
8	Programming (C++, Java, Python)	1,440
9	Software Engineering	1,440
10	Web Development	1,530

nsctprep.dev

This dataset is created and compiled by Muhammad Abdullah Awais

© 2026 NSCT Prep. All rights reserved.

SUBJECT

AI / Machine Learning & Data Analytics

1,530 Multiple Choice Questions

AI / Machine Learning & Data Analytics — Easy

510 questions

Q1. What does AI stand for?

- A. Advanced Iteration
- B. Artificial Intelligence
- C. Automated Integration
- D. Applied Informatics

Answer: B

Q2. Which of the following is a type of machine learning?

- A. Manual Learning
- B. Supervised Learning
- C. Static Learning
- D. Hardware Learning

Answer: B

Q3. Who is considered the father of Artificial Intelligence?

- A. John McCarthy
- B. Charles Babbage
- C. Alan Turing
- D. Tim Berners-Lee

Answer: A

Q4. What is the primary goal of machine learning?

- A. To design efficient computer hardware
- B. To replace human decision-making entirely
- C. To write self-modifying operating systems
- D. To enable computers to learn from data

Answer: D

Q5. Which of these is an example of AI in daily life?

- A. Saving a file to disk
- B. Using a basic calculator
- C. Printing a text document
- D. Spam email filtering

Answer: D

Q6. Data Analytics primarily deals with:

- A. Building and assembling computer hardware
- B. Writing low-level operating system code
- C. Examining data sets to draw conclusions
- D. Designing graphical user interface layouts

Answer: C

Q7. What does ML stand for in the context of AI?

- A. Machine Learning
- B. Meta Language
- C. Memory Linking
- D. Macro Logic

Answer: A

Q8. Which is NOT a type of machine learning?

- A. Reinforcement Learning
- B. Compiled Learning
- C. Supervised Learning
- D. Unsupervised Learning

Answer: B

Q9. A chatbot is an example of:

- A. Artificial Intelligence
- B. Database Management
- C. Operating System Design
- D. Computer Architecture

Answer: A

Q10. What is a dataset in data analytics?

- A. A sequential optimization algorithm
- B. A hardware peripheral device
- C. A compiled programming language
- D. A collection of related data points

Answer: D

Q11. What is a scalar in linear algebra?

- A. A higher-order tensor
- B. A multi-row data matrix
- C. A single numerical value
- D. A multi-element vector

Answer: C

Q12. What is the mean of the numbers 2, 4, 6, 8, 10?

- A. 7
- B. 8
- C. 5
- D. 6

Answer: D

Q13. A matrix with equal rows and columns is called:

- A. Rectangular matrix
- B. Diagonal matrix
- C. Square matrix
- D. Identity matrix

Answer: C

Q14. What does probability measure?

- A. The computational speed of an algorithm
- B. The overall size of a given dataset
- C. The likelihood of an event occurring
- D. The hardware accuracy of a processor

Answer: D

Q15. The derivative of a constant is:

- A. 1
- B. The constant itself
- C. 0
- D. Infinity

Answer: C

Q16. What is the dot product of vectors [1,2] and [3,4]?

- A. 7
- B. 14
- C. 11
- D. 10

Answer: C

Q17. What is the median of {1, 3, 5, 7, 9}?

- A. 3
- B. 5
- C. 4
- D. 7

Answer: B

Q18. An identity matrix has:

- A. All entries set to the value of zero
- B. Entries filled with random real numbers
- C. 1s on the diagonal and 0s elsewhere
- D. All entries set to the value of one

Answer: C

Q19. The range of a probability value is:

- A. -1 to 1
- B. Negative infinity to positive infinity
- C. 0 to 1
- D. 0 to 100

Answer: C

Q20. What is the transpose of a matrix?

- A. The matrix is numerically inverted
- B. The matrix entries are all squared
- C. The matrix is multiplied by itself
- D. Rows and columns are interchanged

Answer: D

Q21. Which Python library is most commonly used for numerical computations in AI?

- A. Tkinter
- B. Flask
- C. NumPy
- D. Django

Answer: C

Q22. What does pandas primarily provide?

- A. Network socket programming APIs
- B. Data structures for data analysis
- C. Web development framework tools
- D. Game development engine support

Answer: B

Q23. Which library is used for plotting graphs in Python?

- A. Pandas
- B. NumPy
- C. Matplotlib
- D. Scikit-learn

Answer: C

Q24. How do you import NumPy with an alias?

- A. include numpy
- B. import np
- C. import numpy as num
- D. import numpy as np

Answer: D

Q25. What is a Jupyter Notebook?

- A. An interactive computing environment for code and visualizations
- B. A production web server for hosting static site content
- C. A relational database for storing structured table records
- D. A lightweight plain text editor without any execution support

Answer: A

Q26. Which function creates a NumPy array?

- A. np.make()
- B. np.list()
- C. np.array()
- D. np.create()

Answer: C

Q27. What does df.head() do in pandas?

- A. Sorts the DataFrame by column values
- B. Returns the last 5 rows of the frame
- C. Deletes the first row of the DataFrame
- D. Returns the first 5 rows of the DataFrame

Answer: D

Q28. Which library provides machine learning algorithms in Python?

- A. Pillow
- B. Scikit-learn
- C. Seaborn
- D. Matplotlib

Answer: B

Q29. What data type does pandas use for tabular data?

- A. List
- B. Dictionary
- C. DataFrame
- D. Array

Answer: C

Q30. Which operator is used for element-wise multiplication in NumPy?

- A. *
- B. //
- C. &
- D. @

Answer: A

Q31. What is data preprocessing?

- A. Deploying trained models into production systems
- B. Writing detailed summary reports on raw findings
- C. Transforming raw data into a clean and usable format
- D. Creating interactive charts to visualize raw data

Answer: C

Q32. What is a missing value in a dataset?

- A. A data point that equals exactly zero
- B. A data point that is absent or null
- C. An extreme statistical outlier point
- D. A strongly negative numerical value

Answer: B

Q33. What is data normalization?

- A. Duplicating data across tables
- B. Sorting data in alphabetical order
- C. Scaling data to a standard range
- D. Removing all data from storage

Answer: C

Q34. Which of these is a common data format?

- A. SYS
- B. EXE
- C. CSV
- D. DLL

Answer: C

Q35. What does data cleaning involve?

- A. Making the dataset physically larger in size
- B. Adding more synthetic records to the data
- C. Encrypting data columns for secure storage
- D. Removing errors and inconsistencies from data

Answer: D

Q36. What is a categorical variable?

- A. A variable that can only hold binary data
- B. A variable with randomly assigned values
- C. A variable with continuous numerical values
- D. A variable with discrete categories or labels

Answer: D

Q37. What is a data source?

- A. A trained machine learning model file
- B. An interactive data visualization chart
- C. A computational optimization algorithm
- D. The origin from which data is collected

Answer: D

Q38. What is the purpose of removing duplicates?

- A. To add new input features
- B. To change column data types
- C. To eliminate redundant records
- D. To increase the dataset size

Answer: C

Q39. What is a structured dataset?

- A. Continuous raw audio waveforms
- B. Data organized in rows and columns
- C. Unorganized raw free-form text
- D. Unstructured pixel image data

Answer: B

Q40. Which pandas method checks for null values?

- A. fillna()
- B. isnull()
- C. notnull()
- D. dropna()

Answer: B

Q41. What does EDA stand for?

- A. Extended Data Algorithm
- B. External Data Analysis
- C. Effective Data Assessment
- D. Exploratory Data Analysis

Answer: D

Q42. Which chart is used to show the distribution of a single variable?

- A. Line chart
- B. Scatter plot
- C. Histogram
- D. Pie chart

Answer: C

Q43. What does a box plot display?

- A. The five-number summary: min, Q1, median, Q3, max
- B. Only the arithmetic mean of the distribution values
- C. Only the overall range from minimum to maximum
- D. Only the mode or most frequent value observed

Answer: A

Q44. Which pandas method gives summary statistics?

- A. describe()
- B. info()
- C. head()
- D. shape

Answer: A

Q45. A scatter plot shows the relationship between:

- A. Two numerical variables
- B. Categories data only
- C. Time series data only
- D. Textual data only

Answer: D

Q46. What is the purpose of a correlation matrix?

- A. To clean and preprocess the raw input data
- B. To deploy and manage software applications
- C. To train and evaluate machine learning models
- D. To show relationships between multiple variables

Answer: D

Q47. What does df.info() display in pandas?

- A. Only the numerical summary statistics values
- B. Only the count of missing or null values
- C. Column names, data types, and non-null counts
- D. Only the first five rows of the DataFrame

Answer: C

Q48. A bar chart is best used for:

- A. Showing data distributions
- B. Comparing quantities across categories
- C. Displaying feature correlations
- D. Plotting time series trends

Answer: D

Q49. What is the mode of a dataset?

- A. The full range of the values
- B. The most frequently occurring value
- C. The middle value of the data
- D. The average value of the data

Answer: C

Q50. Which plot is best for showing proportions of a whole?

- A. Histogram
- B. Pie chart
- C. Scatter plot
- D. Box plot

Answer: B

Q51. In supervised learning, the model learns from:

- A. Labeled data with known outputs
- B. Randomly generated fake data
- C. Unlabeled unstructured data only
- D. No data of any kind at all

Answer: A

Q52. Linear regression is used for:

- A. Classification tasks only
- B. Reducing dataset dimensionality
- C. Predicting continuous numerical values
- D. Clustering data points

Answer: D

Q53. What is classification in supervised learning?

- A. Grouping similar data together
- B. Predicting a discrete category or class
- C. Predicting a continuous output value
- D. Reducing feature dimensions

Answer: C

Q54. Which algorithm is commonly used for classification?

- A. K-Means method
- B. Decision Tree
- C. DBSCAN method
- D. PCA method

Answer: D

Q55. What is the target variable in supervised learning?

- A. The output the model predicts
- B. An internal training weight
- C. A tunable hyperparameter
- D. An individual input feature

Answer: A

Q56. Logistic regression is used for:

- A. Dimensionality reduction
- B. Linear regression
- C. Binary classification
- D. Clustering

Answer: C

Q57. What is overfitting?

- A. Model performs well on training data but poorly on new data
- B. Model is too simple and has very few learned parameters
- C. Model performs poorly on both training and testing data
- D. Model has absolutely no trainable weight parameters

Answer: A

Q58. K-Nearest Neighbors (KNN) classifies based on:

- A. A low-rank matrix decomposition factorization
- B. The steepest gradient descent optimization step
- C. An entirely random selection of a class label
- D. The majority class among the K nearest data points

Answer: D

Q59. What is a feature in machine learning?

- A. The model architecture itself
- B. The optimization training algorithm
- C. An input variable used for prediction
- D. The output target variable label

Answer: C

Q60. What is underfitting?

- A. Model has an excessive number of parameters
- B. Model has been trained for far too many epochs
- C. Model is too simple to capture patterns in data
- D. Model has memorized all the training data points

Answer: C

Q61. What is ensemble learning?

- A. Using a single standalone model for output
- B. Selecting the most relevant input features
- C. Combining multiple models to improve prediction
- D. Cleaning and preprocessing raw input data

Answer: C

Q62. Random Forest is an ensemble of:

- A. Neural networks
- B. Linear models
- C. Decision trees
- D. SVMs

Answer: C

Q63. What is bagging in ensemble learning?

- A. Training models on random subsets of data with replacement
- B. Cleaning and imputing missing data point values
- C. Training a single model on the complete full dataset
- D. Removing irrelevant features from the training data

Answer: A

Q64. What is boosting?

- A. Sequentially training models where each focuses on previous errors
- B. Reducing the overall size of the training dataset
- C. Engineering new features from the existing raw data
- D. Training all individual models simultaneously in parallel

Answer: A

Q65. The final prediction in a classification ensemble is typically made by:

- A. Using only the first model
- B. Random class selection
- C. Majority voting
- D. Using the weakest model

Answer: C

Q66. Which of these is a popular boosting algorithm?

- A. DBSCAN
- B. K-Means
- C. PCA
- D. XGBoost

Answer: D

Q67. How does Random Forest handle overfitting?

- A. By removing all features from the training input
- B. By using a single very deep decision tree only
- C. By only adding more raw data samples to train
- D. By averaging predictions from many diverse trees

Answer: D

Q68. What is a base learner in ensemble methods?

- A. A tunable model hyperparameter value
- B. An individual model within the ensemble
- C. A data preprocessing pipeline step
- D. The final aggregated combined model output

Answer: B

Q69. AdaBoost stands for:

- A. Automated Boosting
- B. Adaptive Boosting
- C. Advanced Boosting
- D. Additional Boosting

Answer: B

Q70. In Random Forest, what is feature randomness?

- A. Features are randomly generated from scratch each time
- B. Each tree considers a random subset of features at each split
- C. Features are removed permanently from the entire dataset
- D. All decision trees use all available features every time

Answer: A

Q71. Unsupervised learning works with:

- A. Unlabeled data
- B. Numerical data only
- C. Textual data only
- D. Labeled output data

Answer: A

Q72. K-Means is a type of:

- A. Regression algorithm
- B. Clustering algorithm
- C. Classification algorithm
- D. Sorting algorithm

Answer: B

Q73. What does clustering do?

- A. Classifies data with known labels
- B. Generates entirely new data points
- C. Predicts continuous numerical values
- D. Groups similar data points together

Answer: D

Q74. PCA stands for:

- A. Principal Component Analysis
- B. Primary Cluster Algorithm
- C. Probability Calculation Approach
- D. Pattern Classification Analysis

Answer: A

Q75. In K-Means, what is a centroid?

- A. A class target label
- B. An extreme outlier point
- C. The center point of a cluster
- D. A single input feature

Answer: C

Q76. What is dimensionality reduction?

- A. Adding more derived features to expand the feature space
- B. Removing all data from the storage system entirely
- C. Reducing the number of features while preserving important information
- D. Increasing the overall size of the raw training dataset

Answer: C

Q77. Which is an example of unsupervised learning?

- A. Predicting house prices
- B. Labeled image classification
- C. Customer segmentation
- D. Email spam detection

Answer: C

Q78. Association rule mining is used to:

- A. Find relationships between items in transactions
- B. Classify documents into predefined categories
- C. Predict future numerical stock price values
- D. Train multi-layer deep neural network models

Answer: A

Q79. What does the 'K' in K-Means represent?

- A. The number of iterations
- B. The number of features
- C. The learning rate
- D. The number of clusters

Answer: D

Q80. What is an unsupervised learning application?

- A. Labeled prediction
- B. Anomaly detection
- C. Supervised classification
- D. Regression

Answer: B

Q81. Accuracy is defined as:

- A. Number of correct predictions divided by total predictions
- B. Only the count of true negative classification results
- C. The total number of features used in the model
- D. Only the count of true positive classification results

Answer: A

Q82. What is a confusion matrix?

- A. A Python data visualization chart library
- B. A type of multi-layer deep neural network
- C. A preprocessing data cleaning utility tool
- D. A table showing predicted vs actual classifications

Answer: D

Q83. What is precision in classification?

- A. The recall evaluation metric
- B. Overall model accuracy metric
- C. True Positives divided by the total count
- D. $\text{True Positives} / (\text{True Positives} + \text{False Positives})$

Answer: A

Q84. What is recall (sensitivity)?

- A. $\text{True Positives} / (\text{True Positives} + \text{False Positives})$
- B. $\text{True Positives} / (\text{True Positives} + \text{False Negatives})$
- C. The specificity metric
- D. The overall accuracy metric

Answer: D

Q85. Mean Squared Error (MSE) is used to evaluate:

- A. Regression models
- B. Classification models only
- C. Clustering algorithms
- D. Association rule mining

Answer: A

Q86. What is a training set used for?

- A. Deploying the model to production
- B. Training the model to learn patterns
- C. Evaluating the final performance
- D. Tuning the model hyperparameters

Answer: B

Q87. What is a test set used for?

- A. Training the model on labeled training samples
- B. Tuning the model's hyperparameter settings
- C. Cleaning and preprocessing the raw data
- D. Evaluating the model's performance on unseen data

Answer: D

Q88. What is a false positive?

- A. Making a correct and accurate model prediction
- B. Predicting positive when the actual value is negative
- C. Encountering a missing value in the raw data
- D. Predicting negative when the actual value is positive

Answer: B

Q89. R-squared measures:

- A. The proportion of variance in the target explained by the model
- B. The total number of data samples in the dataset
- C. The total number of input features in the training dataset
- D. The total elapsed time required for model training

Answer: A

Q90. What is a validation set?

- A. The same set as the independent held-out test set
- B. The entire full dataset including all partitions
- C. A preprocessing data cleaning utility pipeline
- D. A subset used to tune hyperparameters during training

Answer: D

Q91. What is feature engineering?

- A. Building custom computer hardware for faster processing
- B. Creating new features from existing data to improve model performance
- C. Deleting all features from the dataset before model training
- D. Only using the original raw data without any transformations

Answer: B

Q92. What is a feature?

- A. A specific type of learning algorithm
- B. A particular model training method
- C. An input variable used by the model for prediction
- D. The output produced by a trained model

Answer: A

Q93. What is feature selection?

- A. Training the model on the full dataset
- B. Choosing the most relevant features for the model
- C. Deleting the entire dataset from storage
- D. Creating new derived features from raw data

Answer: B

Q94. What is a binary feature?

- A. A continuous numeric feature
- B. A missing feature value
- C. A feature with many distinct categories
- D. A feature with only two possible values (0 or 1)

Answer: C

Q95. What is a numerical feature?

- A. A feature with categorical text categories
- B. A feature containing image data
- C. A feature containing audio data
- D. A feature with continuous or discrete numeric values

Answer: A

Q96. Feature extraction involves:

- A. Deleting features from a dataset
- B. Deriving new features from raw data
- C. Renaming features in a table
- D. Copying features between datasets

Answer: B

Q97. What does dropping a feature mean?

- A. Removing a feature from the dataset
- B. Transforming a feature with a log
- C. Adding a brand new feature column
- D. Scaling a feature to unit range

Answer: A

Q98. Why might you create interaction features?

- A. To remove extreme outlier values from data
- B. To normalize data to a standard range
- C. To capture relationships between two or more features
- D. To reduce the overall size of the dataset

Answer: C

Q99. What is binning?

- A. Removing selected features from the dataset
- B. Adding new synthetic features to the data
- C. Converting continuous features into discrete intervals
- D. Sorting data records by a column value

Answer: C

Q100. What is a datetime feature?

- A. A feature representing dates and times
- B. A purely categorical feature only
- C. A strictly binary feature only
- D. A purely numerical feature only

Answer: B

Q101. What is a neural network?

- A. A network routing protocol for data packets
- B. A computing system inspired by biological neural networks
- C. A comparison-based sorting algorithm for arrays
- D. A relational database for storing structured records

Answer: B

Q102. What is a neuron (node) in a neural network?

- A. A basic computational unit that receives inputs and produces an output
- B. A physical data storage unit for persisting files on disk
- C. A specific type of input feature used in the dataset
- D. A gradient-based training algorithm for optimization

Answer: A

Q103. What is an activation function?

- A. A technique for initializing random weight parameters
- B. A function that introduces non-linearity into the network
- C. A function for loading data from external file sources
- D. A method for calculating the total loss of the model

Answer: B

Q104. What is a hidden layer?

- A. A data preprocessing transformation layer
- B. The very last output layer of the network
- C. The very first input layer of the network
- D. A layer between the input and output layers

Answer: D

Q105. What is the ReLU activation function?

- A. $f(x) = x^2$
- B. $f(x) = 1/(1+e^{-x})$
- C. $f(x) = \max(0, x)$
- D. $f(x) = \tanh(x)$

Answer: C

Q106. What is backpropagation?

- A. A type of non-linear activation function for neurons
- B. A specific deep neural network architecture design
- C. A data preprocessing technique for cleaning raw inputs
- D. An algorithm for computing gradients to update network weights

Answer: D

Q107. Which framework is commonly used for deep learning?

- A. NumPy only
- B. Flask
- C. TensorFlow
- D. Django

Answer: C

Q108. What is an epoch in training?

- A. A type of non-linear activation function node
- B. A single individual data point in the training set
- C. A single hidden layer within the neural network
- D. One complete pass through the entire training dataset

Answer: D

Q109. What is a loss function?

- A. An individual input feature column used for training
- B. A non-linear activation function applied at hidden layers
- C. A structured collection of labeled data for evaluation
- D. A function measuring how well predictions match actual values

Answer: D

Q110. What is gradient descent?

- A. A type of multi-layer feedforward neural network architecture with connections
- B. An optimization algorithm minimizing loss by updating weights via steepest descent
- C. A tree-based hierarchical data structure for nearest-neighbor lookups
- D. A non-linear activation function that maps inputs to the range zero to one

Answer: B

Q111. What is a Convolutional Neural Network (CNN)?

- A. A neural network designed for processing grid-like data such as images
- B. A relational database system for storing structured records
- C. A recurrent network designed exclusively for text sequence data
- D. A comparison-based sorting algorithm for numerical arrays

Answer: A

Q112. What is a Recurrent Neural Network (RNN)?

- A. A neural network for sequential data with memory of previous inputs
- B. A convolutional network designed exclusively for image pixel data
- C. An unsupervised density-based clustering algorithm for grouping
- D. A static feedforward network with no temporal connections

Answer: A

Q113. What is a GAN?

- A. An unsupervised clustering method based on density estimation
- B. A specialized variant of convolutional neural networks for images
- C. Generative Adversarial Network - two networks competing to generate realistic data
- D. A supervised regression algorithm for predicting numeric values

Answer: C

Q114. What is an autoencoder?

- A. A neural network that learns to compress and reconstruct its input
- B. A comparison-based sorting algorithm for numerical arrays
- C. A data loading utility for reading files into memory
- D. A supervised classifier that predicts categorical labels from data

Answer: A

Q115. What is LSTM?

- A. Long Short-Term Memory - an RNN variant handling long-term dependencies
- B. A cross-entropy loss function for multi-class classification
- C. A specialized type of convolutional neural network for images
- D. A structured data format for storing tabular information

Answer: A

Q116. What is a pooling layer in a CNN?

- A. A data loading layer for reading external files
- B. The final output layer for class predictions
- C. A layer that increases the total number of neurons
- D. A layer that reduces spatial dimensions by down-sampling

Answer: D

Q117. What is the purpose of a convolutional filter (kernel)?

- A. To add random noise for data augmentation purposes
- B. To store training data on disk for later retrieval access
- C. To increase the spatial resolution of the input image
- D. To detect specific features like edges or textures in input data

Answer: D

Q118. What is a pre-trained model?

- A. A randomly initialized model with no learned representations
- B. A model with absolutely no prior training on any data
- C. A model that has been permanently deleted from storage
- D. A model already trained on a large dataset for use on new tasks

Answer: D

Q119. What is a generative model?

- A. A model that only performs supervised categorical classification
- B. A model that only performs numerical regression prediction
- C. A model that learns to generate new data similar to training data
- D. A model that only performs unsupervised data point clustering

Answer: C

Q120. What framework was developed by Facebook (Meta) for deep learning?

- A. Caffe
- B. Keras
- C. PyTorch
- D. TensorFlow

Answer: C

Q121. What does NLP stand for?

- A. Numerical Linear Processing
- B. Network Logic Protocol
- C. Natural Language Processing
- D. Neural Language Programming

Answer: C

Q122. Tokenization in NLP means:

- A. Encrypting text for secure data transmission
- B. Splitting text into individual words or subwords
- C. Deleting text from the data pipeline
- D. Compressing text to reduce storage size

Answer: B

Q123. What is sentiment analysis?

- A. Summarizing long documents into shorter text passages
- B. Translating text from one natural language to another language
- C. Generating entirely new text from a given prompt input
- D. Determining the emotional tone of text (positive, negative, neutral)

Answer: D

Q124. Stop words are:

- A. Important domain-specific keywords carrying high information
- B. Common words like 'the' and 'is' often removed in processing
- C. Rare words that appear only once in the entire corpus
- D. Technical terms from a specialized subject vocabulary

Answer: B

Q125. What is text classification?

- A. Generating entirely new text from a prompt
- B. Translating text between different languages
- C. Editing and correcting text for grammar
- D. Assigning predefined categories to text documents

Answer: D

Q126. What is a corpus in NLP?

- A. A large collection of text documents
- B. A specific grammar rule definition
- C. A single word in a sentence
- D. A type of neural network model

Answer: A

Q127. What is stemming?

- A. Reducing words to their root form by removing suffixes
- B. Adding prefixes to words for morphological expansion
- C. Counting the total frequency of words in text
- D. Translating words between two different languages

Answer: A

Q128. What is lemmatization?

- A. Counting the total characters in a word
- B. Removing all vowels from the input text
- C. Adding suffixes for morphological inflection
- D. Reducing words to their base dictionary form

Answer: D

Q129. Named Entity Recognition (NER) identifies:

- A. Grammar errors found throughout the given text
- B. Spelling mistakes identified within the text body
- C. The overall length of each sentence in the document
- D. Named entities like people, organizations, and locations in text

Answer: A

Q130. What is machine translation?

- A. Performing manual text translation
- B. Generating new text from a prompt
- C. Summarizing text into shorter form
- D. Automatically translating text from one language to another

Answer: C

Q131. Computer vision enables computers to:

- A. Interpret and understand visual information from images and videos
- B. Design and build responsive web page user interfaces
- C. Manage and query relational database table records
- D. Process only natural language text data and documents

Answer: A

Q132. What is image classification?

- A. Segmenting an image into parts
- B. Assigning a label to an entire image
- C. Detecting objects in a given image
- D. Generating an image from scratch

Answer: C

Q133. What is a pixel?

- A. An input feature column name
- B. A gradient training algorithm
- C. A type of deep neural network
- D. The smallest unit of a digital image

Answer: D

Q134. What is object detection?

- A. Removing detected objects from an existing image frame
- B. Classifying the entire image into a single category label
- C. Identifying and locating objects within an image with bounding boxes
- D. Generating entirely new photorealistic synthetic images

Answer: C

Q135. RGB stands for:

- A. Random, Generated, Binary
- B. Recursive, Generative, Base
- C. Real, Gradient, Bias
- D. Red, Green, Blue

Answer: D

Q136. What is image segmentation?

- A. Resizing an image to different specified dimensions
- B. Partitioning an image into meaningful regions at the pixel level
- C. Rotating an image by a specified number of degrees
- D. Cropping an image to a smaller rectangular area

Answer: C

Q137. What does a convolution operation do to an image?

- A. Applies a filter to extract features like edges and textures
- B. Increases the image resolution to a higher pixel count
- C. Changes the file format of the image container file
- D. Deletes the image entirely from disk storage space

Answer: A

Q138. What is a grayscale image?

- A. An image with only shades of gray (single channel)
- B. A volumetric three-dimensional image stack
- C. A binary black and white image with two values
- D. A full-color image with three RGB channels

Answer: A

Q139. What is image resizing?

- A. Rotating the image by a specified angle value
- B. Changing the dimensions (width and height) of an image
- C. Deleting the image from the disk file system
- D. Changing the color space of the image channels

Answer: B

Q140. Face detection is an example of:

- A. Reinforcement learning
- B. Computer vision
- C. Data analytics
- D. Natural language processing

Answer: B

Q141. Big Data is characterized by:

- A. Only the fast speed of arrival
- B. Only the diverse data types
- C. Volume, Velocity, Variety (3 Vs)
- D. Only the large volume of data

Answer: C

Q142. What is Hadoop?

- A. A standalone relational database for transactional queries
- B. A multi-layer deep neural network architecture for AI
- C. An open-source framework for distributed storage and processing of big data
- D. A general-purpose compiled programming language for applications

Answer: C

Q143. What is MapReduce?

- A. A comparison-based sorting algorithm for small arrays
- B. A programming model for processing large datasets in parallel across a cluster
- C. A structured query language for relational database management
- D. A supervised machine learning algorithm for classification

Answer: B

Q144. What is Apache Spark?

- A. A production web server for hosting applications
- B. A standalone relational database management system
- C. A general-purpose programming language runtime
- D. A fast distributed computing engine for big data processing

Answer: D

Q145. What is a data warehouse?

- A. A centralized repository for structured data used for analysis and reporting
- B. A multi-layer deep neural network trained on data
- C. A physical room for storing server hardware equipment
- D. A transactional database used for daily operational queries

Answer: A

Q146. What does ETL stand for?

- A. Edit, Transfer, Link
- B. Extract, Transform, Load
- C. Encrypt, Transmit, Log
- D. Evaluate, Test, Learn

Answer: B

Q147. What is a data lake?

- A. A cleaned and normalized relational database table store
- B. An interactive dashboard for data visualization tools
- C. A storage repository holding raw data in its native format
- D. A small curated dataset for specific model training

Answer: C

Q148. What is distributed computing?

- A. Cloud storage of files without any computation tasks
- B. Spreading computation across multiple machines working together
- C. Using a single powerful dedicated computer for processing
- D. Interactive data visualization and charting tools

Answer: B

Q149. What is HDFS?

- A. A general-purpose interpreted programming language runtime
- B. A desktop operating system for personal computers
- C. A multi-layer deep neural network for classification
- D. Hadoop Distributed File System for storing data across a cluster

Answer: D

Q150. What is batch processing?

- A. Processing data interactively with user input prompts
- B. Processing data manually without any automation
- C. Processing large volumes of data as a group at scheduled intervals
- D. Processing data in real-time as each record arrives

Answer: C

Q151. What is MLOps?

- A. Only the model training phase of the ML lifecycle
- B. Only the data visualization phase of the lifecycle
- C. Practices for deploying and maintaining ML models in production
- D. Only the data collection phase of the ML lifecycle

Answer: C

Q152. What is model deployment?

- A. Collecting and gathering raw data from sources
- B. Making a trained model available for use in production
- C. Training a model on labeled training data samples
- D. Preprocessing and cleaning the raw input data

Answer: B

Q153. What is an API in the context of model deployment?

- A. A gradient-based training algorithm for optimizing model weights
- B. An interactive visualization tool for exploring model results
- C. An interface allowing applications to request predictions from the model
- D. A relational database for storing structured training data

Answer: C

Q154. What is model versioning?

- A. Tracking different versions of a model over time
- B. Deleting old models from the storage system
- C. Reducing model size through weight pruning
- D. Training models to converge significantly faster

Answer: A

Q155. What is a Docker container?

- A. A physical rack-mounted server housed in a data center
- B. A standalone relational database for transactional queries
- C. A multi-layer deep neural network model architecture
- D. A lightweight portable package containing everything to run an application

Answer: D

Q156. What is CI/CD in MLOps?

- A. An interactive data visualization and dashboard charting tool
- B. A standalone relational database management query system
- C. A general-purpose compiled programming language for building apps
- D. Continuous Integration and Continuous Deployment for automating ML pipelines

Answer: D

Q157. What is model monitoring?

- A. Collecting and gathering new raw data from external sources
- B. Training new models from scratch on fresh training data
- C. Building automated data preprocessing pipeline workflows
- D. Tracking model performance and behavior in production over time

Answer: D

Q158. What is a REST API?

- A. A gradient-based model training optimization method
- B. A type of relational database storage engine format
- C. An API architecture using HTTP methods for communication
- D. A general-purpose compiled programming language

Answer: C

Q159. What does reproducibility mean in ML?

- A. Using less training data to save storage and costs
- B. Running trained models faster by using better hardware
- C. The ability to recreate the same results with the same data and code
- D. Making models larger by adding more network layers

Answer: C

Q160. What is a model artifact?

- A. The saved output of training including model weights and configuration
- B. Only the source code used to build and train the model
- C. Only the runtime logs generated during model training
- D. Only the raw training data used during model training

Answer: A

Q161. What is AI bias?

- A. Systematic unfairness in AI predictions due to biased data or algorithms
- B. Slow processing speed during model inference time
- C. Random errors that occur during the model prediction process
- D. Hardware failures that cause incorrect output computations

Answer: A

Q162. What is data privacy?

- A. Encrypting only the physical hardware server equipment
- B. Making all collected data publicly available and open
- C. Deleting all collected data from every storage system
- D. Protecting personal information from unauthorized access and use

Answer: D

Q163. What is informed consent in AI?

- A. Sharing collected data freely with any third party entity
- B. Ignoring user preferences about their data usage rights
- C. Ensuring people understand and agree to how their data will be used
- D. Collecting personal data without any prior user permission

Answer: C

Q164. What is transparency in AI?

- A. Making AI decision-making processes understandable and open
- B. Making AI systems run faster and more efficiently
- C. Hiding how AI systems work from all stakeholders
- D. Using significantly more data for model training

Answer: A

Q165. What is the GDPR?

- A. A European regulation for data protection and privacy
- B. A deep neural network architecture for images
- C. A supervised machine learning training algorithm
- D. A general-purpose compiled programming language

Answer: A

Q166. Why is fairness important in AI?

- A. To use less training data during the model building phase
- B. To ensure AI systems treat all groups equitably without discrimination
- C. To make AI systems process data faster and more efficiently
- D. To reduce the financial costs of developing AI systems

Answer: B

Q167. What is accountability in AI?

- A. Automating every aspect without any human involvement
- B. Ensuring responsible parties can be identified when AI causes harm
- C. Making AI systems completely autonomous without oversight
- D. Removing all human oversight from AI decision processes

Answer: B

Q168. What is an adversarial attack on AI?

- A. Deliberately crafted inputs designed to fool AI systems
- B. A physical attack on the server hardware equipment
- C. A network attack on the communication infrastructure
- D. A SQL injection attack on the database system

Answer: A

Q169. What is algorithmic fairness?

- A. Making algorithms process data faster through hardware optimization
- B. Reducing the overall storage size of trained model files
- C. Using more diverse and larger training datasets for models
- D. Ensuring algorithms make decisions without discrimination on protected attributes

Answer: D

Q170. What is the purpose of AI regulations?

- A. To ensure safe fair and responsible development and use of AI
- B. To make AI systems significantly more expensive to build
- C. To deliberately slow down all AI research and development
- D. To limit all AI capabilities to only narrow simple tasks

Answer: A

Q171. What type of AI is designed to perform a single specific task?

- A. Meta AI
- B. Narrow AI
- C. General AI
- D. Super AI

Answer: B

Q172. Which term describes a machine's ability to imitate human conversation?

- A. Chatbot
- B. Compiler
- C. Web Scraping
- D. Data Mining

Answer: A

Q173. What does the term 'training data' refer to in machine learning?

- A. Data transmitted between two network endpoint nodes
- B. Data stored in a database for backup recovery use
- C. Data used to teach a model to recognize patterns
- D. Data used to evaluate final model accuracy scores

Answer: C

Q174. Which field combines statistics, algorithms, and computing to extract insights?

- A. Data Science
- B. Web Design
- C. Compiling
- D. Networking

Answer: A

Q175. What is reinforcement learning primarily based on?

- A. Clustering of similar items
- B. Labeled training datasets
- C. Manual rule programming
- D. Reward and penalty signals

Answer: D

Q176. Which of the following best describes deep learning?

- A. Learning restricted to tabular structured datasets
- B. Learning without any computational hardware need
- C. Learning from only a single data sample at once
- D. Learning using neural networks with many layers

Answer: D

Q177. What is the Turing Test designed to evaluate?

- A. A machine's intelligent behavior
- B. Memory capacity of a system
- C. Energy efficiency of a chip
- D. Processing speed of hardware

Answer: A

Q178. Which AI application converts spoken language into written text?

- A. Data Warehousing
- B. Image Segmentation
- C. Speech Recognition
- D. Load Balancing

Answer: C

Q179. What does the abbreviation ML stand for in the context of AI?

- A. Machine Learning
- B. Meta Language
- C. Model Linking
- D. Memory Loading

Answer: A

Q180. Which type of learning uses labeled examples to train a model?

- A. Unsupervised learning
- B. Supervised learning
- C. Reinforcement learning
- D. Transfer learning

Answer: B

Q181. What is a scalar in linear algebra?

- A. A vector of values
- B. A tensor of arrays
- C. A matrix of numbers
- D. A single numerical value

Answer: D

Q182. What does the dot product of two vectors produce?

- A. A new tensor
- B. A new matrix
- C. A single scalar
- D. A new vector

Answer: C

Q183. What is the mean of the numbers 2, 4, 6, 8, and 10?

- A. 6
- B. 7
- C. 4
- D. 5

Answer: A

Q184. Which operation flips a matrix over its diagonal?

- A. Transpose
- B. Rotation
- C. Inversion
- D. Reduction

Answer: A

Q185. What does a probability value of zero represent?

- A. A random event
- B. Impossible event
- C. A likely event
- D. A certain event

Answer: B

Q186. What is the derivative of a constant value?

- A. Infinity
- B. Undefined
- C. One
- D. Zero

Answer: D

Q187. Which term describes a rectangular array of numbers in rows and columns?

- A. Scalar
- B. Matrix
- C. Vector
- D. String

Answer: B

Q188. What is the range of a standard probability value?

- A. Negative one to one
- B. Minus ten to ten
- C. Zero to one only
- D. Zero to infinity

Answer: C

Q189. What does the summation symbol sigma represent?

- A. Ratio of values
- B. Product of terms
- C. Limit of terms
- D. Sum of a series

Answer: D

Q190. What is an identity matrix?

- A. A matrix with all zeros in every position
- B. A matrix with all elements equal to two
- C. A square matrix with ones on diagonal only
- D. A rectangular matrix with random entries

Answer: C

Q191. Which Python library is primarily used for numerical array computations?

- A. Flask
- B. NumPy
- C. Django
- D. Tkinter

Answer: B

Q192. What does the pandas function read_csv() do?

- A. Compresses CSV into zip
- B. Writes data to a CSV file
- C. Reads CSV into a DataFrame
- D. Deletes a CSV from disk

Answer: C

Q193. Which library is most commonly used for creating plots in Python?

- A. Logging
- B. Subprocess
- C. Requests
- D. Matplotlib

Answer: D

Q194. What data structure does pandas use to represent tabular data?

- A. Binary Tree
- B. DataFrame
- C. Linked List
- D. Dictionary

Answer: B

Q195. Which keyword is used to define a function in Python?

- A. function
- B. func
- C. def
- D. define

Answer: C

Q196. What does the shape attribute of a NumPy array return?

- A. The dimensions of array
- B. Total memory in bytes
- C. The maximum array value
- D. The data type of elements

Answer: A

Q197. Which Python library provides algorithms like SVM and Random Forest?

- A. Beautiful Soup
- B. SQLAlchemy
- C. Pillow Library
- D. Scikit-learn

Answer: D

Q198. What is a Python list comprehension used for?

- A. Deleting items from dictionary
- B. Creating lists with concise syntax
- C. Handling file system exceptions
- D. Connecting to remote databases

Answer: B

Q199. Which method checks for missing values in a pandas DataFrame?

- A. dropna()
- B. notna()
- C. fillna()
- D. isnull()

Answer: D

Q200. What does pip install do in Python?

- A. Runs Python unit tests
- B. Creates virtual environments
- C. Installs Python packages
- D. Compiles Python source code

Answer: C

Q201. What is data cleaning in the context of data preprocessing?

- A. Compressing data for storage
- B. Adding noise to the data
- C. Encrypting data for security
- D. Fixing errors and removing inconsistencies

Answer: D

Q202. Which technique fills missing values with the column average?

- A. Normalization step
- B. Mode imputation
- C. Deletion method
- D. Mean imputation

Answer: D

Q203. What does one-hot encoding convert?

- A. Categories into binary vectors
- B. Numbers into text labels
- C. Images into pixel arrays
- D. Audio into text transcripts

Answer: A

Q204. What is the purpose of removing duplicate records from a dataset?

- A. To add more noise and improve model robustness
- B. To reduce the number of features in the dataset
- C. To avoid bias from repeated identical observations
- D. To increase the total number of training samples

Answer: C

Q205. Which method scales features to a range between zero and one?

- A. Standardization
- B. Log transformation
- C. Min-max normalization
- D. Binning technique

Answer: C

Q206. What is a common file format for storing tabular datasets?

- A. CSV
- B. PNG
- C. MP3
- D. EXE

Answer: A

Q207. What does the term 'missing data' refer to in a dataset?

- A. Data stored in a different file location
- B. Values that are absent for some observations
- C. Values that are larger than the average value
- D. Data that has been encrypted for privacy

Answer: B

Q208. What is the first step typically performed in any data preprocessing pipeline?

- A. Data collection
- B. Hyperparameter tuning
- C. Model training
- D. Feature selection

Answer: A

Q209. Which type of data contains categories like color or country name?

- A. Numerical data
- B. Sequential data
- C. Continuous data
- D. Categorical data

Answer: D

Q210. What does the term 'outlier' refer to in a dataset?

- A. A value significantly different from others
- B. The average value of all observations
- C. A missing entry in a data column
- D. The most common value in the data

Answer: A

Q211. What does EDA stand for in data science?

- A. Extended Data Algorithm
- B. Extra Data Analysis
- C. External Data Approach
- D. Exploratory Data Analysis

Answer: D

Q212. Which type of chart is best for showing the distribution of a single variable?

- A. Pie chart
- B. Line chart
- C. Histogram
- D. Gantt chart

Answer: C

Q213. What does a scatter plot visualize?

- A. The hierarchical structure of a dataset
- B. The proportion of categories in a dataset
- C. The trend of a variable over time periods
- D. The relationship between two numerical variables

Answer: D

Q214. What is the median of a sorted dataset?

- A. The middle value in order
- B. The range of the data
- C. The arithmetic average
- D. The most frequent value

Answer: A

Q215. Which plot type shows the five-number summary of a dataset?

- A. Bar chart
- B. Area chart
- C. Box plot
- D. Scatter plot

Answer: C

Q216. What does a correlation coefficient of +1 indicate?

- A. No relationship at all
- B. Perfect positive linear
- C. Random relationship
- D. Perfect negative linear

Answer: B

Q217. Which visualization is best for comparing quantities across categories?

- A. Line chart
- B. Bar chart
- C. Heat map
- D. Histogram

Answer: B

Q218. What is the mode of a dataset?

- A. The most frequent value
- B. The largest value
- C. The average value
- D. The middle value

Answer: A

Q219. What does a line chart typically display?

- A. Trends over time
- B. Data distribution
- C. Category proportions
- D. Feature correlations

Answer: A

Q220. What is the purpose of a heatmap in data analysis?

- A. To compress images into smaller sizes
- B. To display magnitude of values using colors
- C. To show geographic locations on a map
- D. To play audio data as visual patterns

Answer: B

Q221. What is the goal of a supervised learning algorithm?

- A. Generate new data from scratch
- B. Learn a mapping from inputs to outputs
- C. Find hidden patterns without labels
- D. Group similar data points together

Answer: B

Q222. Which of these is a classification algorithm?

- A. Principal Component
- B. Linear Regression
- C. K-Means Clustering
- D. Logistic Regression

Answer: D

Q223. What does a decision tree use to split data at each node?

- A. Feature thresholds
- B. Random selection
- C. Output labels only
- D. Time stamps alone

Answer: A

Q224. What type of problem predicts a continuous numerical value?

- A. Clustering task
- B. Association task
- C. Classification task
- D. Regression task

Answer: D

Q225. Which algorithm finds the best-fit line through data points?

- A. Random Forest model
- B. K-Nearest Neighbors
- C. Apriori Algorithm
- D. Linear Regression

Answer: D

Q226. What does KNN stand for in machine learning?

- A. K-Null Notation
- B. K-Nearest Neighbors
- C. K-Node Navigation
- D. K-Normal Networks

Answer: B

Q227. In binary classification, how many possible output classes exist?

- A. Three
- B. Two
- C. One
- D. Four

Answer: B

Q228. What is a label in supervised learning?

- A. An input feature column
- B. The correct output answer
- C. A training epoch count
- D. A model parameter value

Answer: B

Q229. Which algorithm is often used for predicting house prices?

- A. Linear Regression model
- B. Naive Bayes classifier
- C. K-Means algorithm
- D. Apriori association rule

Answer: A

Q230. What does the term 'training' mean in supervised learning?

- A. Visualizing data in chart format
- B. Manually coding rules for predictions
- C. Adjusting model parameters using data
- D. Removing features from the dataset

Answer: C

Q231. What is an ensemble method in machine learning?

- A. Combining multiple models
- B. A data cleaning technique
- C. A single complex model
- D. A visualization approach

Answer: A

Q232. Which ensemble method trains models on random data subsets?

- A. Stacking
- B. Bagging
- C. Pruning
- D. Boosting

Answer: B

Q233. What is Random Forest built upon?

- A. Multiple logistic regression models combined together
- B. Multiple SVM models with different kernel functions
- C. Multiple neural networks of varying layer depths
- D. Multiple decision trees with random feature subsets

Answer: D

Q234. How does a majority voting ensemble make predictions?

- A. Using the most common prediction
- B. Using the maximum prediction
- C. Using the average prediction value
- D. Using the minimum prediction

Answer: A

Q235. What is the main advantage of ensemble methods over single models?

- A. They need less data input
- B. They always train faster
- C. They use less memory
- D. They reduce overall error

Answer: D

Q236. Which technique trains models sequentially to correct prior errors?

- A. Bagging
- B. Boosting
- C. Blending
- D. Stacking

Answer: B

Q237. What does the term 'base learner' mean in ensemble learning?

- A. The evaluation metric used
- B. An individual model in ensemble
- C. The training data subset used
- D. The final combined model

Answer: B

Q238. Which of these is a popular gradient boosting library?

- A. XGBoost
- B. NumPy
- C. Matplotlib
- D. Pandas

Answer: A

Q239. What is the purpose of bootstrap sampling in bagging?

- A. To select the best features
- B. To normalize feature values
- C. To create diverse training sets
- D. To remove outliers from data

Answer: C

Q240. How does Random Forest handle the prediction for classification tasks?

- A. Uses the prediction of the last tree only
- B. Uses the prediction of the first tree only
- C. Takes majority vote from all the trees
- D. Averages all tree predictions numerically

Answer: C

Q241. What is unsupervised learning?

- A. Learning from labeled data
- B. Learning with human feedback
- C. Learning from test data only
- D. Learning without any labels

Answer: D

Q242. What is the goal of clustering in unsupervised learning?

- A. Predicting a target value
- B. Generating synthetic data
- C. Grouping similar data points
- D. Reducing training time

Answer: C

Q243. Which algorithm groups data into K predefined number of clusters?

- A. K-Nearest Neighbors
- B. K-Fold Validation
- C. K-Means Clustering
- D. K-Mode Analysis

Answer: C

Q244. What is dimensionality reduction used for?

- A. Increasing the sample count
- B. Reducing the number of features
- C. Converting data to text format
- D. Adding more features to data

Answer: B

Q245. Which technique is a common dimensionality reduction method?

- A. Principal Component Analysis
- B. Logistic Regression
- C. Random Forest model
- D. Gradient Boosting

Answer: A

Q246. What is the centroid in K-Means clustering?

- A. The smallest value in a cluster
- B. The label assigned to a cluster
- C. The data point farthest from cluster
- D. The center point of a cluster

Answer: D

Q247. Which type of learning is used for customer segmentation?

- A. Supervised learning
- B. Reinforcement learning
- C. Unsupervised learning
- D. Semi-supervised learning

Answer: C

Q248. What does anomaly detection aim to find?

- A. The most common patterns
- B. The best classification model
- C. Unusual or rare data points
- D. The optimal feature weights

Answer: C

Q249. What is a dendrogram used to visualize?

- A. Neural network architecture
- B. Linear regression coefficients
- C. Hierarchical clustering results
- D. Confusion matrix of a model

Answer: C

Q250. Which algorithm does not require specifying the number of clusters beforehand?

- A. Mini-Batch K-Means
- B. K-Medoids
- C. K-Means
- D. DBSCAN

Answer: D

Q251. What is a confusion matrix used for?

- A. Visualizing data distributions
- B. Evaluating classification results
- C. Preprocessing raw data
- D. Training neural networks

Answer: B

Q252. What does accuracy measure in classification?

- A. The number of features used
- B. The proportion of correct predictions
- C. The size of the training data
- D. The speed of model training

Answer: B

Q253. What is the purpose of splitting data into training and test sets?

- A. To remove outliers from the original dataset
- B. To increase the training time of the model
- C. To reduce the total amount of data available
- D. To evaluate how well the model generalizes

Answer: D

Q254. What does precision measure in a classification model?

- A. The overall accuracy of the model on all classes
- B. The speed of prediction for each data sample
- C. The proportion of actual positives correctly identified
- D. The proportion of positive predictions that are correct

Answer: D

Q255. What does recall measure in a classification model?

- A. The proportion of actual positives correctly identified
- B. The number of features used by the trained model
- C. The computational time required for each prediction
- D. The proportion of positive predictions that are correct

Answer: A

Q256. What is cross-validation used for?

- A. Generating synthetic data
- B. Reducing the number of features
- C. Assessing model performance robustly
- D. Training the final production model

Answer: C

Q257. What is overfitting in the context of model evaluation?

- A. The model performs well on training but poorly on test data
- B. The model performs poorly on training data
- C. The model has too few parameters
- D. The model takes too long to train

Answer: A

Q258. What does the F1 score combine?

- A. Bias and variance
- B. Training and test error
- C. Precision and recall
- D. Accuracy and speed

Answer: C

Q259. What is a validation set used for?

- A. Testing the final model performance
- B. Training the model parameters
- C. Cleaning and preprocessing the data
- D. Tuning hyperparameters during development

Answer: D

Q260. What does underfitting indicate about a model?

- A. The model is too simple to capture patterns
- B. The model uses too many features
- C. The model has been trained too long
- D. The model is too complex for the data

Answer: A

Q261. What is feature engineering in machine learning?

- A. Evaluating model performance
- B. Training a model automatically
- C. Creating informative input variables
- D. Deploying models to production

Answer: C

Q262. What is feature selection?

- A. Adding all available features to the model without any filtering
- B. Converting feature names from one language to another one
- C. Generating entirely new features from random noise data
- D. Choosing the most relevant features for model training use

Answer: D

Q263. What does one-hot encoding create from a categorical feature?

- A. A single numerical column
- B. A text description field
- C. Multiple binary columns
- D. A sorted ordinal value

Answer: C

Q264. What is the purpose of scaling numerical features?

- A. To add missing value indicators
- B. To bring features to a similar range
- C. To change the data type of features
- D. To remove features from the data

Answer: B

Q265. Which technique creates new features by combining existing ones?

- A. Feature deletion
- B. Feature encryption
- C. Feature interaction
- D. Feature imputation

Answer: C

Q266. What is binning used for in feature engineering?

- A. Encrypting sensitive data for privacy
- B. Converting categories into continuous values
- C. Removing duplicate values from dataset
- D. Converting continuous variables into categories

Answer: D

Q267. What does a polynomial feature transformation create?

- A. Compressed versions of original data
- B. Lower dimensional representations
- C. Higher order combinations of features
- D. Random noise features for testing

Answer: C

Q268. What is the difference between ordinal and nominal encoding?

- A. Ordinal preserves category order while nominal treats categories as equal
- B. Ordinal only works with numbers while nominal only works with strings
- C. Both encodings are identical and always produce the same output values
- D. Ordinal treats categories as equal while nominal preserves their ordering

Answer: A

Q269. What is a derived feature?

- A. An original feature from raw data
- B. A new feature created from existing ones
- C. A feature with many missing values
- D. A feature removed during cleaning

Answer: B

Q270. Why is feature engineering important for model performance?

- A. It always makes models train faster
- B. It removes the requirement for training data
- C. It eliminates the need for model tuning
- D. Good features can significantly improve accuracy

Answer: D

Q271. What is a neural network inspired by?

- A. Database systems
- B. Computer circuits
- C. The human brain
- D. Mathematical proofs

Answer: C

Q272. What is an activation function in a neural network?

- A. A function that loads training data
- B. A function that saves the model
- C. A function that introduces non-linearity
- D. A function that cleans input data

Answer: C

Q273. What is a hidden layer in a neural network?

- A. A layer between input and output
- B. The output layer of the network
- C. The first layer receiving raw data
- D. The layer that stores data on disk

Answer: A

Q274. Which activation function outputs values between 0 and 1?

- A. Linear function
- B. Tanh function
- C. ReLU function
- D. Sigmoid function

Answer: D

Q275. What does the term 'epoch' mean in deep learning training?

- A. The initialization of model weight values
- B. One complete pass through all training data
- C. A single forward pass of one sample
- D. The final test evaluation of the model

Answer: B

Q276. What is the purpose of backpropagation in neural networks?

- A. To calculate and propagate error gradients
- B. To make forward predictions on data
- C. To load training data into memory
- D. To visualize the network architecture

Answer: A

Q277. What is a weight in a neural network?

- A. The number of layers in the network
- B. The learning rate of the optimizer
- C. A learnable parameter connecting neurons
- D. The size of the training dataset

Answer: C

Q278. What is the most commonly used activation function in hidden layers today?

- A. Sigmoid
- B. Softmax
- C. Linear
- D. ReLU

Answer: D

Q279. What does the output layer of a classification network produce?

- A. Hidden representations
- B. Class probability scores
- C. Gradient computations
- D. Raw pixel values

Answer: B

Q280. What is a bias term in a neural network?

- A. The rate at which the model learns from training data
- B. The preference of the model for certain training samples
- C. The error in predictions caused by insufficient data size
- D. An additional learnable parameter added to weighted sums

Answer: D

Q281. What does CNN stand for in deep learning?

- A. Central Neural Network
- B. Circular Neuron Network
- C. Convolutional Neural Network
- D. Connected Node Network

Answer: C

Q282. What does RNN stand for in deep learning?

- A. Recursive Node Network
- B. Recurrent Neural Network
- C. Regional Neuron Network
- D. Random Neural Network

Answer: B

Q283. What is a convolutional filter used for in CNNs?

- A. Storing training labels
- B. Detecting patterns in input
- C. Compressing the model size
- D. Removing noise from data

Answer: B

Q284. What is the purpose of a pooling layer in a CNN?

- A. To add more features to data
- B. To increase image resolution
- C. To reduce spatial dimensions
- D. To convert images to text

Answer: C

Q285. What type of data are RNNs primarily designed to handle?

- A. Tabular data
- B. Image data
- C. Graph data
- D. Sequential data

Answer: D

Q286. What is a GAN in deep learning?

- A. Grouped Attention Network
- B. Gradient Adjusted Network
- C. General Analysis Network
- D. Generative Adversarial Network

Answer: D

Q287. What is max pooling in a CNN?

- A. Taking the average value from a region
- B. Taking the median value from a region
- C. Taking the maximum value from a region
- D. Taking the minimum value from a region

Answer: C

Q288. What is an autoencoder used for?

- A. Generating adversarial examples
- B. Learning compressed data representations
- C. Tuning model hyperparameters
- D. Supervised classification tasks

Answer: B

Q289. What does LSTM stand for?

- A. Layered Spatial Tensor Model
- B. Local Statistical Testing Method
- C. Long Short-Term Memory
- D. Linear Sequential Training Model

Answer: C

Q290. What is the Transformer architecture primarily known for?

- A. Reinforcement learning games
- B. Image classification tasks
- C. Self-attention mechanism for sequences
- D. Database query optimization

Answer: C

Q291. What does NLP stand for?

- A. Neural Logic Processing
- B. Numerical Linear Programming
- C. Natural Language Processing
- D. Network Layer Protocol

Answer: C

Q292. What is tokenization in NLP?

- A. Encrypting text for security use
- B. Compressing text to save space
- C. Splitting text into smaller units
- D. Translating text between languages

Answer: C

Q293. What is sentiment analysis used for?

- A. Detecting the emotion or opinion in text
- B. Counting the words in a document
- C. Correcting grammar errors in writing
- D. Translating text to other languages

Answer: A

Q294. What is a stop word in natural language processing?

- A. A common word often removed in preprocessing
- B. A word that stops program execution
- C. A word that contains a typing error
- D. A word used only in technical documents

Answer: A

Q295. What is named entity recognition?

- A. Identifying names of programming functions in code
- B. Identifying proper nouns like people and places in text
- C. Identifying the most frequent word in a paragraph text
- D. Identifying syntax errors in written documents for fixes

Answer: B

Q296. What is stemming in text preprocessing?

- A. Reducing words to their root or base form
- B. Sorting words in alphabetical sequence
- C. Adding prefixes to expand word meanings
- D. Translating words to another language

Answer: A

Q297. What is a corpus in NLP?

- A. A neural network architecture
- B. A large collection of text data
- C. A visualization chart format
- D. A machine learning algorithm type

Answer: B

Q298. What does a language model predict?

- A. The next word in a sequence
- B. Hardware specifications
- C. Database query results
- D. Image classification labels

Answer: A

Q299. What is text classification in NLP?

- A. Translating text between two languages
- B. Converting speech to written text
- C. Generating images from descriptions
- D. Assigning categories to text documents

Answer: D

Q300. What is a word embedding?

- A. A word encrypted for secure storage
- B. A word stored in a database table
- C. A word highlighted in a text editor
- D. A dense vector representation of a word

Answer: D

Q301. What is computer vision?

- A. A computer's screen resolution settings
- B. AI field enabling machines to interpret images
- C. A method for compressing video file data
- D. A type of display monitor hardware

Answer: B

Q302. What is image classification?

- A. Converting images to text format
- B. Sorting images by file size on disk
- C. Assigning a label to an entire image
- D. Splitting an image into two halves

Answer: C

Q303. What does object detection identify in an image?

- A. The resolution of the image file
- B. The file format of the image data
- C. The color profile of the image data
- D. Objects and their bounding box locations

Answer: D

Q304. What is a pixel in digital image processing?

- A. The smallest unit of a digital image
- B. A color space conversion algorithm
- C. A type of image compression format
- D. A filter used to sharpen images

Answer: A

Q305. What does image segmentation do?

- A. Converts color images to grayscale
- B. Increases image resolution quality
- C. Divides image into meaningful regions
- D. Reduces the file size of images

Answer: C

Q306. What is data augmentation in computer vision?

- A. Reducing the resolution of all images
- B. Converting all images to black and white
- C. Deleting duplicate images from dataset
- D. Creating modified copies of training images

Answer: D

Q307. Which deep learning architecture is most associated with image processing?

- A. Long Short-Term Memory
- B. Generative Adversarial Network
- C. Recurrent Neural Network
- D. Convolutional Neural Network

Answer: D

Q308. What is face recognition technology used for?

- A. Converting faces into cartoon drawings
- B. Identifying or verifying a person's identity
- C. Compressing facial images to save space
- D. Measuring the resolution of face images

Answer: B

Q309. What is optical character recognition?

- A. Recognizing and extracting text from images
- B. Compressing text before saving to disk
- C. Converting text files to images
- D. Encrypting text within image files

Answer: A

Q310. What does RGB stand for in digital images?

- A. Random Generated Bits
- B. Red Green Blue
- C. Regular Grid Base
- D. Rapid Graphics Buffer

Answer: B

Q311. What are the three Vs of big data?

- A. Vector, Variable, Vault
- B. Version, Vendor, Virtue
- C. Vision, Value, Validity
- D. Volume, Velocity, Variety

Answer: D

Q312. What is Apache Hadoop used for?

- A. Building web applications
- B. Distributed storage and processing
- C. Designing user interfaces
- D. Creating mobile applications

Answer: B**Q313. What is Apache Spark?**

- A. A database management system
- B. A fast distributed computing engine
- C. A web development framework
- D. A machine learning algorithm

Answer: B**Q314. What does HDFS stand for in big data?**

- A. High Data File System
- B. Hardware Distribution Framework
- C. Hybrid Data Format Standard
- D. Hadoop Distributed File System

Answer: D**Q315. What is a data lake?**

- A. A small SQL database
- B. A machine learning algorithm
- C. A centralized repository for raw data
- D. A data visualization tool

Answer: C**Q316. What is MapReduce in big data processing?**

- A. A database query language
- B. A machine learning framework
- C. A parallel processing programming model
- D. A data visualization library

Answer: C**Q317. What is streaming data?**

- A. Data generated and processed continuously
- B. Data archived for long-term storage
- C. Data stored in static files
- D. Data compressed for efficient transfer

Answer: A**Q318. What is a data warehouse?**

- A. An unstructured data lake for raw files
- B. A tool for creating data visualizations
- C. A physical storage building for servers
- D. A structured repository for analyzed data

Answer: D**Q319. What is the purpose of data partitioning in big data systems?**

- A. To convert data from one format to another
- B. To divide data across nodes for parallel processing
- C. To encrypt data for security purposes
- D. To delete old data from the storage system

Answer: B

Q320. What is Apache Kafka primarily used for?

- A. Real-time data streaming and messaging
- B. Image processing tasks
- C. Creating web applications
- D. Training deep learning models

Answer: A

Q321. What does MLOps stand for?

- A. Machine Learning Operations
- B. Multi-Layer Optimization
- C. Machine Logic Operations
- D. Model Linear Operations

Answer: A

Q322. What is model deployment in machine learning?

- A. Training a model on data
- B. Collecting new training data
- C. Making a trained model available for use
- D. Deleting an old model version

Answer: C

Q323. What is a REST API commonly used for in ML deployment?

- A. Storing training data in databases
- B. Serving model predictions over HTTP
- C. Training models faster
- D. Visualizing model architecture

Answer: B

Q324. What is model versioning?

- A. Converting models between different formats
- B. Deleting older models from storage
- C. Tracking different versions of a trained model
- D. Training multiple models simultaneously

Answer: C

Q325. What is CI/CD in the context of MLOps?

- A. Calculated Input and Computed Derivation
- B. Continuous Integration and Continuous Deployment
- C. Computer Intelligence and Cloud Delivery
- D. Central Interface and Core Distribution

Answer: B

Q326. What is model monitoring in production?

- A. Training the model with new data
- B. Deleting unused model files from disk
- C. Tracking model performance after deployment
- D. Creating documentation for the model

Answer: C

Q327. What is a Docker container used for in MLOps?

- A. Storing large datasets for training
- B. Visualizing model predictions on charts
- C. Packaging applications with all dependencies
- D. Training deep learning models faster

Answer: C

Q328. What is an ML pipeline?

- A. A visualization of model accuracy metrics
- B. A single model training step
- C. A database for storing predictions
- D. An automated sequence of ML workflow steps

Answer: D

Q329. What is A/B testing used for in model deployment?

- A. Training two models at the same time
- B. Splitting data into training and test sets
- C. Merging two models into a single model
- D. Comparing two model versions with real users

Answer: D

Q330. What is a feature store in MLOps?

- A. A retail store that sells ML hardware
- B. A centralized repository for ML features
- C. A visualization tool for model features
- D. A database for storing raw training data

Answer: B

Q331. What is AI bias?

- A. A technique for improving model accuracy
- B. A type of neural network optimization method
- C. A method for collecting training data faster
- D. Systematic unfairness in AI system outputs

Answer: D

Q332. What is data privacy in the context of AI?

- A. Making all data publicly available online
- B. Sharing data freely between all companies
- C. Protecting personal information in AI systems
- D. Storing data without any encryption method

Answer: C

Q333. What does explainable AI aim to provide?

- A. Larger training datasets for models
- B. More complex model architectures
- C. Faster model training speeds
- D. Understandable AI decision reasoning

Answer: D

Q334. What is fairness in AI systems?

- A. Making AI models run on all hardware
- B. Minimizing the size of the trained model
- C. Ensuring equal treatment across different groups
- D. Maximizing accuracy on the training dataset

Answer: C

Q335. What is the GDPR in relation to AI and data?

- A. A European data protection regulation
- B. A programming language for AI development
- C. A type of neural network architecture
- D. A machine learning algorithm for privacy

Answer: A

Q336. What is an adversarial attack on an AI system?

- A. A method for training models faster now
- B. A technique for compressing model size
- C. Deliberately manipulating input to fool AI
- D. Improving model accuracy by adding data

Answer: C

Q337. What does transparency mean in AI ethics?

- A. Making AI code run more efficiently
- B. Reducing the size of AI model files
- C. Hiding model details from all users
- D. Being open about how AI systems work

Answer: D

Q338. What is informed consent in AI data collection?

- A. Automatically opting in all users to data sharing
- B. Getting user agreement before collecting data
- C. Collecting data without user knowledge
- D. Deleting all collected data after processing

Answer: B

Q339. What is the purpose of AI regulation?

- A. To prevent all AI research and development
- B. To restrict AI use to government agencies only
- C. To ensure AI systems are safe and responsible
- D. To make AI technology more expensive to build

Answer: C

Q340. What is algorithmic transparency?

- A. Making algorithms run faster on hardware
- B. Using transparent display screens for output
- C. Making AI decision processes understandable
- D. Hiding algorithm details from competitors

Answer: C

Q341. Which branch of AI focuses on enabling machines to learn from data?

- A. Robotics
- B. Machine Learning
- C. Computer Graphics
- D. Database Management

Answer: B

Q342. What is a robot in the context of AI?

- A. A software virus
- B. A physical or virtual agent that can perform tasks autonomously
- C. A type of database
- D. A programming language

Answer: B

Q343. Which of the following is an application of natural language processing?

- A. Image compression
- B. Voice assistants like Siri
- C. Video rendering
- D. Battery management

Answer: B

Q344. What does the term 'intelligent agent' mean in AI?

- A. A human expert
- B. An entity that perceives its environment and takes actions to achieve goals
- C. A computer virus
- D. A data table

Answer: B

Q345. Which of the following is NOT an AI application?

- A. Spam email filtering
- B. Self-driving cars
- C. Manual data entry on paper
- D. Recommendation systems

Answer: C

Q346. What is computer vision in AI?

- A. A way to improve monitor resolution
- B. The field enabling computers to interpret visual information from the world
- C. A type of encryption
- D. A database query language

Answer: B

Q347. What type of data does a supervised learning algorithm require?

- A. Only images
- B. Labeled data with input-output pairs
- C. Random noise
- D. Encrypted data

Answer: B

Q348. What is an AI model?

- A. A physical robot
- B. A mathematical representation learned from data to make predictions
- C. A type of hard drive
- D. A monitor display

Answer: B

Q349. Which industry commonly uses AI for fraud detection?

- A. Agriculture
- B. Banking and Finance
- C. Construction
- D. Textile manufacturing

Answer: B

Q350. What is the purpose of a recommendation system?

- A. To delete user accounts
- B. To suggest relevant items to users based on their preferences
- C. To compress files
- D. To format documents

Answer: B

Q351. What is a vector in mathematics?

- A. A single number
- B. An ordered list of numbers with magnitude and direction
- C. A type of graph
- D. A random variable

Answer: B

Q352. What does the term 'dimension' refer to in a dataset?

- A. The physical size of a hard drive
- B. The number of features or variables in the data
- C. The color depth of images
- D. The speed of computation

Answer: B

Q353. What is the formula for calculating the mean?

- A. Sum of values divided by count of values
- B. Product of all values
- C. Largest value minus smallest value
- D. Square root of variance

Answer: A

Q354. What is variance in statistics?

- A. The middle value of a dataset
- B. A measure of how spread out data points are from the mean
- C. The most common value
- D. The total count of data points

Answer: B

Q355. What is a function in mathematics?

- A. A random process
- B. A rule that assigns each input exactly one output
- C. A type of matrix
- D. A data storage format

Answer: B

Q356. What does the slope of a line represent?

- A. The y-intercept
- B. The rate of change between two variables
- C. The color of the line
- D. The length of the line

Answer: B

Q357. What is the difference between a row vector and a column vector?

- A. There is no difference
- B. A row vector is horizontal while a column vector is vertical
- C. Row vectors are larger
- D. Column vectors cannot be used in calculations

Answer: B

Q358. What is a conditional probability?

- A. The probability of an impossible event
- B. The probability of an event occurring given that another event has already occurred
- C. A probability greater than one
- D. The total probability of all events

Answer: B

Q359. What is an exponent in mathematics?

- A. A subtraction operation
- B. The number of times a base number is multiplied by itself
- C. A type of average
- D. A statistical test

Answer: B

Q360. What is a dictionary in Python?

- A. An ordered list of elements
- B. A collection of key-value pairs
- C. A type of loop
- D. A mathematical function

Answer: B

Q361. What does the len() function return in Python?

- A. The data type of an object
- B. The number of items in an object
- C. The memory size in bytes
- D. The last element of a list

Answer: B

Q362. Which Python library is used for deep learning with dynamic computation graphs?

- A. pandas
- B. matplotlib
- C. PyTorch
- D. scikit-learn

Answer: C

Q363. What does the range() function generate in Python?

- A. Random numbers
- B. A sequence of numbers
- C. A list of strings
- D. Boolean values

Answer: B

Q364. What is a tuple in Python?

- A. A mutable list
- B. An immutable ordered sequence of elements
- C. A type of dictionary
- D. A function definition

Answer: B

Q365. What does the type() function do in Python?

- A. Converts data types
- B. Returns the data type of an object
- C. Creates a new variable
- D. Deletes an object

Answer: B

Q366. What is the purpose of the import statement in Python?

- A. To export data to files
- B. To bring external modules and libraries into the current script
- C. To delete variables
- D. To create loops

Answer: B

Q367. What does df.describe() do in pandas?

- A. Deletes the DataFrame
- B. Generates descriptive statistics like mean, std, min, max for numerical columns
- C. Sorts the DataFrame
- D. Renames columns

Answer: B

Q368. What is a for loop used for in Python?

- A. Defining functions
- B. Iterating over a sequence of elements
- C. Importing libraries
- D. Handling exceptions

Answer: B

Q369. What does np.zeros() create in NumPy?

- A. An array filled with ones
- B. An array filled with zeros
- C. An empty array
- D. A random array

Answer: B

Q370. What is data transformation in preprocessing?

- A. Deleting all data
- B. Converting data from one format or structure to another suitable for analysis
- C. Printing data on paper
- D. Sending data over the internet

Answer: B

Q371. What is the purpose of splitting data into training and testing sets?

- A. To reduce dataset size
- B. To evaluate model performance on unseen data and detect overfitting
- C. To make training faster
- D. To remove missing values

Answer: B

Q372. What is a numerical feature in a dataset?

- A. A feature containing text labels
- B. A feature represented by numbers that can be measured or counted
- C. A feature with only two values
- D. A feature with missing values

Answer: B

Q373. What does it mean to drop a column in data preprocessing?

- A. Moving a column to a different table
- B. Removing an entire feature from the dataset
- C. Adding new data to a column
- D. Sorting a column in descending order

Answer: B

Q374. What is label encoding used for?

- A. Adding labels to charts
- B. Converting categorical text values into numerical codes
- C. Encrypting data for security
- D. Splitting data into training sets

Answer: B

Q375. What is the purpose of data validation in preprocessing?

- A. Making data look prettier
- B. Checking that data meets expected quality standards and constraints before processing
- C. Encrypting sensitive data
- D. Compressing data files

Answer: B

Q376. What is a feature vector?

- A. A type of virus
- B. A numerical representation of an object's features used as input to an ML model
- C. A graph showing feature importance
- D. A physical measurement tool

Answer: B

Q377. What does data integration involve in preprocessing?

- A. Splitting data into smaller pieces
- B. Combining data from multiple sources into a unified dataset
- C. Deleting duplicate records only
- D. Changing data types

Answer: B

Q378. Why might you convert text to lowercase during preprocessing?

- A. To save storage space
- B. To ensure consistent treatment of words regardless of capitalization
- C. To make text unreadable
- D. To increase processing speed

Answer: B

Q379. What is the purpose of exploratory data analysis?

- A. To deploy models to production
- B. To understand data patterns, distributions, and relationships before modeling
- C. To encrypt sensitive data
- D. To compress data files

Answer: B

Q380. What does a histogram display?

- A. The relationship between two variables
- B. The frequency distribution of a single variable
- C. A geographical map
- D. A network diagram

Answer: B

Q381. What is the range of a dataset?

- A. The average value
- B. The difference between the maximum and minimum values
- C. The most frequent value
- D. The middle value

Answer: B

Q382. What does a negative correlation between two variables mean?

- A. Both variables increase together
- B. As one variable increases the other tends to decrease
- C. The variables are identical
- D. There is no relationship

Answer: B

Q383. What is the purpose of a count plot?

- A. To show the count of numerical values
- B. To show the frequency of each category in a categorical variable
- C. To display time series data
- D. To plot mathematical functions

Answer: B

Q384. What does df.shape return in pandas?

- A. The data types of columns
- B. A tuple showing the number of rows and columns in the DataFrame
- C. The column names
- D. The sum of all values

Answer: B

Q385. What is a missing value indicator in EDA?

- A. A value of zero
- B. A marker like NaN or null that indicates data was not recorded
- C. A negative number
- D. The most frequent value

Answer: B

Q386. What is a stacked bar chart used for?

- A. Showing 3D data
- B. Comparing the composition of categories across different groups
- C. Displaying continuous distributions
- D. Plotting geographic data

Answer: B

Q387. What is the mean absolute deviation?

- A. The square of variance
- B. The average of the absolute differences between each value and the mean
- C. The maximum deviation from the mean
- D. The standard deviation squared

Answer: B

Q388. Why is it important to check the data types of columns during EDA?

- A. Data types affect font size in visualizations
- B. Incorrect data types can cause errors in computations and analysis, such as numbers stored as strings
- C. Data types are irrelevant for analysis
- D. All columns should always be strings

Answer: B

Q389. What is the difference between classification and regression?

- A. They are the same thing
- B. Classification predicts categories while regression predicts continuous numerical values
- C. Regression predicts categories while classification predicts numbers
- D. Neither involves prediction

Answer: B

Q390. What is a decision boundary in classification?

- A. The edge of a dataset
- B. The boundary separating different class regions in feature space
- C. A programming error
- D. The maximum number of features allowed

Answer: B

Q391. What does the term 'prediction' mean in supervised learning?

- A. Guessing without any basis
- B. The output generated by a trained model for a given input
- C. Storing data in a database
- D. Visualizing data

Answer: B

Q392. What is a hyperparameter in a machine learning model?

- A. A parameter learned from training data
- B. A configuration set before training that controls the learning process
- C. A type of feature in the dataset
- D. The model's final prediction

Answer: B

Q393. What is the goal of minimizing a loss function?

- A. To increase errors
- B. To find model parameters that produce predictions closest to actual values
- C. To maximize data size
- D. To remove features

Answer: B

Q394. What is a leaf node in a decision tree?

- A. The root of the tree
- B. A node at the bottom that contains the final prediction or class label
- C. An intermediate splitting node
- D. The input layer

Answer: B

Q395. What is gradient descent used for in supervised learning?

- A. Sorting data
- B. Optimizing model parameters by iteratively moving in the direction of steepest decrease in the loss function
- C. Creating visualizations
- D. Splitting datasets

Answer: B

Q396. What is multiclass classification?

- A. A problem with exactly two output classes
- B. A classification task where the model predicts one of three or more possible classes
- C. A regression problem
- D. An unsupervised learning task

Answer: B

Q397. What is the purpose of the intercept term in linear regression?

- A. To multiply features together
- B. To represent the predicted value when all features are zero
- C. To remove outliers
- D. To normalize the data

Answer: B

Q398. What is a training error?

- A. An error in the code
- B. The difference between the model's predictions and actual values on the training data
- C. A hardware malfunction
- D. A missing value in the dataset

Answer: B

Q399. What is the wisdom of crowds principle in ensemble learning?

- A. Only the best model matters
- B. Combining multiple diverse models often produces better predictions than any single model alone
- C. Crowds always make wrong decisions
- D. Ensembles use only one model

Answer: B

Q400. What does the term 'weak learner' mean in boosting?

- A. A model that never makes predictions
- B. A model that performs only slightly better than random guessing
- C. A model with perfect accuracy
- D. A model that uses no data

Answer: B

Q401. What is the primary purpose of averaging predictions in a regression ensemble?

- A. To increase variance
- B. To reduce prediction errors by smoothing out individual model mistakes
- C. To make predictions more extreme
- D. To delete outlier predictions

Answer: B

Q402. Which of the following is an example of a bagging algorithm?

- A. AdaBoost
- B. XGBoost
- C. Random Forest
- D. Logistic Regression

Answer: C

Q403. What happens to overfitting when you add more trees to a Random Forest?

- A. Overfitting always increases dramatically
- B. Adding more trees generally does not increase overfitting, it tends to stabilize performance
- C. The model becomes simpler
- D. All trees become identical

Answer: B

Q404. What is the role of the number of estimators hyperparameter in ensemble methods?

- A. It controls the number of features used
- B. It specifies how many base models are combined in the ensemble
- C. It determines the learning rate
- D. It sets the maximum depth of trees

Answer: B

Q405. What is a simple averaging ensemble?

- A. A single model that averages features
- B. An ensemble that takes the arithmetic mean of predictions from multiple models
- C. A feature engineering technique
- D. A data preprocessing step

Answer: B

Q406. Why are decision trees commonly used as base learners in ensemble methods?

- A. They are the only algorithm available
- B. They are fast to train, can model nonlinear relationships, and their high variance is reduced by ensemble aggregation
- C. They never overfit
- D. They always produce identical results

Answer: B

Q407. What is weighted voting in ensemble classification?

- A. All models have equal influence
- B. Each model's vote is weighted by its estimated quality or confidence
- C. Only the best model votes
- D. Weights are applied to features

Answer: B

Q408. What is gradient boosting?

- A. An algorithm that uses random sampling like bagging
- B. A boosting method that sequentially fits models to the residual errors of previous models
- C. A single decision tree algorithm
- D. An unsupervised clustering method

Answer: B

Q409. What is the difference between clustering and classification?

- A. They are the same task
- B. Clustering groups data without predefined labels while classification assigns predefined class labels
- C. Classification does not use labels
- D. Clustering requires labeled data

Answer: B

Q410. What is the purpose of K-Means initialization?

- A. To normalize the data
- B. To set the initial positions of cluster centroids before the iterative algorithm begins
- C. To determine the number of features
- D. To clean the data

Answer: B

Q411. What is market basket analysis?

- A. Analyzing shopping cart designs
- B. Using association rules to find products frequently purchased together
- C. A type of regression
- D. A supervised learning technique

Answer: B

Q412. What is the objective of PCA?

- A. To increase data dimensions
- B. To find directions of maximum variance and reduce dimensionality while preserving information
- C. To classify data into groups
- D. To remove outliers

Answer: B

Q413. What is the elbow in the elbow method for K-Means?

- A. The first cluster
- B. The point on the inertia curve where adding more clusters gives diminishing improvement, resembling an elbow shape
- C. The last data point
- D. The maximum number of features

Answer: B

Q414. What is feature reduction?

- A. Adding more features to a dataset
- B. Decreasing the number of features while retaining important information
- C. Deleting all features
- D. Normalizing feature values

Answer: B

Q415. What is a cluster center?

- A. The first data point in a cluster
- B. The representative point of a cluster, often the mean of all points in the cluster
- C. The farthest point from other clusters
- D. A data point that belongs to no cluster

Answer: B

Q416. What is the difference between PCA and feature selection?

- A. They are identical techniques
- B. PCA creates new composite features from original features while feature selection chooses a subset of original features
- C. PCA only removes features
- D. Feature selection creates new features

Answer: B

Q417. What is an autoencoder's bottleneck layer?

- A. The output layer
- B. A hidden layer with fewer neurons that forces the network to learn compressed representations of the input
- C. The input layer
- D. A dropout layer

Answer: B

Q418. What is the purpose of the MiniBatchKMeans algorithm?

- A. To increase clustering accuracy
- B. To provide a faster approximation of K-Means by using random subsets of data for each iteration
- C. To handle text data only
- D. To replace DBSCAN

Answer: B

Q419. What is a true positive in classification?

- A. The model incorrectly predicts the positive class
- B. The model correctly predicts the positive class
- C. The model correctly predicts the negative class
- D. The model incorrectly predicts the negative class

Answer: B

Q420. What is a true negative in classification?

- A. The model correctly predicts the positive class
- B. The model incorrectly predicts the positive class
- C. The model correctly predicts the negative class
- D. The model incorrectly predicts the negative class

Answer: C

Q421. What is a false negative in classification?

- A. The model correctly predicts positive
- B. The model correctly predicts negative
- C. The model incorrectly predicts negative when the actual class is positive
- D. The model incorrectly predicts positive when the actual class is negative

Answer: C

Q422. What is the purpose of a holdout test set?

- A. To train the model
- B. To provide a final unbiased evaluation of the model's performance on unseen data
- C. To tune hyperparameters
- D. To clean the data

Answer: B

Q423. What does a higher accuracy mean for a classification model?

- A. The model makes more errors
- B. A larger proportion of all predictions are correct
- C. The model is more complex
- D. The model trains faster

Answer: B

Q424. What is root mean squared error?

- A. The mean of errors
- B. The square root of the average of squared differences between predicted and actual values
- C. The maximum error
- D. The median error

Answer: B

Q425. What does the R-squared value tell us about a regression model?

- A. The number of features used
- B. The proportion of variance in the target variable explained by the model
- C. The training time
- D. The number of data points

Answer: B

Q426. Why is it important to evaluate a model on data it has not been trained on?

- A. It is not important
- B. To check if the model generalizes to new data rather than just memorizing the training data
- C. To make training faster
- D. To increase the dataset size

Answer: B

Q427. What is a baseline model in model evaluation?

- A. The most complex model possible
- B. A simple reference model used as a minimum performance benchmark for comparison
- C. A model with zero accuracy
- D. The final deployed model

Answer: B

Q428. What is the specificity of a classifier?

- A. The same as recall
- B. The proportion of actual negatives that are correctly identified as negative
- C. The total number of predictions
- D. The model's complexity level

Answer: B

Q429. What is an interaction feature?

- A. A feature from user interactions with a website
- B. A new feature created by combining two or more existing features, such as multiplying them together
- C. A feature that interacts with the model
- D. A type of categorical feature

Answer: B

Q430. What is feature extraction?

- A. Removing all features from data
- B. Creating new informative features from raw data like extracting day of week from a date
- C. A type of model training
- D. Deleting columns from a dataset

Answer: B

Q431. Why is domain knowledge valuable in feature engineering?

- A. It is not valuable at all
- B. Experts can identify meaningful transformations and combinations that algorithms alone might miss
- C. Domain knowledge slows down the process
- D. It only helps with data collection

Answer: B

Q432. What is a ratio feature?

- A. A feature with a fixed ratio
- B. A new feature created by dividing one feature by another to capture proportional relationships
- C. A feature that counts occurrences
- D. A binary feature

Answer: B

Q433. What is the purpose of extracting date components as features?

- A. To encrypt date information
- B. To create separate features like year, month, day, and day of week that capture temporal patterns
- C. To delete date columns
- D. To format dates for display

Answer: B

Q434. What is text vectorization?

- A. Converting vectors to text
- B. Converting text data into numerical feature vectors that ML algorithms can process
- C. A text editing operation
- D. A type of font transformation

Answer: B

Q435. What is the purpose of indicator (dummy) variables?

- A. To indicate errors in data
- B. To represent categorical variables as binary columns where each category gets its own column
- C. To point to other datasets
- D. To indicate missing values only

Answer: B

Q436. What is feature aggregation?

- A. Splitting features into smaller parts
- B. Summarizing multiple values into a single statistic like count, mean, or sum for grouped data
- C. Deleting aggregated data
- D. A type of normalization

Answer: B

Q437. What is a count feature?

- A. A feature that counts model parameters
- B. A feature representing the frequency or count of occurrences of a particular event or category
- C. A feature with only integer values
- D. The total number of features

Answer: B

Q438. What is the benefit of creating squared features?

- A. They reduce all values to zero
- B. They allow linear models to capture quadratic (U-shaped or inverted-U) relationships between features and the target
- C. They make features easier to read
- D. They always improve accuracy

Answer: B

Q439. What is a perceptron in deep learning?

- A. A type of dataset
- B. The simplest form of neural network with a single neuron that computes a weighted sum of inputs
- C. A loss function
- D. A type of optimizer

Answer: B

Q440. What is the purpose of the input layer in a neural network?

- A. To process data through complex computations
- B. To receive and pass the raw input features to the next layer
- C. To produce final predictions
- D. To apply activation functions

Answer: B

Q441. What is the difference between a deep and a shallow neural network?

- A. Deep networks use more data
- B. Deep networks have many hidden layers while shallow networks have one or two hidden layers
- C. Shallow networks are more accurate
- D. There is no difference

Answer: B

Q442. What is the tanh activation function?

- A. A linear function
- B. A function that outputs values between -1 and 1, shaped like a stretched sigmoid
- C. A function that outputs only 0 or 1
- D. A function that outputs only positive values

Answer: B

Q443. What is a batch size in neural network training?

- A. The total dataset size
- B. The number of training samples processed before the model's weights are updated
- C. The number of layers in the network
- D. The number of output classes

Answer: B

Q444. What is a dense layer in a neural network?

- A. A layer with no connections
- B. A fully connected layer where every neuron is connected to every neuron in the previous layer
- C. A layer that reduces dimensions
- D. A layer that only processes images

Answer: B

Q445. What is the purpose of a learning rate schedule?

- A. To set a fixed learning rate
- B. To adjust the learning rate during training, typically decreasing it to allow finer convergence
- C. To increase the learning rate continuously
- D. To determine the batch size

Answer: B

Q446. What is the output layer of a neural network?

- A. The first layer that receives input
- B. The final layer that produces the network's prediction or output
- C. Any hidden layer
- D. The layer that stores training data

Answer: B

Q447. What is the role of a GPU in deep learning?

- A. It stores the training data
- B. It accelerates matrix computations needed for training neural networks through parallel processing
- C. It provides internet connectivity
- D. It manages the operating system

Answer: B

Q448. What is model inference in deep learning?

- A. The process of training a model
- B. Using a trained model to make predictions on new data
- C. Designing the model architecture
- D. Cleaning the training data

Answer: B

Q449. What is a stride in CNN convolution?

- A. The size of the filter
- B. The number of pixels the filter moves across the input for each step
- C. The depth of the network
- D. The number of filters used

Answer: B

Q450. What is padding in CNNs?

- A. Adding noise to images
- B. Adding extra pixels around the input borders to control the spatial dimensions of the output
- C. Removing pixels from images
- D. A type of activation function

Answer: B

Q451. What is average pooling?

- A. Computing the average of all network weights
- B. Taking the average of values in each pooling window to reduce spatial dimensions
- C. Averaging predictions from multiple models
- D. Computing the average loss

Answer: B

Q452. What is the vanishing gradient problem specifically in RNNs?

- A. Gradients become too large
- B. Gradients diminish exponentially over long sequences, making it hard for the network to learn long-range dependencies
- C. RNNs do not have gradients
- D. Gradients are always stable in RNNs

Answer: B

Q453. What is the generator in a GAN?

- A. The part that classifies real from fake
- B. The neural network that creates synthetic data samples from random noise
- C. A data preprocessing tool
- D. The loss function

Answer: B

Q454. What is the purpose of flattening in a CNN?

- A. Making images flat
- B. Converting multi-dimensional feature maps into a one-dimensional vector to feed into fully connected layers
- C. Reducing the number of filters
- D. Removing padding from feature maps

Answer: B

Q455. What is global average pooling?

- A. Averaging all model parameters
- B. Taking the average of each entire feature map to produce a single value per map, replacing fully connected layers
- C. A global training strategy
- D. Averaging across all training epochs

Answer: B

Q456. What is a feature map in a CNN?

- A. A geographical map of features
- B. The output produced when a convolutional filter is applied to the input, showing where patterns are detected
- C. A list of feature names
- D. The model architecture diagram

Answer: B

Q457. What is sequence-to-sequence learning?

- A. Learning from one data point to another
- B. A framework where the model reads an input sequence and produces an output sequence, used in translation and summarization
- C. A type of sorting algorithm
- D. Learning sequences in random order

Answer: B

Q458. What is fine-tuning a pre-trained model?

- A. Training from scratch on new data
- B. Taking a model trained on one task and continuing training on a new related task with a smaller learning rate
- C. Removing all layers of a model
- D. Only using the model for inference

Answer: B

Q459. What is a vocabulary in NLP?

- A. A grammar textbook
- B. The set of all unique words or tokens recognized by an NLP model
- C. A list of synonyms
- D. A type of neural network

Answer: B

Q460. What is text preprocessing in NLP?

- A. Writing text for a publication
- B. Cleaning and transforming raw text into a format suitable for analysis, including lowercasing, removing punctuation, and tokenizing
- C. Compressing text files
- D. Printing text documents

Answer: B

Q461. What is a bigram in NLP?

- A. A type of neural network
- B. A sequence of two consecutive words or characters in text
- C. A binary encoding scheme
- D. A grammatical rule

Answer: B

Q462. What is the bag-of-words representation?

- A. A physical bag containing word cards
- B. A text representation that counts word occurrences regardless of order or grammar
- C. A neural network for word generation
- D. A method for sorting words alphabetically

Answer: B

Q463. What is question answering in NLP?

- A. A quiz game format
- B. The task of automatically finding or generating answers to questions posed in natural language
- C. A type of data collection
- D. A text formatting technique

Answer: B

Q464. What is text summarization?

- A. Making text longer and more detailed
- B. Automatically creating a shorter version of text that retains the key information
- C. Translating text to another language
- D. Correcting grammar errors

Answer: B

Q465. What is a token in NLP?

- A. A coin-like object
- B. A basic unit of text such as a word, subword, or character that the model processes
- C. A model parameter
- D. A type of loss function

Answer: B

Q466. What is spam detection in NLP?

- A. Finding physical spam mail
- B. Using text classification to automatically identify unwanted or unsolicited messages
- C. A network security tool
- D. A data cleaning technique

Answer: B

Q467. What is text generation in NLP?

- A. Manually typing text
- B. Automatically producing coherent and contextually appropriate text using language models
- C. Copying text from one file to another
- D. Formatting text for printing

Answer: B

Q468. What is image preprocessing in computer vision?

- A. Taking photographs
- B. Preparing images for analysis through operations like resizing, normalization, and color conversion
- C. Printing images
- D. Compressing image file sizes only

Answer: B

Q469. What is a bounding box in object detection?

- A. A box used for shipping
- B. A rectangle drawn around a detected object specifying its location in the image
- C. A type of convolutional filter
- D. A border around the entire image

Answer: B

Q470. What is the difference between a color image and a grayscale image?

- A. They have the same number of channels
- B. Color images have three channels (RGB) while grayscale images have a single channel representing brightness
- C. Grayscale images are larger in file size
- D. Color images have only one channel

Answer: B

Q471. What is image annotation in computer vision?

- A. Writing captions for social media
- B. Labeling images with class labels, bounding boxes, or pixel masks to create training data for CV models
- C. Editing image colors
- D. Compressing images

Answer: B

Q472. What is transfer learning commonly used for in computer vision?

- A. Transferring images between devices
- B. Using a model pre-trained on a large dataset and adapting it to a new visual task with less data
- C. Transferring data between databases
- D. Moving files between folders

Answer: B

Q473. What does image normalization do?

- A. Makes all images the same size
- B. Scales pixel values to a standard range like 0 to 1 or standardizes them using mean and standard deviation
- C. Removes noise from images
- D. Converts images to text

Answer: B

Q474. What is edge detection in image processing?

- A. Detecting the physical edges of a monitor
- B. Identifying boundaries between different regions in an image where pixel intensity changes sharply
- C. Finding the borders of the image file
- D. A type of data augmentation

Answer: B

Q475. What is image cropping in computer vision?

- A. Growing images like crops
- B. Selecting and extracting a rectangular region of interest from an image
- C. Deleting all images
- D. Rotating images

Answer: B

Q476. What is image flipping as a data augmentation technique?

- A. Inverting pixel values
- B. Mirroring an image horizontally or vertically to create new training samples
- C. Converting the image format
- D. Changing image resolution

Answer: B

Q477. What is a pre-trained model like ResNet or VGG used for?

- A. Only for the task it was originally trained on
- B. As a feature extractor or starting point for transfer learning on new visual tasks
- C. It cannot be reused
- D. Only for generating images

Answer: B

Q478. What is a cluster in distributed computing?

- A. A group of similar data points
- B. A collection of interconnected computers working together to process data
- C. A type of database index
- D. A sorting algorithm

Answer: B

Q479. What is the Map phase in MapReduce?

- A. Creating geographic maps
- B. Processing input data and emitting key-value pairs for each record
- C. Reducing data to a single value
- D. Storing data in memory

Answer: B

Q480. What is the Reduce phase in MapReduce?

- A. Deleting unnecessary data
- B. Aggregating all intermediate values associated with the same key to produce final results
- C. Creating new data points
- D. Mapping data to categories

Answer: B

Q481. What is a NoSQL database?

- A. A database that does not store data
- B. A non-relational database designed for flexible schemas, horizontal scaling, and specific data models like document, key-value, or graph
- C. A SQL database with extra features
- D. A database only for numerical data

Answer: B

Q482. What is real-time data processing?

- A. Processing data once a year
- B. Processing data immediately or within milliseconds as it arrives, enabling instant insights and actions
- C. Processing data in monthly batches
- D. Storing data for future processing

Answer: B

Q483. What is data replication in distributed systems?

- A. Deleting duplicate data
- B. Storing copies of data on multiple nodes to ensure availability and fault tolerance
- C. Compressing data to save space
- D. Encrypting data for security

Answer: B

Q484. What is a data pipeline?

- A. A physical pipe for data cables
- B. A series of automated steps that move and transform data from source systems to destination systems
- C. A single database query
- D. A type of network cable

Answer: B

Q485. What is horizontal scaling?

- A. Making a single server larger
- B. Adding more machines to a system to distribute the workload across multiple nodes
- C. Increasing CPU speed
- D. Adding more RAM to one server

Answer: B

Q486. What is a Spark RDD?

- A. A type of hard drive
- B. A Resilient Distributed Dataset, an immutable distributed collection of objects that can be processed in parallel
- C. A relational database design
- D. A random data distribution

Answer: B

Q487. What is cloud computing in the context of big data?

- A. Computing in cloudy weather
- B. Using remote servers hosted on the internet to store, manage, and process data instead of local infrastructure
- C. A type of encryption
- D. A programming language

Answer: B

Q488. What is the goal of MLOps?

- A. To replace data scientists
- B. To streamline the process of deploying, monitoring, and maintaining ML models in production
- C. To eliminate the need for testing
- D. To make models train faster only

Answer: B

Q489. What is a model endpoint in deployment?

- A. The end of model training
- B. A URL or API where a deployed model receives input data and returns predictions
- C. The last layer of a neural network
- D. The final evaluation metric

Answer: B

Q490. What is containerization in MLOps?

- A. Putting computers in containers
- B. Packaging a model with its dependencies and environment into a container for consistent deployment across different platforms
- C. A type of data storage
- D. A compression technique

Answer: B

Q491. What is the purpose of logging in ML systems?

- A. Cutting trees for lumber
- B. Recording system events, predictions, inputs, and errors to enable debugging, monitoring, and auditing of ML systems
- C. Slowing down the system
- D. Only tracking training loss

Answer: B

Q492. What is batch inference?

- A. Training in batches
- B. Running predictions on a large collection of data at once, typically on a schedule, rather than one at a time
- C. Real-time prediction for each request
- D. A type of model training

Answer: B

Q493. What is the purpose of health checks for deployed ML models?

- A. Checking the health of data scientists
- B. Automated checks that verify the model service is running, responsive, and functioning correctly
- C. Only checking hardware temperature
- D. Measuring model accuracy once

Answer: B

Q494. What is a rollback in model deployment?

- A. Rolling the model forward
- B. Reverting to a previous model version when the new version performs poorly or causes issues
- C. Restarting the entire system
- D. Deleting all model versions

Answer: B

Q495. What is an ML experiment?

- A. A chemistry lab experiment
- B. A specific model training run with defined hyperparameters, data, and code that produces measurable results
- C. A random guess
- D. A data collection process

Answer: B

Q496. What is latency in the context of model serving?

- A. The accuracy of the model
- B. The time it takes from receiving an input to returning a prediction
- C. The size of the model
- D. The training time

Answer: B

Q497. What is the purpose of automated testing in MLOps?

- A. To test hardware components
- B. To automatically verify that ML code, data, and models meet quality standards before deployment
- C. Testing is not important for ML
- D. To slow down the deployment process

Answer: B

Q498. What is responsible AI?

- A. AI that responds quickly
- B. The practice of developing and deploying AI systems that are fair, transparent, accountable, and safe
- C. AI that costs less money
- D. AI that only works for governments

Answer: B

Q499. What is an example of AI bias in hiring?

- A. AI that works faster than humans
- B. An AI system that unfairly favors candidates of a particular gender or race based on biased training data
- C. AI that processes more applications
- D. AI that costs less than human recruiters

Answer: B

Q500. Why is human oversight important in AI decision-making?

- A. Humans are always faster than AI
- B. To ensure AI decisions are reviewed, correctable, and accountable, especially in high-stakes situations
- C. Human oversight is not necessary
- D. AI should make all decisions independently

Answer: B

Q501. What is data consent in AI ethics?

- A. Agreeing to pay for data
- B. Obtaining permission from individuals before collecting, using, or sharing their personal data for AI systems
- C. Data consent is not required
- D. Only companies need to give consent

Answer: B

Q502. What is the digital divide in the context of AI?

- A. A type of data splitting technique
- B. The gap between those who have access to AI technology and those who do not, often along socioeconomic or geographic lines
- C. A coding technique
- D. A security vulnerability

Answer: B

Q503. What does the term 'black box' mean in AI ethics?

- A. A physical black box device
- B. An AI system whose internal decision-making process is not understandable or explainable to users
- C. A box for storing AI hardware
- D. A type of test environment

Answer: B

Q504. What is the potential impact of AI on employment?

- A. AI has no effect on jobs
- B. AI can automate certain tasks, potentially displacing some jobs while creating new roles that require human-AI collaboration skills
- C. AI will eliminate all human jobs
- D. AI only creates new jobs

Answer: B

Q505. What is the purpose of AI ethics guidelines?

- A. To slow down AI development
- B. To provide principles and frameworks for developing AI systems that are beneficial, fair, and minimize harm to society
- C. To make AI more expensive
- D. To prevent all AI research

Answer: B

Q506. What is surveillance using AI technology?

- A. AI monitoring computer performance
- B. Using AI systems to monitor, track, or analyze people's activities, raising privacy and civil liberties concerns
- C. A type of AI training method
- D. A data backup process

Answer: B

Q507. Why is diversity important in AI development teams?

- A. It has no impact on AI quality
- B. Diverse teams are more likely to identify potential biases, consider varied perspectives, and build AI systems that work fairly for all populations
- C. Only technical skills matter
- D. Diversity slows down development

Answer: B

Q508. What is the absolute value of a number?

- A. The number multiplied by itself
- B. The non-negative distance of a number from zero on the number line
- C. The reciprocal of the number
- D. The square root of the number

Answer: B

Q509. What is the purpose of data type conversion in preprocessing?

- A. To delete data
- B. To change values from one data type to another so they can be processed correctly by algorithms
- C. To increase file size
- D. To add new columns

Answer: B

Q510. What is a chatbot in NLP?

- A. A physical robot that chats
- B. A software application that simulates human conversation using natural language processing
- C. A chat room for programmers
- D. A tool for compressing text files

Answer: B

AI / Machine Learning & Data Analytics — Medium

510 questions

Q511. The Turing Test is used to evaluate:

- A. Whether a machine can exhibit intelligent behavior
- B. The total memory capacity of a given system
- C. How quickly a processor handles computations
- D. How much bandwidth a network can support

Answer: A

Q512. Which of the following best describes Deep Learning?

- A. A relational system for storing structured data
- B. A paradigm for writing event-driven programs
- C. A subset of ML using multi-layer neural networks
- D. A protocol for routing data between networks

Answer: C

Q513. In reinforcement learning, an agent learns by:

- A. Memorizing all possible training inputs
- B. Receiving rewards or penalties for actions
- C. Reading through a set of labeled data points
- D. Clustering together groups of similar data

Answer: B

Q514. Which of the following is a descriptive analytics technique?

- A. Predicting future trends
- B. Real-time optimization
- C. Prescribing actions
- D. Summarizing historical data

Answer: D

Q515. What is the difference between AI and ML?

- A. AI is a subset of ML
- B. They are the same thing
- C. They are unrelated fields
- D. ML is a subset of AI

Answer: D

Q516. Which type of analytics answers 'What will happen?'

- A. Prescriptive Analytics
- B. Descriptive Analytics
- C. Predictive Analytics
- D. Diagnostic Analytics

Answer: C

Q517. An expert system in AI uses:

- A. A knowledge base and inference engine
- B. Only probabilistic fuzzy logic rules
- C. Only evolutionary genetic algorithms
- D. Only deep artificial neural networks

Answer: A

Q518. Which of these is a weak AI system?

- A. A chess-playing program
- B. A general-purpose intellect
- C. A conscious thinking machine
- D. A fully self-aware robot

Answer: A

Q519. What is the role of a training set in ML?

- A. To deploy the model into production
- B. To visualize the model's final results
- C. To test the final trained model output
- D. To train the model to learn patterns

Answer: D

Q520. Prescriptive analytics is used to:

- A. Describe past events only
- B. Store data more efficiently
- C. Recommend actions to take
- D. Detect anomalies in data

Answer: C

Q521. The gradient of a function points in the direction of:

- A. Steepest descent
- B. Random direction
- C. Steepest ascent
- D. Zero change

Answer: C

Q522. Bayes' theorem relates:

- A. Prior and posterior probabilities
- B. Only matrix decompositions
- C. Only the mean and variance
- D. Only derivative computations

Answer: A

Q523. The eigenvalue equation is $Av = \lambda v$. What is λ ?

- A. Eigenvector
- B. Trace
- C. Determinant
- D. Eigenvalue

Answer: D

Q524. What is the chain rule used for in calculus?

- A. Computing joint probabilities
- B. Adding matrices element-wise
- C. Differentiating composite functions
- D. Sorting data sequentially

Answer: C

Q525. Standard deviation measures:

- A. Spread of data around the mean
- B. The central tendency of values
- C. The minimum observed value
- D. The maximum observed value

Answer: A

Q526. A positive definite matrix has:

- A. All negative eigenvalues
- B. Mixed sign eigenvalues
- C. All zero eigenvalues
- D. All positive eigenvalues

Answer: D

Q527. In probability, two events are independent if:

- A. $P(A|B) = P(B)$
- B. $P(A \cap B) = P(A) * P(B)$
- C. $P(A \cup B) = P(A) + P(B)$
- D. $P(A \cap B) = 0$

Answer: B

Q528. The determinant of a 2x2 matrix $\begin{bmatrix} a & b \\ c & d \end{bmatrix}$ is:

- A. $ab + cd$
- B. $ab - cd$
- C. $ad - bc$
- D. $ad + bc$

Answer: C

Q529. A convex function has the property that:

- A. It is always an increasing function
- B. It has no minimum point at all
- C. Any local minimum is also a global minimum
- D. It has many distinct local minima

Answer: C

Q530. The covariance between two identical variables equals:

- A. A value of exactly zero
- B. The variance of that variable
- C. A value of negative one
- D. A value of exactly one

Answer: B

Q531. What is broadcasting in NumPy?

- A. A networking feature for sending data to multiple receivers
- B. Automatic expansion of arrays with different shapes for arithmetic
- C. A type of sorting algorithm for ordering array elements
- D. A logging mechanism for recording runtime error messages

Answer: B

Q532. Which scikit-learn class is used for train-test splitting?

- A. SplitData
- B. DataSplitter
- C. train_test_split
- D. TestTrainDivide

Answer: C

Q533. What does `df.groupby()` do in pandas?

- A. Renames columns using a provided mapping dict
- B. Sorts the DataFrame by index values in order
- C. Filters rows based on a boolean condition mask
- D. Groups data by one or more columns for aggregation

Answer: D

Q534. Which pandas method fills missing values?

- A. fillna()
- B. isna()
- C. notna()
- D. dropna()

Answer: A

Q535. What is a Python generator?

- A. A built-in immutable data type like a list
- B. A function that yields values lazily using yield
- C. A special class constructor for initialization
- D. A comparison-based sorting algorithm method

Answer: B

Q536. How do you perform matrix multiplication in NumPy?

- A. np.multiply(A, B) function
- B. A + B element-wise addition
- C. A * B element-wise product
- D. np.dot(A, B) or A @ B operator

Answer: A

Q537. What does the pandas method .apply() do?

- A. Deletes specified rows from the DataFrame index
- B. Merges two DataFrames on a common column key
- C. Adds a new column to the end of a DataFrame
- D. Applies a function along an axis of a DataFrame

Answer: D

Q538. Which visualization library is built on top of Matplotlib?

- A. Pygal
- B. Bokeh
- C. Seaborn
- D. Plotly

Answer: C

Q539. What is the purpose of np.reshape()?

- A. To sort the elements of an array in ascending order
- B. To change the shape of an array without changing its data
- C. To filter elements of an array by a condition mask
- D. To delete selected elements from an existing array

Answer: B

Q540. What does pickle do in Python?

- A. Manages relational database connections
- B. Performs numerical matrix calculations
- C. Creates interactive charts and plots
- D. Serializes and deserializes Python objects

Answer: D

Q541. What is one-hot encoding?

- A. Encrypting data columns for secure transmission
- B. Compressing data files for reduced disk usage
- C. Converting categorical variables into binary vectors
- D. Normalizing continuous numerical data columns

Answer: C

Q542. What is the difference between normalization and standardization?

- A. Normalization scales to [0,1], standardization scales to zero mean and unit variance
- B. Normalization is always the superior approach for every ML algorithm
- C. They are identical processes producing the exact same transformed outputs
- D. Standardization is only applicable to integer-valued feature columns

Answer: A

Q543. What is label encoding?

- A. Assigning numerical values to categorical labels
- B. Encrypting labels for secure data storage
- C. Adding new labels to the feature space
- D. Removing labels from a dataset entirely

Answer: A

Q544. What is imputation?

- A. Adding new columns to expand the feature set
- B. Replacing missing values with estimated values
- C. Deleting every row containing missing values
- D. Sorting data by ascending numerical order

Answer: B

Q545. Why is feature scaling important for ML algorithms?

- A. It automatically removes all outlier data point values
- B. It adds new synthetic features to expand the model
- C. It makes the overall dataset physically larger in storage
- D. It ensures all features contribute equally regardless of their scale

Answer: D

Q546. What is data augmentation?

- A. Encrypting data columns for privacy and security
- B. Deleting redundant records from the training set
- C. Compressing data files to reduce storage overhead
- D. Creating new training data by modifying existing data

Answer: D

Q547. What is the purpose of the train-test split?

- A. To evaluate model performance on unseen data
- B. To increase the total size of the full dataset
- C. To clean and preprocess the raw input data
- D. To visualize the distribution of data points

Answer: A

Q548. What is an outlier?

- A. A value that is completely missing or absent
- B. A feature that stores categorical text labels
- C. A data point significantly different from others
- D. A record that has been duplicated in error

Answer: C

Q549. What does the z-score standardization formula compute?

- A. $(x - \min) / (\max - \min)$ range
- B. x divided by the max value
- C. $(x - \text{mean}) / \text{standard deviation}$
- D. log of the value x

Answer: C

Q550. What is web scraping in data collection?

- A. Designing interactive web page layouts
- B. Testing web application load performance
- C. Building responsive websites from scratch
- D. Extracting data from websites programmatically

Answer: D

Q551. What is a heatmap used for in EDA?

- A. Showing temperature measurements from a sensor only
- B. Visualizing the magnitude of values in a matrix using colors
- C. Building multi-layer neural network architectures
- D. Creating three-dimensional surface projection plots

Answer: B

Q552. Skewness in a distribution refers to:

- A. Asymmetry of the distribution around the mean
- B. The total number of distribution peaks
- C. The height of the data distribution
- D. The overall width of the distribution

Answer: C

Q553. What is a pair plot (pairplot) in Seaborn?

- A. A single scatter plot between exactly two variables
- B. A grid of plots showing pairwise relationships between variables
- C. A grouped bar chart comparing category counts
- D. A circular pie chart showing value proportions

Answer: B

Q554. Kurtosis measures:

- A. The arithmetic mean
- B. The tailedness of a distribution
- C. The center of the distribution
- D. The measure of skewness

Answer: C

Q555. What is the Interquartile Range (IQR)?

- A. Mean minus Median
- B. Max - Min range
- C. Standard deviation * 2
- D. Q3 - Q1 difference

Answer: C

Q556. A violin plot combines which two visualizations?

- A. Box plot and KDE plot
- B. Pie chart and histogram
- C. Scatter plot and bar chart
- D. Line chart and area chart

Answer: A

Q557. What does the value_counts() method do in pandas?

- A. Returns the frequency of unique values in a Series
- B. Sums together all values stored in a Series
- C. Counts the total number of rows in a Series
- D. Counts the number of null values in a Series

Answer: B

Q558. A QQ plot is used to:

- A. Build linear regression prediction models
- B. Plot quarterly financial data over time
- C. Create grouped and stacked bar charts
- D. Check if data follows a particular distribution

Answer: D

Q559. What is a kernel density estimation (KDE) plot?

- A. A circular chart showing proportions
- B. A variant of the standard scatter plot
- C. A smoothed continuous version of a histogram
- D. A type of stacked or grouped bar chart

Answer: C

Q560. What is the purpose of a log transformation in EDA?

- A. To reduce skewness and make data more normally distributed
- B. To change the data types of columns in the table
- C. To encrypt data columns for compliance and security
- D. To completely delete all outlier values from data

Answer: A

Q561. What is the cost function in linear regression?

- A. Hinge loss function
- B. Log loss function
- C. Cross-entropy loss
- D. Mean Squared Error (MSE)

Answer: B

Q562. The sigmoid function maps values to:

- A. Range between 0 and 1
- B. Only integers
- C. Range between -1 and 1
- D. Any real number

Answer: A

Q563. What is regularization in supervised learning?

- A. Increasing overall model complexity further
- B. Adding more training data to the model
- C. Adding a penalty term to prevent overfitting
- D. Removing unnecessary features from the model

Answer: C

Q564. What is the difference between L1 and L2 regularization?

- A. L1 uses absolute values and can produce sparse models; L2 uses squared values
- B. L1 and L2 are completely identical techniques
- C. L2 regularization always produces sparse model weights
- D. L1 regularization is always the better choice overall

Answer: D

Q565. What is a Support Vector Machine (SVM)?

- A. An algorithm that finds the optimal hyperplane to separate classes
- B. A linear dimensionality reduction projection method
- C. A type of deep feedforward artificial neural network
- D. An unsupervised density-based clustering algorithm

Answer: A

Q566. What is the bias-variance tradeoff?

- A. Balancing model simplicity (bias) against sensitivity to training data (variance)
- B. Choosing between running computation on a CPU versus a GPU device
- C. Selecting the right proportions for training and test data splits
- D. Choosing an appropriate learning rate for gradient optimization

Answer: A

Q567. In decision trees, what is information gain?

- A. The reduction in entropy after splitting on a feature
- B. The overall accuracy of the trained model
- C. The total number of leaf nodes in the tree
- D. The current depth of the decision tree

Answer: B

Q568. What is cross-validation?

- A. Evaluating using only the training data subset
- B. Splitting data into multiple folds to train and validate the model
- C. Training the model on all available data at once
- D. Evaluating using only the test data subset

Answer: D

Q569. What is the purpose of the softmax function?

- A. Converting raw scores into probabilities for multi-class classification
- B. Performing binary classification tasks on labeled data only
- C. Performing unsupervised clustering of unlabeled data points
- D. Performing regression to predict continuous numeric values

Answer: A

Q570. What is Naive Bayes based on?

- A. A deep learning based neural network model
- B. Bayes' theorem with an assumption of feature independence
- C. An iterative gradient descent optimization method
- D. An ensemble of many decision tree classifiers

Answer: B

Q571. What is the difference between bagging and boosting?

- A. Bagging and boosting are completely identical approaches
- B. Boosting is always faster than bagging in every scenario
- C. Bagging always performs better than boosting overall
- D. Bagging trains models independently in parallel; boosting trains them sequentially on errors

Answer: C

Q572. In AdaBoost, how are misclassified samples handled?

- A. They are completely removed from the data
- B. Their weights are increased for the next iteration
- C. Their weights are decreased in the next round
- D. They are duplicated in the training dataset

Answer: C

Q573. What is stacking in ensemble learning?

- A. Using only a single standalone model for all predictions
- B. Random feature selection from the available input columns
- C. Using a meta-model to combine predictions from multiple base models
- D. Data augmentation to increase the training set size

Answer: C

Q574. What is the out-of-bag (OOB) error in Random Forest?

- A. The overall training error computed on the full training dataset
- B. The test error computed on a held-out independent test set
- C. The validation error from a separate cross-validation fold
- D. Error estimated using samples not included in each tree's bootstrap sample

Answer: D

Q575. Gradient Boosting minimizes the loss function by:

- A. Removing the least informative features iteratively
- B. Increasing the overall size of the training dataset
- C. Adding trees that fit the negative gradient of the loss
- D. Random bootstrap sampling of the original dataset

Answer: C

Q576. What is the max_features parameter in Random Forest?

- A. The minimum number of samples per leaf node
- B. The number of features to consider at each split
- C. The maximum depth of each individual tree
- D. The maximum total number of trees in the forest

Answer: C

Q577. Why does ensemble learning generally outperform single models?

- A. It reduces variance and/or bias by combining diverse models
- B. It always uses significantly less training data overall
- C. It is always computationally faster than a single model
- D. It always requires fewer input features to train on

Answer: A

Q578. What is a weak learner?

- A. A model that performs slightly better than random guessing
- B. A model that consistently achieves 100% accuracy
- C. A model that never makes any classification errors
- D. A model with absolutely no trainable parameters

Answer: A

Q579. In XGBoost, what is the purpose of the learning rate?

- A. It controls the contribution of each tree to shrink step size
- B. It directly determines the maximum depth of each tree
- C. It selects which features are used at each tree split
- D. It explicitly sets the total number of trees in the forest

Answer: A

Q580. What is voting in ensemble methods?

- A. Training a single standalone model on the complete dataset
- B. Engineering new derived features from the raw input data
- C. Combining predictions by having each model vote on the outcome
- D. Removing outlier data points from the training samples

Answer: C

Q581. How does the elbow method work for K-Means?

- A. Select the value of K completely at random from a uniform range
- B. Use the largest possible K equal to the number of data points
- C. Plot inertia vs K and find where adding clusters gives diminishing returns
- D. Always use K equal to three regardless of the dataset size

Answer: C

Q582. What is the silhouette score?

- A. The Euclidean distance between every pair of cluster centroids
- B. A measure of how similar a point is to its own cluster vs neighboring clusters
- C. The total number of clusters found by the clustering algorithm
- D. The total number of iterations before the algorithm converges

Answer: B

Q583. DBSCAN stands for:

- A. Data-Based Statistical Clustering Algorithm
- B. Database Scan Algorithm Notation
- C. Density-Based Spatial Clustering of Applications with Noise
- D. Deep Binary Sorting Classification Algorithm

Answer: D

Q584. What advantage does DBSCAN have over K-Means?

- A. It only works with fully labeled and supervised training data
- B. It can find arbitrarily shaped clusters and doesn't require specifying K
- C. It always gives strictly better results than all alternatives
- D. It is always computationally faster than every other algorithm

Answer: B

Q585. In PCA, principal components are:

- A. Orthogonal directions of maximum variance
- B. Cluster center points
- C. Randomly chosen directions
- D. Features that contain missing values

Answer: B

Q586. What is hierarchical clustering?

- A. A supervised regression technique for numerical prediction
- B. A deep neural network architecture for image recognition
- C. Building a tree of clusters by merging or splitting iteratively
- D. Using K-Means clustering multiple times in succession

Answer: C

Q587. What is the Apriori algorithm used for?

- A. Frequent itemset mining and association rules
- B. Dimensionality reduction projections
- C. Numerical regression for predictions
- D. Supervised classification with labeled data

Answer: A

Q588. What is t-SNE used for?

- A. Visualizing high-dimensional data in 2D or 3D
- B. Imputing missing data point values
- C. Training deep neural network models
- D. Selecting the most relevant features

Answer: A

Q589. In association rules, what does 'support' measure?

- A. The conditional confidence level of the rule
- B. The overall predictive accuracy of the rule
- C. The proportion of transactions containing an itemset
- D. The total number of features in the dataset

Answer: C

Q590. What is the difference between agglomerative and divisive hierarchical clustering?

- A. Neither method uses a hierarchy
- B. Agglomerative is bottom-up; divisive is top-down
- C. Agglomerative is top-down; divisive is bottom-up
- D. The two methods are completely identical

Answer: C

Q591. What is the F1-score?

- A. The arithmetic mean of precision and recall
- B. The product of precision and recall values
- C. The harmonic mean of precision and recall
- D. The square of the model accuracy

Answer: A

Q592. What does the ROC curve plot?

- A. Training error values plotted against validation error values
- B. True Positive Rate vs False Positive Rate at various thresholds
- C. Precision values plotted against corresponding recall values
- D. Training accuracy values plotted against the loss values

Answer: B

Q593. What does AUC-ROC measure?

- A. Only the precision metric of the trained classification model
- B. Only the accuracy metric of the trained classification model
- C. Only the recall metric of the trained classification model
- D. The overall ability of the model to discriminate between classes

Answer: A

Q594. What is K-fold cross-validation?

- A. Dividing data into K folds and using each fold as validation once
- B. Training the model completely without any validation step
- C. Using all available data solely for the training process
- D. Using only two simple train-test splits for evaluation

Answer: B

Q595. When is recall more important than precision?

- A. When simple accuracy alone is a sufficient metric
- B. When the dataset classes are perfectly balanced overall
- C. When false positives are very costly to the outcome
- D. When missing positive cases is costly, like disease detection

Answer: C

Q596. What is stratified cross-validation?

- A. Performing completely random splitting of the data
- B. Ignoring overall class balance during data splitting
- C. Using only the majority class for model training purposes
- D. K-fold CV that preserves the class distribution in each fold

Answer: B

Q597. What is the log loss (cross-entropy loss)?

- A. A feature selection method based on mutual information scores
- B. The exact same metric as mean squared error for regression
- C. A metric specifically designed for evaluating cluster quality
- D. A loss function that penalizes confident wrong predictions heavily

Answer: D

Q598. What is the purpose of a learning curve?

- A. To select the most important features from the training data
- B. To diagnose overfitting or underfitting by plotting performance vs training size
- C. To clean and preprocess the raw input data before training
- D. To learn about new machine learning algorithms from documentation

Answer: B

Q599. What is the difference between micro and macro averaging?

- A. Micro averaging is always a better choice than macro averaging
- B. Micro aggregates all instances globally; macro averages per-class metrics equally
- C. Micro and macro averaging are completely identical approaches
- D. Macro averaging ignores the overall class distribution entirely

Answer: A

Q600. What is Mean Absolute Error (MAE)?

- A. The single maximum prediction error value
- B. The median of all prediction errors
- C. The average of absolute differences between predicted and actual values
- D. The average of all squared differences between predictions

Answer: A

Q601. What is the purpose of polynomial features?

- A. To remove unnecessary features
- B. To normalize the data values
- C. To handle all missing values
- D. To capture non-linear relationships by creating powers and interactions of features

Answer: A

Q602. What is the Variance Inflation Factor (VIF)?

- A. An indicator for missing feature values
- B. A measure of multicollinearity between features
- C. A metric for overall model accuracy
- D. A particular feature scaling method

Answer: D

Q603. What is target encoding?

- A. Encoding categories using one-hot vector representation
- B. Replacing categories with the mean of the target variable for that category
- C. Removing all categorical features from dataset
- D. Encoding using binary representation of categories

Answer: A

Q604. Why is feature scaling important for KNN?

- A. KNN only works well with categorical type features
- B. KNN uses distance calculations that are affected by feature magnitudes
- C. Feature scaling actually slows down KNN
- D. KNN actually ignores all feature scales

Answer: A

Q605. What is a lag feature in time series?

- A. A feature with all of its values missing
- B. A feature created from previous time steps' values
- C. A feature that holds only binary zero-one
- D. A feature that is computationally slow to run

Answer: B

Q606. What is the purpose of log transformation on features?

- A. To deliberately add more noise to features
- B. To reduce skewness and handle multiplicative relationships
- C. To intentionally increase the skewness of data
- D. To remove all features entirely from the dataset

Answer: B

Q607. What is feature hashing?

- A. Encrypting feature values for secure storage and privacy compliance
- B. Mapping high-dimensional features to a fixed-size vector using a hash function
- C. Sorting features in the dataset by their statistical importance
- D. Removing duplicate feature entries from the data table columns

Answer: B

Q608. What are rolling window features?

- A. Features that are created without any time component at all
- B. Simple binary indicator features only
- C. Static features with no temporal aspect
- D. Features computed from a sliding window over time series data like rolling mean

Answer: C

Q609. What is the purpose of the chi-squared test in feature selection?

- A. To test the independence between categorical features and the target
- B. To normalize all features to a common range
- C. To impute missing values in the features
- D. To create entirely new features from scratch

Answer: B

Q610. What is domain knowledge in feature engineering?

- A. Using only fully automated feature methods
- B. Deliberately ignoring the problem context entirely
- C. Creating features completely at random each time
- D. Using expertise about the problem area to create meaningful features

Answer: C

Q611. What is the vanishing gradient problem?

- A. Gradients become very small in early layers, making training slow
- B. Gradients become extremely large causing numerical overflow
- C. The model has far too many hidden layers to train
- D. The training process converges far too quickly to minimum

Answer: A

Q612. What is dropout?

- A. Deleting training data samples to reduce dataset size
- B. Randomly deactivating neurons during training to prevent overfitting
- C. Reducing the learning rate schedule over the epochs
- D. Removing network layers permanently from the architecture

Answer: B

Q613. What is batch normalization?

- A. A loss function for measuring classification prediction error
- B. A preprocessing step applied only to the raw input data
- C. Normalizing the inputs of each layer to stabilize and accelerate training
- D. A type of non-linear activation function for hidden layers

Answer: C

Q614. What is the difference between a batch and a mini-batch?

- A. Batch and mini-batch gradient descent are completely identical
- B. A batch is the full dataset; a mini-batch is a subset used per iteration
- C. A mini-batch is always larger than a full batch of data
- D. A batch always contains exactly one single data sample only

Answer: A

Q615. What is transfer learning?

- A. Moving files between different computer file systems
- B. Transferring data records between separate databases
- C. Using a pre-trained model on a new but related task
- D. Training every model from scratch on random weights

Answer: C

Q616. What is the purpose of the learning rate?

- A. It determines the total number of hidden layers in the model
- B. It explicitly sets the mini-batch size during each epoch
- C. It defines the specific neural network architecture layout
- D. It controls the step size of weight updates during optimization

Answer: D

Q617. What is the exploding gradient problem?

- A. The overall training process runs far too slowly
- B. Gradients become extremely small, causing no weight changes
- C. The model severely underfits the training data distribution
- D. Gradients become extremely large, causing unstable weight updates

Answer: D

Q618. What is weight initialization and why is it important?

- A. Setting initial weights to enable effective training and avoid gradient problems
- B. Weights are assigned their final values only after training
- C. Weights are always initialized to exactly zero for every neuron
- D. The initialization strategy has no effect on the training process

Answer: A

Q619. What is the Adam optimizer?

- A. A non-linear activation function like sigmoid or ReLU
- B. An adaptive learning rate optimizer combining momentum and RMSProp
- C. A cross-entropy loss function for classification training
- D. A type of multi-layer feedforward neural network architecture

Answer: B

Q620. What is a softmax output layer used for?

- A. Feature extraction from the input data
- B. Binary classification with only two output classes
- C. Regression for continuous value prediction
- D. Multi-class classification producing a probability distribution

Answer: D

Q621. What is the Transformer architecture?

- A. A model based on self-attention mechanisms without recurrence
- B. A generative adversarial network for data synthesis
- C. A density-based clustering algorithm for grouping
- D. A standard convolutional neural network for image tasks

Answer: A

Q622. What is self-attention?

- A. A cross-entropy loss function that measures classification prediction error
- B. A type of spatial pooling operation that reduces feature map dimensions
- C. A mechanism relating different positions within a sequence to compute representations
- D. A data augmentation method that generates synthetic training examples

Answer: C

Q623. What is a Variational Autoencoder (VAE)?

- A. A linear regression model for numeric outputs
- B. A generative model that learns a probabilistic latent space
- C. A standard deterministic autoencoder without sampling
- D. A supervised classifier for categorical prediction

Answer: B

Q624. What is the difference between GRU and LSTM?

- A. GRU is always more accurate than LSTM on all tasks
- B. GRU and LSTM are completely identical architectures
- C. GRU has fewer gates (2 vs 3) and is computationally simpler
- D. LSTM has fewer learnable parameters than the GRU

Answer: D

Q625. What is 1x1 convolution used for?

- A. Increasing the spatial dimensions of the feature maps in the network
- B. Changing the number of channels and adding non-linearity without changing spatial dimensions
- C. Performing spatial pooling to reduce the feature map dimensions
- D. Applying batch normalization across all the feature map channels

Answer: A

Q626. What is depthwise separable convolution?

- A. Performing a pooling operation to reduce the spatial feature dimensions
- B. Splitting convolution into depthwise and pointwise operations to reduce computation
- C. Applying a normalization technique to stabilize the training process
- D. Performing a standard full convolution operation on the input feature maps

Answer: A

Q627. What is the encoder-decoder architecture?

- A. A cross-entropy loss function for measuring prediction quality
- B. A type of generative adversarial network with two competing agents
- C. A structured data format for storing model configuration files
- D. A structure where an encoder compresses input and a decoder generates output

Answer: D

Q628. What is data parallelism in deep learning?

- A. A type of batch normalization applied across distributed nodes
- B. Distributing training data across multiple GPUs, each running a model copy
- C. Reducing the overall size of the training data by sampling
- D. Using a single GPU for all training and inference computation

Answer: B

Q629. What is a residual network (ResNet)?

- A. A very shallow network with only a single hidden layer
- B. A standard recurrent neural network for sequences
- C. A deep network using skip connections that add input to output of layers
- D. A network architecture with no inter-layer connections

Answer: C

Q630. What is the purpose of attention mechanisms?

- A. To allow models to focus on relevant parts of the input
- B. To increase the mini-batch size during training
- C. To initialize the network weights before training
- D. To reduce the learning rate over training epochs

Answer: A

Q631. What is TF-IDF?

- A. A recurrent neural network variant designed for sequential text modeling
- B. A structured query language for relational database data retrieval
- C. A multi-layer deep feedforward neural network architecture for classification
- D. A weighting scheme reflecting word importance in a document relative to a corpus

Answer: D

Q632. What is Word2Vec?

- A. A simple word frequency counter for bag-of-words features
- B. A technique learning dense vector representations of words from context
- C. A language dictionary for looking up word definitions
- D. A rule-based grammar checker for text correction

Answer: B

Q633. What is the bag-of-words model?

- A. Representing text with full word order preserved
- B. A syntactic parsing technique for sentences
- C. Representing text as the frequency of words, ignoring order
- D. A neural language model using word embeddings

Answer: C

Q634. What is BERT?

- A. A rule-based expert system for language processing
- B. A generative adversarial network designed for text data
- C. A pre-trained transformer for bidirectional language understanding
- D. A type of recurrent neural network for sequential processing

Answer: C

Q635. What is attention in sequence-to-sequence models?

- A. A data augmentation method for creating synthetic text training data
- B. A mechanism allowing the decoder to focus on relevant parts of the input sequence
- C. A cross-entropy loss function for measuring sequence prediction accuracy
- D. A regularization technique for preventing overfitting in sequence models

Answer: B

Q636. What is text embedding?

- A. Converting text into dense numerical vectors
- B. Removing text from the raw dataset
- C. Encrypting text for secure storage
- D. Converting numbers into readable text

Answer: A

Q637. What is part-of-speech (POS) tagging?

- A. Removing individual words from the input text sequence
- B. Translating text between two different natural languages
- C. Assigning grammatical categories (noun, verb, etc.) to each word
- D. Counting the number of syllables in each word

Answer: C

Q638. What is sequence-to-sequence (Seq2Seq) modeling?

- A. Mapping a single input word to a single corresponding output word
- B. Sorting input sequences into a predefined correct order
- C. Clustering multiple sequences into meaningful related groups
- D. Mapping an input sequence to an output sequence of potentially different length

Answer: C

Q639. What is cosine similarity used for in NLP?

- A. Counting the total frequency of words in text
- B. Parsing sentences into syntax tree structures
- C. Generating new text from a trained model
- D. Measuring the similarity between two text vectors

Answer: D

Q640. What is a language model?

- A. A static vocabulary dictionary for word lookup
- B. A rule-based grammar checker for syntax only
- C. A database of aligned translation memory pairs
- D. A model that predicts the probability of a sequence of words

Answer: D

Q641. What is the difference between semantic and instance segmentation?

- A. Instance segmentation completely ignores the object class labels
- B. Semantic segmentation is always the better approach overall
- C. Semantic labels every pixel by class; instance distinguishes individual objects of the same class
- D. Semantic and instance segmentation are completely identical approaches

Answer: B

Q642. What is data augmentation in computer vision?

- A. Deleting images from the dataset to make it smaller and more manageable
- B. Reducing the overall image quality to speed up the training process
- C. Applying transformations like rotation, flipping, and cropping to increase training data
- D. Converting all images to grayscale only without any other transformation

Answer: D

Q643. What is the YOLO algorithm?

- A. A pixel-level semantic image segmentation algorithm
- B. A supervised multi-class image classification algorithm
- C. You Only Look Once - a real-time object detection algorithm
- D. An unsupervised K-means image clustering algorithm

Answer: C

Q644. What is image feature extraction?

- A. Adding textual annotation overlays onto the image pixels
- B. Compressing images into smaller file archive formats
- C. Identifying and representing distinctive visual patterns in images
- D. Deleting specific features or layers from the image data

Answer: C

Q645. What is the purpose of max pooling?

- A. Adding random noise to the feature map pixel values
- B. Changing the color representation of the image data
- C. Reducing spatial dimensions while retaining the most prominent features
- D. Increasing the spatial size of the feature map dimensions

Answer: C

Q646. What is optical flow?

- A. A type of spatial convolutional image filter kernel operation
- B. A standard color model representation for digital images
- C. The pattern of apparent motion of objects between consecutive frames
- D. A container file format for storing compressed video data

Answer: C

Q647. What is the IoU (Intersection over Union) metric?

- A. A metric measuring overlap between predicted and ground truth bounding boxes
- B. A metric that measures overall image quality and clarity
- C. A metric that assesses the overall color accuracy of images
- D. A metric for evaluating the spatial resolution of images

Answer: C

Q648. What are anchor boxes in object detection?

- A. The pixel coordinate system used for image representation
- B. The individual color channels within the input image data
- C. The border regions around the edges of an input image
- D. Predefined bounding box shapes that help detect objects of various sizes and ratios

Answer: C

Q649. What is non-maximum suppression (NMS)?

- A. A batch normalization method for stabilizing training
- B. A post-processing step that removes overlapping duplicate detections
- C. A gradient-based model training optimization technique
- D. A non-linear activation function for hidden layers

Answer: B

Q650. What is the difference between object detection and image classification?

- A. Image classification actually locates and draws bounding boxes around objects
- B. Detection locates multiple objects with bounding boxes; classification assigns one label to the whole image
- C. Object detection assigns a single label to the entire image without any localization
- D. Object detection and image classification are completely identical tasks

Answer: B

Q651. What is stream processing?

- A. Only storing data without processing it
- B. Processing data in real-time as it arrives
- C. Processing data in large scheduled batches
- D. Only visualizing data without analysis

Answer: B

Q652. What is Apache Kafka?

- A. A distributed event streaming platform for real-time data pipelines
- B. A standalone relational database for transactional workloads
- C. A supervised machine learning framework for model training
- D. An interactive data visualization and charting dashboard

Answer: A

Q653. What is the difference between Spark and Hadoop MapReduce?

- A. They are completely identical technologies with no performance differences
- B. Spark only works effectively on small datasets under one gigabyte
- C. MapReduce is significantly faster than Spark for all data workloads
- D. Spark processes in-memory making it faster; MapReduce writes to disk between steps

Answer: D

Q654. What is NoSQL?

- A. An interactive data visualization and charting dashboard
- B. Non-relational databases designed for flexible schemas and horizontal scaling
- C. A general-purpose compiled programming language for apps
- D. A newer version of the standard SQL query language syntax

Answer: B

Q655. What is data partitioning?

- A. Encrypting data at rest for privacy and compliance
- B. Dividing data into smaller pieces distributed across multiple nodes
- C. Merging all distributed data together into a single node
- D. Deleting old data records to free up storage space

Answer: B

Q656. What is Apache Hive?

- A. A distributed event streaming platform for real-time pipelines
- B. A graph database for storing relationship network data
- C. A data warehouse tool providing SQL-like queries over Hadoop data
- D. A supervised machine learning library for model training

Answer: C

Q657. What is a Spark DataFrame?

- A. A simple NumPy array structure
- B. A distributed collection of data organized into named columns
- C. A Python pandas DataFrame structure only
- D. A traditional database table structure only

Answer: A

Q658. What is the CAP theorem?

- A. It is primarily about effective data visualization techniques
- B. It applies only to single standalone machines and servers
- C. All three properties can always be fully achieved simultaneously
- D. A distributed system can guarantee at most two of: Consistency, Availability, and Partition tolerance

Answer: B

Q659. What is data sharding?

- A. Encrypting data at rest for privacy regulation compliance
- B. Copying the entire database to every single server node
- C. Splitting a database into smaller pieces across multiple servers
- D. Deleting old data records to reclaim disk storage space

Answer: C

Q660. What is the role of a scheduler in big data systems?

- A. Training supervised machine learning classification models
- B. Visualizing data trends using interactive chart dashboards
- C. Managing and coordinating the execution of distributed tasks
- D. Storing and persisting data on distributed file systems

Answer: C

Q661. What is model drift?

- A. Model training getting significantly faster with each iteration
- B. Degradation in model performance over time due to changes in data
- C. Model file size increasing automatically without configuration
- D. Model accuracy consistently increasing without retraining

Answer: B

Q662. What is A/B testing in model deployment?

- A. Comparing two model versions by serving them to different user groups
- B. Running computations on two GPUs for faster processing speed
- C. Training two different models simultaneously on the same data
- D. Using two separate datasets for training and evaluation only

Answer: A

Q663. What is a feature store?

- A. A model registry for tracking model versions and status
- B. A code repository for managing source version control
- C. A centralized repository for storing managing and serving ML features
- D. A data warehouse for storing raw analytical data records

Answer: C

Q664. What is model serialization?

- A. Deploying a model to a production serving endpoint
- B. Saving a trained model to a file format that can be loaded later
- C. Training a model on a labeled dataset from scratch
- D. Evaluating a model's accuracy on a held-out test set

Answer: B

Q665. What is canary deployment?

- A. Rolling back a deployed model to an earlier stable version
- B. Gradually rolling out a new model to a small percentage of users first
- C. Deploying the new model to all users immediately without rollout
- D. Training a brand new model from scratch on latest data

Answer: B

Q666. What is the purpose of a model registry?

- A. Storing and managing raw training data in a data lake
- B. Centrally managing model versions metadata and deployment status
- C. Visualizing model performance results on a dashboard
- D. Writing and version-controlling source code for models

Answer: B

Q667. What is data drift?

- A. Changes in the statistical distribution of input data over time
- B. Changes in the source code of the training pipeline
- C. Hardware failures in the production serving cluster
- D. Changes in the internal model architecture configuration

Answer: A

Q668. What is the ONNX format?

- A. A markup text file format for web page content
- B. An open format for ML models enabling interoperability between frameworks
- C. A proprietary relational database file storage format
- D. A compressed video streaming container file format

Answer: B

Q669. What is a ML pipeline?

- A. A database query for extracting records from tables
- B. An automated sequence of steps from data processing to model deployment
- C. An interactive data visualization charting dashboard
- D. A single isolated model training step without automation

Answer: B

Q670. What is the role of Kubernetes in MLOps?

- A. Visualizing model performance on dashboard charts
- B. Orchestrating containerized ML services for scaling and management
- C. Storing raw data files on a distributed file system
- D. Training machine learning models on training data

Answer: B

Q671. What is the difference between individual and group fairness?

- A. Group fairness ignores the outcomes of demographic groups
- B. Individual fairness treats similar people similarly; group fairness ensures equal outcomes
- C. They are completely identical fairness concepts with no differences
- D. Individual fairness ignores the treatment of specific individuals

Answer: B

Q672. What is differential privacy?

- A. A firewall system for blocking unauthorized network access
- B. A process of permanently deleting all personal data records
- C. A mathematical framework providing privacy guarantees when analyzing data
- D. A standard encryption method for securing data files at rest

Answer: C

Q673. What is explainable AI (XAI)?

- A. AI systems that require significantly less training data
- B. AI systems that use significantly larger model architectures
- C. AI systems that process data at significantly faster speeds
- D. AI systems that provide interpretable explanations for their decisions

Answer: D

Q674. What is model interpretability?

- A. The total file size in megabytes of the serialized model
- B. The inference speed at which the model generates predictions
- C. The overall accuracy metric of the trained classification model
- D. The degree to which humans can understand a model's predictions

Answer: D

Q675. What is the black box problem in AI?

- A. A hardware issue where the server casing is colored black
- B. A storage problem where disk capacity runs out unexpectedly
- C. Complex models making decisions that cannot be easily understood or explained
- D. A networking error where connections time out intermittently

Answer: C

Q676. What are deepfakes?

- A. Corrupted files that cannot be opened or processed
- B. AI-generated synthetic media that convincingly replaces a person's likeness
- C. Encrypted videos that require a decryption key to play
- D. Low-quality images with significant compression artifacts

Answer: B

Q677. What is federated learning?

- A. Sharing all raw training data openly across all participants
- B. Training ML models across decentralized devices without sharing raw data
- C. Training models exclusively on a single centralized server machine
- D. Using only publicly available open-source datasets for training

Answer: B

Q678. What is the right to explanation under GDPR?

- A. Individuals have the right to receive explanations for automated decisions
- B. AI companies can keep their algorithms and models completely secret
- C. No explanation is ever needed for any automated AI decisions
- D. Only government agencies can request explanations of AI systems

Answer: A

Q679. What is data anonymization?

- A. Making all personal data publicly available and fully open
- B. Encrypting data with standard symmetric encryption only
- C. Deleting all data from every storage system permanently
- D. Removing or modifying personal identifiers so individuals cannot be identified

Answer: D

Q680. What is the trolley problem in AI ethics?

- A. A moral dilemma about how autonomous systems should handle potential harm
- B. A transportation route optimization problem for logistics planning
- C. A data storage capacity issue for managing large files
- D. A network routing problem for directing data packets

Answer: A

Q681. Which concept refers to an AI that can learn any intellectual task a human can?

- A. Artificial Narrow Intelligence
- B. Artificial Specific Intelligence
- C. Artificial General Intelligence
- D. Artificial Basic Intelligence

Answer: C

Q682. What distinguishes descriptive analytics from predictive analytics?

- A. Descriptive forecasts future events while predictive summarizes past data
- B. Descriptive needs real-time data while predictive works with batch data only
- C. Descriptive uses deep learning while predictive relies on simple statistics
- D. Descriptive summarizes past data while predictive forecasts future outcomes

Answer: D

Q683. In machine learning, what does the term 'generalization' mean?

- A. Converting categorical variables into numerical representation
- B. Memorizing every training sample exactly as it is provided
- C. Performing well on new and previously unseen data points
- D. Reducing the total number of features in a given dataset

Answer: C

Q684. Which of the following is an example of unsupervised learning?

- A. Stock price prediction model
- B. Customer segmentation clustering
- C. Spam email classification
- D. Handwriting digit recognition

Answer: B

Q685. What role does a loss function play in training a machine learning model?

- A. It measures how far predictions deviate from actual values
- B. It specifies the data format used for input preprocessing step
- C. It determines the total memory required for model storage
- D. It controls the number of features selected for training data

Answer: A

Q686. What is the difference between structured and unstructured data?

- A. Structured data is always larger in file size than unstructured data is
- B. Structured data requires GPUs while unstructured data uses only CPU cores
- C. Structured data is only images while unstructured data is found in tables
- D. Structured data uses rows and columns while unstructured data does not

Answer: D

Q687. Which concept describes a model that learns noise instead of true patterns?

- A. Underfitting the training dataset
- B. Overfitting the training data
- C. Normalizing input features
- D. Standardizing output labels

Answer: B

Q688. What is transfer learning in modern AI systems?

- A. Transferring model weights from GPU memory to CPU memory
- B. Converting one programming language into a different language
- C. Moving data between two different storage systems for backup
- D. Reusing a pre-trained model on a new but related task domain

Answer: D

Q689. Which metric is commonly used to evaluate classification model performance?

- A. Mean Squared Error
- B. Root Mean Absolute
- C. Adjusted R-squared
- D. Accuracy and F1-score

Answer: D

Q690. What is a 'feature' in machine learning terminology?

- A. A visualization analysis tool
- B. An individual measurable property
- C. A type of network architecture
- D. A bug found in software testing

Answer: B

Q691. What does the gradient of a function represent in optimization?

- A. The maximum value the function can possibly achieve overall
- B. The direction of steepest ascent at a given point in space
- C. The total area under the curve of the function on a plot
- D. The number of local minima in the function search space

Answer: B

Q692. What does Bayes' theorem allow us to compute?

- A. The posterior probability given prior and likelihood
- B. The determinant of any square matrix efficiently
- C. The exact maximum value of a continuous function
- D. The shortest path between nodes in a graph

Answer: A

Q693. What is the purpose of the chain rule in calculus for neural networks?

- A. It selects which activation function to use in each layer
- B. It determines the optimal number of hidden layers needed
- C. It controls the learning rate during the training process
- D. It enables gradient computation through composed functions

Answer: D

Q694. What does the covariance between two variables measure?

- A. The exact ratio of one variable value to another variable
- B. The causal relationship from one variable to the other one
- C. The joint variability and direction of relationship between them
- D. The maximum possible value either variable can take at all

Answer: C

Q695. What is the role of eigenvalues in Principal Component Analysis?

- A. They determine the optimal number of training epochs to use
- B. They specify the learning rate for gradient descent processes
- C. They represent the variance amount captured by each component
- D. They control the dropout probability in neural network layers

Answer: C

Q696. What property must a matrix have to be invertible?

- A. All its elements must be positive real numbers
- B. It must have more rows than total columns
- C. Its determinant must be non-zero valued
- D. All diagonal elements must equal exactly one

Answer: C

Q697. What does the standard deviation tell us about a dataset?

- A. The spread of values around the mean value
- B. The difference between maximum and minimum
- C. The middle value when data is sorted in order
- D. The most frequently occurring value in the data

Answer: A

Q698. In ML, what is a convex function?

- A. A function that always produces negative output values only
- B. A function with exactly two critical points at most always
- C. A function only defined on integer input values overall
- D. A function where any local minimum is also a global minimum

Answer: D

Q699. What does the L2 norm of a vector measure?

- A. The Euclidean length or magnitude of vector
- B. The largest absolute element in entire vector
- C. The sum of all elements without absolute value
- D. The count of non-zero elements in the vector

Answer: A

Q700. What is a probability distribution function?

- A. A function describing likelihood of possible outcome values
- B. A function that sorts data points in descending order value
- C. A function removing outlier values from a given data set
- D. A function converting categorical data into numerical codes

Answer: A

Q701. What is the difference between a shallow copy and a deep copy of a Python list?

- A. Shallow copy works only with integers while deep copy handles all data types
- B. Shallow copy duplicates all nested objects while deep copy copies references only
- C. Shallow copy copies references to nested objects while deep copy duplicates all
- D. Shallow copy is faster because it uses more memory than a deep copy overall

Answer: C

Q702. How does NumPy broadcasting work for arrays of different shapes?

- A. It stretches smaller arrays to match larger ones following compatibility rules
- B. It pads smaller arrays with zeros until both arrays have identical dimensions
- C. It randomly samples elements from larger arrays to reduce them to smaller size
- D. It raises an error whenever array shapes do not match exactly in all dimensions

Answer: A

Q703. What advantage does vectorized NumPy code have over standard Python loops?

- A. Vectorized code uses more memory but produces more accurate results overall
- B. Vectorized code executes in optimized C making it significantly much faster
- C. Vectorized code only works with integer data types unlike Python loop code
- D. Vectorized code automatically parallelizes across multiple machine clusters

Answer: B

Q704. What does the pandas groupby() method enable you to do?

- A. Remove duplicate rows based on all column values in the frame
- B. Sort DataFrame rows in ascending order by their index value names
- C. Split data into groups and apply aggregate functions to each one
- D. Merge two DataFrames horizontally using a common column as key

Answer: C

Q705. What is the purpose of Python's yield keyword in generator functions?

- A. It terminates the function and returns the final value permanently
- B. It creates a new thread for parallel function execution runs
- C. It imports external modules required by function at call time
- D. It pauses execution and produces a value without losing state

Answer: D

Q706. How do you create a virtual environment in Python for project isolation?

- A. Using the command `python -m venv env_name` in the terminal window
- B. Using the command `pip create environment env_name` from terminal
- C. Using the command `python --create-env env_name` in the settings
- D. Using the command `python --isolate project env_name` from shell

Answer: A

Q707. What is the role of the fit_transform() method in scikit-learn preprocessing?

- A. It only transforms data using previously learned parameters from data
- B. It fits preprocessor to data and transforms it in a single step
- C. It evaluates model performance on the validation set after train
- D. It splits the dataset into separate training and testing subsets

Answer: B

Q708. What does the lambda keyword create in Python?

- A. A small anonymous function with a single expression body
- B. A new class definition with automatic initialization method
- C. A named function stored permanently in the global namespace
- D. A decorator that modifies behavior of existing functions

Answer: A

Q709. What is the purpose of the pickle module in Python for ML workflows?

- A. It provides functions for creating interactive data visualizations
- B. It serializes Python objects to save and load trained ML models
- C. It manages database connections for storing large training data
- D. It handles HTTP requests for downloading datasets from the web

Answer: B

Q710. How does pandas merge() differ from concat() for combining DataFrames?

- A. merge always produces fewer rows while concat always produces more rows
- B. merge requires sorted data while concat works only with unsorted frames
- C. merge joins on common columns while concat stacks along an axis direction
- D. merge only works with numeric data while concat handles all data types

Answer: C

Q711. What is the difference between normalization and standardization?

- A. Normalization works only with text while standardization works only with numeric data
- B. Normalization removes outliers while standardization only handles missing values in data
- C. Normalization scales to 0-1 range while standardization centers with zero mean unit variance
- D. Normalization increases feature count while standardization decreases feature count total

Answer: C

Q712. Why might you use stratified sampling instead of simple random sampling?

- A. It guarantees that the sample will be larger than simple random samples
- B. It preserves the proportion of each class in the sample matching population
- C. Stratified sampling is always faster than simple random sampling methods
- D. Stratified sampling eliminates the need for any data cleaning steps after

Answer: B

Q713. What is the purpose of applying a log transformation to skewed data?

- A. It converts categorical variables into numerical features for models
- B. It reduces skewness making the distribution closer to normal shape
- C. It removes all outliers automatically from the dataset completely
- D. It increases the total number of samples available for training

Answer: B

Q714. When should you use label encoding versus one-hot encoding for categories?

- A. Label encoding is used for numerical features while one-hot is for text only
- B. One-hot encoding should be used only when the dataset has fewer than ten rows
- C. Label encoding for ordinal categories and one-hot encoding for nominal ones
- D. Label encoding is always superior to one-hot encoding in every scenario

Answer: C

Q715. What problem does the SMOTE technique address in preprocessing?

- A. It handles missing values by predicting them using regression models
- B. It removes highly correlated features from the dataset efficiently
- C. It standardizes all features to have the same scale and distribution
- D. It generates synthetic minority class samples to balance class distribution

Answer: D

Q716. What is the purpose of binning continuous variables into discrete groups?

- A. To increase the total number of features in the dataset significantly
- B. To convert numerical data into completely random categorical labels
- C. To remove all missing values from the continuous variable columns
- D. To reduce noise and capture non-linear patterns in the data values

Answer: D

Q717. How does the Interquartile Range method detect outliers in a dataset?

- A. Values above the mean are always considered outliers by the IQR method
- B. Only the single maximum and minimum values are considered outliers here
- C. Values outside 1.5 times the IQR from Q1 and Q3 are flagged as outliers
- D. Values within one standard deviation of median are flagged as outlier

Answer: C

Q718. What is data augmentation and when is it commonly used?

- A. Encrypting sensitive data fields before storing them in the database
- B. Removing features to reduce dimensionality in high-dimensional datasets
- C. Converting unstructured data into structured tabular format for use
- D. Creating modified copies of existing data to increase training set size

Answer: D

Q719. Why is it important to handle missing values before training a model?

- A. Missing values always indicate that the entire dataset is unreliable here
- B. Missing values only affect visualization and have no impact on models
- C. Most ML algorithms cannot process datasets that contain any missing values
- D. Handling missing values is optional since models automatically ignore them

Answer: C

Q720. What is feature scaling and why is it important for distance-based algorithms?

- A. It converts all features to categorical type for better model interpretability
- B. It duplicates important features to give them more weight during training
- C. It ensures features contribute equally so larger-scale features do not dominate
- D. It removes features with low variance to reduce the dimensionality of data

Answer: C

Q721. What does a right-skewed distribution indicate about the data?

- A. Most values are clustered on the left with a long tail extending right
- B. Most values are clustered on the right with a long tail extending left
- C. The distribution has exactly two peaks at equal distances from center
- D. Values are perfectly symmetrically distributed around the central mean

Answer: A

Q722. How does the Pearson correlation differ from the Spearman correlation?

- A. Pearson requires sorted data while Spearman works with unsorted data only
- B. Pearson only works with categorical data while Spearman works with numbers
- C. Pearson measures linear relationships while Spearman measures monotonic ones
- D. Pearson is always larger than Spearman for any given pair of variables

Answer: C

Q723. What is a pair plot used for in exploratory data analysis?

- A. Visualizing the decision boundary of a trained classification model
- B. Displaying all pairwise feature relationships and distributions simultaneously
- C. Comparing model performance metrics across different algorithm choices
- D. Showing only the relationship between the target variable and one feature

Answer: B

Q724. Why is it important to check for multicollinearity during EDA?

- A. Checking for multicollinearity is optional and never affects final results
- B. Multicollinearity only affects visualization and has no impact on any model
- C. Multicollinearity always improves model accuracy and prediction performance
- D. Highly correlated features can cause instability in regression model coefficients

Answer: D

Q725. What information does a Variance Inflation Factor provide?

- A. The optimal number of clusters for a K-means clustering algorithm
- B. The degree to which a feature is explained by other features in model
- C. The total number of missing values in each feature of the dataset
- D. The probability that a feature follows a normal distribution pattern

Answer: B

Q726. What is the purpose of a QQ plot in statistical analysis?

- A. To compare model predictions against actual observed target values
- B. To display the learning curve of a model during training process
- C. To visualize the decision boundary of a classification model used
- D. To assess whether data follows a specific theoretical distribution

Answer: D

Q727. How do violin plots improve upon standard box plots?

- A. They only display the mean and are simpler to interpret than box plots
- B. They show the actual probability density distribution shape of the data
- C. They remove all outlier information to simplify the visualization look
- D. They work only with categorical data unlike box plots which need numbers

Answer: B

Q728. What does the kurtosis of a distribution describe?

- A. The asymmetry of the distribution around its mean value point
- B. The tailedness and peakedness of the distribution compared to normal
- C. The number of distinct modes or peaks present in the distribution
- D. The total range between the minimum and maximum values observed

Answer: B

Q729. Why is it important to visualize data before applying machine learning algorithms?

- A. Visualization is only for presentation and does not affect modeling choices
- B. It is a mandatory step that all ML frameworks enforce before model training
- C. Visualization automatically selects the best algorithm for the given dataset
- D. It reveals patterns, outliers, and distributions that guide preprocessing decisions

Answer: D

Q730. What does a correlation matrix heatmap help identify?

- A. The exact prediction accuracy of a trained machine learning model
- B. Pairs of features with strong positive or negative linear relationships
- C. The best train-test split ratio for the given dataset size overall
- D. The optimal hyperparameters for a gradient boosting model training

Answer: B

Q731. How does regularization help prevent overfitting in linear regression models?

- A. It adds a penalty term for large coefficients constraining model complexity
- B. It increases the number of features to capture more complex data patterns
- C. It multiplies all input features by a constant scaling factor for balance
- D. It removes the least important training samples from the dataset entirely

Answer: A

Q732. What is the key assumption of Naive Bayes classification?

- A. The decision boundary must be linear in the original feature space always
- B. All features must be continuous numerical values with normal distribution
- C. The dataset must have exactly equal numbers of samples in every class
- D. Features are conditionally independent given the class label of the sample

Answer: D

Q733. How does the C parameter in SVM affect the decision boundary?

- A. C determines the number of support vectors regardless of the data layout
- B. Higher C creates a tighter margin penalizing misclassification more strictly
- C. C only affects the kernel type used and has no impact on the margin width
- D. Higher C creates a wider margin allowing more misclassification of points

Answer: B

Q734. What advantage does a decision tree have over logistic regression for certain tasks?

- A. Decision trees always achieve higher accuracy than logistic regression models
- B. Decision trees can capture non-linear relationships without feature engineering
- C. Decision trees require less training data than logistic regression to converge
- D. Decision trees never overfit regardless of the depth or number of features

Answer: B

Q735. What is the difference between hard and soft margin classification in SVM?

- A. Hard margin is faster to train while soft margin requires exponentially more time
- B. Hard margin uses linear kernels while soft margin always uses non-linear kernels
- C. Hard margin allows no violations while soft margin permits some misclassification
- D. Hard margin permits misclassification while soft margin requires perfect separation

Answer: C

Q736. Why might logistic regression outperform complex models on small datasets?

- A. Its simplicity and fewer parameters make it less prone to overfitting on small data
- B. Logistic regression always converges in fewer epochs than any other algorithm does
- C. It automatically augments small datasets to increase the number of training samples
- D. Logistic regression ignores the training data and uses only prior knowledge

Answer: A

Q737. What does the kernel trick enable SVM to do?

- A. Map data to a higher-dimensional space to find non-linear decision boundaries
- B. Reduce training time by eliminating the need for any optimization process
- C. Convert regression problems into classification problems without data changes
- D. Automatically select the best features from the dataset for classification

Answer: A

Q738. How does K in KNN affect the bias-variance tradeoff of the classifier?

- A. Both small and large K always produce identical bias and variance values here
- B. K has no effect on the bias-variance tradeoff of the KNN classifier at all
- C. Small K has low bias and high variance while large K has high bias low variance
- D. Small K has high bias and low variance while large K has low bias high variance

Answer: C

Q739. What is the purpose of the sigmoid function in logistic regression?

- A. It normalizes input features to have zero mean and unit variance values
- B. It maps raw output values to probabilities between zero and one range
- C. It determines the optimal learning rate for gradient descent optimization
- D. It selects the most important features for the model automatically here

Answer: B

Q740. Why is feature importance a useful output of tree-based models?

- A. It guarantees that the model will achieve the highest possible accuracy
- B. It automatically removes unimportant features from the training dataset
- C. It converts categorical features into numerical features for processing
- D. It reveals which features contribute most to the model prediction decisions

Answer: D

Q741. How does feature randomness in Random Forest improve model performance?

- A. It guarantees each tree will have identical accuracy on test data set
- B. It eliminates the need for any hyperparameter tuning in the model
- C. It automatically removes all irrelevant features from the dataset used
- D. It reduces correlation between trees leading to better ensemble diversity

Answer: D

Q742. What is the key difference between bagging and boosting ensemble techniques?

- A. Bagging trains models sequentially while boosting trains all models in parallel
- B. Bagging trains models in parallel independently while boosting trains sequentially
- C. Bagging always outperforms boosting on every possible dataset and problem type
- D. Bagging only works with neural networks while boosting only works with trees

Answer: B

Q743. What does the learning rate parameter control in gradient boosting?

- A. The contribution of each tree to the final ensemble prediction output
- B. The percentage of features randomly selected at each tree split
- C. The number of trees to include in the final ensemble model overall
- D. The maximum depth allowed for each individual decision tree used

Answer: A

Q744. What is stacking in the context of ensemble methods?

- A. Using a meta-learner to combine predictions from diverse base models
- B. Removing the worst-performing model from the ensemble iteratively
- C. Training the same model architecture on different random data subsets
- D. Simply averaging the predictions from all models in the ensemble

Answer: A

Q745. Why might an ensemble of weak learners outperform a single strong learner?

- A. Combining diverse weak learners reduces variance without increasing bias much
- B. Weak learners always have lower bias than any single strong learner model
- C. Weak learners are always faster to train than any single strong learner is
- D. An ensemble of weak learners never overfits regardless of ensemble size used

Answer: A

Q746. How does AdaBoost assign weights to training samples?

- A. All samples receive equal weights throughout the entire training process
- B. Misclassified samples receive higher weights so subsequent models focus on them
- C. Weights are randomly assigned at each iteration without considering any errors
- D. Correctly classified samples receive higher weights for reinforcement purposes

Answer: B

Q747. What is out-of-bag error estimation in Random Forest?

- A. Evaluating each tree on samples not included in its bootstrap training set
- B. Using a separate test set that was manually created before model training
- C. Measuring training error on the full dataset after all trees are built
- D. Running cross-validation with exactly five folds on the entire dataset used

Answer: A

Q748. What is the purpose of subsampling in Stochastic Gradient Boosting?

- A. To convert the boosting algorithm into a bagging algorithm automatically
- B. To ensure every sample is used exactly once across all trees in ensemble
- C. To increase training time by processing every sample multiple times each
- D. To reduce overfitting and training time by using random data subsets per tree

Answer: D

Q749. How does the max_depth parameter affect individual trees in an ensemble?

- A. Deeper trees always improve ensemble accuracy without any risk of overfitting
- B. Max depth has no effect on individual tree complexity or ensemble performance
- C. Deeper trees have lower bias and higher variance capturing more complex splits
- D. Deeper trees have higher bias and lower variance across ensemble predictions

Answer: C

Q750. Why is diversity important among base learners in an ensemble?

- A. Diversity ensures all models make exactly the same predictions on every input
- B. Diversity is irrelevant since all models should be identical for best results
- C. Diverse models make different errors that can cancel out when combined together
- D. Diverse models always have individually higher accuracy than identical models

Answer: C

Q751. How does DBSCAN determine cluster membership compared to K-Means?

- A. DBSCAN requires the number of clusters while K-Means determines it automatically
- B. DBSCAN groups points by density connectivity while K-Means uses centroid distance
- C. DBSCAN uses distance to centroids while K-Means uses density-based grouping
- D. Both algorithms use identical methods to determine cluster membership always

Answer: B

Q752. What is the elbow method used for in K-Means clustering?

- A. To decide the maximum number of iterations for the K-Means convergence
- B. To determine the optimal number of features to use for the clustering task
- C. To evaluate whether K-Means or DBSCAN is the better algorithm for the data
- D. To select the optimal K by finding the point where inertia reduction slows

Answer: D

Q753. What does the silhouette score measure in clustering evaluation?

- A. How well each point fits its cluster versus the nearest neighboring cluster
- B. The accuracy of cluster labels compared to true class labels in dataset
- C. The total computational time required to train the clustering algorithm
- D. The number of outliers present in each cluster after model convergence

Answer: A

Q754. How does hierarchical clustering differ from K-Means in its approach?

- A. Hierarchical builds a tree of nested clusters while K-Means uses flat partitions
- B. Hierarchical requires specifying K while K-Means does not need any parameters
- C. Hierarchical only works with text data while K-Means works with numbers only
- D. Hierarchical always produces fewer clusters than K-Means on the same data set

Answer: A

Q755. What is the role of the epsilon parameter in DBSCAN?

- A. It defines the minimum number of clusters the algorithm must discover
- B. It sets the maximum distance for two points to be considered neighbors
- C. It controls the learning rate for optimizing the cluster assignments
- D. It specifies the dimensionality reduction ratio before clustering starts

Answer: B

Q756. Why might PCA lose important information during dimensionality reduction?

- A. PCA always preserves all information regardless of components selected count
- B. It projects onto directions of maximum variance and may discard informative ones
- C. PCA can only reduce to exactly two dimensions and no other number of them
- D. It requires labeled data and fails when applied to unsupervised datasets

Answer: B

Q757. What is the difference between agglomerative and divisive hierarchical clustering?

- A. Agglomerative starts with individual points and merges while divisive starts whole and splits
- B. Both approaches always produce identical clustering results on any dataset given to them
- C. Agglomerative only works with numerical data while divisive handles categorical data only
- D. Agglomerative starts with one cluster and splits while divisive starts with individual points

Answer: A

Q758. How does t-SNE differ from PCA for dimensionality reduction?

- A. Both methods produce identical visualizations on any dataset given to them always
- B. PCA is non-linear while t-SNE is a purely linear dimensionality reduction method
- C. PCA preserves global linear structure while t-SNE preserves local non-linear structure
- D. PCA can only reduce to two dimensions while t-SNE works with any target dimension

Answer: C

Q759. What challenge does the curse of dimensionality pose for clustering algorithms?

- A. High-dimensional data always forms perfectly separated clusters automatically here
- B. It makes clustering faster because more dimensions provide more information now
- C. The curse of dimensionality only affects supervised learning and never clustering
- D. In high dimensions distance metrics become less meaningful reducing cluster quality

Answer: D

Q760. What is Gaussian Mixture Model clustering based on?

- A. A density-based approach identical to DBSCAN using epsilon neighborhoods
- B. Probabilistic assignment assuming data comes from a mixture of Gaussian distributions
- C. Hard assignment of each point to exactly one cluster using closest centroid
- D. A hierarchical approach that builds a dendrogram of cluster relationships

Answer: B

Q761. Why is accuracy a poor metric for imbalanced classification problems?

- A. Accuracy requires balanced classes and cannot be computed on imbalanced data sets
- B. Accuracy is only defined for regression problems and not for classification tasks
- C. A model predicting only the majority class achieves high accuracy despite being useless
- D. Accuracy always gives a value of zero when classes are imbalanced in the dataset

Answer: C

Q762. What is the ROC curve and what does it visualize?

- A. A plot of training loss versus validation loss over each training epoch number
- B. A plot of precision versus recall for different classification threshold values
- C. A plot of model complexity versus generalization error for model selection use
- D. A plot of true positive rate versus false positive rate at different thresholds

Answer: D

Q763. How does K-fold cross-validation improve upon a simple train-test split?

- A. It randomly assigns labels to the test set to simulate real-world data distribution
- B. It always uses exactly two folds regardless of the K value specified by the user
- C. It uses all data for both training and testing reducing variance of performance estimate
- D. It eliminates the need for any test data by using only training data for evaluation

Answer: C

Q764. What is the difference between micro and macro averaging for multi-class metrics?

- A. Micro averages per class then takes mean while macro aggregates all predictions globally
- B. Micro aggregates predictions globally while macro averages per class metrics equally
- C. Micro averaging is used for binary classification while macro is for multi-class only
- D. Both micro and macro averaging always produce exactly identical metric values here

Answer: B

Q765. What does the Area Under the ROC Curve represent?

- A. The probability a random positive sample is ranked higher than a random negative one
- B. The exact accuracy of the model at the optimal classification threshold selected
- C. The maximum number of features the model can effectively use for prediction
- D. The total training time required to achieve the best validation set performance

Answer: A

Q766. When should you use stratified K-fold cross-validation instead of standard K-fold?

- A. When class distributions are imbalanced to ensure each fold represents all classes
- B. When all features are categorical and need special handling during evaluation
- C. When the dataset is very large and standard K-fold is computationally too expensive
- D. When the dataset has perfectly balanced class proportions across all groups

Answer: A

Q767. What is the bias-variance tradeoff in model evaluation?

- A. Bias and variance are independent and changing one has no effect on the other one
- B. Higher bias always leads to higher variance and both increase simultaneously always
- C. The tradeoff only applies to neural networks and not to any traditional ML algorithms
- D. Simpler models have high bias low variance while complex models have low bias high variance

Answer: D

Q768. What is the purpose of a learning curve in model evaluation?

- A. To visualize the feature importance rankings from a trained random forest model
- B. To display the optimal hyperparameters found during grid search tuning process
- C. To show how training and validation performance change with training set size
- D. To plot the distribution of prediction errors across all test set observations

Answer: C

Q769. How does the log loss metric differ from accuracy for evaluation?

- A. Log loss and accuracy always produce identical rankings of model performance
- B. Log loss can only be used with regression models while accuracy is for classifiers
- C. Log loss measures prediction speed while accuracy measures correctness of results
- D. Log loss penalizes confident wrong predictions more heavily than accuracy does

Answer: D

Q770. What is the purpose of a calibration curve in model evaluation?

- A. To assess whether predicted probabilities match actual observed frequencies
- B. To determine the optimal number of features to include in the final model
- C. To measure how fast the model makes predictions on new test data samples
- D. To visualize the convergence of the training loss during model optimization

Answer: A

Q771. What is the wrapper method for feature selection?

- A. Selecting features alphabetically by their column name in the original data frame
- B. Using a model to evaluate different feature subsets and selecting the best one found
- C. Selecting features based only on their individual correlation with the target variable
- D. Removing all features with missing values regardless of their predictive importance

Answer: B

Q772. How does target encoding work for high-cardinality categorical features?

- A. It replaces each category with the mean of the target variable for that group
- B. It replaces each category with a randomly assigned numerical identifier value
- C. It creates one binary column for each unique category in the feature column
- D. It removes all categories that appear fewer than ten times in the dataset

Answer: A

Q773. What is the purpose of creating interaction features between variables?

- A. To convert all numerical features into categorical features for tree models
- B. To reduce the total number of features in the dataset to save memory space
- C. To capture combined effects that individual features alone cannot represent
- D. To remove multicollinearity between features by merging them into one column

Answer: C

Q774. Why is the filter method for feature selection computationally efficient?

- A. It requires training the model multiple times with different feature subsets each
- B. It evaluates features independently of any model using statistical measures only
- C. It uses a complex model to evaluate every possible subset of feature combinations
- D. It only works with datasets that have fewer than one hundred features total

Answer: B

Q775. What is the benefit of creating date-based features from timestamps?

- A. Machine learning algorithms automatically extract date features without engineering
- B. Date-based features always decrease model accuracy compared to raw timestamps
- C. Extracting components like day of week and month reveals cyclical patterns in data
- D. Timestamps are already in the optimal format for all machine learning algorithms

Answer: C

Q776. How does the Variance Threshold method work for feature selection?

- A. It increases the variance of low-variance features through data augmentation
- B. It selects features with the highest variance above the target variable mean
- C. It calculates the variance between features to detect multicollinearity issues
- D. It removes features with variance below a specified threshold as uninformative

Answer: D

Q777. What is feature hashing and when is it useful?

- A. It sorts features by importance and removes the least important ones from data
- B. It encrypts feature values for security during model training in production use
- C. It maps high-dimensional categorical features to fixed-size vectors using a hash
- D. It creates polynomial features from the original feature set for better models

Answer: C

Q778. Why might you apply a log transformation to a highly skewed feature?

- A. It removes all outliers from the feature by compressing them to zero values
- B. Log transformation converts categorical features to numerical features directly
- C. Log transformation always improves model accuracy regardless of distribution shape
- D. It reduces skewness and helps satisfy normality assumptions of some algorithms

Answer: D

Q779. What is the embedded method for feature selection?

- A. Selecting features by running the model on every possible subset combination
- B. Selecting features before training any model using only statistical measures
- C. Feature selection performed as part of the model training process itself directly
- D. Removing all features except the one with the highest correlation to target

Answer: C

Q780. How do you handle cyclical features like hour of day or day of week?

- A. Treat them as regular continuous numerical features without any transformation
- B. Simply remove cyclical features since they do not contain useful information
- C. Encode them using sine and cosine transformations to preserve cyclic continuity
- D. Convert them to categorical features and apply one-hot encoding to each value

Answer: C

Q781. Why does the vanishing gradient problem occur in deep networks?

- A. Because the learning rate is always set too high for deep architectures
- B. The problem only occurs in shallow networks and never in deep architectures
- C. Deep networks always have too few parameters to compute meaningful gradients
- D. Gradients shrink exponentially through many layers making early layers learn slowly

Answer: D

Q782. How does batch normalization improve deep neural network training?

- A. It normalizes layer inputs reducing internal covariate shift for faster training
- B. It removes the need for any activation functions in the neural network layers
- C. It replaces the optimizer and directly updates weights without gradient descent
- D. It increases the batch size automatically to improve training speed for models

Answer: A

Q783. What is the purpose of dropout regularization in neural networks?

- A. It increases the number of neurons in each layer to improve model capacity
- B. It drops entire training samples that the model already classifies correctly
- C. It randomly deactivates neurons during training to prevent co-adaptation of features
- D. It permanently removes neurons that have zero weight values in the network

Answer: C

Q784. How does the Adam optimizer combine momentum and adaptive learning rates?

- A. Adam uses only momentum without any adaptive learning rate adjustment at all
- B. It maintains running averages of both gradients and squared gradients for updates
- C. It uses a fixed learning rate that never changes throughout the training process
- D. Adam ignores gradient history and uses only the current gradient for each update

Answer: B

Q785. What is the difference between a dense layer and a convolutional layer?

- A. Dense layers connect every neuron while convolutional layers use local receptive fields
- B. Dense layers only work with images while convolutional layers work with tabular data
- C. Dense layers use local connections while convolutional layers connect every neuron pair
- D. Both layer types are identical and can be used interchangeably in any architecture

Answer: A

Q786. Why is the learning rate one of the most important hyperparameters?

- A. The learning rate only affects the final accuracy and not the training speed at all
- B. It only matters for the first epoch and is automatically adjusted after that point
- C. The learning rate has no effect on training dynamics and can be set to any value
- D. Too high causes divergence and too low causes extremely slow convergence of loss

Answer: D

Q787. What is transfer learning in the context of deep learning?

- A. Transferring data from one storage system to another for model training use
- B. Using a pre-trained network and fine-tuning it for a different related task
- C. Training a model from random initialization on a completely new dataset only
- D. Converting a classification model into a regression model without retraining

Answer: B

Q788. How does gradient clipping prevent training instability in deep networks?

- A. It caps gradient values at a maximum threshold preventing exploding gradients
- B. It removes all gradients entirely preventing any weight updates from occurring
- C. It increases gradient values when they are too small to accelerate convergence
- D. It replaces all negative gradients with zero to ensure only positive updates

Answer: A

Q789. What is the purpose of an embedding layer in deep learning?

- A. To reduce the learning rate during training to prevent overfitting of model
- B. To map discrete tokens to dense continuous vector representations efficiently
- C. To normalize the input data before it enters the first hidden layer values
- D. To compute the final loss value at the output of the neural network model

Answer: B

Q790. Why do deep learning models typically require GPU acceleration?

- A. GPUs have more storage space than CPUs for holding the training dataset files
- B. GPUs excel at parallel matrix operations which dominate deep learning computation
- C. GPUs are required by all deep learning frameworks and cannot use CPU at all
- D. GPUs automatically optimize hyperparameters while CPUs cannot perform this task

Answer: B

Q791. How does the LSTM cell solve the vanishing gradient problem of standard RNNs?

- A. It processes sequences in reverse order to reduce the effective sequence length
- B. It uses a much larger learning rate than standard RNN implementations typically use
- C. It uses gates to control information flow allowing gradients to persist long-term
- D. It removes all activation functions to ensure gradients never shrink in magnitude

Answer: C

Q792. What is the purpose of the discriminator in a GAN?

- A. To distinguish between real data samples and generated fake data samples
- B. To preprocess input data before it is fed to the generator network model
- C. To generate realistic synthetic data samples from random noise input vectors
- D. To evaluate the final quality of the trained model on a held-out test set

Answer: A

Q793. How does dilated convolution differ from standard convolution in CNNs?

- A. Dilated convolution uses smaller filters than standard convolution always does
- B. Standard convolution has larger receptive fields than dilated convolution does
- C. Dilated convolution inserts gaps between filter elements expanding receptive field
- D. Dilated convolution only works with one-dimensional input data like text only

Answer: C

Q794. What is the key innovation of the self-attention mechanism in Transformers?

- A. It eliminates the need for any training data by using pre-defined attention maps
- B. It requires data to be processed in strict sequential order one step at a time
- C. It only attends to the immediately preceding element in the input sequence
- D. It computes relationships between all positions in a sequence simultaneously

Answer: D

Q795. What problem does mode collapse cause in GAN training?

- A. The discriminator becomes too weak to provide any useful training signal at all
- B. The generator produces diverse outputs but none of them look like real data
- C. The generator produces realistic but limited variety of outputs ignoring modes
- D. The training process converges too slowly requiring excessive computation time

Answer: C

Q796. How do 1x1 convolutions serve as a dimensionality reduction technique in deep networks?

- A. They only work with one-dimensional data and cannot be applied to image features
- B. They remove the need for any pooling layers in the convolutional neural network
- C. They increase the spatial dimensions of the feature maps while reducing channels
- D. They reduce the number of channels by combining feature maps with learned weights

Answer: D

Q797. What is the purpose of positional encoding in Transformer architectures?

- A. It replaces the need for any attention mechanism in the Transformer model
- B. It determines the optimal position for inserting new layers during training
- C. It encodes the position of the model parameters in memory for faster access
- D. It injects sequence position information since self-attention has no order notion

Answer: D

Q798. How does depthwise separable convolution reduce computational cost?

- A. It removes all bias terms from the convolutional layers to reduce parameters
- B. It skips every other spatial position during convolution to halve computation
- C. It uses only the first channel of input and ignores all other channels entirely
- D. It separates spatial and channel-wise filtering reducing multiply-add operations

Answer: D

Q799. What is the role of the encoder and decoder in a sequence-to-sequence model?

- A. The encoder handles training while the decoder is only used during evaluation
- B. The encoder generates output and decoder processes the raw input sequence data
- C. Both encoder and decoder perform identical operations on the same input data set
- D. The encoder compresses input to a representation and decoder generates the output

Answer: D

Q800. What advantage does the GRU have over the standard LSTM architecture?

- A. GRU uses fewer gates and parameters achieving similar performance more efficiently
- B. GRU can process images while LSTM is limited to only text and audio data types
- C. GRU does not use any gating mechanism unlike LSTM which uses three gates total
- D. GRU always achieves higher accuracy than LSTM on every sequential task given

Answer: A

Q801. How does TF-IDF improve upon simple bag-of-words representation?

- A. It counts words more quickly than bag-of-words without changing the representation
- B. It weights words by importance reducing the influence of common uninformative terms
- C. It only considers the first and last words in each document ignoring all others
- D. TF-IDF and bag-of-words produce identical representations for all text documents

Answer: B

Q802. What is the key innovation of Word2Vec compared to one-hot encoding?

- A. Word2Vec only works with very small vocabularies of less than one hundred words total
- B. Word2Vec produces identical representations as one-hot encoding for all vocabulary items
- C. Word2Vec requires labeled data for every word pair relationship in the corpus used
- D. Word2Vec creates dense vectors capturing semantic relationships between words efficiently

Answer: D

Q803. What is the difference between stemming and lemmatization?

- A. Stemming applies crude rules while lemmatization uses vocabulary and morphological analysis
- B. Stemming uses dictionary-based root finding while lemmatization uses rule-based cutting
- C. Stemming is more accurate than lemmatization for all natural language processing tasks
- D. Both methods always produce identical output for every word in all languages used

Answer: A

Q804. How does the attention mechanism in seq2seq models improve machine translation?

- A. It allows the decoder to focus on relevant source words for each output word generated
- B. Attention removes the need for both encoder and decoder in the translation pipeline
- C. It forces the decoder to translate each word independently without any context used
- D. It limits the model to translating only between English and one other target language

Answer: A

Q805. What is the purpose of byte-pair encoding in modern NLP tokenizers?

- A. It encrypts text using byte-level operations for secure communication purposes
- B. It splits every word into individual characters without any merging of subwords
- C. It converts all text to lowercase and removes all punctuation marks from strings
- D. It iteratively merges frequent character pairs creating a subword vocabulary set

Answer: D

Q806. How does the BERT model differ from earlier language models in its training approach?

- A. BERT uses bidirectional context by masking random tokens and predicting them
- B. BERT uses left-to-right unidirectional context like all previous language models did
- C. BERT processes text character by character unlike word-level previous models
- D. BERT requires no pre-training and is trained only on downstream task data sets

Answer: A

Q807. What is the purpose of the CLS token in BERT?

- A. It separates individual characters within each word of the input sequence
- B. It indicates missing or unknown words in the vocabulary of the tokenizer
- C. It provides a pooled representation for classification tasks using the model
- D. It marks the end of every sentence in the input text sequence provided

Answer: C

Q808. What is perplexity and how is it used to evaluate language models?

- A. It measures the vocabulary size of the language model's training corpus data
- B. It measures the training time required to converge to the optimal solution found
- C. It counts the number of parameters in the language model architecture structure
- D. It measures how well the model predicts text with lower values being better scores

Answer: D

Q809. How does transfer learning with pre-trained models benefit NLP tasks with limited data?

- A. Transfer learning only works when the new task has more data than the pre-training set
- B. Pre-trained models must always be retrained from scratch on new task specific data
- C. Pre-trained models provide rich language understanding that transfers to new tasks well
- D. Pre-trained models can only be used for the exact task they were originally trained on

Answer: C

Q810. What is the difference between extractive and abstractive text summarization?

- A. Extractive selects existing sentences while abstractive generates new sentences from scratch
- B. Extractive only works with short texts while abstractive only works with long documents
- C. Both methods always produce identical summaries for any given input document provided
- D. Extractive generates new sentences while abstractive selects existing sentences from input

Answer: A

Q811. How does transfer learning with pre-trained CNNs benefit computer vision tasks with limited data?

- A. Transfer learning only works for text classification and not for any image-based tasks
- B. Early layers learn general features that transfer well reducing data needs for new tasks
- C. Pre-trained models always perform worse than randomly initialized models on new tasks
- D. Pre-trained models must be completely retrained from scratch for every new CV task

Answer: B

Q812. What is the difference between semantic and instance segmentation?

- A. Semantic segmentation only works on grayscale images while instance works on color
- B. Semantic labels each pixel by class while instance distinguishes individual objects too
- C. Both segmentation types always produce identical pixel-level output maps on images
- D. Instance labels each pixel by class while semantic distinguishes individual objects only

Answer: B

Q813. How does the YOLO architecture achieve real-time object detection?

- A. It only detects objects in the center of the image ignoring all border regions always
- B. It frames detection as a single regression problem processing the entire image at once
- C. It processes each potential object region separately using a sliding window approach
- D. YOLO requires multiple passes through the image increasing detection time significantly

Answer: B

Q814. What is the purpose of non-maximum suppression in object detection?

- A. To maximize the size of all bounding boxes to ensure objects are fully contained
- B. To increase the number of detected bounding boxes for better detection coverage
- C. To remove redundant overlapping bounding boxes keeping only the best predictions
- D. To suppress all detections with confidence scores above a certain threshold level

Answer: C

Q815. How does image augmentation help prevent overfitting in computer vision models?

- A. It creates diverse variations forcing the model to learn invariant features robustly
- B. It reduces the training set size to prevent the model from memorizing all samples
- C. Augmentation always causes more overfitting because it adds artificial noise to data
- D. It only works with very large datasets and has no benefit for small training sets

Answer: A

Q816. What is the Intersection over Union metric used for in object detection?

- A. Measuring the total number of pixels in the input image being processed by model
- B. Determining the optimal learning rate for the detection model training run
- C. Measuring the overlap between predicted and ground truth bounding box regions
- D. Calculating the training time for the object detection model on GPU hardware

Answer: C

Q817. How do feature pyramid networks improve multi-scale object detection?

- A. They only detect objects at a single fixed scale ignoring all other size variations
- B. They reduce all objects to the same size before detection to simplify the algorithm
- C. They build a pyramid of feature maps at different scales for detecting objects of various sizes
- D. Feature pyramids increase detection time without improving accuracy at any scale value

Answer: C

Q818. What is the role of anchor boxes in modern object detection architectures?

- A. Anchor boxes are only used during evaluation and have no effect during training
- B. They provide predefined reference boxes of various sizes and ratios for detection
- C. They replace the non-maximum suppression step with a simpler averaging method
- D. They fix the position of all detected objects to the center of the input image always

Answer: B

Q819. How does a U-Net architecture achieve precise segmentation of medical images?

- A. It uses skip connections between encoder and decoder to preserve spatial details
- B. U-Net only works with natural images and cannot process medical imaging data well
- C. It requires pixel-level annotations for only ten percent of the training images used
- D. It only uses downsampling layers without any upsampling to produce segmentation maps

Answer: A

Q820. What is the purpose of batch normalization in deep CNN architectures?

- A. It normalizes layer activations stabilizing training and allowing higher learning rates
- B. Batch normalization increases training time without providing any accuracy benefit
- C. It only works with fully connected layers and cannot be used in convolutional layers
- D. It removes all color information from input images before processing them further

Answer: A

Q821. How does Apache Spark improve upon Hadoop MapReduce for iterative ML algorithms?

- A. Spark only works with small datasets while MapReduce handles large-scale data
- B. Spark is always slower than MapReduce but provides better fault tolerance here
- C. Spark writes intermediate results to disk just like MapReduce for reliability
- D. Spark keeps data in-memory across iterations avoiding costly disk read-write cycles

Answer: D

Q822. What is the difference between batch processing and stream processing?

- A. Both approaches process data identically and the terms are completely interchangeable here
- B. Batch processing is always faster than stream processing for all data processing workloads
- C. Batch processes data one item at a time while streaming processes all data at once together
- D. Batch processes all data at once while streaming processes data as it arrives continuously

Answer: D

Q823. What is data sharding and why is it used in distributed databases?

- A. It encrypts data using a shared key distributed across all machines in cluster
- B. It horizontally partitions data across machines enabling scalability and performance
- C. It compresses data into smaller files to reduce storage costs on each machine
- D. It creates backup copies of all data on a single machine for redundancy purposes

Answer: B

Q824. How does Apache Spark SQL enable big data analytics for users familiar with SQL?

- A. It converts all SQL queries into Python code before executing them on the cluster
- B. It replaces Apache Spark entirely with a traditional database management system
- C. Spark SQL only works with small datasets that fit in a single machine memory
- D. It provides a SQL interface over distributed DataFrames for familiar query syntax

Answer: D

Q825. What is the CAP theorem and why is it relevant to distributed data systems?

- A. It states distributed systems can have consistency, availability, and partition tolerance all at once
- B. CAP theorem only applies to single-machine databases and not to distributed systems at all
- C. It guarantees all distributed databases will always be faster than centralized database systems
- D. It states only two of consistency, availability, and partition tolerance can be guaranteed together

Answer: D

Q826. What role does Apache Airflow play in big data engineering workflows?

- A. It orchestrates and schedules complex data pipeline workflows as directed acyclic graphs
- B. It trains machine learning models on big data using distributed gradient descent
- C. It provides real-time stream processing capabilities for high-velocity data sources
- D. It stores large datasets in distributed file systems across multiple cluster nodes

Answer: A

Q827. How does data replication improve fault tolerance in distributed systems?

- A. Replication always degrades read performance to improve write speed of the system
- B. It compresses data before storing to reduce the total storage space required
- C. It copies data across multiple nodes so the system survives individual node failures
- D. Replication stores data on a single machine to reduce network overhead costs

Answer: C

Q828. What is the difference between a data lake and a data warehouse?

- A. Data lakes only support SQL queries while warehouses support all query languages used
- B. Data lakes store raw multi-format data while warehouses store structured processed data
- C. Data lakes store structured processed data while warehouses store raw unprocessed data
- D. Both store data identically and the terms are completely interchangeable in practice

Answer: B

Q829. How does columnar storage format benefit analytical queries on big data?

- A. Columnar storage is slower than row storage for all types of queries without exception
- B. It stores data row by row which is optimal for reading entire records efficiently
- C. It stores data column by column enabling efficient reads of only needed columns quickly
- D. It only works with numerical data and cannot handle text or categorical columns well

Answer: C

Q830. What is the role of a cluster manager like YARN in distributed computing?

- A. It converts distributed queries into single-machine queries for faster processing
- B. It allocates resources and schedules tasks across nodes in the computing cluster
- C. It handles data encryption and security for all files in the distributed system
- D. It manages the final output visualization of processed big data results on screen

Answer: B

Q831. What is model drift and why does it require monitoring?

- A. Model drift means the model is getting faster over time needing no intervention at all
- B. Model performance degrades over time as data distributions change requiring retraining
- C. Model drift only affects models deployed for less than one week in production systems
- D. It refers to the physical movement of servers causing hardware failures in the system

Answer: B

Q832. How does canary deployment reduce risk when releasing new ML models?

- A. It deploys the new model to all users simultaneously without any gradual rollout plan
- B. It routes a small percentage of traffic to the new model while monitoring for issues
- C. It only works for web applications and cannot be used with machine learning models
- D. Canary deployment prevents any model from being deployed to production at all ever

Answer: B

Q833. What is the purpose of model serialization in deployment?

- A. It deletes model weights that are no longer needed after training is completed
- B. It visualizes the model architecture as a diagram for documentation purposes only
- C. It converts trained models into a portable format for loading in production systems
- D. It trains multiple models in sequence one after another on the same training data

Answer: C

Q834. How does infrastructure as code benefit MLOps workflows?

- A. It eliminates the need for any cloud computing resources for ML model deployment
- B. It requires all infrastructure to be manually configured through web interfaces only
- C. It defines infrastructure in version-controlled code enabling reproducible environments
- D. Infrastructure as code only applies to web development and not to ML workflows

Answer: C

Q835. What is the role of experiment tracking in the ML development lifecycle?

- A. Experiment tracking is optional and provides no value to the ML development process
- B. It logs parameters, metrics, and artifacts enabling comparison and reproducibility
- C. It only tracks the final model and ignores all intermediate experiments and results
- D. It automatically selects the best model without any human review or intervention

Answer: B

Q836. What is the shadow deployment pattern for ML models?

- A. Shadow deployment completely ignores the old model predictions for all decisions
- B. The new model replaces the old model immediately without any parallel operation
- C. It deploys models only during nighttime hours when user traffic is at its lowest
- D. The new model runs alongside the old one comparing outputs without serving users

Answer: D

Q837. How does model compression benefit edge deployment scenarios?

- A. Model compression always improves accuracy while reducing the model size significantly
- B. It reduces model size and latency enabling deployment on resource-constrained devices
- C. Compression eliminates the need for any preprocessing of input data at inference
- D. It only works with linear models and cannot be applied to deep neural networks

Answer: B

Q838. What is the purpose of data validation in an ML pipeline?

- A. Data validation is unnecessary since all incoming data is always perfectly clean
- B. To verify data quality and schema consistency before model training or inference
- C. To train the model faster by removing data that takes long to process in pipeline
- D. To convert all data types to strings for uniform processing in the ML pipeline

Answer: B

Q839. How does blue-green deployment work for ML model updates?

- A. It trains the model in the blue environment and evaluates in the green one only
- B. It gradually increases traffic to the new model over several weeks of monitoring
- C. It maintains two identical environments and switches traffic between them instantly
- D. Blue-green deployment only works with batch processing and not real-time serving

Answer: C

Q840. What is the benefit of using ONNX format for ML model interoperability?

- A. ONNX always increases model inference time compared to native framework formats
- B. ONNX only supports models trained with PyTorch and no other ML framework at all
- C. ONNX provides a standard format allowing models to run across different frameworks
- D. ONNX can only represent simple linear models and not complex neural networks

Answer: C

Q841. How can training data bias lead to discriminatory AI outcomes?

- A. Training data bias always improves model performance and never causes discrimination
- B. Bias in training data is automatically detected and removed by all ML algorithms used
- C. If training data reflects historical discrimination the model learns and perpetuates it
- D. Training data bias only affects model speed and has no impact on fairness of outputs

Answer: C

Q842. What is the difference between individual fairness and group fairness in AI?

- A. Individual fairness focuses on groups while group fairness focuses on each individual person specifically
- B. Both concepts are identical and always produce the same fairness outcomes for all situations
- C. Individual fairness only applies to classification while group fairness only applies to regression
- D. Individual fairness treats similar individuals similarly while group fairness ensures equal outcomes across groups

Answer: D

Q843. What is the right to explanation under GDPR for automated decision-making?

- A. It requires companies to share all source code of their AI models publicly with everyone
- B. Individuals have the right to obtain meaningful information about the logic of automated decisions
- C. The right to explanation only applies to decisions made by human workers not AI systems
- D. GDPR does not address automated decision-making at all and has no such requirement

Answer: B

Q844. How does differential privacy protect individual data in ML training?

- A. Differential privacy removes all privacy protections to improve model accuracy scores
- B. It adds calibrated noise to queries or gradients making individual data points unidentifiable
- C. It requires all training data to be publicly available for anyone to inspect at any time
- D. It encrypts the entire training dataset using standard encryption algorithms for security

Answer: B

Q845. What is the problem of automation bias in AI-assisted decision-making?

- A. It refers to the bias introduced by automating the data collection process for ML
- B. Automation bias only occurs in fully autonomous systems and never in AI-assisted ones
- C. Users tend to over-rely on AI recommendations even when the AI may be incorrect here
- D. Users never trust AI recommendations and always override them with their own judgment

Answer: C

Q846. How can AI systems inadvertently create feedback loops that amplify bias?

- A. AI feedback loops always reduce bias over time through self-correcting learning loops
- B. Biased predictions influence future data collection which further reinforces the bias
- C. Feedback loops only occur in recommendation systems and never in other AI applications
- D. AI systems never create feedback loops because they process each prediction independently

Answer: B

Q847. What is the dual-use dilemma in AI technology development?

- A. AI technology can only be used for beneficial purposes and never for harmful applications
- B. Dual-use only refers to AI running on two different hardware platforms simultaneously here
- C. AI technology developed for harmful purposes can never be repurposed for beneficial uses
- D. The same AI technology can be used for both beneficial and harmful purposes creating tension

Answer: D

Q848. What role does model interpretability play in building trust with AI stakeholders?

- A. Model interpretability is only relevant for developers and has no value for end users
- B. Interpretability reduces model accuracy so it should be avoided for all production models
- C. All AI models are inherently interpretable and no additional effort is needed for trust
- D. It enables stakeholders to understand and verify AI reasoning building confidence in system

Answer: D

Q849. How does the concept of AI safety differ from AI ethics?

- A. AI safety is relevant only for AGI systems while AI ethics applies only to narrow AI systems
- B. AI safety focuses on preventing AI from causing harm while ethics addresses broader societal values
- C. AI safety and AI ethics are identical concepts with no meaningful difference between them
- D. AI safety only concerns physical robots while AI ethics only concerns software applications

Answer: B

Q850. What is the challenge of measuring fairness across multiple protected attributes simultaneously?

- A. Intersectional groups create complexity where satisfying fairness for one may violate another
- B. Only one protected attribute can be considered at a time and others must be fully ignored
- C. Multiple protected attributes have no interaction and can be handled completely independently
- D. Fairness across multiple attributes is trivial to achieve by optimizing for single attribute

Answer: A

Q851. What is the difference between narrow AI and general AI?

- A. Narrow AI is cheaper while general AI is expensive
- B. Narrow AI handles specific tasks while general AI can handle any intellectual task
- C. Narrow AI uses more data than general AI
- D. There is no difference between them

Answer: B

Q852. What role does a reward signal play in reinforcement learning?

- A. It provides labeled examples for classification
- B. It tells the agent how good or bad its action was
- C. It compresses the training data
- D. It removes outliers from data

Answer: B

Q853. Which type of machine learning discovers hidden patterns in data without labels?

- A. Supervised learning
- B. Reinforcement learning
- C. Unsupervised learning
- D. Semi-supervised learning

Answer: C

Q854. What is the role of a validation set in the machine learning workflow?

- A. To train the model parameters
- B. To tune hyperparameters and prevent overfitting during development
- C. To collect new data from users
- D. To deploy the model to production

Answer: B

Q855. What is semi-supervised learning?

- A. A method that uses only unlabeled data
- B. A method that combines a small amount of labeled data with a large amount of unlabeled data
- C. A method that only uses reinforcement signals
- D. A method that does not use any data

Answer: B

Q856. What is the explore-exploit tradeoff in reinforcement learning?

- A. Choosing between training and testing
- B. Balancing trying new actions versus using known rewarding actions
- C. Deciding between CPU and GPU usage
- D. Selecting between classification and regression

Answer: B

Q857. What is a knowledge graph in the context of AI?

- A. A chart showing training loss
- B. A structured representation of real-world entities and their relationships
- C. A neural network architecture
- D. A type of decision tree

Answer: B

Q858. Why is data quality important for machine learning models?

- A. Models always work regardless of data quality
- B. Poor data quality leads to inaccurate or biased model predictions
- C. Data quality only matters for image data
- D. High quality data slows down training

Answer: B

Q859. What is an AI pipeline?

- A. A physical pipe used in factories
- B. A sequence of data processing and modeling steps from raw data to predictions
- C. A single algorithm that solves all problems
- D. A hardware component in computers

Answer: B

Q860. What is the difference between a parametric and non-parametric model?

- A. Parametric models are always more accurate
- B. Parametric models have a fixed number of parameters while non-parametric models grow with data
- C. Non-parametric models cannot make predictions
- D. There is no practical difference

Answer: B

Q861. What is the purpose of partial derivatives in machine learning?

- A. To sort data alphabetically
- B. To measure the rate of change of a function with respect to one variable while holding others constant
- C. To create random numbers
- D. To delete features from a dataset

Answer: B

Q862. What is a probability distribution?

- A. A sorted list of numbers
- B. A function that describes the likelihood of different outcomes for a random variable
- C. A type of neural network
- D. A data storage method

Answer: B

Q863. What is matrix multiplication and when is it defined?

- A. Element-wise multiplication of any two matrices
- B. Multiplying rows of the first matrix by columns of the second, defined when the first matrix's columns equal the second's rows
- C. Always defined for any two matrices
- D. Adding corresponding elements of two matrices

Answer: B

Q864. What is the normal distribution and why is it important in ML?

- A. A distribution where all values are equal
- B. A bell-shaped symmetric distribution that naturally arises in many real-world phenomena due to the central limit theorem
- C. A uniform distribution over integers
- D. A distribution only used in physics

Answer: B

Q865. What is the inverse of a matrix used for?

- A. Transposing the matrix
- B. Solving systems of linear equations and computing certain ML closed-form solutions
- C. Increasing matrix dimensions
- D. Sorting matrix elements

Answer: B

Q866. What is the central limit theorem?

- A. A theorem about the center of a dataset
- B. The sampling distribution of the mean approaches a normal distribution as sample size increases, regardless of the population distribution
- C. All distributions are normal
- D. Larger samples are always biased

Answer: B

Q867. What is the Euclidean distance between two points?

- A. The Manhattan distance
- B. The straight-line distance computed as the square root of the sum of squared differences
- C. The cosine of the angle between them
- D. The absolute difference of their means

Answer: B

Q868. What is the difference between a discrete and continuous random variable?

- A. They are the same thing
- B. A discrete variable takes countable values while a continuous variable can take any value in a range
- C. Discrete variables are always larger
- D. Continuous variables can only be integers

Answer: B

Q869. Why is the logarithm function commonly used in machine learning?

- A. It makes all values negative
- B. It converts multiplicative relationships to additive ones and helps with numerical stability
- C. It is only used for display purposes
- D. Logarithms are not used in ML

Answer: B

Q870. What is a correlation coefficient?

- A. A measure of causation between variables
- B. A value between -1 and 1 that measures the linear relationship strength between two variables
- C. The mean of two variables
- D. The number of data points

Answer: B

Q871. What is the difference between `append()` and `extend()` for Python lists?

- A. They are identical
- B. `append()` adds a single element while `extend()` adds all elements from an iterable individually
- C. `extend()` adds a single element while `append()` adds multiple
- D. Neither modifies the original list

Answer: B

Q872. What is the purpose of the `with` statement in Python?

- A. Defining classes
- B. Managing resources with automatic cleanup via context managers
- C. Creating loops
- D. Importing modules

Answer: B

Q873. How does pandas handle missing data by default?

- A. It deletes the entire dataset
- B. It represents missing values as NaN and provides methods like fillna() and dropna() to handle them
- C. It replaces missing values with zeros automatically
- D. It crashes when encountering missing data

Answer: B

Q874. What is slicing in NumPy arrays?

- A. Deleting arrays
- B. Extracting subarrays using index ranges with start:stop:step syntax
- C. Sorting arrays
- D. Concatenating arrays

Answer: B

Q875. What is the purpose of sklearn.preprocessing.StandardScaler?

- A. It creates standard Python variables
- B. It standardizes features by removing the mean and scaling to unit variance
- C. It sorts data in standard order
- D. It converts data to string format

Answer: B

Q876. What does the map() function do in Python?

- A. Creates geographic maps
- B. Applies a given function to each item of an iterable and returns the results
- C. Maps keys to values in a dictionary
- D. Measures memory allocation

Answer: B

Q877. How do you handle exceptions in Python?

- A. Using if-else statements only
- B. Using try-except blocks to catch and handle errors gracefully
- C. Exceptions cannot be handled in Python
- D. Using import statements

Answer: B

Q878. What is the difference between loc and iloc in pandas?

- A. They are identical
- B. loc uses label-based indexing while iloc uses integer position-based indexing
- C. loc is faster than iloc always
- D. iloc uses label-based indexing

Answer: B

Q879. What is a decorator in Python?

- A. A way to add colors to output
- B. A function that modifies the behavior of another function without changing its source code
- C. A type of loop
- D. A data visualization tool

Answer: B

Q880. What does the pandas pivot_table() method do?

- A. Rotates a table physically
- B. Creates a spreadsheet-style pivot table for summarizing data by grouping and aggregating
- C. Deletes pivot columns
- D. Transposes the DataFrame

Answer: B

Q881. What is the difference between min-max scaling and z-score standardization?

- A. They produce identical results
- B. Min-max scales to a fixed range like 0-1 while z-score centers data to zero mean and unit variance
- C. Min-max is always better
- D. Z-score only works with integers

Answer: B

Q882. What is multivariate imputation and when is it preferred?

- A. Replacing missing values with zeros only
- B. Using relationships between multiple features to estimate missing values, preferred when features are correlated
- C. It only works for time series data
- D. It is never preferred over simple imputation

Answer: B

Q883. Why is stratified sampling important when splitting imbalanced datasets?

- A. It makes processing faster
- B. It ensures each split maintains the same class distribution as the original dataset
- C. It removes all minority class samples
- D. It is only used for text data

Answer: B

Q884. What is the purpose of a pipeline in scikit-learn for preprocessing?

- A. To create data visualizations
- B. To chain multiple preprocessing steps and the model into a single object ensuring consistent application
- C. To store data in databases
- D. To generate synthetic data

Answer: B

Q885. What is the effect of class imbalance on model training?

- A. It has no effect on model performance
- B. Models tend to be biased toward the majority class, performing poorly on minority class predictions
- C. It always improves accuracy
- D. It only affects unsupervised learning

Answer: B

Q886. What is feature binarization?

- A. Converting all features to text
- B. Converting numerical features to binary values using a threshold
- C. Deleting binary features
- D. Doubling all feature values

Answer: B

Q887. How does the Robust Scaler differ from Standard Scaler?

- A. They are identical
- B. Robust Scaler uses median and IQR instead of mean and std, making it resistant to outliers
- C. Robust Scaler is always slower
- D. Standard Scaler handles outliers better

Answer: B

Q888. What is the purpose of winsorization in data preprocessing?

- A. Deleting all outliers permanently
- B. Capping extreme values at a specified percentile to reduce the impact of outliers
- C. Converting data to Windows format
- D. Increasing the number of outliers

Answer: B

Q889. Why should you encode categorical variables before feeding data to most ML algorithms?

- A. Categorical data takes more storage
- B. Most ML algorithms perform mathematical operations and cannot process text categories directly
- C. Encoding is only needed for neural networks
- D. Categorical variables are always irrelevant

Answer: B

Q890. What is the difference between MCAR, MAR, and MNAR missing data?

- A. They are all the same type of missing data
- B. MCAR is completely random, MAR depends on observed data, and MNAR depends on the missing values themselves
- C. These are types of data encoding
- D. They describe data augmentation methods

Answer: B

Q891. What is the difference between a kernel density plot and a histogram?

- A. They are identical visualizations
- B. A KDE provides a smooth continuous estimate of the probability density while a histogram uses discrete bins
- C. KDE only works for categorical data
- D. Histograms are always more accurate

Answer: B

Q892. How does a parallel coordinates plot help with multivariate EDA?

- A. It only shows two variables at once
- B. It visualizes each observation as a line across parallel axes for each feature, revealing clusters and patterns in high-dimensional data
- C. It is a type of bar chart
- D. It can only show categorical data

Answer: B

Q893. What does the chi-squared test assess in EDA?

- A. The mean of a distribution
- B. Whether there is a statistically significant association between two categorical variables
- C. The slope of a regression line
- D. The normality of data

Answer: B

Q894. Why is it useful to create a missing value heatmap during EDA?

- A. To make the dataset look colorful
- B. To visually identify patterns in missing data across features and observations
- C. To remove all missing values automatically
- D. To encrypt sensitive columns

Answer: B

Q895. What does a jointplot in Seaborn display?

- A. Only a histogram
- B. A bivariate plot showing the relationship between two variables along with their marginal distributions
- C. A 3D surface plot
- D. A pie chart with labels

Answer: B

Q896. What is the significance of the interquartile range in identifying outliers?

- A. IQR identifies the mean value
- B. Points falling below $Q1 - 1.5 \times IQR$ or above $Q3 + 1.5 \times IQR$ are considered potential outliers
- C. IQR only applies to categorical data
- D. IQR measures correlation

Answer: B

Q897. What is the purpose of a lag plot in time series EDA?

- A. To create time delays in processing
- B. To check for autocorrelation by plotting each observation against a lagged version of itself
- C. To display GPS coordinates
- D. To compress time series data

Answer: B

Q898. How does the Shapiro-Wilk test complement visual normality assessment?

- A. It replaces all visualizations
- B. It provides a formal statistical test for normality to supplement visual methods like QQ plots and histograms
- C. It only tests for uniformity
- D. It is identical to the chi-squared test

Answer: B

Q899. What does a swarm plot show that a box plot might miss?

- A. The mean value
- B. Individual data points showing the actual distribution density and potential clusters within groups
- C. The standard deviation
- D. Only outliers

Answer: B

Q900. What is the purpose of profiling tools like pandas-profiling during EDA?

- A. To slow down analysis
- B. To automatically generate comprehensive data quality reports including statistics, correlations, and missing value analysis
- C. To convert data to PDF format
- D. To deploy ML models

Answer: B

Q901. What is elastic net regularization and when is it preferred?

- A. A type of neural network
- B. A combination of L1 and L2 regularization that balances feature selection and coefficient shrinkage
- C. A data preprocessing technique
- D. A visualization method

Answer: B

Q902. How does a decision tree handle continuous features for splitting?

- A. It ignores continuous features
- B. It evaluates all possible threshold values to find the split that maximizes information gain or minimizes impurity
- C. It converts them to binary first
- D. It only uses the mean value

Answer: B

Q903. What is the difference between batch gradient descent and stochastic gradient descent?

- A. They are identical algorithms
- B. Batch uses the entire dataset per update while SGD uses one sample, trading accuracy per step for faster iterations
- C. SGD uses the entire dataset
- D. Batch gradient descent is always faster

Answer: B

Q904. Why is polynomial regression considered a linear model despite fitting curves?

- A. It cannot fit curves
- B. It is linear in its parameters even though it fits nonlinear relationships by using polynomial features of the input
- C. It uses neural networks internally
- D. Polynomial regression is not a linear model

Answer: B

Q905. What is the difference between one-vs-rest and one-vs-one for multiclass classification?

- A. They produce identical results always
- B. One-vs-rest trains K binary classifiers while one-vs-one trains $K*(K-1)/2$ classifiers for K classes
- C. One-vs-rest only works for two classes
- D. One-vs-one is always less accurate

Answer: B

Q906. What is the purpose of cost-sensitive learning?

- A. To reduce the monetary cost of computing
- B. To assign different penalties for different types of misclassification errors based on their real-world impact
- C. To make models train faster
- D. To reduce the number of features

Answer: B

Q907. How does early stopping work as a regularization technique?

- A. It stops data collection early
- B. It monitors validation performance during training and stops when performance starts degrading, preventing overfitting
- C. It removes features early in the process
- D. It only works for linear models

Answer: B

Q908. What is the role of maximum margin in SVM classification?

- A. To minimize the margin between classes
- B. To find the hyperplane that maximizes the distance between itself and the nearest data points from each class
- C. To maximize training errors
- D. To reduce the number of features

Answer: B

Q909. What is the difference between parametric and non-parametric supervised learning methods?

- A. They are identical in approach
- B. Parametric methods assume a fixed functional form while non-parametric methods make fewer assumptions and grow in complexity with data
- C. Non-parametric methods always perform worse
- D. Parametric methods have no parameters

Answer: B

Q910. What is the purpose of grid search in hyperparameter tuning?

- A. To create data grids for visualization
- B. To exhaustively search through a specified subset of hyperparameter combinations to find the best model configuration
- C. To split data into grids
- D. To remove outliers from grid-shaped data

Answer: B

Q911. How does the subsample parameter in gradient boosting affect the model?

- A. It has no effect
- B. Using a fraction of training data per tree introduces randomness that reduces overfitting similar to bagging
- C. It increases training data size
- D. It controls the number of features per split

Answer: B

Q912. What is the difference between soft and hard voting in ensemble classification?

- A. They always produce identical results
- B. Hard voting uses majority class votes while soft voting averages predicted probabilities before making the decision
- C. Soft voting only works for regression
- D. Hard voting uses probabilities

Answer: B

Q913. How does feature importance work in Random Forest?

- A. It ranks features alphabetically
- B. It measures how much each feature contributes to reducing impurity across all trees in the forest
- C. It only considers the first tree
- D. It is based solely on feature correlations

Answer: B

Q914. What is the role of the max_depth parameter in tree-based ensembles?

- A. It controls the number of trees
- B. It limits how deep each individual tree can grow, trading expressiveness for regularization
- C. It sets the maximum number of features
- D. It determines the learning rate

Answer: B

Q915. How does LightGBM achieve faster training than XGBoost?

- A. It uses fewer trees
- B. It uses histogram-based splitting and leaf-wise tree growth instead of level-wise, reducing computation
- C. It skips the boosting step
- D. It uses a simpler loss function

Answer: B

Q916. What is the purpose of regularization parameters like lambda and alpha in XGBoost?

- A. To increase training speed only
- B. To penalize complex trees by adding L1 and L2 penalties on leaf weights, reducing overfitting
- C. To add more features
- D. To skip cross-validation

Answer: B

Q917. What is blending in ensemble learning?

- A. Mixing data from different datasets
- B. A simplified stacking approach where base model predictions on a holdout set are used to train a meta-learner
- C. Averaging all features together
- D. Removing outliers from predictions

Answer: B

Q918. Why does boosting typically have higher bias than bagging at the start but lower at the end?

- A. Boosting always has low bias
- B. Boosting starts with a weak learner (high bias) and sequentially reduces bias by correcting errors, while bagging combines full models from the start
- C. Bagging always has higher bias
- D. Bias is irrelevant in ensemble methods

Answer: B

Q919. What is the purpose of column subsampling in gradient boosting?

- A. To reduce the number of rows
- B. To randomly select a subset of features for each tree, increasing diversity and reducing overfitting
- C. To change the number of output classes
- D. To modify the learning rate

Answer: B

Q920. How do you determine the optimal number of boosting rounds?

- A. Always use exactly 100 rounds
- B. Use early stopping on a validation set, stopping when validation performance stops improving
- C. More rounds are always better
- D. Use exactly as many rounds as features

Answer: B

Q921. What is the difference between K-Means and K-Medoids clustering?

- A. They are identical algorithms
- B. K-Means uses mean centroids while K-Medoids uses actual data points as cluster centers, making it more robust to outliers
- C. K-Medoids is always faster
- D. K-Means uses actual data points as centers

Answer: B

Q922. How does the explained variance ratio help in choosing the number of PCA components?

- A. It is not useful for PCA
- B. It shows the proportion of total variance captured by each component, helping choose enough components to retain a desired percentage of information
- C. It measures clustering quality
- D. It indicates the number of outliers

Answer: B

Q923. What is the difference between DBSCAN's epsilon and minPoints parameters?

- A. They control the same thing
- B. Epsilon defines the neighborhood radius while minPoints sets the minimum neighbors needed to form a dense core point
- C. Epsilon counts points and minPoints measures distance
- D. They are both optional parameters

Answer: B

Q924. What is the Calinski-Harabasz index for clustering evaluation?

- A. A measure of classification accuracy
- B. The ratio of between-cluster dispersion to within-cluster dispersion, where higher values indicate better-defined clusters
- C. A feature selection metric
- D. A dimensionality reduction technique

Answer: B

Q925. How does Gaussian Mixture Model clustering assign soft cluster memberships?

- A. It uses hard assignments like K-Means
- B. Each data point receives a probability of belonging to each cluster based on the Gaussian component likelihoods
- C. It only assigns points to the nearest center
- D. It does not produce cluster assignments

Answer: B

Q926. What is the difference between complete linkage and single linkage in hierarchical clustering?

- A. They produce identical results
- B. Complete linkage uses maximum inter-cluster distance while single linkage uses minimum, leading to different cluster shapes
- C. Single linkage always produces better clusters
- D. Complete linkage is single linkage with extra steps

Answer: B

Q927. What is the reconstruction error in autoencoders and why is it important?

- A. It measures classification accuracy
- B. It measures how well the autoencoder's output matches its input, indicating the quality of the learned representation
- C. It counts the number of layers
- D. It measures training speed

Answer: B

Q928. What is the role of the Expectation-Maximization algorithm in GMM?

- A. It is a clustering evaluation metric
- B. EM iteratively estimates cluster parameters (E-step: compute responsibilities, M-step: update parameters) until convergence
- C. It only initializes clusters
- D. It is used for supervised learning only

Answer: B

Q929. How does t-SNE preserve local structure during dimensionality reduction?

- A. It preserves global structure instead
- B. It converts pairwise distances to conditional probabilities and minimizes KL divergence between high and low-dimensional distributions
- C. It uses linear transformations only
- D. It discards all distance information

Answer: B

Q930. What is the purpose of the perplexity parameter in t-SNE?

- A. It measures classification accuracy
- B. It controls the effective number of nearest neighbors considered, balancing attention between local and global aspects of the data
- C. It sets the output dimensions
- D. It determines the number of clusters

Answer: B

Q931. What is the precision-recall tradeoff and why does it exist?

- A. They always improve together
- B. Increasing the decision threshold increases precision but decreases recall because fewer positive predictions are made but they are more confident
- C. There is no tradeoff between them
- D. They are the same metric

Answer: B

Q932. What is the advantage of using mean absolute percentage error for regression evaluation?

- A. It is always the best metric
- B. MAPE expresses errors as percentages making it scale-independent and intuitive for stakeholders to understand
- C. It works well when actual values are near zero
- D. MAPE is identical to RMSE

Answer: B

Q933. How does leave-one-out cross-validation work?

- A. It removes one feature at a time
- B. Each data point is used as a single test set while all remaining points are used for training, repeated for every point
- C. It splits data into two halves
- D. It only uses the first data point for testing

Answer: B

Q934. What is the adjusted R-squared and why is it better than regular R-squared?

- A. They are identical metrics
- B. Adjusted R-squared penalizes for the number of features, preventing artificial inflation from adding irrelevant variables
- C. Adjusted R-squared is always lower
- D. Regular R-squared accounts for model complexity

Answer: B

Q935. What is a classification report in scikit-learn?

- A. A report about the dataset
- B. A summary table showing precision, recall, F1-score, and support for each class along with averages
- C. A list of all model parameters
- D. A data quality report

Answer: B

Q936. When is the F-beta score more appropriate than the F1 score?

- A. Never, F1 is always best
- B. When you want to weight precision and recall differently, with $\beta > 1$ favoring recall and $\beta < 1$ favoring precision
- C. When you have exactly two classes
- D. When datasets are perfectly balanced

Answer: B

Q937. What is the purpose of a cumulative gains chart in model evaluation?

- A. To show cumulative training loss
- B. To show how much of the positive class is captured by targeting the top-scored percentage of predictions
- C. To display feature importance
- D. To measure computational gains

Answer: B

Q938. What is the difference between macro and weighted averaging for multi-class metrics?

- A. They always produce the same result
- B. Macro averaging treats all classes equally regardless of size while weighted averaging accounts for class frequency
- C. Macro averaging is always preferred
- D. Weighted averaging ignores class sizes

Answer: B

Q939. What is the Huber loss and when is it preferred over MSE?

- A. They are identical losses
- B. Huber loss is quadratic for small errors and linear for large errors, making it robust to outliers while still being differentiable
- C. Huber loss is always worse than MSE
- D. It only works for classification

Answer: B

Q940. What is the concept of an evaluation metric being proper?

- A. It means the metric is correctly coded
- B. A proper scoring rule is maximized when the predicted probabilities match the true probabilities, encouraging honest probability estimates
- C. It means the metric is simple
- D. Proper metrics are only for regression

Answer: B

Q941. How does frequency encoding work for categorical variables?

- A. It converts categories to sound frequencies
- B. It replaces each category with its frequency of occurrence in the dataset, capturing popularity information
- C. It only works for time series
- D. It is identical to one-hot encoding

Answer: B

Q942. What is the purpose of creating ratio and proportion features?

- A. To normalize all features to 1
- B. To capture relative relationships between variables that raw values alone may not express
- C. To divide the dataset into portions
- D. To calculate model accuracy

Answer: B

Q943. How do sine and cosine transformations encode cyclical features?

- A. They convert cycles to straight lines
- B. They map cyclical values to a circle using sin and cos, preserving the fact that values like 23:00 and 01:00 are close
- C. They only work for audio data
- D. They remove cyclical patterns from data

Answer: B

Q944. What is the difference between backward elimination and forward selection for feature selection?

- A. They always select the same features
- B. Backward elimination starts with all features and removes the least important, while forward selection starts empty and adds the most important
- C. They are identical algorithms
- D. Forward selection removes features

Answer: B

Q945. What is weight of evidence encoding for categorical variables?

- A. It measures the physical weight of data
- B. WoE measures the predictive power of each category by computing the log ratio of the distribution of events to non-events within each category
- C. It only works for numerical features
- D. It assigns random weights to categories

Answer: B

Q946. How does the correlation-based filter method work for feature selection?

- A. It selects features randomly
- B. It selects features that have high correlation with the target variable but low correlation with each other, reducing redundancy
- C. It only removes constant features
- D. It requires model training

Answer: B

Q947. What is the purpose of creating window-based features for time series data?

- A. To resize data visualizations
- B. To capture temporal patterns by computing statistics like rolling mean, sum, or standard deviation over fixed time windows
- C. To remove time information
- D. To encrypt temporal data

Answer: B

Q948. What is the Recursive Feature Elimination method?

- A. A method that recursively adds features
- B. RFE trains a model, removes the least important feature, retrains, and repeats until the desired number of features remains
- C. It eliminates all features recursively
- D. It is a type of data augmentation

Answer: B

Q949. Why might you create statistical aggregate features from grouped data?

- A. To reduce the number of rows only
- B. Aggregates like group means, counts, medians, and standard deviations capture entity-level behavioral patterns from transactional data
- C. Aggregation always loses useful information
- D. Statistical aggregates are only for visualization

Answer: B

Q950. What is the difference between one-hot encoding and binary encoding for high cardinality features?

- A. They produce identical outputs
- B. One-hot creates one column per category while binary encoding uses $\log_2(n)$ columns with binary representation, greatly reducing dimensionality
- C. Binary encoding creates more columns
- D. One-hot always uses fewer columns

Answer: B

Q951. How does the Leaky ReLU activation function address the dying ReLU problem?

- A. It has the same issue as ReLU
- B. Instead of outputting zero for negative inputs, Leaky ReLU outputs a small negative slope, keeping neurons active and allowing gradient flow
- C. It only works for positive inputs
- D. It eliminates all negative values

Answer: B

Q952. What is the purpose of weight regularization in neural networks?

- A. To make weights larger
- B. To penalize large weight values in the loss function, constraining model complexity and reducing overfitting
- C. To initialize weights to zero
- D. To speed up inference

Answer: B

Q953. What is the difference between SGD with momentum and vanilla SGD?

- A. They are identical
- B. Momentum accumulates past gradients to smooth updates and accelerate convergence, especially in directions with consistent gradient sign
- C. Vanilla SGD is faster
- D. Momentum only works for small datasets

Answer: B

Q954. What is the Xavier/Glorot initialization strategy?

- A. Setting all weights to one
- B. Initializing weights from a distribution scaled by the number of input and output neurons to maintain signal variance across layers
- C. Using random integers as weights
- D. Copying weights from another model

Answer: B

Q955. How does data augmentation help in deep learning?

- A. It reduces the dataset size
- B. It artificially increases training data diversity by applying transformations, helping the model generalize and reducing overfitting
- C. It cleans the data
- D. It only works for text data

Answer: B

Q956. What is the difference between online learning and batch learning in neural networks?

- A. They are the same approach
- B. Online learning updates weights after each sample while batch learning processes the entire dataset before updating
- C. Online learning is always better
- D. Batch learning processes one sample at a time

Answer: B

Q957. What is the purpose of the He initialization for ReLU networks?

- A. It is identical to Xavier initialization
- B. He initialization scales weights by $\sqrt{2/\text{fan\_in}}$ to account for ReLU's zeroing of negative values, maintaining proper signal variance
- C. It initializes all weights to zero
- D. It only works for sigmoid networks

Answer: B

Q958. What is label smoothing and how does it help training?

- A. Making labels blurry in images
- B. Replacing hard 0/1 labels with soft values like 0.1/0.9, reducing model overconfidence and improving generalization
- C. It removes labels from the dataset
- D. It only works for regression

Answer: B

Q959. What is the warmup strategy in learning rate scheduling?

- A. Preheating the GPU hardware
- B. Starting with a very small learning rate and gradually increasing it to the target rate over initial training steps
- C. Using a constant high learning rate
- D. Warmup is only used during inference

Answer: B

Q960. What is the purpose of gradient accumulation in deep learning?

- A. To store gradients permanently
- B. To simulate larger batch sizes by accumulating gradients over multiple mini-batches before updating weights, useful when GPU memory is limited
- C. To remove gradients
- D. It slows down training intentionally

Answer: B

Q961. How does the forget gate in an LSTM cell work?

- A. It forgets the model parameters
- B. It decides what information to discard from the cell state using a sigmoid gate that outputs values between 0 and 1
- C. It removes training data
- D. It deletes gradient values

Answer: B

Q962. What is the discriminator's role in GAN training?

- A. It generates new data
- B. It classifies inputs as real or fake, providing training signal to the generator to improve its outputs
- C. It preprocesses data
- D. It sets the learning rate

Answer: B

Q963. What is batch renormalization and why was it introduced?

- A. It is identical to batch normalization
- B. Batch renormalization corrects batch normalization's issues with small or non-iid mini-batches by using running statistics with trainable correction parameters
- C. It removes normalization entirely
- D. It only works for very large batches

Answer: B

Q964. What is the attention mechanism's key-query-value framework?

- A. A database query system
- B. Queries determine what to attend to, keys determine relevance scores, and values carry the information to be aggregated based on attention weights
- C. Keys contain the input data only
- D. Queries and keys are always identical

Answer: B

Q965. What is the purpose of the bottleneck layer in a ResNet block?

- A. To slow down training
- B. To reduce computational cost by using 1x1 convolutions to decrease and then increase the channel dimensions around the expensive 3x3 convolution
- C. To create an autoencoder
- D. To remove skip connections

Answer: B

Q966. How does teacher forcing work in training sequence generation models?

- A. Using a teacher to label data
- B. During training, the model receives the actual ground truth output token as input for the next step rather than its own predicted token
- C. It only works during inference
- D. It replaces the decoder entirely

Answer: B

Q967. What is group normalization and when is it preferred over batch normalization?

- A. They are the same technique
- B. Group normalization normalizes over groups of channels within each sample, independent of batch size, preferred when batch sizes are too small for stable batch statistics
- C. Group norm requires large batches
- D. It only works for RNNs

Answer: B

Q968. What is the difference between transposed convolution and upsampling in CNNs?

- A. They are identical operations
- B. Transposed convolution learns upsampling parameters while standard upsampling uses fixed interpolation methods like bilinear or nearest neighbor
- C. Upsampling always produces better results
- D. Transposed convolution reduces dimensions

Answer: B

Q969. What is the concept of feature reuse in DenseNet architecture?

- A. Reusing training data
- B. Each layer receives feature maps from all preceding layers as input, enabling direct feature reuse, gradient flow improvement, and parameter efficiency
- C. Only the last layer reuses features
- D. Feature reuse causes overfitting

Answer: B

Q970. What is knowledge distillation's dark knowledge?

- A. Secret information hidden in models
- B. The soft probability distributions over non-target classes produced by the teacher model, containing inter-class similarity information
- C. Dark knowledge does not exist
- D. It refers to unlabeled data

Answer: B

Q971. How does the Skip-gram model in Word2Vec work?

- A. It skips words during processing
- B. It predicts surrounding context words given a center word, learning word embeddings that capture semantic relationships
- C. It classifies sentences into categories
- D. It only works with consecutive words

Answer: B

Q972. What is the difference between GloVe and Word2Vec embeddings?

- A. They produce identical embeddings
- B. Word2Vec uses local context windows with neural networks while GloVe uses global word co-occurrence statistics with matrix factorization
- C. GloVe is always better
- D. Word2Vec uses global statistics

Answer: B

Q973. What is the masked language modeling objective used in BERT?

- A. Removing all words from text
- B. Randomly masking some input tokens and training the model to predict the masked tokens from bidirectional context
- C. Translating masked text
- D. Only predicting the next word

Answer: B

Q974. What is the role of the encoder in a machine translation system?

- A. It translates the text directly
- B. It processes the input sentence into a contextual representation that captures its meaning for the decoder to use
- C. It only tokenizes the input
- D. It generates the output translation

Answer: B

Q975. What is the difference between autoregressive and autoencoding language models?

- A. They are identical approaches
- B. Autoregressive models predict tokens left-to-right (GPT-style) while autoencoding models use bidirectional context with masked tokens (BERT-style)
- C. Autoregressive models are always better
- D. Autoencoding models generate text one token at a time

Answer: B

Q976. What is named entity recognition used for in practical applications?

- A. Naming files on a computer
- B. Identifying and classifying entities like person names, organizations, locations, and dates in text for information extraction
- C. Renaming variables in code
- D. Creating new entity types

Answer: B

Q977. How does beam search differ from greedy decoding for text generation?

- A. They produce identical outputs
- B. Beam search maintains multiple candidate sequences at each step, exploring a broader search space than greedy decoding which only keeps the single best token
- C. Greedy decoding is always better
- D. Beam search uses random sampling

Answer: B

Q978. What is the purpose of the attention mask in Transformer models?

- A. To hide the model from users
- B. To prevent the model from attending to padding tokens or future positions, ensuring valid attention computation
- C. To mask training errors
- D. To speed up computation without any functional purpose

Answer: B

Q979. What is text embedding and how does it differ from word embedding?

- A. They are identical concepts
- B. Text embedding produces a single vector for an entire sentence or document while word embedding produces individual vectors for each word
- C. Text embedding only works for short texts
- D. Word embedding handles entire documents

Answer: B

Q980. What is the purpose of the temperature parameter in language model generation?

- A. It controls GPU temperature
- B. It scales the logits before softmax, with lower values making output more focused and deterministic and higher values increasing randomness and diversity
- C. It has no effect on output
- D. It controls the learning rate

Answer: B

Q981. How does the region proposal network work in Faster R-CNN?

- A. It manually selects regions
- B. RPN slides a small network over the feature map to predict object proposals and objectness scores simultaneously
- C. It uses random cropping for proposals
- D. RPN is a separate model from the detector

Answer: B

Q982. What is the difference between top-1 and top-5 accuracy in image classification?

- A. They measure the same thing
- B. Top-1 requires the correct class to be the single highest prediction while top-5 counts it as correct if the true class is among the five highest predictions
- C. Top-5 is always lower than top-1
- D. They are used for different types of images

Answer: B

Q983. What is mean Average Precision in object detection evaluation?

- A. The average image resolution
- B. mAP averages the precision across different recall levels and across all object classes, providing a single performance metric for detection
- C. It counts the number of detected objects
- D. It measures processing speed

Answer: B

Q984. How does the U-Net skip connections help with image segmentation?

- A. They skip the encoder entirely
- B. Skip connections concatenate encoder features with decoder features, providing high-resolution spatial information that helps precise boundary localization
- C. They remove low-level features
- D. Skip connections only connect the first and last layers

Answer: B

Q985. What is the purpose of image augmentation techniques like random erasing?

- A. To corrupt data permanently
- B. Random erasing masks random rectangular patches in training images, forcing the model to rely on diverse features rather than specific local patterns
- C. To reduce dataset size
- D. To remove objects from images permanently

Answer: B

Q986. What is the difference between single-shot and multi-scale object detection?

- A. They detect objects at the same scale
- B. Multi-scale detection processes images at multiple resolutions or uses feature pyramids to detect objects of various sizes, while single-scale may miss small or large objects
- C. Single-shot always detects more objects
- D. Multi-scale only works for large objects

Answer: B

Q987. What is semantic segmentation's pixel-wise classification?

- A. Classifying entire images into categories
- B. Assigning a class label to every individual pixel in the image, creating a dense prediction map
- C. Detecting objects with bounding boxes
- D. Counting pixels in each color

Answer: B

Q988. How does Mixup data augmentation work for training CV models?

- A. Mixing different datasets together
- B. Creating new training samples by linearly interpolating between pairs of images and their labels, encouraging smoother decision boundaries
- C. Mixing color channels randomly
- D. Combining two models into one

Answer: B

Q989. What is the role of backbone networks in object detection architectures?

- A. They provide structural support for servers
- B. Backbone networks extract hierarchical visual features from the input image that are shared by detection head and other task-specific modules
- C. They generate bounding boxes directly
- D. Backbones only process text data

Answer: B

Q990. What is the difference between instance segmentation and panoptic segmentation?

- A. They are identical tasks
- B. Instance segmentation only handles countable objects while panoptic segmentation unifies both instance segmentation for things and semantic segmentation for stuff regions
- C. Panoptic only detects backgrounds
- D. Instance segmentation covers everything in the image

Answer: B

Q991. How does Apache Spark's lazy evaluation improve performance?

- A. It makes Spark slower
- B. Spark builds an execution plan without computing results until an action is called, enabling query optimization and avoiding unnecessary computations
- C. It processes data immediately
- D. Lazy evaluation has no effect on performance

Answer: B

Q992. What is the difference between a data lake and a data lakehouse?

- A. They are identical concepts
- B. A data lakehouse combines the flexibility of data lakes with the management and performance features of data warehouses, supporting both analytics and ML workloads
- C. Data lakehouses only store structured data
- D. Data lakes are always better

Answer: B

Q993. What is the purpose of data serialization formats like Parquet and Avro?

- A. They are programming languages
- B. They efficiently encode structured data for storage and transmission, with Parquet optimized for columnar analytics and Avro for row-based serialization
- C. They only work with text data
- D. They replace databases entirely

Answer: B

Q994. How does Apache Kafka enable real-time data streaming?

- A. It stores data in batches only
- B. Kafka provides a distributed publish-subscribe messaging system where producers publish events to topics and consumers process them in near real-time
- C. It replaces all databases
- D. Kafka only works with text data

Answer: B

Q995. What is the role of a distributed file system in big data?

- A. It only stores small files
- B. It stores large datasets across multiple machines, providing fault tolerance through replication and enabling parallel access to data
- C. It replaces local file systems entirely
- D. It only works with structured data

Answer: B

Q996. What is the purpose of a message queue in big data architectures?

- A. To send emails between users
- B. To decouple data producers from consumers, buffering messages and enabling asynchronous processing at different rates
- C. To queue database queries
- D. To sort messages alphabetically

Answer: B

Q997. How does data partitioning improve query performance in big data systems?

- A. It always slows down queries
- B. Partitioning divides data into segments based on key columns, allowing queries to scan only relevant partitions instead of the entire dataset
- C. It creates duplicate data
- D. Partitioning only works for small datasets

Answer: B

Q998. What is the difference between structured, semi-structured, and unstructured data in big data?

- A. They are all the same format
- B. Structured has fixed schema (tables), semi-structured has flexible schema (JSON/XML), and unstructured has no schema (images/text)
- C. Only structured data is useful
- D. Unstructured data cannot be processed

Answer: B

Q999. What is the purpose of Apache Airflow in big data workflows?

- A. It provides air conditioning for data centers
- B. Airflow is a workflow orchestration platform that schedules, monitors, and manages complex data pipeline dependencies as directed acyclic graphs
- C. It processes data in real-time
- D. It replaces Apache Spark

Answer: B

Q1000. What is the concept of data locality in distributed computing?

- A. Storing all data in one location
- B. Moving computation to where the data resides rather than moving large datasets to the compute nodes, minimizing network transfer
- C. Data must always be moved to a central server
- D. Data locality only matters for small datasets

Answer: B

Q1001. What is the difference between concept drift and data drift?

- A. They are identical phenomena
- B. Data drift is a change in input feature distributions while concept drift is a change in the relationship between features and the target variable
- C. Concept drift only affects deep learning
- D. Data drift does not require monitoring

Answer: B

Q1002. How does the shadow deployment pattern work for ML models?

- A. Models run in the dark
- B. The new model receives production traffic and generates predictions alongside the existing model, but only the old model's predictions are served while the new model's performance is monitored
- C. It replaces the model immediately
- D. Shadow models never see real data

Answer: B

Q1003. What is the purpose of a model registry in MLOps?

- A. A government registration system
- B. A centralized repository that stores, versions, tracks, and manages ML model artifacts, metadata, and lifecycle stages
- C. A list of model architectures
- D. A training data storage system

Answer: B

Q1004. What is infrastructure as code and why is it important for MLOps?

- A. Writing code on physical infrastructure
- B. Defining and managing infrastructure through code files, enabling reproducible, version-controlled, and automated provisioning of ML environments
- C. It only applies to web development
- D. IaC replaces all manual processes

Answer: B

Q1005. How does model compression benefit edge deployment?

- A. It makes models larger
- B. Techniques like pruning, quantization, and distillation reduce model size and computation, enabling deployment on resource-constrained edge devices like mobile phones and IoT sensors
- C. Edge devices have unlimited resources
- D. Compression always reduces accuracy unacceptably

Answer: B

Q1006. What is the role of feature flags in ML model deployment?

- A. Flags used in feature engineering
- B. Feature flags enable gradually enabling new models for specific user segments, allowing controlled rollout and quick disabling if issues arise
- C. They only work for software features
- D. Feature flags are identical to A/B tests

Answer: B

Q1007. What is the difference between online and offline model evaluation?

- A. They produce identical results
- B. Offline evaluation uses held-out historical data while online evaluation measures performance on live production traffic with real user interactions
- C. Online evaluation is always more accurate
- D. Offline evaluation uses production data

Answer: B

Q1008. What is the purpose of model explainability tools in production?

- A. To make models more complex
- B. To provide human-understandable explanations for individual predictions, supporting debugging, regulatory compliance, and user trust
- C. Explainability is only for research
- D. All models are inherently explainable

Answer: B

Q1009. What is the concept of model staleness in production?

- A. Models that smell bad
- B. A model becoming less accurate over time as the real-world data distribution changes from the data it was trained on
- C. Models always improve with age
- D. Staleness only affects neural networks

Answer: B

Q1010. What is the concept of algorithmic recourse and why is it important?

- A. Recourse means repeating algorithms
- B. Algorithmic recourse provides individuals with actionable steps to change an unfavorable AI decision, ensuring people are not permanently harmed by automated systems
- C. It is a type of optimization algorithm
- D. Recourse is only relevant for recommendation systems

Answer: B

Q1011. How can feedback loops amplify bias in AI systems?

- A. Feedback loops always reduce bias
- B. Biased AI decisions influence future data collection which reinforces the original bias, creating a cycle where the system becomes increasingly biased over time
- C. Feedback loops do not exist in AI
- D. They only affect recommendation systems

Answer: B

Q1012. What is the difference between explainability and interpretability in AI?

- A. They are identical concepts
- B. Interpretability means a model is inherently understandable while explainability provides post-hoc explanations for complex models that are not inherently transparent
- C. Explainability is always preferred
- D. Interpretability only applies to neural networks

Answer: B

Q1013. What is the concept of privacy-preserving machine learning?

- A. Making ML models private property
- B. Techniques that enable training and inference on sensitive data while protecting individual privacy, including federated learning, differential privacy, and secure computation
- C. Privacy cannot be preserved in ML
- D. It only applies to healthcare data

Answer: B

Q1014. What are model cards and why are they important?

- A. Physical cards with model images
- B. Standardized documentation describing a model's intended use, performance metrics, limitations, and ethical considerations, promoting transparency and responsible deployment
- C. They are used for trading models
- D. Model cards replace all other documentation

Answer: B

Q1015. What is the problem of proxy discrimination in AI systems?

- A. Using proxy servers for AI
- B. When AI uses seemingly neutral features that are highly correlated with protected attributes like race or gender, resulting in discriminatory outcomes despite not explicitly using those attributes
- C. Proxy discrimination cannot occur in AI
- D. It only happens with simple models

Answer: B

Q1016. How does the concept of AI safety differ from AI security?

- A. They are identical concerns
- B. AI safety focuses on preventing unintended harmful behaviors from AI systems themselves, while AI security protects AI systems from external malicious attacks
- C. Safety is only about physical robots
- D. Security is not relevant to AI

Answer: B

Q1017. What is the right to erasure and how does it apply to AI systems?

- A. Erasing AI code permanently
- B. Under GDPR, individuals can request deletion of their personal data, which creates challenges for AI systems trained on that data as the model may retain learned patterns
- C. It does not apply to AI
- D. Only applies to social media

Answer: B

Q1018. What is the environmental impact of training large AI models?

- A. AI training has zero environmental impact
- B. Training large models consumes significant energy and produces carbon emissions, with some large model training runs emitting as much CO2 as multiple transatlantic flights
- C. Environmental impact is negligible
- D. Only GPU manufacturing has environmental impact

Answer: B

Q1019. What is the concept of AI value alignment?

- A. Aligning data values in a database
- B. Ensuring AI systems' goals, behaviors, and decisions are consistent with human values, intentions, and ethical principles
- C. Aligning model parameters to specific values
- D. Value alignment is not important for AI

Answer: B

Q1020. What is the purpose of a staging environment in ML deployment?

- A. A theater stage for AI presentations
- B. A pre-production environment that mirrors production for testing models with realistic data and infrastructure before live deployment
- C. A storage area for old models
- D. A training environment for data scientists

Answer: B

AI / Machine Learning & Data Analytics — Hard

510 questions

Q1021. The Chinese Room argument by John Searle challenges:

- A. Standard statistical data analytics methods
- B. Strong AI - that machines can truly understand
- C. Modern machine learning training algorithms
- D. Weak AI - that machines can mimic intelligence

Answer: B

Q1022. Which of the following problems is considered AI-complete?

- A. Sorting a list of numbers
- B. Computing matrix multiplication
- C. Performing a binary search
- D. Natural language understanding

Answer: D

Q1023. The frame problem in AI refers to:

- A. Challenges of reducing network latency across distributed systems
- B. Complications of managing memory allocation in modern computers
- C. Difficulty in determining what changes and what stays the same after an action
- D. Problems with optimizing database indexing for query performance

Answer: C

Q1024. Which approach to AI attempts to mimic biological neural structures?

- A. Symbolism
- B. Evolutionary computation
- C. Bayesian methods
- D. Connectionism

Answer: D

Q1025. In the context of AI, what is the combinatorial explosion?

- A. Rapid growth of possible solutions making brute force infeasible
- B. An advanced technique for compressing large data files
- C. A sudden failure in hardware caused by electrical overload
- D. A specialized type of deep neural network architecture

Answer: A

Q1026. What distinguishes Artificial General Intelligence (AGI) from Narrow AI?

- A. AGI consumes significantly less working memory
- B. AGI requires absolutely no input training data
- C. AGI can perform any intellectual task a human can
- D. AGI is faster at one specific computational task

Answer: C

Q1027. The symbol grounding problem in AI concerns:

- A. How to optimize the speed of sorting algorithms
- B. How to efficiently compress large data archives
- C. How to increase the clock speed of a processor
- D. How symbols in an AI system get their meaning

Answer: D

Q1028. Which of the following is a characteristic of a multi-agent system?

- A. Only one autonomous decision-maker
- B. Multiple interacting intelligent agents
- C. A single centralized processing unit
- D. No communication between components

Answer: B

Q1029. What is the knowledge representation bottleneck in AI?

- A. Slow data transfer speeds across network connections
- B. Difficulty in encoding real-world knowledge into a formal system
- C. Insufficient physical storage space on server hardware
- D. Limited pixel resolution on visual display monitors

Answer: B

Q1030. In the context of data analytics, what is the CRISP-DM model?

- A. A standard process model for data mining projects
- B. A multi-layer neural network architecture design
- C. A general-purpose interpreted programming language
- D. A scalable relational database management system

Answer: A

Q1031. The Hessian matrix contains:

- A. Eigenvalues of the matrix only
- B. The determinant values only
- C. Second-order partial derivatives
- D. First-order derivatives only

Answer: C

Q1032. The Kullback-Leibler divergence measures:

- A. The spread or variance of a single distribution
- B. The central tendency or mean of a distribution
- C. How one probability distribution differs from another
- D. The linear correlation between two variables

Answer: C

Q1033. A Jacobian matrix represents:

- A. First-order partial derivatives of a vector-valued function
- B. Only the output values of a single scalar function
- C. Only the second-order derivatives of a scalar function
- D. Only the eigenvalues of a square transformation matrix

Answer: A

Q1034. In the context of ML, why is the log-likelihood used instead of likelihood?

- A. It requires significantly less memory during training
- B. It converts products to sums, making computation easier
- C. It produces more statistically accurate final estimates
- D. It runs faster on modern parallel computing hardware

Answer: B

Q1035. The singular value decomposition (SVD) decomposes a matrix into:

- A. U, Sigma, and V-transpose matrices
- B. Only the eigenvector components
- C. Exactly two factor matrices
- D. Only the eigenvalue components

Answer: A

Q1036. A saddle point in optimization is where:

- A. The objective function reaches its overall global minimum
- B. The gradient of the function is completely undefined
- C. The objective function reaches its overall global maximum
- D. The gradient is zero but it is neither a minimum nor maximum

Answer: D

Q1037. The moment generating function uniquely determines:

- A. Only the median statistic
- B. Only the expected mean value
- C. The probability distribution
- D. Only the variance estimate

Answer: C

Q1038. In matrix calculus, the gradient of $x^T A x$ with respect to x is:

- A. $2x$
- B. Ax
- C. A^T
- D. $(A + A^T)x$

Answer: D

Q1039. The condition number of a matrix indicates:

- A. The trace or diagonal sum of the matrix
- B. Sensitivity of the solution to small changes in input
- C. The total number of columns in the matrix
- D. The total number of rows in the matrix

Answer: B

Q1040. Jensen's inequality states that for a convex function f :

- A. $f(E[X]) \leq E[f(X)]$
- B. $f(E[X]) \geq E[f(X)]$
- C. $f(E[X]) = 0$
- D. $f(E[X]) = E[f(X)]$

Answer: A

Q1041. What is the Global Interpreter Lock (GIL) in Python?

- A. A file system locking mechanism for preventing concurrent write access
- B. A garbage collection tool for automated memory resource management
- C. A cryptographic security feature for encrypting sensitive runtime data
- D. A mutex preventing multiple threads from executing Python bytecode simultaneously

Answer: D

Q1042. Which tool is best for parallel processing of large datasets in Python?

- A. Flask
- B. Django
- C. Tkinter
- D. Dask

Answer: D

Q1043. What is the difference between deepcopy and copy in Python?

- A. deepcopy creates copies of nested objects recursively, copy does not
- B. copy is faster and creates a more thorough deep duplication
- C. deepcopy only works on list objects and not dictionaries
- D. They are completely identical in behavior and performance

Answer: A

Q1044. In NumPy, what does np.einsum() do?

- A. Performs only element-wise operations on flat one-D arrays
- B. Performs Einstein summation for multi-dimensional array operations
- C. Creates identity matrices of a specified dimension size
- D. Calculates eigenvalues of a square matrix decomposition

Answer: B

Q1045. What is vectorization in the context of NumPy?

- A. Converting natural language text into numeric vectors
- B. Performing operations on entire arrays instead of loops
- C. Creating scalable vector graphics for visualization
- D. A lossless data compression encoding technique

Answer: B

Q1046. What does the __slots__ attribute do in Python classes?

- A. Defines abstract methods for interface contracts
- B. Enables multiple inheritance across class chains
- C. Restricts instance attributes and reduces memory usage
- D. Creates new class-level methods from descriptors

Answer: C

Q1047. Which library provides GPU-accelerated computing for Python ML?

- A. CuPy
- B. Pandas
- C. Matplotlib
- D. SciPy

Answer: A

Q1048. What is a context manager in Python and why is it useful in ML?

- A. It creates persistent database connections for data loading
- B. It manages resources using with statements ensuring proper cleanup
- C. It handles asynchronous network requests for API access
- D. It manages the training context of machine learning models

Answer: B

Q1049. What is the purpose of __call__ in a Python class?

- A. Copies the object into a new reference
- B. Initializes a new class from a template
- C. Destroys the instance and frees memory
- D. Makes an instance callable like a function

Answer: D

Q1050. How does memory mapping (np.memmap) help in handling large datasets?

- A. It maps files to memory allowing access without loading the entire file
- B. It encrypts data at rest to ensure privacy and compliance
- C. It compresses data into smaller archive files for disk savings
- D. It duplicates data across servers for redundancy and backups

Answer: A

Q1051. What is the SMOTE technique used for?

- A. Reducing the dimensionality of the feature space with PCA
- B. Oversampling the minority class by creating synthetic examples
- C. Undersampling the majority class by randomly removing records
- D. Selecting only the most informative features from the dataset

Answer: B

Q1052. When should you use target encoding instead of one-hot encoding?

- A. When the overall dataset is small
- B. When there are only two categories present
- C. When categorical variables have high cardinality
- D. When the data is purely numerical

Answer: B

Q1053. What is data leakage in preprocessing?

- A. When data columns are encrypted for secure storage
- B. When data records are accidentally duplicated twice
- C. When data is physically lost during the cleaning step
- D. When information from the test set influences training

Answer: D

Q1054. What is the purpose of Yeo-Johnson transformation?

- A. To detect and remove extreme outlier values from the numerical columns
- B. To encode high-cardinality categorical features into integer representations
- C. To make data more normally distributed, handling both positive and negative values
- D. To fill in missing values using the mean or median of each column

Answer: C

Q1055. In handling missing data, what is MCAR?

- A. Missing data that follows a clear and observable systematic pattern
- B. Missing Completely At Random - missingness is unrelated to any variable
- C. Missing data that occurs in only one specific column of data
- D. Missing data that was caused by systematic data entry errors

Answer: B

Q1056. What is the KNN imputation method?

- A. Replacing missing values with randomly generated numerical estimates
- B. Deleting every row that contains any missing values from the dataset
- C. Filling all missing entries with the single overall global mean value
- D. Using K nearest neighbors to estimate missing values based on similar records

Answer: D

Q1057. Why should standardization be fit only on training data?

- A. Because training data is inherently more important
- B. Because the computation is significantly faster
- C. To prevent data leakage from test set statistics
- D. Because the test data is a much smaller sample

Answer: C

Q1058. What is the Isolation Forest algorithm used for in preprocessing?

- A. Automated feature selection
- B. Anomaly and outlier detection
- C. Data range normalization
- D. Missing value imputation

Answer: B

Q1059. What is ordinal encoding and when is it preferred over one-hot encoding?

- A. It assigns ordered integers to categories with natural ordering
- B. It is exactly the same approach as basic label encoding
- C. It removes categorical variables from the feature set entirely
- D. It always creates separate binary columns for each category

Answer: A

Q1060. What is the effect of multicollinearity on preprocessing?

- A. It has absolutely no effect on preprocessing or model training
- B. Highly correlated features can cause instability in model coefficients
- C. It always improves the overall accuracy of the trained model
- D. It consistently reduces the total required model training time

Answer: B

Q1061. Simpson's paradox in EDA refers to:

- A. A trend that appears in groups but reverses when groups are combined
- B. A visualization error caused by incorrect axis scale calibration
- C. A specific type of missing data pattern observed in large datasets
- D. A random sampling method for selecting representative subsets

Answer: A

Q1062. What is the purpose of the Kolmogorov-Smirnov test?

- A. To calculate the arithmetic mean of a data sample
- B. To test if a sample comes from a specific distribution
- C. To encode categorical data into numerical integers
- D. To remove extreme outlier values from the dataset

Answer: B

Q1063. What is the curse of dimensionality in EDA?

- A. There are too few features for the model to learn
- B. There is simply too much data volume to process
- C. Data becomes increasingly sparse as dimensions increase
- D. The data is already too clean to analyze further

Answer: C

Q1064. Cramér's V statistic measures:

- A. The skewness of a single distribution
- B. Correlation between two numerical variables
- C. The arithmetic mean of a given dataset
- D. Association between two categorical variables

Answer: B

Q1065. What is the difference between correlation and causation in EDA?

- A. Correlation measures association; causation means one variable directly affects another
- B. Causation is a weaker relationship than correlation
- C. Correlation always directly implies a causal link
- D. They are completely identical concepts in statistics

Answer: B

Q1066. What is the Shapiro-Wilk test used for?

- A. Testing for the presence of outlier values
- B. Testing if a dataset is normally distributed
- C. Testing for the existence of duplicate rows
- D. Testing for the count of missing null values

Answer: B

Q1067. In EDA, what is a bimodal distribution?

- A. A completely uniform distribution
- B. A distribution with two distinct peaks
- C. A distribution with no peaks at all
- D. A distribution with one single peak

Answer: B

Q1068. What does the Durbin-Watson statistic test for?

- A. Normality of the data
- B. Homoscedasticity
- C. Multicollinearity issues
- D. Autocorrelation in residuals

Answer: B

Q1069. What is heteroscedasticity and why is it problematic?

- A. Non-constant variance of residuals, violating regression assumptions
- B. A specific type of variable correlation
- C. A pattern of missing data values
- D. Perfectly constant variance of all residuals

Answer: C

Q1070. What is the purpose of the Andrews curves visualization?

- A. To display proportional breakdowns in a circular pie format
- B. To visualize multivariate data as curves, each representing one observation
- C. To create standard grouped and stacked bar charts for categories
- D. To plot time series trend lines with seasonal decomposition

Answer: B

Q1071. What is the kernel trick in SVM?

- A. Mapping data to a higher-dimensional space without explicit computation
- B. A dimensionality reduction visualization projection technique
- C. A wrapper-based method for iterative feature selection
- D. A preprocessing data cleaning and normalization technique

Answer: A

Q1072. What is the VC dimension?

- A. The total count of input features in the dataset
- B. The step size or magnitude of the learning rate
- C. A measure of the capacity and complexity of a classification model
- D. The total number of training samples available

Answer: C

Q1073. In gradient boosting, what does each subsequent tree learn?

- A. The residuals (errors) of the previous ensemble
- B. Random patterns found within the noise
- C. The original target values from the training data
- D. Nothing new beyond the original input features

Answer: D

Q1074. What is the difference between hard and soft margin SVM?

- A. Hard margin and soft margin are completely identical approaches
- B. Hard margin is always the preferred choice over soft margin
- C. Soft margin SVM never works well on real-world datasets
- D. Hard margin allows no misclassification; soft margin allows some with a penalty

Answer: C

Q1075. What is the Representer Theorem?

- A. A foundational theorem about efficient compressed data representation and storage formats
- B. The optimal solution in kernel methods can be expressed as a linear combination of kernel evaluations at training points
- C. A graphical visualization principle for plotting high-dimensional feature space relationships
- D. A data warehousing method for efficiently storing columnar records in structured databases

Answer: B

Q1076. What is Platt scaling?

- A. A method to calibrate SVM outputs into probabilities using a sigmoid function
- B. A standardization feature scaling method using z-score transforms of values
- C. A model combination ensemble technique using weighted averaging of outputs
- D. A min-max data normalization technique for scaling feature ranges to a bound

Answer: A

Q1077. What is the effect of increasing the C parameter in SVM?

- A. The model becomes much simpler and heavily regularized overall
- B. The decision boundary margin becomes significantly wider
- C. The model becomes more sensitive to individual points, risking overfitting
- D. The amount of L2 regularization penalty increases significantly

Answer: C

Q1078. What is the difference between Gini impurity and entropy in decision trees?

- A. Gini impurity is always the better choice for all tree problems
- B. Both measure node impurity but Gini is computationally simpler; entropy uses logarithms
- C. Entropy is always the better choice for all tree problems
- D. They produce completely different tree structures every single time

Answer: B

Q1079. What is the Structural Risk Minimization principle?

- A. Always using the simplest available model regardless of performance
- B. Maximizing model complexity to fit all training data points exactly
- C. Balancing empirical risk and model complexity to minimize generalization error
- D. Minimizing only the training error and ignoring test performance

Answer: B

Q1080. In logistic regression, what does multicollinearity cause?

- A. More accurate and well-calibrated probability estimates
- B. Significantly faster convergence during the training process
- C. Unstable and unreliable coefficient estimates with large standard errors
- D. Consistently better and more accurate model predictions overall

Answer: C

Q1081. What is the bias-variance decomposition of ensemble methods?

- A. Bagging primarily reduces variance; boosting primarily reduces bias
- B. Neither technique affects the bias or variance components
- C. Both techniques only reduce the variance component of error
- D. Both techniques only reduce the bias component of error

Answer: A

Q1082. How does XGBoost handle regularization differently from traditional Gradient Boosting?

- A. XGBoost uses no regularization at all in its training objective
- B. Traditional Gradient Boosting has stronger regularization overall
- C. They handle regularization in the exact same identical manner
- D. XGBoost includes L1 and L2 regularization terms in its objective function

Answer: D

Q1083. What is the difference between XGBoost and LightGBM's tree growing strategy?

- A. Neither of them actually uses tree models
- B. XGBoost grows level-wise; LightGBM grows leaf-wise
- C. They grow trees in an identical manner always
- D. XGBoost uses leaf-wise; LightGBM uses level-wise

Answer: D

Q1084. In CatBoost, how are categorical features handled?

- A. They are silently ignored during the training and inference steps
- B. They are completely removed from the feature set before training
- C. They must be manually one-hot encoded before model training begins
- D. Using ordered target statistics with random permutations to avoid target leakage

Answer: D

Q1085. What is the effect of increasing the number of trees in a Random Forest?

- A. It always causes the model to underfit the training data
- B. Performance improves then plateaus; it generally does not overfit
- C. It always causes the model to overfit the training data
- D. Performance always decreases with each additional tree

Answer: B

Q1086. What is the role of the subsample parameter in Gradient Boosting?

- A. It sets the step size or learning rate of the optimization
- B. It selects which features are considered at each individual tree split
- C. It introduces stochastic gradient boosting by using a fraction of samples per tree
- D. It controls the maximum allowed depth of each individual tree

Answer: C

Q1087. How does the isolation mechanism work in Isolation Forest for ensemble anomaly detection?

- A. Anomalies require significantly more random splits to isolate
- B. The splits are entirely random and have no diagnostic meaning
- C. Anomalies are isolated in fewer splits because they are rare and different
- D. All data points require an exactly equal number of splits

Answer: C

Q1088. What is Bayesian Model Averaging?

- A. Random model selection from the available ensemble
- B. Simple majority voting across all base model predictions
- C. Weighting ensemble models by their posterior probabilities
- D. Using only the single highest-accuracy base model

Answer: C

Q1089. What is negative correlation learning in ensemble methods?

- A. Using the exact same training data for every individual model
- B. Removing highly correlated input features from the dataset
- C. Training all base models to produce completely identical outputs
- D. Encouraging base learners to make diverse errors through a penalty term

Answer: D

Q1090. How does DART (Dropouts meet Multiple Additive Regression Trees) improve boosting?

- A. By using significantly deeper trees for each individual boosting iteration
- B. By removing the least important features entirely from the training data
- C. By randomly dropping trees during boosting iterations to prevent over-specialization
- D. By adding many more trees to the overall boosted ensemble of the model

Answer: B

Q1091. What is the Gaussian Mixture Model (GMM)?

- A. A tree-based decision boundary model for supervised classification prediction
- B. A probabilistic model assuming data comes from a mixture of Gaussian distributions
- C. A linear regression method for predicting continuous numerical target values
- D. A filter-based feature selection technique using statistical hypothesis tests

Answer: B

Q1092. How does spectral clustering work?

- A. Uses deep neural networks to learn cluster assignment mappings
- B. Uses decision tree ensembles to group data into discrete clusters
- C. Uses eigenvalues of a similarity matrix to reduce dimensionality before clustering
- D. Applies K-Means clustering directly without any preprocessing transformation

Answer: C

Q1093. What is the UMAP algorithm?

- A. A dimensionality reduction technique based on manifold learning and topology
- B. A gradient-based numerical regression training technique
- C. A supervised multi-class classification prediction method
- D. An unsupervised density-based spatial clustering algorithm

Answer: A

Q1094. In DBSCAN, what are core points, border points, and noise points?

- A. Noise points are the ones that form the densest most cohesive clusters
- B. Core points have enough neighbors; border points are near core points; noise points are isolated
- C. Core points are the extreme outliers that are furthest from all clusters
- D. All three point types are treated exactly the same by the clustering algorithm

Answer: B

Q1095. What is the difference between hard and soft clustering?

- A. Hard assigns each point to one cluster; soft assigns probabilities of belonging to each cluster
- B. Soft clustering does not exist as a valid machine learning technique
- C. Hard clustering is always the strictly better approach for all problems
- D. They are completely identical approaches with no meaningful differences

Answer: A

Q1096. What is the Information Criterion (BIC/AIC) used for in clustering?

- A. Performing feature selection for supervised model training
- B. Measuring the internal purity of each individual data cluster
- C. Calculating the pairwise distances between all data points
- D. Selecting the optimal number of clusters by balancing fit and complexity

Answer: D

Q1097. What is Self-Organizing Map (SOM)?

- A. An unsupervised neural network that produces a low-dimensional representation of data
- B. A supervised classification model that predicts categorical labels
- C. A gradient boosting algorithm that builds sequential tree models
- D. A linear regression model that predicts continuous target values

Answer: A

Q1098. How does the OPTICS algorithm improve upon DBSCAN?

- A. It requires explicitly specifying the number K of clusters
- B. It only works with fully labeled and supervised training data
- C. It is always computationally faster than every other algorithm
- D. It handles clusters of varying density without requiring a fixed epsilon

Answer: D

Q1099. What is the manifold hypothesis in unsupervised learning?

- A. All real-world data is always uniformly distributed in space
- B. Data always naturally forms spherical well-separated clusters
- C. All feature dimensions are always statistically independent
- D. High-dimensional data often lies on a lower-dimensional manifold

Answer: D

Q1100. What is contrastive learning in the context of unsupervised representation learning?

- A. A linear regression method for predicting numerical outputs
- B. A purely supervised classification technique using labels
- C. A density-based spatial clustering approach for grouping data
- D. Learning representations by contrasting similar and dissimilar pairs

Answer: D

Q1101. What is the Matthews Correlation Coefficient (MCC)?

- A. The exact same metric as standard classification accuracy for balanced data
- B. A metric specifically designed for evaluating unsupervised clustering
- C. A metric that is only applicable to multi-class classification problems
- D. A balanced metric using all four confusion matrix values, ranging from -1 to 1

Answer: D

Q1102. What is the difference between Type I and Type II errors?

- A. The two types are completely identical errors
- B. Type I is a false positive; Type II is a false negative
- C. Type I is a false negative; Type II is a false positive
- D. Both are actually true positive outcomes

Answer: B

Q1103. What is nested cross-validation?

- A. An outer CV loop for evaluation with an inner CV loop for hyperparameter tuning
- B. A single random train-test split without any cross-validation folds
- C. Cross-validation performed entirely without any held-out validation data
- D. The same standard approach as regular single-loop K-fold cross-validation

Answer: A

Q1104. What is the Cohen's Kappa statistic?

- A. A metric that accounts for agreement occurring by chance between predicted and actual
- B. The exact same calculation as standard overall classification accuracy
- C. A regression-only metric for evaluating continuous predictions
- D. A feature importance measure for ranking input variables

Answer: A

Q1105. What is calibration in the context of model evaluation?

- A. Whether predicted probabilities reflect true likelihoods of outcomes
- B. The overall classification accuracy of the trained model
- C. A data preprocessing step for cleaning raw data
- D. The process of selecting the most relevant features

Answer: A

Q1106. What is the Brier score?

- A. A feature importance measure for trees
- B. The mean squared error of probabilistic predictions
- C. A ranking metric for ordering search results
- D. A clustering score for partition quality

Answer: B

Q1107. When should you use the PR-AUC instead of ROC-AUC?

- A. For evaluating unsupervised clustering partition quality
- B. When dealing with highly imbalanced datasets where positive class is rare
- C. When the class distribution is perfectly balanced and equal
- D. For evaluating regression models with continuous outputs

Answer: B

Q1108. What is the purpose of bootstrapping in model evaluation?

- A. Performing automated feature selection and dimensionality reduction
- B. Estimating confidence intervals for metrics by resampling with replacement
- C. Running hyperparameter tuning using grid or random search
- D. Increasing the total amount of available training data samples

Answer: B

Q1109. What is the difference between leave-one-out CV and K-fold CV?

- A. They are exactly identical approaches with no meaningful differences whatsoever
- B. LOOCV is always significantly faster computationally than K-fold approaches
- C. K-fold always produces strictly more accurate results than LOOCV methods
- D. LOOCV uses N-1 samples for training and 1 for testing each time; K-fold uses larger folds

Answer: D

Q1110. What is the expected calibration error (ECE)?

- A. A regression metric measuring the average absolute difference in predictions
- B. The exact same calculation as the standard log loss or cross-entropy loss metric
- C. The weighted average of difference between predicted confidence and actual accuracy across bins
- D. A clustering metric measuring the silhouette score across all partitions

Answer: C

Q1111. What is the Boruta algorithm for feature selection?

- A. A multi-layer feedforward deep neural network architecture for prediction
- B. A gradient-based linear regression method for continuous value estimation
- C. A wrapper method using shadow features and Random Forest to find relevant features
- D. A density-based unsupervised clustering algorithm for grouping similar data

Answer: C

Q1112. What is the difference between wrapper, filter, and embedded feature selection?

- A. Wrappers evaluate subsets with a model; filters use stats; embedded methods learn during training
- B. Filters use trained models to evaluate each candidate feature subset iteratively
- C. Wrappers are always strictly the best approach for every feature selection problem
- D. They are all identical approaches that produce the exact same feature subset selections

Answer: A

Q1113. What is the curse of dimensionality's impact on feature engineering?

- A. More features require exponentially more data to avoid sparsity
- B. The number of dimensions has no effect on model training
- C. Fewer features always cause the model to severely underfit
- D. More features always improve the model performance metrics

Answer: A

Q1114. What is mutual information in feature selection?

- A. A measure of the dependency between two variables based on information theory
- B. The linear correlation coefficient between two variables
- C. The overall variance of a single variable in a dataset
- D. The arithmetic mean of a single variable in a dataset

Answer: A

Q1115. What is feature crossing in deep learning?

- A. Combining two or more features to create a new feature capturing their interaction
- B. Normalizing features to have zero mean and unit variance
- C. Removing individual features from the dataset one at a time
- D. Visualizing features using scatter plots and histograms

Answer: A

Q1116. How does LASSO perform feature selection?

- A. By using only simple correlation analysis
- B. By removing features entirely at random
- C. By ranking all features only by variance
- D. By driving some feature coefficients exactly to zero through L1 regularization

Answer: C

Q1117. What is the difference between forward and backward feature selection?

- A. Forward starts empty and adds features; backward starts with all and removes features
- B. Forward selection actually removes features one at a time from model
- C. Backward selection actually adds features one at a time to the model
- D. Forward and backward selection are completely identical methods

Answer: D

Q1118. What is permutation importance?

- A. Simply ranking features by their alphabetical column order
- B. A method for sorting features by their variance values
- C. A data augmentation technique for expanding training data
- D. Measuring importance by shuffling each feature and observing the drop in performance

Answer: D

Q1119. What is the concept of information leakage in feature engineering?

- A. Features that are redundant copies of other existing features
- B. Features that have missing values in some of their data records
- C. Features containing target info that would not be available at prediction time
- D. Features that contain a large amount of measurement noise

Answer: C

Q1120. What is automated feature engineering (e.g., Featuretools)?

- A. Generating features completely at random without any data context
- B. Deleting features from the dataset to reduce dimensionality
- C. Relying exclusively on manual hand-crafted feature creation processes
- D. Using algorithms to auto-generate features from relational data via deep feature synthesis

Answer: D

Q1121. What is the dying ReLU problem?

- A. ReLU causes gradients to explode during backpropagation
- B. Neurons permanently output zero because they get stuck in the negative region
- C. ReLU increases peak memory usage beyond available capacity
- D. ReLU activation is computationally too slow for large networks

Answer: B

Q1122. What is the Lottery Ticket Hypothesis?

- A. All neural network architectures are equally effective regardless of size
- B. Sparse network architectures never work well for any real task
- C. Dense networks contain sparse subnetworks that can match full network performance
- D. Larger networks are always strictly better than smaller alternatives

Answer: C

Q1123. What is knowledge distillation?

- A. Removing irrelevant features from the input feature space
- B. Training a smaller student model to mimic a larger teacher model
- C. Transferring database records between different storage systems
- D. Augmenting the training data with synthetic generated samples

Answer: B

Q1124. What is the purpose of skip connections (residual connections)?

- A. Removing entire layers from the network to reduce complexity
- B. Skipping the data preprocessing step to speed up training
- C. Reducing the size of the training dataset to save storage
- D. Allowing gradients to flow directly by adding input to output of a block

Answer: D

Q1125. What is gradient clipping?

- A. Removing all gradients entirely from the backpropagation computation
- B. Setting all gradient values to exactly zero after each update step
- C. Increasing the learning rate to speed up convergence time
- D. Limiting gradient values to a maximum threshold to prevent exploding gradients

Answer: D

Q1126. What is the difference between SGD with momentum and Adam?

- A. SGD momentum uses a fixed learning rate with velocity; Adam adapts per parameter
- B. Adam optimizer never converges to a good solution on real data
- C. SGD with momentum is always faster than Adam for every problem
- D. They are completely identical optimization algorithms with no differences

Answer: A

Q1127. What is neural architecture search (NAS)?

- A. Automated methods for finding optimal neural network architectures
- B. Feature selection and dimensionality reduction algorithms
- C. Manual hand-designed network architecture engineering process
- D. Data preprocessing and cleaning pipeline automation tools

Answer: A

Q1128. What is the difference between pre-training and fine-tuning?

- A. They are completely identical stages in the same training process
- B. Fine-tuning always happens before the pre-training phase
- C. Pre-training is designed specifically for very small datasets
- D. Pre-training learns general representations; fine-tuning adapts to a specific task

Answer: D

Q1129. What is the effect of batch size on training?

- A. Smaller batch sizes are always strictly worse for model generalization
- B. Larger batch sizes are always strictly better for all training scenarios
- C. Larger batches give smoother gradients but may generalize worse; smaller batches add helpful noise
- D. Batch size has absolutely no measurable effect on training or generalization

Answer: C

Q1130. What is mixed-precision training?

- A. Using both 16-bit and 32-bit floating point to speed up training while maintaining accuracy
- B. Using only 8-bit integer arithmetic for all network computation steps
- C. Using only 64-bit double-precision floats for maximum numerical precision
- D. Training neural network models entirely without any GPU acceleration

Answer: A

Q1131. What is the difference between causal and bidirectional self-attention?

- A. Causal attention only looks at previous positions; bidirectional looks at all positions
- B. Bidirectional attention only looks backward at previous tokens in sequence
- C. Causal and bidirectional attention are completely identical mechanisms
- D. Causal attention actually looks at all positions in the input sequence

Answer: C

Q1132. What is the mode collapse problem in GANs?

- A. The generator produces limited variety instead of the full data distribution
- B. The model architecture has grown far too large for memory
- C. The discriminator fails completely and cannot distinguish anything
- D. The training process runs far too slowly to converge in time

Answer: A

Q1133. What is the Wasserstein GAN (WGAN) and how does it improve training?

- A. It uses only convolutional layers without any dense layers
- B. It uses the Wasserstein distance for more stable training gradients
- C. It completely removes the discriminator from the architecture
- D. It uses standard cross-entropy loss for the discriminator

Answer: B

Q1134. What is Neural ODE?

- A. Modeling continuous-depth neural networks using ordinary differential equations
- B. A generative adversarial network variant for image synthesis
- C. A standard discrete-layer feedforward neural network architecture
- D. An unsupervised density-based technique for data clustering

Answer: A

Q1135. What is the Flash Attention algorithm?

- A. An IO-aware exact attention algorithm reducing memory from quadratic to linear
- B. A data loading technique for prefetching training batches from disk
- C. An approximate attention method that trades accuracy for modest speed
- D. A new non-linear activation function for transformer hidden layers

Answer: A

Q1136. What is the difference between model parallelism and data parallelism?

- A. Data parallelism splits the model layers across devices for training
- B. Model parallelism splits the model across devices; data parallelism replicates model with split data
- C. Model parallelism splits the data while keeping the full model on one device
- D. They are completely identical distributed training strategies with no differences

Answer: B

Q1137. What is the concept of equivariance in CNNs?

- A. Output is always identical regardless of any transformation of input
- B. All layers in the network have exactly equal learned weight values
- C. Output transforms in the same way as the input (e.g., shift in input causes shift in output)
- D. The neural network architecture is perfectly symmetric across all layers

Answer: A

Q1138. What is a diffusion model?

- A. A recurrent neural network architecture for sequence modeling
- B. An unsupervised density-based clustering algorithm for data
- C. A generative model that learns to reverse a gradual noising process
- D. A generative adversarial network variant using two competing agents

Answer: C

Q1139. What is the Mixture of Experts (MoE) architecture?

- A. A standard model ensemble using simple majority voting
- B. A data preprocessing method for feature normalization
- C. A single monolithic large network without any routing
- D. A model using a gating network to route inputs to specialized sub-networks

Answer: D

Q1140. What is quantization in deep learning?

- A. Increasing the overall model size by adding more trainable parameters
- B. Applying data augmentation transforms to expand the training set
- C. Reducing weight and activation precision to lower bit representations for efficiency
- D. Adding more hidden layers to increase the depth of the network

Answer: C

Q1141. What is the difference between GPT and BERT architectures?

- A. Both use exactly the same attention mechanism and training objective
- B. GPT uses bidirectional attention; BERT uses causal left-to-right attention
- C. They are completely identical transformer architectures with no differences
- D. GPT uses causal attention for generation; BERT uses bidirectional attention for understanding

Answer: D

Q1142. What is byte-pair encoding (BPE) in tokenization?

- A. A tokenization approach that only works at the word level
- B. A tokenization strategy that works at sentence level
- C. A tokenization method operating at character level only
- D. A subword tokenization algorithm that iteratively merges the most frequent character pairs

Answer: B

Q1143. What is the perplexity metric in language models?

- A. The total elapsed time required for training the model
- B. The number of parameters or total size of the model
- C. A measure of how well a model predicts a sample; lower is better
- D. The total number of words in the vocabulary of the model

Answer: C

Q1144. What is retrieval-augmented generation (RAG)?

- A. Combining retrieval with generation to ground outputs in external knowledge
- B. An unsupervised clustering method for grouping similar documents
- C. A purely generative approach relying entirely on internal knowledge
- D. A purely retrieval approach returning documents without generation

Answer: A

Q1145. What is the problem of hallucination in large language models?

- A. Generating plausible-sounding but factually incorrect or fabricated information
- B. Using too much GPU memory during the inference computation
- C. Producing completely empty outputs with no generated tokens
- D. Generating text output at an excessively slow processing speed

Answer: A

Q1146. What is the concept of positional encoding in Transformers?

- A. Normalizing the positional coordinates of data in feature space
- B. Encoding the physical position of GPUs in a server rack cluster
- C. Adding position info to embeddings since Transformers lack inherent sequence order
- D. Sorting the input data samples by their index in the dataset

Answer: C

Q1147. What is cross-lingual transfer learning?

- A. Translating all text first then training on the translated version
- B. Ignoring language differences and treating all text the same
- C. Training on one language and applying knowledge to another language
- D. Training separate independent models for each individual language

Answer: C

Q1148. What is the difference between extractive and abstractive summarization?

- A. Extractive generates completely new text from the source material
- B. Abstractive only selects existing sentences from the original text
- C. Extractive and abstractive summarization are identical methods
- D. Extractive selects existing sentences; abstractive generates new sentences

Answer: C

Q1149. What is prompt engineering?

- A. Cleaning and preprocessing the raw input training data
- B. Modifying the internal neural network model architecture
- C. Training the language model from scratch on new domain data
- D. Designing effective input prompts to guide LLM behavior without fine-tuning

Answer: D

Q1150. What is the RLHF (Reinforcement Learning from Human Feedback) technique?

- A. Unsupervised pre-training using masked language modeling objectives
- B. Standard supervised fine-tuning using labeled input-output training pairs
- C. Data augmentation using paraphrasing and back-translation methods
- D. Fine-tuning language models using human preference judgments as reward signals

Answer: D

Q1151. What is the Feature Pyramid Network (FPN)?

- A. A single-scale object detector that operates at only one fixed resolution level
- B. A standard image classification network without multi-scale feature fusion
- C. A generative adversarial network for producing synthetic realistic images
- D. A multi-scale feature extractor combining low-res semantic and high-res spatial features

Answer: D

Q1152. What is the difference between one-stage and two-stage object detectors?

- A. One-stage (YOLO) detects directly; two-stage (R-CNN) proposes regions then classifies
- B. One-stage detectors are always significantly slower than two-stage methods
- C. They are completely identical detection approaches with no performance differences
- D. Two-stage detectors do not use region proposals at any processing stage

Answer: A

Q1153. What is panoptic segmentation?

- A. A task that performs only semantic segmentation without instances
- B. A task that performs only instance segmentation without semantics
- C. A task that is limited to standard object detection only
- D. A unified task combining semantic segmentation (stuff) and instance segmentation (things)

Answer: D

Q1154. What is the Vision Transformer (ViT)?

- A. A recurrent neural network architecture designed for image sequences
- B. A generative adversarial network for producing synthetic images
- C. A specialized convolutional neural network variant for image processing
- D. Applying the Transformer architecture directly to image patches for classification

Answer: D

Q1155. What is contrastive learning in computer vision (e.g., SimCLR)?

- A. A two-stage region-proposal based object detection algorithm pipeline
- B. Learning visual representations by contrasting augmented views of the same image against others
- C. A fully supervised classification method that requires complete labeled training data
- D. A generative adversarial network training method for synthesizing realistic images

Answer: B

Q1156. What is the concept of deformable convolutions?

- A. Standard fixed-grid convolutions with no adaptive sampling offsets
- B. Convolutions with learnable offsets that adapt the sampling grid to object shapes
- C. One-by-one pointwise convolutions for changing channel depth
- D. Dilated convolutions that increase the receptive field without offsets

Answer: B

Q1157. What is neural style transfer?

- A. Classifying images into predefined categorical labels with CNNs
- B. Applying the artistic style of one image to the content of another using CNNs
- C. Training a brand new neural network from scratch on random data
- D. Applying data augmentation transforms to expand the training set

Answer: B

Q1158. What is depth estimation in computer vision?

- A. Predicting the distance of each pixel from the camera
- B. Measuring the resolution of the image in pixels
- C. Classifying the entire image into a category
- D. Counting the total number of objects in a scene

Answer: A

Q1159. What is the CLIP model?

- A. A model learning visual concepts from language supervision, connecting images and text
- B. A single-modality object detector that only processes image pixel data
- C. A standard supervised image classifier trained only on labeled image categories
- D. A text-only language model without any visual understanding capabilities

Answer: A

Q1160. What is 3D point cloud processing?

- A. Processing natural language text data from document sources
- B. Processing unstructured 3D spatial data points from sensors like LiDAR
- C. Processing continuous audio waveform signals from microphones
- D. Processing standard two-dimensional image pixel data from cameras

Answer: B

Q1161. What is the Lambda architecture?

- A. A single-server architecture running everything on one machine
- B. An architecture supporting only offline batch processing workloads
- C. An architecture supporting only real-time stream processing
- D. A big data architecture combining batch and real-time stream processing

Answer: D

Q1162. What is Apache Flink?

- A. A standalone relational database for transactional query workloads
- B. A stream processing framework with event-time processing and exactly-once semantics
- C. A batch-only data processor without any stream processing capability
- D. An interactive data visualization and dashboard charting tool

Answer: B

Q1163. What is the difference between horizontal and vertical scaling?

- A. Vertical scaling adds additional machines to the cluster pool
- B. They are completely identical approaches to scaling infrastructure
- C. Horizontal scaling upgrades individual machine hardware resources
- D. Horizontal adds more machines; vertical adds more power to existing machines

Answer: D

Q1164. What is the concept of data lineage?

- A. The specific file format of the data such as CSV or JSON
- B. Tracking data from its origin through all transformations to its current state
- C. The chronological age of the data since it was created
- D. The total size of the data measured in bytes or rows

Answer: B

Q1165. What is eventual consistency in distributed systems?

- A. All replicas converge to the same value given enough time without new updates
- B. Consistency is fundamentally impossible in any distributed system
- C. All replicas are immediately and perfectly consistent at all times
- D. Data is never consistent across any of the distributed replicas

Answer: A

Q1166. What is Apache Arrow?

- A. A cross-language platform for in-memory columnar data with zero-copy reads
- B. A distributed event streaming platform for real-time pipelines
- C. A supervised machine learning library for model training
- D. A standalone relational database for transactional query workloads

Answer: A

Q1167. What is the Kappa architecture?

- A. A simplified Lambda alternative using only stream processing for real-time and batch
- B. An architecture designed only for small-scale single-machine data
- C. The exact same architecture as Lambda with no simplification
- D. An architecture supporting only offline batch processing workloads

Answer: A

Q1168. What is data skew in distributed processing?

- A. Perfectly even and balanced distribution of data across all nodes
- B. Uneven data distribution across partitions causing some nodes to be overloaded
- C. Data records that are missing or have null values throughout
- D. Data records that have been accidentally duplicated in storage

Answer: B

Q1169. What is the role of a distributed consensus algorithm like Raft or Paxos?

- A. Visualizing distributed system metrics on dashboards
- B. Sorting large distributed data records across nodes
- C. Training machine learning models on distributed nodes
- D. Ensuring all nodes agree on shared state despite failures

Answer: D

Q1170. What is columnar storage and why is it efficient for analytics?

- A. A type of secondary database index for accelerating query lookups
- B. Storing data by rows as in traditional relational database table layouts
- C. Storing data by columns not rows, enabling efficient compression and selective reads
- D. A data visualization technique for displaying column chart graphics

Answer: C

Q1171. What is the concept of shadow deployment?

- A. Running a new model in parallel with production without serving its predictions
- B. Deploying the model at night during off-peak hours only
- C. A/B testing between two model versions with live traffic
- D. Canary deployment with a small percentage of live users

Answer: A

Q1172. What is concept drift vs data drift?

- A. Data drift only affects the target output variable distribution
- B. Concept drift only affects the distribution of input features alone
- C. Concept drift changes input-target relationship; data drift changes input distribution
- D. They are completely identical phenomena with no meaningful differences

Answer: C

Q1173. What is model explainability in production?

- A. Storing more raw training data for future retraining sessions
- B. Providing interpretable explanations for why a model makes specific predictions
- C. Making the model inference significantly faster and more efficient
- D. Reducing the overall size and memory footprint of the model

Answer: B

Q1174. What is the purpose of experiment tracking tools like MLflow?

- A. Recording parameters metrics code and artifacts for each experiment
- B. Only visualizing final model results on a dashboard
- C. Only storing the raw training data on a file system
- D. Only deploying models to a production serving endpoint

Answer: A

Q1175. What is the blue-green deployment strategy?

- A. A data visualization technique using a blue-green color palette
- B. Running models on GPUs with blue and green LED indicator lights
- C. A model training strategy using two alternating learning rates
- D. Maintaining two identical production environments and switching traffic between them

Answer: D

Q1176. What is model compression for deployment?

- A. Increasing numerical precision of weights from FP16 to FP64
- B. Adding more hidden layers to increase the depth of network
- C. Reducing model size and complexity while maintaining acceptable performance
- D. Making models significantly larger by adding trainable parameters

Answer: C

Q1177. What is the role of infrastructure as code (IaC) in MLOps?

- A. Collecting and gathering raw data from external data sources
- B. Writing machine learning algorithm implementations and model code
- C. Defining and managing ML infrastructure through code for reproducibility
- D. Training machine learning models on labeled training datasets

Answer: C

Q1178. What is a serving graph in ML deployment?

- A. A computation graph used only during the model training process
- B. An interactive data visualization chart for exploring model outputs
- C. A diagram showing the internal layers and neurons of a network
- D. A DAG of preprocessing inference and postprocessing steps during prediction

Answer: D

Q1179. What is the challenge of training-serving skew?

- A. The training dataset being too small for adequate model learning overall
- B. The model training process converging far too slowly to be practical
- C. Models becoming too large to fit within the available GPU memory capacity
- D. Differences between training and serving environments causing inconsistent predictions

Answer: B

Q1180. What is GitOps for ML?

- A. Using Git as single source of truth for ML infrastructure configs and deployments
- B. Only storing the Python source code files in a Git repository
- C. A gradient-based training method for optimizing neural networks
- D. A type of version control system different from standard Git

Answer: A

Q1181. What is the alignment problem in AI?

- A. Ensuring advanced AI goals and behaviors align with human values and intentions
- B. Reducing the computational costs of training large-scale AI models
- C. Making AI systems process data at significantly faster inference speeds
- D. Training machine learning models to achieve the highest accuracy metrics

Answer: A

Q1182. What is membership inference attack?

- A. Deleting training data from the model's storage permanently
- B. Stealing the complete set of model weights from the server
- C. Determining whether a specific data point was used in training a model
- D. Changing the model's predictions by modifying its architecture

Answer: C

Q1183. What is the concept of Fairness through Awareness?

- A. Completely ignoring all protected attributes during model prediction
- B. Making random predictions without considering any input features
- C. Using the exact same model configuration for every single individual
- D. Treating similar individuals similarly using a task-specific similarity metric

Answer: D

Q1184. What is model inversion attack?

- A. Improving the model's overall prediction accuracy
- B. Reconstructing training data features from model outputs
- C. Making the model inference significantly faster
- D. Augmenting the training data with synthetic samples

Answer: B

Q1185. What is the difference between equality and equity in AI fairness?

- A. Equity ignores all differences between individuals and groups
- B. Equality is always strictly the better approach for fairness goals
- C. They are completely identical concepts with no meaningful differences
- D. Equality gives the same treatment; equity adjusts treatment for fair outcomes

Answer: D

Q1186. What is the EU AI Act?

- A. A supervised machine learning algorithm for classification tasks
- B. A structured data format for storing model configurations
- C. A risk-based regulatory framework for AI in the European Union
- D. A general-purpose compiled programming language specification

Answer: C

Q1187. What is algorithmic recourse?

- A. Providing individuals with actionable steps to change an unfavorable AI decision
- B. Completely removing the algorithm from production deployment
- C. Using a completely different dataset for model retraining
- D. Making the algorithm run faster through hardware optimization

Answer: A

Q1188. What is the difference between disparate treatment and disparate impact?

- A. Disparate treatment is the unintentional adverse effect on specific groups
- B. Disparate treatment uses protected attributes intentionally; disparate impact causes unequal effects
- C. Disparate impact is the intentional use of protected attributes in a decision model
- D. They are completely identical legal concepts with no meaningful differences

Answer: C

Q1189. What is homomorphic encryption and how does it relate to AI privacy?

- A. A compression algorithm for reducing data file sizes
- B. A process of permanently deleting data from storage
- C. Encryption allowing computation on encrypted data without decrypting it
- D. A standard symmetric encryption method for securing files

Answer: C

Q1190. What is the concept of AI safety?

- A. Making AI systems process data at significantly faster speeds for better user experience
- B. Ensuring AI systems behave as intended without causing unintended harm as capabilities grow
- C. Marketing and promoting AI products and services to potential commercial customers
- D. Reducing the overall financial costs of training and deploying large AI model systems

Answer: B

Q1191. A model achieves 99% training accuracy but only 55% test accuracy. Which strategy best addresses this?

- A. Increase model complexity by adding many more parameters and layers
- B. Remove all validation data to maximize available training data volume
- C. Increase the learning rate to help the model converge much faster
- D. Apply regularization techniques and collect more diverse training data

Answer: D

Q1192. When labeled data is scarce but unlabeled data is abundant, which learning paradigm is most suitable?

- A. Rule-based expert system with manually crafted knowledge base rules
- B. Pure unsupervised learning ignoring all available labeled data points
- C. Semi-supervised learning combining labeled and unlabeled data sources
- D. Fully supervised learning with extensive data augmentation pipelines

Answer: C

Q1193. Which challenge arises when deploying ML models trained on historical data that no longer reflects reality?

- A. Gradient vanishing causing training to stall at suboptimal solutions
- B. Feature explosion resulting from too many correlated input variables
- C. Memory overflow due to the increasing size of the stored model files
- D. Data drift causing degraded model performance over time in production

Answer: D

Q1194. A startup wants to predict rare diseases from patient records. What is the primary ML challenge?

- A. Class imbalance where positive cases are extremely underrepresented
- B. Feature extraction is impossible from electronic health record data
- C. The data will be too structured for machine learning algorithms used
- D. The model will require too many GPU hours for simple tabular data

Answer: A

Q1195. Why is the No Free Lunch theorem significant in ML algorithm selection?

- A. It states no single algorithm performs best across every possible task
- B. It proves deep learning always outperforms all traditional ML methods
- C. It shows unsupervised learning is superior to supervised approaches
- D. It guarantees ensemble methods will always improve baseline accuracy

Answer: A

Q1196. For autonomous vehicles, which combination of AI subfields is most critical?

- A. Recommendation systems combined with collaborative filtering engines
- B. Computer vision with reinforcement learning and sensor fusion systems
- C. Sentiment analysis combined with topic modeling for understanding
- D. Natural language processing combined with text summarization methods

Answer: B

Q1197. What is the key difference between model-based and instance-based learning?

- A. Model-based is always faster at inference while instance-based is always faster
- B. Model-based builds explicit parameter models while instance-based memorizes examples
- C. Model-based requires no training while instance-based needs extensive training time
- D. Model-based only works with images while instance-based only handles text data

Answer: B

Q1198. Adding more training data yields no accuracy improvement. What does this indicate about the model?

- A. The computing hardware has reached its maximum processing speed limits
- B. The model has reached its capacity and needs a more complex architecture
- C. The learning rate is set too high causing the optimizer to diverge rapidly
- D. The dataset contains no useful patterns for any learning algorithm to find

Answer: B

Q1199. In AI ethics, what does algorithmic accountability primarily require from organizations?

- A. Explaining and taking responsibility for AI-driven decisions outcomes
- B. Publishing all source code as open source for public use and changes
- C. Restricting AI development to only government-approved research labs
- D. Mandating that all AI models achieve above ninety-five percent rate

Answer: A

Q1200. Which approach best addresses catastrophic forgetting in continual learning systems?

- A. Removing all previous training data before introducing new examples
- B. Reducing the model size to prevent storing too much old information
- C. Using elastic weight consolidation to protect important parameters
- D. Retraining the entire model from scratch every time new data arrives

Answer: C

Q1201. Why is the Hessian matrix important in second-order optimization methods?

- A. It captures second-order derivative information to estimate curvature
- B. It stores the training data samples needed for batch processing
- C. It determines the initial random weight values for network layers
- D. It selects the appropriate activation function for each neuron

Answer: A

Q1202. When KL divergence between two distributions is zero, what does this indicate?

- A. One distribution has zero variance while the other has infinite range
- B. The two distributions have no overlapping support in their domains
- C. Both distributions are perfectly negatively correlated with each other
- D. The two probability distributions are completely identical everywhere

Answer: D

Q1203. Why can saddle points be problematic for optimization in high-dimensional spaces?

- A. Saddle points make it impossible to compute any derivative at that point
- B. Saddle points always cause the loss function value to increase unboundedly
- C. Gradients near saddle points approach zero causing very slow convergence
- D. Saddle points only occur in one-dimensional optimization never in higher

Answer: C

Q1204. What is the mathematical justification for using cross-entropy loss in classification?

- A. It guarantees convergence in exactly one epoch of training every time
- B. It derives from maximum likelihood estimation under a Bernoulli model
- C. It eliminates the need for any regularization in the model training
- D. It is computationally simpler than all other available loss functions used

Answer: B

Q1205. Why does the curse of dimensionality degrade distance-based algorithms?

- A. Higher dimensions always require more memory than any hardware provides
- B. Higher dimensions guarantee all data points will be perfectly separable
- C. In high dimensions the class count always exceeds the sample count total
- D. In high dimensions all pairwise distances converge making neighbors meaningless

Answer: D

Q1206. What role does the Jacobian matrix play in training multi-output neural networks?

- A. It contains all first-order partial derivatives of outputs versus inputs
- B. It determines the optimal batch size for stochastic gradient descent
- C. It stores bias terms for each layer of the neural network model used
- D. It specifies architecture including number of layers and their widths

Answer: A

Q1207. In Bayesian inference, how does the prior influence the posterior when data is very limited?

- A. The prior has minimal influence and the posterior matches the likelihood
- B. The prior is completely ignored by all Bayesian inference methods used
- C. The prior dominates the posterior since limited data provides weak signal
- D. The prior and posterior are always identical regardless of observed data

Answer: C

Q1208. What mathematical property makes softmax suitable for multi-class classification output?

- A. It compresses all inputs to a fixed range of negative one to one
- B. It outputs values that always sum to exactly one across all classes
- C. It always produces integer values representing discrete class indices
- D. It ensures every output is either exactly zero or exactly one always

Answer: B

Q1209. Why is the Fisher Information Matrix used in natural gradient descent?

- A. It captures statistical manifold curvature for more efficient update steps
- B. It eliminates the requirement for computing any first-order derivatives
- C. It replaces the need for any loss function during model training completely
- D. It guarantees global convergence regardless of the model architecture used

Answer: A

Q1210. What is the significance of positive semi-definiteness in covariance matrices?

- A. It restricts the matrix to having exactly two distinct eigenvalue values
- B. It ensures the matrix can only contain integer values in all its entries
- C. It means the matrix inverse always exists and is unique for every case
- D. It guarantees all eigenvalues are non-negative representing valid variances

Answer: D

Q1211. For a custom scikit-learn transformer, which methods must be implemented for pipeline compatibility?

- A. Only the predict method is required since transformers always make predictions
- B. The score and evaluate methods are the only ones required for all pipelines
- C. Only the transform method needs to be implemented for pipeline compatibility
- D. The fit and transform methods must both be implemented at minimum for use

Answer: D

Q1212. What is the advantage of multiprocessing over threading for CPU-bound ML tasks in Python?

- A. Multiprocessing creates separate processes bypassing GIL for true parallelism
- B. Threading bypasses the GIL while multiprocessing is limited by the GIL lock
- C. Threading uses less memory because it creates separate isolated memory spaces
- D. Multiprocessing is always slower but provides much better error handling now

Answer: A

Q1213. Why should you use np.allclose() instead of == when comparing floating-point arrays?

- A. The allclose function is significantly faster than equality operator for arrays
- B. The equality operator cannot be applied to any NumPy arrays at all ever
- C. Floating-point arithmetic introduces rounding errors making exact comparison fail
- D. The equality operator only works for integer arrays and raises errors on floats

Answer: C

Q1214. Why is it critical to apply fit() only on training data for preprocessing in production ML?

- A. Fitting on test data would make preprocessing steps run much more slowly
- B. Fitting on all data causes data leakage by incorporating test set statistics
- C. Fitting on training data alone ensures the model always achieves higher loss
- D. Fitting on test data would automatically delete the test set from memory

Answer: B

Q1215. What problem does functools.lru_cache solve in feature engineering pipelines?

- A. It logs all function calls to a database for auditing and tracking uses
- B. It converts all function outputs to lower precision to save system memory
- C. It automatically distributes computation across multiple GPU device units
- D. It caches expensive function results to avoid redundant recomputation work

Answer: D

Q1216. When handling very large datasets in pandas, what approach reduces memory most effectively?

- A. Specifying optimal dtypes and reading data in chunks with iteration
- B. Removing all column names to reduce the metadata overhead in RAM
- C. Sorting the DataFrame by index before performing any analysis steps
- D. Converting all columns to string type to ensure uniform representation

Answer: A

Q1217. What is the purpose of `__getitem__` and `__len__` in a PyTorch custom Dataset class?

- A. They define how to save and load model checkpoints to and from disk
- B. They specify loss function and optimizer used during backpropagation
- C. They enable indexed data access and size reporting for the DataLoader
- D. They control the learning rate schedule during model training loop runs

Answer: C

Q1218. Why is `np.einsum()` considered more flexible than standard NumPy operations for tensors?

- A. It runs exclusively on GPU hardware while standard operations use CPU only
- B. It expresses complex multi-dimensional operations in compact Einstein notation
- C. It automatically converts all inputs to higher precision floating point types
- D. It only works with two-dimensional arrays making the interface much simpler

Answer: B

Q1219. What design pattern does scikit-learn's Pipeline class implement and why is it beneficial?

- A. The observer pattern notifying all components when training data changes
- B. The composite pattern chaining steps to prevent data leakage in CV folds
- C. The singleton pattern ensuring only one model instance exists in memory
- D. The factory pattern creating new model instances based on input data type

Answer: B

Q1220. How does `contextlib.contextmanager` help manage resources in data processing?

- A. It provides automatic type checking for all function arguments runtime
- B. It creates permanent global variables accessible from any module at all
- C. It enables resource management using generator-based context managers
- D. It automatically converts synchronous code into asynchronous coroutines

Answer: C

Q1221. A dataset has 40% missing values in a critical feature. Which imputation strategy is most appropriate?

- A. Ignore the feature entirely and remove it from the training dataset now
- B. Simply delete all rows with missing values to ensure data quality overall
- C. Replace all missing values with zero since it is the simplest approach
- D. Use a predictive model like KNN or regression to impute the missing values

Answer: D

Q1222. Why can applying preprocessing transformations before splitting data cause data leakage?

- A. Preprocessing transformations always modify the target variable directly
- B. Statistics from test data influence the transformation applied to training data
- C. Transformations always remove the most important features from the dataset
- D. Preprocessing makes the model run slower during the inference stage overall

Answer: B

Q1223. When dealing with high-cardinality categorical features, why might target encoding outperform one-hot?

- A. One-hot encoding creates massive sparse matrices while target encoding is compact
- B. Target encoding always produces more accurate results regardless of the data
- C. Target encoding removes the need for any validation set during training time
- D. One-hot encoding automatically introduces data leakage into training process

Answer: A

Q1224. In time-series data, why must preprocessing respect temporal order during train-test splitting?

- A. Temporal order only matters for visualization and does not affect model results
- B. Time-series data must always be converted to images before model training steps
- C. Random splitting can leak future information into training causing unrealistic scores
- D. Time-series data always has fewer missing values than cross-sectional datasets

Answer: C

Q1225. A feature has extreme outliers that cannot be removed because they represent valid rare events. What approach is best?

- A. Use mean imputation to replace all outlier values with the column average
- B. Apply min-max scaling which will properly handle extreme outlier values well
- C. Apply robust scaling using median and IQR which is resistant to outlier values
- D. Delete the entire feature since it contains any outlier values in the data

Answer: C

Q1226. Why is the Yeo-Johnson transformation preferred over Box-Cox for some datasets?

- A. Yeo-Johnson is always faster computationally than Box-Cox transformation
- B. Yeo-Johnson only works with categorical features unlike Box-Cox method
- C. Yeo-Johnson can handle both positive and negative values unlike Box-Cox
- D. Yeo-Johnson produces exactly normal distributions while Box-Cox does not

Answer: C

Q1227. When preprocessing text data for NLP, why might subword tokenization outperform word-level tokenization?

- A. Subword tokenization removes the need for any stopword removal preprocessing
- B. It handles out-of-vocabulary words by breaking them into known subword units
- C. It converts all text into numerical features without any vocabulary mapping
- D. Subword tokenization always produces shorter sequences than word tokenization

Answer: B

Q1228. A dataset has features measured in different units with varying scales and distributions. Which preprocessing combination is most robust?

- A. Apply no preprocessing since modern algorithms handle raw data effectively
- B. Apply only min-max scaling to all features uniformly without further analysis
- C. Use power transformation followed by standardization for each feature column
- D. Use one-hot encoding on all features regardless of their data type values

Answer: C

Q1229. Why should you be cautious when using mean imputation for features with non-normal distributions?

- A. It artificially reduces variance and can distort the true distribution shape
- B. Mean imputation only works with categorical features not numerical ones
- C. Mean imputation always increases the standard deviation of the feature values
- D. It automatically removes all outliers before computing the replacement value

Answer: A

Q1230. In a pipeline processing streaming data, why is incremental preprocessing preferred over batch preprocessing?

- A. It automatically handles all missing values without any configuration needed
- B. It processes data point by point without needing the entire dataset in memory
- C. Incremental preprocessing eliminates the need for any feature engineering step
- D. Incremental preprocessing always produces more accurate transformations overall

Answer: B

Q1231. A dataset shows a Pearson correlation of 0.02 between two features, but the scatter plot reveals a clear U-shaped pattern. What explains this?

- A. The data has too many missing values to compute correlation correctly
- B. The scatter plot is incorrectly rendered due to a visualization error
- C. A correlation of 0.02 actually indicates a very strong relationship here
- D. Pearson only captures linear relationships and misses non-linear patterns

Answer: D

Q1232. When exploring a dataset with 500 features, which dimensionality reduction technique is most useful for initial EDA visualization?

- A. Removing all features except the first two columns from the original data
- B. Manually selecting 2 features based on column name alphabetical ordering
- C. Using t-SNE or UMAP to project high-dimensional data into 2D for plotting
- D. Creating 500 individual histograms and reviewing each one sequentially

Answer: C

Q1233. During EDA you discover that a feature has a bimodal distribution. What might this suggest about the underlying data?

- A. The data may contain two distinct subpopulations mixed into one feature
- B. Bimodal distributions always indicate that the feature is perfectly normal
- C. The feature should be removed because bimodal distributions are always errors
- D. The feature has too many outliers and needs to be replaced with its mean

Answer: A

Q1234. Why might the mean be a misleading measure of central tendency for income data?

- A. Income data always follows a perfectly normal distribution in all nations
- B. The mean cannot be computed for continuous variables like income amounts
- C. Income data is always categorical so the mean has no mathematical meaning
- D. Income data is typically right-skewed so the mean is pulled by high earners

Answer: D

Q1235. How does Simpson's Paradox affect conclusions drawn during exploratory data analysis?

- A. It only occurs in datasets with fewer than one hundred total observations
- B. A trend in subgroups can reverse when groups are combined into aggregate
- C. It guarantees that adding more features will always improve model accuracy
- D. It makes all statistical tests invalid regardless of sample size used here

Answer: B

Q1236. A feature shows near-zero variance in the training set. Why should this be investigated before modeling?

- A. They carry almost no discriminative information and may cause instability
- B. Near-zero variance features provide the strongest predictive signal always
- C. Near-zero variance means the feature is perfectly normally distributed
- D. They always indicate data collection errors that must be manually fixed

Answer: A

Q1237. You observe that two features have a Spearman correlation of 0.95. What preprocessing step might you consider?

- A. Multiply the two features together to create a more powerful interaction term
- B. Remove one of the redundant features or use PCA to reduce the collinearity
- C. Convert both features to categorical variables to eliminate the correlation
- D. Add both features to the model since high correlation always helps performance

Answer: B

Q1238. During EDA of a classification dataset, you notice heavy class overlap in feature space. What does this suggest?

- A. The classification task is inherently difficult and may need complex models
- B. The dataset is perfect for a simple linear classifier without any tuning
- C. The features should all be removed and replaced with randomly generated ones
- D. Class overlap always means the labels are incorrect and need relabeling

Answer: A

Q1239. Why is it important to analyze the relationship between features and the target variable separately for different subgroups?

- A. Analyzing subgroups always produces the same results as aggregate analysis does
- B. Different subgroups may exhibit different patterns that aggregate analysis hides
- C. Subgroup analysis is never useful and only wastes computational resources
- D. Subgroup analysis is only applicable to time-series data and no other types

Answer: B

Q1240. A dataset contains timestamps spanning five years. Which EDA technique best reveals seasonal patterns and trends?

- A. Creating a single histogram of all values regardless of their timestamps
- B. Sorting the data alphabetically by the timestamp column values only
- C. Computing a single mean value across all five years of data points
- D. Time-series decomposition into trend, seasonal, and residual components

Answer: D

Q1241. A linear SVM fails to classify a dataset with circular decision boundaries. Which approach best resolves this?

- A. Switching to an RBF kernel to capture the non-linear circular boundary
- B. Increasing the regularization parameter C to a very large positive value
- C. Reducing the training set size to remove noise from the boundary area
- D. Adding more linear features computed from the existing feature columns

Answer: A

Q1242. Why does Naive Bayes often perform surprisingly well despite its strong independence assumption?

- A. The independence assumption is actually true for nearly all real-world datasets
- B. Classification only needs correct ranking of class probabilities not exact values
- C. Naive Bayes secretly uses feature correlations even though it assumes independence
- D. The algorithm automatically corrects for violated assumptions during training

Answer: B

Q1243. In gradient boosting, how does each subsequent tree improve the overall model?

- A. Each tree duplicates the first tree but with different random weight values
- B. Each tree trains on entirely new data that was not used by any prior tree
- C. Each tree is trained on a random subset of features ignoring previous results
- D. Each tree fits the residual errors of the combined predictions of prior trees

Answer: D

Q1244. When should you prefer Ridge regression over Lasso regression for regularization?

- A. When you need exactly half the features to have zero coefficient values only
- B. When you want to perform automatic feature selection by zeroing coefficients
- C. When most features are relevant and you want to shrink all coefficients evenly
- D. When the dataset has categorical features that need one-hot encoding first

Answer: C

Q1245. A decision tree achieves perfect training accuracy but poor test accuracy. Which pruning strategy is most effective?

- A. Grow the tree to maximum depth then add more features to improve test scores
- B. Increase the minimum samples per leaf to the total training set size value
- C. Remove all leaf nodes and keep only the root node for maximum simplicity
- D. Apply cost-complexity pruning to find the subtree with best cross-validated score

Answer: D

Q1246. Why is the choice of distance metric critical for KNN performance on high-dimensional data?

- A. In high dimensions most distance metrics lose discriminative power between points
- B. High-dimensional data always has the same distances regardless of metric used
- C. KNN only works with Euclidean distance so the choice is never actually needed
- D. Distance metrics have no effect on KNN in any number of dimensions at all

Answer: A

Q1247. How does Platt scaling calibrate the output probabilities of an SVM classifier?

- A. It fits a logistic regression on the SVM decision values to produce probabilities
- B. It trains a separate neural network to convert SVM outputs into probabilities
- C. It replaces the SVM kernel with a probability density estimation function
- D. It applies min-max normalization directly to the raw SVM decision values

Answer: A

Q1248. In multi-class classification, what advantage does one-vs-rest have over one-vs-one strategy?

- A. Both strategies always produce identical results regardless of the dataset
- B. One-vs-rest trains fewer classifiers making it computationally more efficient
- C. One-vs-one trains fewer classifiers making it computationally more efficient
- D. One-vs-rest always achieves higher accuracy than one-vs-one on all datasets

Answer: B

Q1249. Why might a well-tuned logistic regression outperform a deep neural network for tabular data?

- A. Logistic regression always converges to the global optimum unlike neural networks
- B. Logistic regression can learn non-linear patterns better than neural networks
- C. Deep neural networks cannot process numerical features in tabular format data
- D. Deep networks need massive data and tabular datasets are often small with noise

Answer: D

Q1250. What is the effect of class imbalance on the decision boundary of a standard SVM?

- A. Class imbalance has no effect on SVM decision boundaries in any scenario ever
- B. The boundary shifts toward the minority class increasing its misclassification rate
- C. The boundary always shifts toward the majority class regardless of data layout
- D. SVM automatically handles class imbalance without any parameter adjustments

Answer: B

Q1251. In a gradient boosting ensemble, how does early stopping prevent overfitting?

- A. It reduces the learning rate to exactly zero after a fixed number of trees
- B. It removes the first few trees which are assumed to have the highest error
- C. It monitors validation loss and stops adding trees when performance degrades
- D. It limits the maximum number of features each tree can use at any split

Answer: C

Q1252. Why does XGBoost use regularization terms in its objective function unlike traditional GBDT?

- A. Regularization slows training to allow more careful optimization at each step
- B. It penalizes tree complexity reducing overfitting and improving generalization
- C. XGBoost regularization only affects the first tree and has no effect on others
- D. Regularization in XGBoost replaces the need for any learning rate parameter

Answer: B

Q1253. When would you prefer LightGBM's leaf-wise growth over XGBoost's level-wise growth?

- A. When all features are categorical and no numerical features exist in data
- B. When the dataset is very small and overfitting is a major concern for accuracy
- C. When you need the simplest possible model with minimal hyperparameters to set
- D. When training speed is critical and the dataset is large with many features

Answer: D

Q1254. How does the DART algorithm modify standard gradient boosting to reduce overfitting?

- A. It randomly drops previously built trees during each boosting iteration step
- B. It uses only the first and last trees and ignores all trees built in between
- C. It doubles the learning rate at each iteration to converge more quickly here
- D. It trains all trees simultaneously rather than sequentially to save total time

Answer: A

Q1255. In a production system, why might model ensembles be problematic despite higher accuracy?

- A. Ensembles are impossible to train on any modern GPU or cloud computing platform
- B. They increase inference latency, memory usage, and deployment complexity significantly
- C. Ensembles always produce less accurate predictions than single model systems
- D. They require manual human review for every single prediction made by the system

Answer: B

Q1256. What is the theoretical basis for why bagging reduces variance in predictions?

- A. The theoretical basis is unknown and bagging works only through empirical results
- B. Bagging increases variance because it uses more models than a single learner
- C. Bagging reduces bias not variance since it trains on different data subsets each
- D. Averaging independent estimates reduces variance by a factor of the ensemble size

Answer: D

Q1257. How does CatBoost handle categorical features differently from XGBoost and LightGBM?

- A. CatBoost requires all features to be one-hot encoded before training can begin
- B. CatBoost converts all categorical features into random numerical values first
- C. CatBoost uses ordered target statistics to encode categoricals avoiding leakage
- D. CatBoost cannot process any categorical features and only handles numerical ones

Answer: C

Q1258. In a stacking ensemble, why should base model predictions be generated using cross-validation?

- A. It prevents data leakage since base models should not predict on their training data
- B. It ensures that all base models have identical performance before stacking them
- C. Cross-validation makes base models train faster than using the full training set
- D. Cross-validation automatically selects the best base models for the final stack

Answer: A

Q1259. What is the relationship between ensemble size and the law of diminishing returns?

- A. Smaller ensembles always outperform larger ensembles due to reduced complexity
- B. Adding more models always linearly improves accuracy without any upper bound limit
- C. After a point additional models contribute marginal gains while increasing costs
- D. Ensemble size has no effect on performance and only affects training speed value

Answer: C

Q1260. Why is Bayesian optimization particularly effective for tuning ensemble hyperparameters?

- A. It only works with ensemble methods and cannot tune single model parameters
- B. It models the objective function to intelligently select promising configurations
- C. Bayesian optimization always finds the global optimum in exactly one iteration
- D. Bayesian optimization randomly samples hyperparameters like grid search does

Answer: B

Q1261. Why might K-Means converge to a suboptimal solution and how is this typically addressed?

- A. Suboptimal convergence only occurs when the dataset has fewer than ten samples
- B. K-Means always finds the globally optimal solution regardless of initialization
- C. It can converge to local optima so K-Means++ initialization is used for better starts
- D. K-Means is guaranteed to diverge without specialized GPU hardware for computation

Answer: C

Q1262. How does spectral clustering handle non-convex cluster shapes that K-Means cannot?

- A. It constructs a similarity graph and uses eigenvectors of its Laplacian for clustering
- B. It converts all non-convex clusters into convex ones before applying K-Means again
- C. It simply increases the number of clusters K until the shapes are captured correctly
- D. It uses the same distance metric as K-Means but with more iterations for accuracy

Answer: A

Q1263. What is the key advantage of UMAP over t-SNE for dimensionality reduction?

- A. UMAP is a linear method while t-SNE is non-linear which makes UMAP faster here
- B. UMAP requires labeled data while t-SNE works in a fully unsupervised manner only
- C. UMAP better preserves global structure and scales to much larger datasets efficiently
- D. UMAP always produces better visualizations than t-SNE on every possible dataset

Answer: C

Q1264. In topic modeling, how does Latent Dirichlet Allocation discover topics from documents?

- A. It models documents as mixtures of topics and topics as distributions over words
- B. It uses supervised labels to assign documents to predefined topic categories only
- C. It requires human experts to manually define all topics before running analysis
- D. It clusters documents using K-Means on word frequency vectors directly always

Answer: A

Q1265. Why is the choice of linkage criterion important in agglomerative hierarchical clustering?

- A. All linkage criteria produce identical dendrograms regardless of the input dataset
- B. Linkage criterion has no effect on the final clustering result in any case ever
- C. Linkage criterion only affects the visualization of dendrogram and not the clusters
- D. Different linkage criteria produce different cluster shapes and handle noise differently

Answer: D

Q1266. How does the Variational Autoencoder differ from a standard autoencoder for unsupervised learning?

- A. Both models are identical in architecture and only differ in their training data used
- B. VAE produces only deterministic encodings while standard autoencoders are probabilistic
- C. VAE learns a probabilistic latent space enabling meaningful data generation sampling
- D. Standard autoencoders can generate new data while VAEs can only compress existing data

Answer: C

Q1267. When evaluating clustering without ground truth labels, why is the Davies-Bouldin Index useful?

- A. It can only evaluate K-Means clustering and does not work with any other algorithms
- B. It always produces a score between zero and one like the accuracy metric in classification
- C. It measures the ratio of within-cluster to between-cluster distances with lower being better
- D. It requires ground truth labels to compute making it a supervised evaluation metric

Answer: C

Q1268. What is the fundamental limitation of using reconstruction error alone to evaluate autoencoders?

- A. Low reconstruction error may indicate memorization rather than useful representation learning
- B. Reconstruction error can only be computed for image data and not for tabular data
- C. Reconstruction error is too computationally expensive to compute on modern hardware
- D. Reconstruction error always perfectly reflects the quality of the learned latent space

Answer: A

Q1269. How does contrastive learning create useful representations without labels?

- A. It uses K-Means clustering as an intermediate step before training the final model
- B. It requires ground truth cluster assignments to define the contrastive loss function
- C. It trains models using manually annotated labels for every single data sample used
- D. It learns to pull augmented views of same samples together and push different apart

Answer: D

Q1270. Why is the Information Criterion approach preferred over the elbow method for determining optimal K?

- A. Information criteria require no data and determine K from theoretical analysis alone here
- B. The elbow method and information criteria always agree on the same optimal K value
- C. The elbow method is always more accurate than any information criterion approach used
- D. Information criteria like BIC provide a principled statistical framework for model selection

Answer: D

Q1271. In medical diagnosis, why is recall often prioritized over precision?

- A. Recall is always higher than precision so it is the easier metric to optimize for
- B. Missing a positive diagnosis is more dangerous than a false alarm in many cases
- C. Medical models always achieve perfect precision making recall the only variable metric
- D. Precision is only relevant for non-medical applications and cannot be used in health

Answer: B

Q1272. Why can nested cross-validation be necessary for unbiased model selection and evaluation?

- A. Nested cross-validation is only needed for datasets with fewer than fifty observations
- B. Hyperparameter tuning on the same CV folds used for evaluation causes optimistic bias
- C. Standard cross-validation already provides completely unbiased model selection results
- D. It is a deprecated technique that has been replaced by simple holdout validation now

Answer: B

Q1273. A model achieves 0.95 AUC but poor performance at the deployed threshold. What is the likely issue?

- A. The model was trained on too much data causing it to generalize too well overall
- B. The model is well-calibrated but the AUC metric was computed with an error here
- C. AUC summarizes all thresholds but the operational threshold may fall in a poor region
- D. An AUC of 0.95 guarantees excellent performance at every possible threshold value

Answer: C

Q1274. When comparing models, why is statistical significance testing important beyond comparing mean metrics?

- A. Statistical testing is only applicable to models trained on datasets exceeding millions
- B. Performance differences may be due to random variation rather than true model superiority
- C. Mean metrics always provide sufficient information to select the best model reliably
- D. It is never needed because cross-validation already accounts for all statistical variance

Answer: B

Q1275. What is the Matthews Correlation Coefficient and why is it preferred for imbalanced datasets?

- A. It is identical to accuracy and provides no additional information for imbalanced data
- B. It only considers true positives and ignores all other confusion matrix cell values
- C. It uses all confusion matrix values providing balanced evaluation even with skewed classes
- D. It requires balanced classes to compute and returns undefined for imbalanced datasets

Answer: C

Q1276. How does the precision-recall curve complement the ROC curve for rare event detection?

- A. Both curves always provide identical information regardless of the class balance ratio
- B. PR curves better expose model performance on the minority class that ROC may obscure
- C. ROC curves are always superior to PR curves in every evaluation scenario encountered
- D. PR curves can only be computed when classes are perfectly balanced in the dataset

Answer: B

Q1277. Why might a model perform well in cross-validation but poorly in real-world deployment?

- A. Distribution shift between training data and production data degrades model performance
- B. Cross-validation uses test data for training which inflates deployment expectations
- C. Models that perform well in CV are guaranteed to perform well in any deployment setup
- D. Cross-validation always provides perfectly accurate estimates of deployment performance

Answer: A

Q1278. What is the purpose of the DeLong test in comparing ROC curves of two models?

- A. It statistically tests whether two AUC values are significantly different from each other
- B. It measures the calibration quality of both models probability estimates simultaneously
- C. It determines whether one model trains faster than the other on the same dataset
- D. It automatically selects the optimal threshold for deploying both models in production

Answer: A

Q1279. In time-series model evaluation, why is standard K-fold cross-validation inappropriate?

- A. Standard K-fold is actually the recommended approach for all time-series evaluations
- B. Random fold assignment violates temporal ordering and leaks future data into training
- C. Time-series data always has too few samples for any cross-validation method to work
- D. K-fold cross-validation can only be applied to classification not regression problems

Answer: B

Q1280. How does Bayesian evaluation of classifiers improve upon frequentist hypothesis testing?

- A. It provides probability distributions over performance differences enabling richer analysis
- B. Bayesian evaluation only works with neural network classifiers and no other methods
- C. It eliminates the need for any test data by computing performance from prior alone
- D. Bayesian methods always agree with frequentist tests making them completely redundant

Answer: A

Q1281. Why can target encoding cause data leakage and how is it typically mitigated?

- A. It uses target values during encoding so leave-one-out or fold-based encoding prevents leak
- B. Target encoding never causes data leakage so no mitigation is ever required
- C. Data leakage only occurs with one-hot encoding and never with target encoding method
- D. Target encoding automatically prevents leakage through its internal smoothing mechanism

Answer: A

Q1282. When engineering features for a gradient boosted tree model, why is binning continuous features usually unnecessary?

- A. Tree-based models already find optimal split points so binning loses useful granularity
- B. Gradient boosted trees cannot process binned features and require raw continuous input
- C. Binning always improves gradient boosted tree performance without any exceptions
- D. Binning is only useful for neural networks and cannot be applied to any tree models

Answer: A

Q1283. How does autoML automate feature engineering and what are its limitations?

- A. AutoML perfectly replaces all manual feature engineering without any limitations ever
- B. AutoML only works with image data and cannot handle tabular or text feature engineering
- C. It explores transformations automatically but may miss domain-specific meaningful features
- D. It always produces fewer features than manual engineering making it strictly inferior

Answer: C

Q1284. Why might feature engineering differ between linear models and tree-based models?

- A. Tree-based models require normalized features while linear models handle raw values
- B. Both model types benefit from exactly identical feature engineering approaches always
- C. Linear models need non-linear transformations while trees handle non-linearity natively
- D. Feature engineering is only needed for tree-based models and never for linear models

Answer: C

Q1285. What is the risk of using highly engineered features that are specific to the training distribution?

- A. Engineered features are always more robust than raw features in all scenarios tested
- B. The risk is minimal because feature engineering always reduces overfitting in models
- C. They may not generalize to production data if the distribution shifts over time period
- D. Highly engineered features always generalize perfectly to any future data distribution

Answer: C

Q1286. How does mutual information differ from Pearson correlation for feature selection?

- A. Pearson captures non-linear relationships while mutual information is limited to linear
- B. Both methods measure identical relationships and always produce the same feature ranking
- C. Mutual information only works with binary features while Pearson works with all types
- D. Mutual information captures any dependency while Pearson only detects linear relationships

Answer: D

Q1287. In NLP, why might TF-IDF features outperform raw word count features for text classification?

- A. Raw word counts always outperform TF-IDF so this premise is incorrect in all cases
- B. TF-IDF downweights common words and upweights distinctive words improving discrimination
- C. TF-IDF converts words to their synonyms before counting improving vocabulary coverage
- D. TF-IDF always produces fewer features than raw word counts making models train faster

Answer: B

Q1288. What is the purpose of creating lag features in time-series feature engineering?

- A. They capture temporal dependencies by using past values as features for current predictions
- B. Lag features convert time-series data into image format for convolutional neural networks
- C. They eliminate seasonality from time-series data making it stationary for all models
- D. Lag features remove temporal dependencies to make data suitable for standard ML models

Answer: A

Q1289. How does the permutation importance method evaluate feature relevance after model training?

- A. It sorts features alphabetically and assigns importance based on their position order
- B. It shuffles each feature's values and measures how much model performance degrades
- C. It adds random noise to each feature and measures how much training accuracy improves
- D. It removes each feature entirely and retrains the model from scratch to compare results

Answer: B

Q1290. Why should you be cautious about using future information when engineering features for time-series models?

- A. Using future values creates temporal leakage leading to unrealistically optimistic evaluation
- B. Future information is always available at prediction time so there is no risk in using it
- C. Future information always improves time-series model accuracy without any negative effects
- D. Time-series models automatically detect and exclude future information during training

Answer: A

Q1291. Why do residual connections in ResNet help train very deep networks effectively?

- A. They allow gradients to flow directly through skip connections preventing vanishing
- B. Residual connections double the number of parameters in each layer of network
- C. Residual connections remove the need for any activation functions in the network
- D. They automatically reduce the learning rate as the network depth increases here

Answer: A

Q1292. What is the dying ReLU problem and how is Leaky ReLU a solution?

- A. Dying ReLU only affects the output layer while Leaky ReLU only affects hidden layers
- B. ReLU never has this problem and Leaky ReLU was created for a different purpose
- C. Dying ReLU means the function grows too large while Leaky ReLU caps the output
- D. Neurons with negative inputs permanently output zero and Leaky ReLU allows small negatives

Answer: D

Q1293. How does the attention mechanism in neural networks improve sequence modeling?

- A. It forces the model to process all elements equally regardless of their relevance
- B. It removes the need for any training data by using pre-computed attention weights
- C. It limits the model to processing only the first and last elements of sequences
- D. It allows the model to dynamically focus on relevant parts of the input sequence

Answer: D

Q1294. Why is the choice of weight initialization strategy critical for deep network training?

- A. Weight initialization has no effect on training since optimizers always find the same solution
- B. Weight initialization only matters for the first layer and has no effect on deeper layers
- C. Poor initialization can cause vanishing or exploding activations preventing effective learning
- D. All initialization strategies produce identical results after sufficient training epochs pass

Answer: C

Q1295. What is the lottery ticket hypothesis and its implications for neural network pruning?

- A. It suggests that random network architectures always outperform carefully designed ones
- B. It proves that larger networks always outperform smaller networks without any exceptions
- C. Dense networks contain sparse subnetworks that achieve comparable accuracy when trained alone
- D. It states that all neural network weights are equally important and none can be removed

Answer: C

Q1296. How does knowledge distillation compress a large model into a smaller one?

- A. It copies all weights directly from the large model into the smaller model architecture
- B. It trains the small model only on data that the large model classified incorrectly
- C. A smaller student model learns to mimic the soft output probabilities of a larger teacher
- D. It removes random layers from the large model until the desired size is reached

Answer: C

Q1297. Why do neural architecture search methods face a massive computational challenge?

- A. The challenge is purely theoretical and modern hardware makes NAS trivially inexpensive
- B. NAS only searches over learning rates not architectures so the space is very small
- C. The search space of possible architectures is exponentially large requiring many evaluations
- D. Neural architecture search always converges in a single iteration without much computation

Answer: C

Q1298. What role does the temperature parameter play in softmax for knowledge distillation?

- A. Temperature only affects training speed and has no effect on the output distribution
- B. Temperature is fixed at one and cannot be changed in any softmax implementation ever
- C. Lower temperature always produces better distillation results regardless of the task
- D. Higher temperature produces softer probabilities revealing more about learned similarities

Answer: D

Q1299. How does mixed precision training accelerate deep learning while maintaining accuracy?

- A. It halves the training data to reduce computation time at the cost of some accuracy
- B. It skips every other layer during forward pass to achieve faster computation speed
- C. It uses float16 for most operations and float32 for critical ones reducing memory and time
- D. It uses only 8-bit integers for all computations throughout the entire training process

Answer: C

Q1300. What is the neural tangent kernel framework and why is it theoretically important?

- A. It is a specific type of convolutional kernel used only for image recognition tasks
- B. It is a pruning technique for removing unnecessary connections from trained networks
- C. It is an optimizer that replaces gradient descent for all neural network training
- D. It shows infinite-width networks behave like kernel methods enabling theoretical analysis

Answer: D

Q1301. Why do Vision Transformers require large-scale pre-training to outperform CNNs?

- A. Vision Transformers have fewer parameters than CNNs and need more data to compensate
- B. CNNs always outperform Vision Transformers regardless of the training data scale used
- C. Vision Transformers can only process text data and need pre-training to handle images
- D. ViTs lack inductive biases like locality and translation equivariance that CNNs have

Answer: D

Q1302. How does the Wasserstein loss improve GAN training stability compared to the original formulation?

- A. It eliminates the need for a discriminator network in the adversarial training setup
- B. It provides smoother gradients based on earth mover distance even when distributions differ
- C. Wasserstein loss only works with image generation and not with other data modalities
- D. It uses the same loss as the original GAN but with a different optimizer algorithm

Answer: B

Q1303. What is the computational complexity of self-attention and how do efficient variants address it?

- A. Efficient variants make attention slower but produce more accurate attention weights
- B. Self-attention is $O(n)$ and efficient variants increase it to $O(n^2)$ for accuracy
- C. Self-attention is $O(n^2)$ and variants like Linformer use projections for $O(n)$
- D. Self-attention has constant complexity regardless of the sequence length provided

Answer: C

Q1304. How does the diffusion model generate high-quality images compared to GANs?

- A. They require significantly less training data than GANs to produce comparable results
- B. Diffusion models train adversarially like GANs but with more discriminator networks
- C. Diffusion models generate images in a single forward pass without any iterative steps
- D. They learn to reverse a gradual noising process producing diverse high-quality samples

Answer: D

Q1305. What is the significance of the universal approximation theorem for neural networks?

- A. It guarantees that deeper networks always outperform shallower networks on all tasks
- B. A sufficiently wide single-layer network can approximate any continuous function closely
- C. It states that neural networks can only approximate linear functions and nothing else
- D. It proves that neural networks always find the global optimum during training runs

Answer: B

Q1306. How does neural architecture search with weight sharing reduce computational cost?

- A. It trains each candidate architecture from scratch independently on separate hardware
- B. It searches only over learning rates not architectures making the space much smaller
- C. Multiple architectures share weights from a supernet avoiding full retraining each time
- D. Weight sharing eliminates the need for any gradient computation during architecture search

Answer: C

Q1307. Why is training stability a major challenge in GAN optimization?

- A. GAN training is always stable and converges to the optimal solution every single time
- B. The minimax game between generator and discriminator can oscillate without convergence
- C. The generator and discriminator always converge at the exact same rate making it easy
- D. Stability issues only occur with very small datasets of fewer than hundred total samples

Answer: B

Q1308. How does the mixture of experts architecture achieve computational efficiency in large models?

- A. It reduces model size by removing most parameters from each expert subnetwork
- B. It uses a single large expert for all inputs and ignores the routing mechanism
- C. It routes each input to only a subset of experts reducing computation per sample
- D. It trains all expert networks on every input sample to ensure comprehensive coverage

Answer: C

Q1309. What is the role of contrastive loss in self-supervised representation learning?

- A. It replaces all other loss functions and is the only one used in deep learning
- B. It only works with generative models and cannot be used with discriminative ones
- C. It pulls similar sample representations together and pushes dissimilar ones apart
- D. It requires labeled data pairs to compute the loss function during training phase

Answer: C

Q1310. How does the Flash Attention algorithm improve Transformer training efficiency?

- A. It uses tiling and recomputation to reduce memory IO making attention faster on GPU
- B. Flash Attention only works on CPU hardware and provides no benefit on GPU systems
- C. It reduces attention accuracy to save computation time during the training process
- D. It removes the attention mechanism entirely replacing it with simple averaging only

Answer: A

Q1311. Why do large language models sometimes generate factually incorrect but plausible-sounding text?

- A. LLMs verify all facts against a database before generating any output text content
- B. LLMs always generate perfectly accurate text since they are trained on factual data
- C. Factual errors only occur when the model has fewer than one million total parameters
- D. LLMs learn statistical patterns rather than factual knowledge causing hallucination errors

Answer: D

Q1312. How does the Transformer's multi-head attention provide an advantage over single-head attention?

- A. Multi-head attention is identical to single-head attention and provides no real advantage
- B. Multiple heads only work with image data and provide no benefit for text processing
- C. Multiple heads attend to different representation subspaces capturing diverse relationships
- D. Multi-head attention uses fewer parameters than single-head attention in all cases

Answer: C

Q1313. What is the fundamental challenge of evaluating open-ended text generation models?

- A. Open-ended generation models always produce identical outputs for the same input text
- B. Evaluation is trivial because BLEU score perfectly captures all aspects of text quality
- C. There is always exactly one correct output for any given text generation prompt used
- D. Multiple valid outputs exist making automatic metrics insufficient for quality assessment

Answer: D

Q1314. How does reinforcement learning from human feedback improve language model alignment?

- A. RLHF only affects the tokenizer and has no impact on the language model weights
- B. It fine-tunes the model using a reward signal trained on human preference judgments
- C. It replaces the Transformer architecture with a recurrent neural network approach
- D. RLHF removes all pre-training knowledge and trains the model entirely from scratch

Answer: B

Q1315. Why is the position of negation words critical for sentiment analysis accuracy?

- A. Negation words have no effect on sentiment and can be safely removed during preprocessing
- B. Negation reverses sentiment of subsequent words and mishandling it causes incorrect polarity
- C. Negation only appears in formal text and never occurs in informal social media content
- D. All negation words are stop words and are automatically removed by all NLP pipelines used

Answer: B

Q1316. What is the catastrophic forgetting problem when fine-tuning pre-trained language models?

- A. Catastrophic forgetting only occurs in computer vision models and never in NLP models
- B. The model loses general pre-trained knowledge when heavily adapted to a specific new task
- C. Fine-tuning always improves performance on all tasks simultaneously without any tradeoffs
- D. Fine-tuning adds new knowledge without affecting any previously learned representations

Answer: B

Q1317. How do retrieval-augmented generation models reduce hallucination in language models?

- A. They remove the generative component entirely and only return retrieved document text
- B. They retrieve relevant documents and condition generation on factual source material
- C. RAG models eliminate hallucination completely with a one hundred percent success rate
- D. They increase the model size to store more factual information in the network parameters

Answer: B

Q1318. Why is cross-lingual transfer learning effective for low-resource languages?

- A. Multilingual models learn shared representations that transfer knowledge between languages
- B. Cross-lingual transfer degrades performance on all languages including high-resource ones
- C. Each language requires a completely separate model with no knowledge sharing possible
- D. Cross-lingual transfer only works between languages that share the same script system

Answer: A

Q1319. What is the computational challenge of scaling attention to very long document contexts?

- A. Long documents always produce better results so there is no computational challenge here
- B. Long documents can be processed in constant time regardless of their total length overall
- C. Attention memory and compute scale quadratically with sequence length becoming prohibitive
- D. Attention complexity is independent of sequence length and depends only on model depth

Answer: C

Q1320. How does chain-of-thought prompting improve reasoning in large language models?

- A. It breaks complex problems into intermediate steps allowing systematic reasoning output
- B. It reduces the model's accuracy but makes the output more readable for human reviewers
- C. Chain-of-thought only works for arithmetic and cannot improve any other reasoning tasks
- D. It forces the model to generate the answer in a single token without any explanation

Answer: A

Q1321. Why do adversarial examples pose a serious threat to deployed computer vision systems?

- A. Adversarial examples only affect models during training and never at inference time
- B. All modern CNN architectures are completely immune to any adversarial perturbations
- C. Adversarial examples only exist in theory and cannot be generated for real images
- D. Imperceptible perturbations can cause confident misclassification in production models

Answer: D

Q1322. How does contrastive learning for visual representations reduce dependence on labeled data?

- A. It completely replaces supervised fine-tuning and achieves perfect accuracy on all tasks
- B. Contrastive learning only works with text data and cannot be applied to images ever
- C. It requires more labeled data than supervised learning to achieve comparable accuracy
- D. It learns representations by contrasting augmented views of images without any labels

Answer: D

Q1323. What is the domain adaptation challenge when deploying CV models across different environments?

- A. Distribution differences between training and deployment environments degrade performance
- B. Domain adaptation is only relevant for NLP tasks and does not apply to computer vision
- C. All visual environments produce identical image statistics so adaptation is unnecessary
- D. Models trained in one visual domain always generalize perfectly to all other domains

Answer: A

Q1324. How does the panoptic segmentation task unify semantic and instance segmentation?

- A. It assigns every pixel both a semantic class label and an instance identifier value
- B. It replaces both tasks with a simpler bounding box detection approach for efficiency
- C. Panoptic segmentation only works on outdoor scenes and cannot process indoor images
- D. It only performs semantic segmentation and completely ignores instance-level information

Answer: A

Q1325. Why is 3D object detection from point clouds fundamentally different from 2D image detection?

- A. Point clouds are identical to images and can be processed by standard 2D CNN methods
- B. 3D detection is always easier than 2D detection because more information is available
- C. Point clouds are sparse, unordered, and require specialized architectures like PointNet
- D. Point clouds can only represent indoor environments and not outdoor driving scenarios

Answer: C

Q1326. What is the role of the feature extractor backbone in modern object detection frameworks?

- A. It handles data augmentation and preprocessing before the actual detection begins
- B. It generates the final bounding box predictions and class labels for each object found
- C. The backbone is only used during training and is removed completely during inference
- D. It extracts hierarchical visual features that the detection head uses for predictions

Answer: D

Q1327. How does self-supervised pre-training with masked image modeling work for vision models?

- A. It only works with text-image pairs and cannot be applied to images alone at all
- B. It masks random image patches and trains the model to reconstruct the missing content
- C. It requires fully labeled segmentation maps for every image in the training dataset
- D. It trains models to classify images into predefined categories using labeled datasets

Answer: B

Q1328. Why is temporal consistency important in video object detection and how is it maintained?

- A. Temporal consistency is irrelevant because each video frame should be processed independently
- B. Video frames are always identical so temporal consistency is automatically guaranteed here
- C. Flickering predictions between frames indicate instability and tracking methods enforce consistency
- D. Temporal consistency only matters for video compression and has no effect on detection

Answer: C

Q1329. What is the advantage of vision-language models like CLIP for zero-shot image classification?

- A. They require separate fine-tuning for every possible class label before making predictions
- B. Vision-language models can only classify images into the exact categories seen during training
- C. They always perform worse than supervised models regardless of the number of training classes
- D. They match images to text descriptions enabling classification without task-specific training

Answer: D

Q1330. How does neural radiance field technology enable novel view synthesis from images?

- A. It requires thousands of input images from every possible angle of the target scene
- B. It simply stitches existing images together without any three-dimensional understanding
- C. It learns a volumetric scene representation mapping 3D coordinates to color and density
- D. NeRF only works with synthetic computer-generated scenes and not real-world photographs

Answer: C

Q1331. Why is exactly-once processing semantics challenging to achieve in distributed streaming systems?

- A. Network failures and node crashes can cause duplicates or losses requiring complex protocols
- B. Exactly-once processing is trivial and all modern streaming systems implement it easily
- C. Streaming systems never experience failures so exactly-once is guaranteed automatically
- D. Exactly-once only matters for batch processing and is irrelevant for streaming systems

Answer: A

Q1332. How does the Lambda architecture address the needs of both batch and real-time analytics?

- A. Lambda architecture only supports batch processing and has no real-time component
- B. It maintains separate batch and speed layers that serve a combined serving layer
- C. It uses only a single processing layer for both batch and real-time data analysis
- D. It replaces all batch processing with stream processing for unified data handling

Answer: B

Q1333. What is data skew and why does it degrade performance in distributed processing?

- A. Data skew only affects storage and has no impact on query processing performance
- B. Data skew always improves performance by concentrating work on the fastest machines
- C. All distributed systems automatically handle data skew without any intervention
- D. Uneven data distribution causes some nodes to be overloaded while others sit idle

Answer: D

Q1334. How does the Kappa architecture simplify the Lambda architecture?

- A. Kappa uses only a stream processing layer treating everything as streams eliminating batch
- B. Kappa removes the speed layer and relies solely on batch processing for all analytics
- C. Kappa adds a third processing layer on top of Lambda's existing batch and speed layers
- D. Kappa and Lambda architectures are identical and the terms are completely interchangeable

Answer: A

Q1335. Why is shuffle operation one of the most expensive operations in distributed data processing?

- A. Shuffle operations only move data within a single machine which is very fast operation
- B. Shuffles only occur in batch processing and never happen in stream processing pipelines
- C. Shuffles move data across the network between nodes causing significant IO and latency
- D. Shuffle operations are optimized by all frameworks and never cause performance issues

Answer: C

Q1336. How does delta lake improve upon traditional data lake architectures?

- A. It replaces the distributed file system with a single centralized database for storage
- B. Delta lake removes all structure and stores data as completely unformatted text files
- C. It adds ACID transactions and schema enforcement to data lake storage for reliability
- D. Delta lake only works with structured data and cannot store any unstructured data files

Answer: C

Q1337. What are the tradeoffs between using a wide table versus a normalized schema in big data?

- A. There are no tradeoffs as both approaches produce identical query performance results
- B. Wide tables are always superior to normalized schemas for every possible query pattern
- C. Wide tables reduce joins but increase redundancy while normalized schemas minimize redundancy
- D. Normalized schemas always perform better than wide tables in distributed query engines

Answer: C

Q1338. How does approximate query processing enable faster analytics on massive datasets?

- A. It only works for count queries and cannot approximate any other aggregation types
- B. It produces exact results faster by using more efficient algorithms on all the data present
- C. Approximate processing always produces completely wrong results with no useful insights
- D. It trades some accuracy for speed using sampling and sketches on subsets of the data

Answer: D

Q1339. Why is data lineage tracking important in complex big data pipelines?

- A. Data lineage is only useful for documentation and has no practical operational value
- B. It traces data transformations enabling debugging, compliance, and impact analysis
- C. Data lineage tracking always degrades pipeline performance by more than fifty percent
- D. It only applies to batch processing and cannot track streaming data transformations

Answer: B

Q1340. How does federated learning enable ML on distributed data without centralizing it?

- A. It only works with image data and cannot be applied to tabular or text data formats
- B. Models train locally on each node and only share model updates not raw data for privacy
- C. Federated learning always produces worse models than centralized training on all tasks
- D. It requires all data to be collected in a central location before any training can begin

Answer: B

Q1341. How does a feature store ensure consistency between training and serving in production ML?

- A. Feature stores only handle training features and have no role in serving predictions
- B. It provides a single source of truth for feature definitions and computation logic used
- C. Feature stores always cause training-serving skew due to different code paths in them
- D. It stores raw data only and requires separate feature computation for training and serving

Answer: B

Q1342. Why is reproducibility challenging in ML systems and how does MLOps address it?

- A. ML systems are perfectly reproducible by default and require no additional tooling at all
- B. MLOps makes all models deterministic by fixing random seeds and nothing else is needed
- C. Reproducibility is only important in academic research and irrelevant for production ML
- D. Non-determinism in training, data changes, and environment drift require systematic tracking

Answer: D

Q1343. What is the role of a model registry in enterprise MLOps workflows?

- A. Model registries are only useful for small teams and provide no value at enterprise scale
- B. It automatically trains new models without any human oversight or approval workflows
- C. It manages model lifecycle including versioning, staging, approval, and deployment tracking
- D. It only stores model code and has no capability to track model metadata or lineage

Answer: C

Q1344. How does data-centric AI differ from model-centric AI in improving ML system performance?

- A. Data-centric AI only applies to unstructured data and model-centric to structured data
- B. Data-centric and model-centric approaches are identical and interchangeable in all cases
- C. Data-centric focuses on improving data quality while model-centric focuses on architecture
- D. Model-centric AI never requires any training data and relies solely on model architecture

Answer: C

Q1345. What are the challenges of implementing real-time ML inference at scale?

- A. Real-time inference is trivial to implement and never poses any engineering challenges
- B. All models inherently provide sub-millisecond inference without any optimization needed
- C. Latency requirements, throughput scaling, and model optimization create significant complexity
- D. Real-time inference only works with batch predictions and cannot handle individual requests

Answer: C

Q1346. How does GitOps extend traditional DevOps practices for ML workflow management?

- A. It eliminates the need for any CI/CD pipelines by manually deploying all changes to prod
- B. It uses Git as the single source of truth for declarative infrastructure and ML pipeline config
- C. GitOps replaces Git with a custom version control system designed only for ML artifacts
- D. GitOps only works with data preprocessing and cannot manage model training or deployment

Answer: B

Q1347. Why is testing ML systems more complex than testing traditional software?

- A. Testing ML systems requires only standard unit tests identical to traditional software ones
- B. ML systems are simpler to test because they only require checking if the code compiles
- C. ML involves data dependencies, model behavior, and performance thresholds beyond unit tests
- D. ML systems never need testing because model accuracy guarantees correct system behavior

Answer: C

Q1348. What is the purpose of a model serving framework like TensorFlow Serving or Triton?

- A. They replace the need for any monitoring or logging of deployed model performance metrics
- B. They optimize and serve models efficiently handling batching, versioning, and concurrency
- C. They are only used for training models and have no capability for serving predictions
- D. Model serving frameworks only work with a single model and cannot manage multiple models

Answer: B

Q1349. How does ML governance address regulatory requirements for AI systems in production?

- A. Governance eliminates the need for any model monitoring after initial deployment is done
- B. It only applies to models making financial decisions and not to any other domain areas
- C. It provides audit trails, documentation, fairness monitoring, and compliance frameworks
- D. ML governance is only relevant for academic research and has no production implications

Answer: C

Q1350. What is the concept of ML technical debt and how does it accumulate in production systems?

- A. Technical debt only applies to traditional software and is irrelevant for ML systems always
- B. ML technical debt is automatically resolved by retraining the model on newer data only
- C. ML systems never accumulate technical debt because they are self-maintaining by design
- D. Data dependencies, pipeline complexity, and configuration debt compound over time in systems

Answer: D

Q1351. Why is it mathematically impossible to satisfy all fairness criteria simultaneously in most real-world scenarios?

- A. All fairness criteria can be satisfied simultaneously with enough training data available
- B. Different fairness metrics are fundamentally incompatible when base rates differ across groups
- C. Only linear models face this limitation while neural networks can satisfy all criteria easily
- D. Impossibility only applies to theoretical scenarios and never to practical AI applications

Answer: B

Q1352. How does federated learning attempt to address privacy concerns while enabling collaborative ML?

- A. It trains models locally sharing only model updates while keeping raw data decentralized
- B. It requires centralizing all data from all participants before any model training begins
- C. Federated learning provides no privacy benefits compared to centralized training approaches
- D. It encrypts all data using standard AES encryption and sends it to a central server only

Answer: A

Q1353. What is the value alignment problem in the context of advanced AI systems?

- A. Ensuring AI systems pursue goals consistent with complex and sometimes conflicting human values
- B. Value alignment is only relevant for recommendation systems and not for other AI applications
- C. It refers to aligning model weights with optimal values during the training optimization
- D. It guarantees AI systems will always act in humanity's best interest without any intervention

Answer: A

Q1354. How can model cards and datasheets improve responsible AI development practices?

- A. They replace the need for any testing or evaluation of AI models before deployment occurs
- B. Model cards are only useful for marketing purposes and provide no technical or ethical value
- C. They provide structured documentation about model capabilities, limitations, and intended use
- D. They automatically fix all bias issues in AI models without any human intervention needed

Answer: C

Q1355. What is the tension between model accuracy and fairness in high-stakes AI applications?

- A. Higher accuracy always guarantees better fairness outcomes for all groups simultaneously
- B. Optimizing for fairness constraints may reduce accuracy and vice versa requiring tradeoffs
- C. There is never any tension because accurate models are always fair to all groups by nature
- D. Fairness constraints always improve model accuracy making the tension nonexistent overall

Answer: B

Q1356. How do deepfakes pose unique challenges for AI ethics and society?

- A. They enable realistic fabrication of content undermining trust in media and enabling fraud
- B. Deepfakes are easily detectable by all humans and pose no significant societal challenges
- C. Deepfakes are only used for entertainment purposes and have no potential for misuse ever
- D. Deepfake technology can only produce low-quality content that no one would believe is real

Answer: A

Q1357. Why is the concept of meaningful human oversight important for autonomous AI systems?

- A. Meaningful oversight only applies to AI systems in military applications and no other fields
- B. Humans must retain ability to understand, intervene, and override AI decisions when needed
- C. Human oversight should be completely removed to maximize the efficiency of AI systems now
- D. Autonomous AI systems never make errors so human oversight is unnecessary and wasteful

Answer: B

Q1358. What are the environmental concerns associated with training large AI models?

- A. Training large AI models has no environmental impact and uses negligible energy resources
- B. Large model training consumes significant energy and computing resources increasing carbon footprint
- C. Environmental concerns only apply to models trained on renewable energy sources specifically
- D. Training small models produces more emissions than large models due to lower computational efficiency

Answer: B

Q1359. How does the EU AI Act categorize AI systems based on risk levels?

- A. The EU AI Act bans all AI systems regardless of their risk level or application domain
- B. The EU AI Act only applies to military AI systems and exempts all commercial applications
- C. It treats all AI systems identically without any risk-based categorization or differentiation
- D. It classifies systems into unacceptable, high, limited, and minimal risk with corresponding rules

Answer: D

Q1360. What is the challenge of ensuring AI accountability when AI systems make autonomous decisions?

- A. AI systems can be held legally accountable themselves without any human responsibility needed
- B. Accountability is only relevant when AI causes physical harm and not for any other type of impact
- C. AI accountability is straightforward since the developer is always clearly responsible for everything
- D. Complex AI systems make it difficult to attribute responsibility across developers, deployers, and users

Answer: D

Q1361. What is the Moravec paradox in AI research?

- A. Simple computational tasks are harder for AI than complex reasoning
- B. High-level reasoning requires less computation than sensorimotor skills
- C. AI systems excel at tasks humans find easy but struggle with tasks humans find hard
- D. Tasks easy for humans like perception and mobility are extremely hard for AI while abstract reasoning is comparatively easier

Answer: D

Q1362. What is the difference between deductive and inductive reasoning in AI systems?

- A. They are identical approaches
- B. Deductive reasoning goes from general rules to specific conclusions while inductive reasoning infers general rules from specific observations
- C. Deductive reasoning is always better than inductive
- D. Inductive reasoning does not use data

Answer: B

Q1363. What is the computational theory of mind and how does it relate to AI?

- A. Computers are conscious beings
- B. The mind operates like a computational system processing information, providing theoretical basis for AI
- C. All computations are intelligent
- D. Minds cannot be modeled computationally

Answer: B

Q1364. What is the Bitter Lesson identified by Rich Sutton in AI research?

- A. AI research is always unsuccessful
- B. General methods leveraging computation scale better than methods leveraging human domain knowledge
- C. Human knowledge is always superior to machine learning
- D. Smaller models always outperform larger ones

Answer: B

Q1365. What is the concept of emergence in complex AI systems?

- A. AI systems always behave predictably
- B. Complex behaviors and capabilities arising from simpler components that were not explicitly programmed
- C. All AI capabilities must be manually coded
- D. Emergence only occurs in biological systems

Answer: B

Q1366. How does the concept of bounded rationality apply to AI agent design?

- A. AI agents always have unlimited resources
- B. Agents must make satisficing decisions due to limited computational resources and information
- C. Bounded rationality means AI cannot learn
- D. It only applies to human decision-making

Answer: B

Q1367. What is the difference between symbolic AI and connectionist AI?

- A. They are the same approach
- B. Symbolic AI uses explicit rules and logic while connectionist AI uses neural networks to learn representations
- C. Symbolic AI is always superior
- D. Connectionist AI does not use mathematics

Answer: B

Q1368. What is the scaling hypothesis in modern AI research?

- A. Smaller models are always better
- B. Increasing model size, data, and compute will continue to yield improved capabilities and potentially lead to more general intelligence
- C. Scaling always leads to diminishing returns
- D. AI models cannot be scaled beyond a fixed size

Answer: B

Q1369. What is the concept of an AI winter and what historically caused them?

- A. A season when AI labs close for vacation
- B. Periods of reduced funding and interest in AI due to unmet expectations and hype cycles
- C. AI systems that only work in cold temperatures
- D. A type of cooling system for AI hardware

Answer: B

Q1370. What is multi-task learning and why can it improve generalization?

- A. Training separate models for each task independently
- B. Training a single model on multiple related tasks simultaneously so shared representations improve performance on each task
- C. It always reduces model accuracy
- D. Multi-task learning only works for image tasks

Answer: B

Q1371. Why is the softmax function a generalization of the sigmoid function for multiple classes?

- A. They are completely unrelated functions
- B. Softmax reduces to sigmoid for the two-class case and extends the logistic function to produce a probability distribution over K classes
- C. Sigmoid is more general than softmax
- D. Softmax only works with two classes

Answer: B

Q1372. What is the connection between maximum likelihood estimation and cross-entropy loss?

- A. They are unrelated concepts
- B. Minimizing cross-entropy loss is mathematically equivalent to maximizing the log-likelihood of the data under the model
- C. Cross-entropy always equals zero
- D. MLE does not apply to classification

Answer: B

Q1373. What is the significance of the spectral norm of a matrix in deep learning?

- A. It only applies to image processing
- B. The largest singular value that controls the Lipschitz constant and affects training stability in neural networks
- C. It is the sum of all matrix elements
- D. It measures matrix sparsity

Answer: B

Q1374. How does the reparameterization trick enable backpropagation through stochastic nodes?

- A. It removes all randomness from the model
- B. It separates the stochastic component from learnable parameters so gradients can flow through the deterministic path
- C. It doubles the learning rate
- D. It is only used in supervised learning

Answer: B

Q1375. What is the role of the Lagrangian in constrained optimization for ML?

- A. It is a type of neural network layer
- B. It converts constrained optimization into an unconstrained problem by incorporating constraints as penalty terms with multipliers
- C. It only applies to linear regression
- D. It eliminates the need for gradient descent

Answer: B

Q1376. Why is the positive semi-definiteness of a kernel matrix required in kernel methods?

- A. It makes computation faster
- B. It ensures the kernel implicitly computes inner products in a valid feature space per Mercer's theorem
- C. It is not actually required
- D. It guarantees the matrix is invertible

Answer: B

Q1377. What is the information-theoretic interpretation of entropy in machine learning?

- A. It measures dataset size
- B. It quantifies the expected amount of information or uncertainty in a random variable's outcomes
- C. It is the same as variance
- D. It only applies to text data

Answer: B

Q1378. How does the matrix rank relate to the information content of a dataset?

- A. Rank always equals the number of rows
- B. The rank indicates the number of linearly independent dimensions, revealing the true dimensionality and potential redundancy in the data
- C. Rank is always equal to one
- D. It measures the speed of matrix operations

Answer: B

Q1379. What is the relationship between the trace of a matrix and eigenvalues?

- A. They are unrelated
- B. The trace equals the sum of eigenvalues, connecting matrix operations to spectral properties
- C. The trace is always zero
- D. Eigenvalues cannot be summed

Answer: B

Q1380. Why does stochastic gradient descent converge despite using noisy gradient estimates?

- A. It does not actually converge
- B. Under conditions of decreasing learning rates and bounded variance, the noise averages out and SGD converges to a local minimum
- C. SGD always finds the global minimum
- D. Noise prevents any form of convergence

Answer: B

Q1381. How does Python's garbage collector handle circular references?

- A. It cannot handle them at all
- B. It uses a generational garbage collector that detects and collects reference cycles beyond simple reference counting
- C. Circular references never occur in Python
- D. It requires manual memory management

Answer: B

Q1382. What is the purpose of `__init__.py` in Python packages?

- A. It initializes variables to zero
- B. It marks a directory as a Python package, can run initialization code, and controls what is exported via `__all__`
- C. It is required for every Python file
- D. It only contains comments

Answer: B

Q1383. Why is joblib preferred over pickle for serializing scikit-learn models with large NumPy arrays?

- A. Joblib is newer and therefore better
- B. Joblib efficiently handles large NumPy arrays by using memory mapping and optimized serialization for numerical data
- C. Pickle cannot serialize any ML models
- D. Joblib produces smaller files for all data types

Answer: B

Q1384. What is the difference between asyncio and multiprocessing for concurrent ML workloads?

- A. They are identical approaches
- B. asyncio handles I/O-bound concurrency in a single thread while multiprocessing uses separate processes for CPU-bound parallelism
- C. asyncio is always faster
- D. multiprocessing cannot be used for ML

Answer: B

Q1385. How do Python descriptors work and where are they used in ML frameworks?

- A. They describe variable names
- B. Descriptors are objects defining `__get__`, `__set__`, or `__delete__` methods that control attribute access, used in frameworks for parameter validation and lazy loading
- C. They are only used for printing
- D. Descriptors are a type of loop

Answer: B

Q1386. What are metaclasses in Python and how do ML frameworks use them?

- A. They are classes that define how other classes are created, used in frameworks for automatic registration and validation of model components
- B. They are a type of data structure
- C. Metaclasses only exist in Java
- D. They are the same as abstract classes

Answer: A

Q1387. Why is using `numpy.random.Generator` preferred over the legacy `numpy.random` functions?

- A. It produces the same results
- B. Generator provides better statistical properties, reproducibility via explicit seed management, and thread safety compared to the global random state
- C. Legacy functions are faster
- D. There is no difference in practice

Answer: B

Q1388. How does Cython improve performance for computationally intensive ML code?

- A. It replaces Python entirely
- B. It compiles Python-like code to C, enabling static type declarations and direct C-level operations that bypass Python interpreter overhead
- C. Cython only works for web applications
- D. It automatically parallelizes all code

Answer: B

Q1389. What is the purpose of the `__repr__` and `__str__` methods in custom ML classes?

- A. They delete the object
- B. `__repr__` provides an unambiguous developer-facing representation while `__str__` provides a user-friendly readable string
- C. They are identical methods
- D. They are used for encryption

Answer: B

Q1390. What is the advantage of using Python's dataclasses for ML configuration management?

- A. They are faster than regular classes
- B. They auto-generate `__init__`, `__repr__`, `__eq__` and support type hints, reducing boilerplate for configuration objects while maintaining readability
- C. They cannot store numerical data
- D. They replace all other class types

Answer: B

Q1391. How does the iterative imputer (MICE) handle missing data differently from simple imputation?

- A. It deletes rows with missing values
- B. It models each feature with missing values as a function of other features, iterating until convergence
- C. It only replaces with the mean
- D. It requires complete data to begin

Answer: B

Q1392. Why can target encoding with naive implementation cause severe overfitting?

- A. It removes too many features
- B. Target information leaks into the encoded features, and rare categories get unreliable statistics that overfit to training noise
- C. It always underfits
- D. Target encoding is never problematic

Answer: B

Q1393. What is the impact of applying StandardScaler before train-test split versus after?

- A. There is no difference either way
- B. Applying before split causes data leakage because test set statistics influence the scaling parameters
- C. Applying after split is always wrong
- D. It only matters for deep learning

Answer: B

Q1394. How does ADASYN differ from SMOTE for handling class imbalance?

- A. They are identical algorithms
- B. ADASYN adaptively generates more synthetic samples for harder-to-learn minority instances near the decision boundary
- C. ADASYN only works for regression
- D. SMOTE generates more samples than ADASYN

Answer: B

Q1395. Why is preprocessing streaming data fundamentally different from batch preprocessing?

- A. Streaming data is always cleaner
- B. Streaming requires online or incremental algorithms that update statistics without access to the full dataset and must handle concept drift
- C. There is no practical difference
- D. Streaming data does not need preprocessing

Answer: B

Q1396. What is the optimal approach for preprocessing features with mixed distributions in a single dataset?

- A. Apply the same transformation to all features
- B. Use column transformers to apply different preprocessing pipelines to different feature subsets based on their characteristics
- C. Always use StandardScaler for everything
- D. Mixed distributions cannot be preprocessed

Answer: B

Q1397. How does quantile transformation make features follow a uniform or normal distribution?

- A. It multiplies all values by a constant
- B. It maps values to their cumulative distribution percentiles, then applies the inverse CDF of the target distribution
- C. It sorts values alphabetically
- D. It only works for binary features

Answer: B

Q1398. What challenges arise when preprocessing data with concept drift in production systems?

- A. No challenges arise since preprocessing is static
- B. Preprocessing statistics become stale as data distributions shift, requiring adaptive monitoring and periodic recalibration of transformers
- C. Concept drift only affects model weights
- D. Production systems never have concept drift

Answer: B

Q1399. Why is entity resolution important when integrating data from multiple sources?

- A. It improves data encryption
- B. It identifies and merges records referring to the same real-world entity despite different representations across sources
- C. It is only relevant for social media data
- D. Entity resolution deletes all duplicate data

Answer: B

Q1400. How does differential privacy noise injection during preprocessing protect sensitive data?

- A. It makes data completely useless
- B. It adds calibrated random noise that provides mathematical privacy guarantees while preserving aggregate statistical properties
- C. It encrypts individual records
- D. Noise injection always destroys data utility

Answer: B

Q1401. How can principal component analysis be used as an EDA tool for high-dimensional data?

- A. PCA is only for model training
- B. PCA reduces dimensions while preserving variance, enabling 2D/3D visualization that reveals clusters, outliers, and structure invisible in individual features
- C. PCA removes all data variation
- D. PCA only works for two features

Answer: B

Q1402. What is the Anscombe quartet and what lesson does it teach about EDA?

- A. It is a type of neural network
- B. Four datasets with identical summary statistics but drastically different distributions, demonstrating the necessity of visualization over summary statistics alone
- C. It is a data compression algorithm
- D. It is a type of hypothesis test

Answer: B

Q1403. How does mutual information provide a more general dependency measure than correlation in EDA?

- A. Mutual information is identical to correlation
- B. Mutual information captures any statistical dependency including nonlinear relationships, while correlation only measures linear association
- C. It only works for binary variables
- D. Mutual information is always zero

Answer: B

Q1404. Why might the Durbin-Watson statistic be crucial for EDA in regression analysis?

- A. It measures data size
- B. It detects autocorrelation in residuals, which violates regression assumptions and can lead to unreliable coefficient estimates and confidence intervals
- C. It only applies to classification
- D. It is a measure of accuracy

Answer: B

Q1405. What is the purpose of Cook's distance in regression EDA?

- A. To measure cooking time for algorithms
- B. To identify influential observations whose removal would substantially change the regression model's fitted values
- C. To measure the distance between clusters
- D. To compute feature importance

Answer: B

Q1406. How does the Kolmogorov-Smirnov test differ from the Anderson-Darling test for distribution testing?

- A. They are identical tests
- B. K-S measures the maximum difference between CDFs (sensitive to the center) while Anderson-Darling gives more weight to tails, better detecting tail deviations
- C. K-S is always more powerful
- D. Anderson-Darling cannot test normality

Answer: B

Q1407. What is the ecological fallacy and how can it mislead EDA conclusions?

- A. It is a type of environmental data
- B. Drawing conclusions about individuals based on aggregate group data, which can be misleading because group-level patterns may not hold for individuals
- C. It only occurs in biological data
- D. It is the same as Simpson's paradox

Answer: B

Q1408. How can SHAP summary plots enhance EDA by revealing feature-target relationships?

- A. SHAP only works after model deployment
- B. SHAP summary plots show each feature's impact on predictions across the dataset, revealing nonlinear relationships and interaction effects invisible in standard correlation analysis
- C. SHAP replaces all other EDA methods
- D. SHAP only works for linear models

Answer: B

Q1409. What is the Mahalanobis distance and why is it superior to Euclidean distance for multivariate outlier detection?

- A. It is the same as Euclidean distance
- B. It accounts for correlations and variances of features, correctly identifying outliers in correlated multivariate data where Euclidean distance fails
- C. It only works for two dimensions
- D. It ignores feature correlations

Answer: B

Q1410. How does the concept of data granularity affect EDA conclusions and subsequent modeling?

- A. Granularity has no effect on analysis
- B. The level of detail in data aggregation affects visible patterns, and analyzing at wrong granularity can hide important trends or create misleading artifacts
- C. Finer granularity is always better
- D. Granularity only matters for images

Answer: B

Q1411. What is the PAC learning framework and what does it guarantee?

- A. It is a data format standard
- B. It provides theoretical guarantees that a learning algorithm will probably approximately correctly learn a concept given sufficient data and computation
- C. It is a type of neural network
- D. PAC stands for Parallel Adaptive Computing

Answer: B

Q1412. How does Bayesian linear regression differ from ordinary least squares?

- A. They are mathematically identical
- B. Bayesian regression treats parameters as distributions rather than point estimates, incorporating prior beliefs and providing uncertainty quantification
- C. Bayesian regression cannot handle multiple features
- D. OLS always provides better predictions

Answer: B

Q1413. What is the dual formulation of SVM and why is it computationally advantageous?

- A. It doubles the training time
- B. The dual formulation expresses the problem in terms of inner products between data points, enabling the kernel trick and being more efficient when features outnumber samples
- C. It is only for linear problems
- D. The dual formulation always gives different results

Answer: B

Q1414. What is the relationship between logistic regression and maximum entropy models?

- A. They are completely unrelated
- B. Logistic regression is equivalent to a maximum entropy model with the same feature functions, both finding the least biased distribution consistent with observed data
- C. Maximum entropy is only for physics
- D. Logistic regression has higher entropy

Answer: B

Q1415. How does the Rademacher complexity bound relate to model generalization?

- A. It measures training speed
- B. It bounds the generalization gap by measuring how well the hypothesis class can fit random noise, with lower complexity indicating better generalization
- C. It is a type of activation function
- D. It only applies to unsupervised learning

Answer: B

Q1416. What is the difference between model stability and model accuracy in supervised learning?

- A. They always improve together
- B. Stability measures how much predictions change with small data perturbations while accuracy measures correctness, and there can be tradeoffs between them
- C. They are identical metrics
- D. Stability is irrelevant for good models

Answer: B

Q1417. Why can high-dimensional feature spaces make linear separability more likely?

- A. They always make classification harder
- B. Cover's theorem states that complex patterns are more likely to be linearly separable when projected into high-dimensional spaces
- C. High dimensions reduce data quality
- D. Linear separability is independent of dimensionality

Answer: B

Q1418. What is conformal prediction and how does it provide prediction guarantees?

- A. It is a type of data augmentation
- B. It produces prediction sets with a guaranteed coverage probability by using nonconformity scores from calibration data, requiring minimal distributional assumptions
- C. It only works for regression
- D. It is another name for Bayesian prediction

Answer: B

Q1419. How does quantile regression differ from ordinary least squares regression?

- A. They produce identical outputs
- B. Quantile regression estimates conditional quantiles rather than the conditional mean, providing a more complete picture of the response distribution
- C. Quantile regression is simpler than OLS
- D. OLS estimates all quantiles simultaneously

Answer: B

Q1420. What is the theory behind margin-based generalization bounds in classification?

- A. Margins are irrelevant to generalization
- B. Larger classification margins imply tighter generalization bounds because the model is more robust to perturbations, connecting SVM margins to PAC-Bayes theory
- C. Smaller margins always generalize better
- D. Margin theory only applies to linear models

Answer: B

Q1421. How does the second-order Taylor expansion in XGBoost's objective improve optimization?

- A. It makes training slower
- B. Using both first and second derivatives of the loss function provides curvature information for more accurate leaf value estimation and better splitting decisions
- C. It eliminates the need for a loss function
- D. Second-order expansion is identical to first-order

Answer: B

Q1422. What is the relationship between ensemble diversity and error reduction?

- A. Diversity is irrelevant to error
- B. The error of an ensemble is bounded by the average individual error minus a term proportional to diversity, so greater diversity leads to larger error reduction
- C. Maximum diversity always gives zero error
- D. Identical models give the best ensemble

Answer: B

Q1423. How does CatBoost's ordered boosting address target leakage in gradient boosting?

- A. It does not address target leakage
- B. It computes residuals using only previously seen observations in a random permutation, preventing the model from using future target information
- C. It removes all categorical features
- D. It uses the same ordering for all iterations

Answer: B

Q1424. What is the effect of correlation between base learners on ensemble variance reduction?

- A. Correlation has no effect
- B. Higher correlation between base learners reduces the variance reduction benefit, as the ensemble variance includes the average pairwise correlation term
- C. High correlation improves ensembles
- D. Correlation is impossible between ensemble members

Answer: B

Q1425. What is Shapley value-based feature importance and how does it improve upon impurity-based importance?

- A. They are identical measures
- B. Shapley values provide theoretically fair attribution by considering all possible feature coalitions, avoiding the bias toward high-cardinality features that affects impurity-based importance
- C. Shapley values are faster to compute
- D. Impurity-based importance is always more accurate

Answer: B

Q1426. Why can gradient boosting be interpreted as gradient descent in function space?

- A. It cannot be interpreted that way
- B. Each boosting iteration adds a function (tree) that approximates the negative gradient of the loss in function space, minimizing the loss by taking steps in the space of functions
- C. It only uses numerical gradients
- D. Function space interpretation is purely theoretical

Answer: B

Q1427. What is the computational complexity advantage of histogram-based gradient boosting?

- A. It has no advantage
- B. Binning continuous features into discrete histograms reduces split finding from $O(n \cdot \text{features})$ to $O(\text{bins} \cdot \text{features})$ per level, dramatically speeding up training on large datasets
- C. It increases memory usage
- D. It only works for categorical features

Answer: B

Q1428. How does the concept of complementarity differ from diversity in ensemble design?

- A. They are identical concepts
- B. Diversity means models make different errors, while complementarity means models' errors compensate each other in a way that benefits the combined prediction
- C. Complementarity is about feature selection
- D. Diversity is always sufficient for good ensembles

Answer: B

Q1429. What is the theoretical justification for why Random Forest's OOB error is a good estimate of generalization error?

- A. OOB error is not a good estimate
- B. Each sample is evaluated only by trees that did not include it in their bootstrap sample, making OOB predictions equivalent to cross-validation predictions
- C. OOB always overestimates error
- D. It is only valid for small datasets

Answer: B

Q1430. What is the multi-armed bandit approach to dynamic ensemble selection?

- A. It is a type of card game
- B. It treats each base model as an arm, dynamically selecting models for each prediction based on their recent performance using exploration-exploitation strategies
- C. It always selects all models
- D. It only works for binary classification

Answer: B

Q1431. What is the theoretical connection between K-Means and Gaussian Mixture Models?

- A. They are completely unrelated
- B. K-Means can be viewed as a special case of GMM where all covariances are equal isotropic matrices and cluster assignments become hard as variance approaches zero
- C. GMM is a simpler version of K-Means
- D. K-Means uses probabilistic assignments

Answer: B

Q1432. How does HDBSCAN improve upon DBSCAN for clusters of varying density?

- A. It uses the same algorithm as DBSCAN
- B. HDBSCAN builds a hierarchy of DBSCAN clusterings over all epsilon values and extracts the most stable clusters, automatically handling varying densities
- C. It requires the user to specify more parameters
- D. It only works for uniform density clusters

Answer: B

Q1433. Why does the curse of dimensionality particularly affect distance-based clustering algorithms?

- A. It has no effect on clustering
- B. In high dimensions, distances between points converge making clusters indistinguishable, and the volume of the space grows exponentially requiring exponentially more data
- C. High dimensions always improve clustering
- D. Distance-based algorithms are immune to dimensionality

Answer: B

Q1434. What is the information bottleneck method in unsupervised learning?

- A. A physical bottleneck in data pipelines
- B. A method that finds compressed representations that retain maximum mutual information with a relevant variable while discarding irrelevant information
- C. It only applies to autoencoders
- D. It is the same as PCA

Answer: B

Q1435. How does the Wasserstein distance provide a more meaningful comparison between distributions than KL divergence for clustering?

- A. They are mathematically equivalent
- B. Wasserstein distance accounts for the geometry of the space and is defined even when distributions have non-overlapping support, unlike KL divergence which becomes infinite
- C. KL divergence is always better
- D. Wasserstein distance cannot be computed

Answer: B

Q1436. What is the relationship between spectral clustering and graph partitioning?

- A. They are unrelated
- B. Spectral clustering constructs a similarity graph, then uses eigenvalues of the graph Laplacian to partition it, making clustering equivalent to the normalized graph cut problem
- C. Spectral clustering does not use graphs
- D. Graph partitioning is a supervised method

Answer: B

Q1437. Why is the Evidence Lower Bound (ELBO) used as the training objective in variational autoencoders?

- A. It is simpler than reconstruction error
- B. The true log-likelihood is intractable, so ELBO provides a tractable lower bound that combines reconstruction quality with regularization of the latent space
- C. ELBO is the same as MSE loss
- D. It is only used for classification

Answer: B

Q1438. What is the stability-based approach to determining the number of clusters?

- A. It always selects $K=2$
- B. It evaluates how consistently clustering solutions are reproduced under data perturbations, with the correct K yielding the most stable results
- C. It measures computational stability
- D. It only works for DBSCAN

Answer: B

Q1439. How does deep embedded clustering combine deep learning with clustering?

- A. It uses a single K-Means step
- B. It jointly optimizes a deep autoencoder's feature representation and cluster assignments, iteratively refining both the embedding and clusters
- C. It only uses linear transformations
- D. It is identical to standard K-Means

Answer: B

Q1440. What is the fundamental problem of identifiability in mixture models?

- A. Mixture models always produce unique solutions
- B. The likelihood function has multiple equivalent optima due to label switching and potential parameter degeneracies, making the solution non-unique
- C. Identifiability only matters for deep learning
- D. Mixture models have no parameters

Answer: B

Q1441. What is the relationship between the Brier score and model calibration?

- A. Brier score only measures discrimination
- B. The Brier score can be decomposed into calibration, refinement, and uncertainty components, making it sensitive to both calibration and discrimination
- C. Brier score ignores calibration
- D. It is only useful for multi-class problems

Answer: B

Q1442. Why might the McNemar test be more appropriate than paired t-test for comparing classifiers?

- A. They test the same thing
- B. McNemar's test compares classifiers based on their disagreement patterns on specific samples rather than aggregate metrics, providing more statistical power
- C. McNemar test requires more data
- D. Paired t-test is always more powerful

Answer: B

Q1443. What is the purpose of permutation testing for model evaluation?

- A. To permute the features
- B. To establish a null distribution by repeatedly shuffling labels and retraining, determining if the model's performance is significantly better than chance
- C. It always validates the model
- D. Permutation tests only work for linear models

Answer: B

Q1444. How does the concept of distribution shift affect model evaluation in production?

- A. Distribution shift never occurs in practice
- B. When production data differs from training/test data distribution, held-out test performance overestimates real-world performance, requiring ongoing monitoring
- C. Distribution shift only affects accuracy
- D. It can be completely prevented by larger test sets

Answer: B

Q1445. What is the purpose of isotonic regression in model calibration?

- A. It is a type of linear regression
- B. Isotonic regression fits a non-decreasing step function to map model scores to calibrated probabilities, more flexible than Platt scaling
- C. It always produces worse calibration
- D. It is only used for regression models

Answer: B

Q1446. What is the effective sample size problem in cross-validation and how does it affect reliability?

- A. Cross-validation always gives reliable estimates
- B. When data is limited, cross-validation folds share training data making fold estimates correlated, violating independence assumptions and causing underestimation of variance
- C. It only affects large datasets
- D. Effective sample size is always equal to dataset size

Answer: B

Q1447. Why is the area under the precision-recall curve more informative than ROC-AUC for highly imbalanced problems?

- A. ROC-AUC is always more informative
- B. PR-AUC is sensitive to the absolute number of false positives which dominates in imbalanced settings, while ROC-AUC can be misleadingly high due to the large number of true negatives
- C. PR-AUC and ROC-AUC always agree
- D. PR-AUC ignores the positive class

Answer: B

Q1448. What is the difference between prospective and retrospective evaluation of ML models?

- A. They are the same approach
- B. Prospective evaluation tests on future data collected after model training while retrospective uses historical held-out data, with prospective being more reliable for deployment
- C. Retrospective is always better
- D. Prospective evaluation uses training data

Answer: B

Q1449. How does the concept of metric elicitation help select the right evaluation metric?

- A. It always selects accuracy
- B. Metric elicitation formally derives the optimal evaluation metric from stakeholder preferences about the relative costs of different types of errors
- C. It selects metrics randomly
- D. It only applies to unsupervised learning

Answer: B

Q1450. What is the limitation of using a single train-test split for model evaluation and how does repeated stratified K-fold address it?

- A. Single splits are always sufficient
- B. A single split gives high variance estimates dependent on the specific random partition, while repeated stratified K-fold averages over multiple splits and preserves class ratios for more stable estimates
- C. Repeated K-fold is always worse
- D. Single splits are only bad for small datasets

Answer: B

Q1451. How does the Boruta algorithm determine feature relevance compared to regular feature importance?

- A. It uses the same approach as Random Forest importance
- B. Boruta creates shadow features by shuffling original features and compares each real feature's importance against the maximum shadow importance over multiple iterations
- C. It only removes constant features
- D. It is faster than simple importance ranking

Answer: B

Q1452. What is the concept of feature store consistency between training and serving?

- A. It means using the same programming language
- B. Feature store consistency ensures the same feature computation logic and transformations are applied during both model training and real-time inference, preventing training-serving skew
- C. It only matters for batch processing
- D. Consistency is guaranteed automatically

Answer: B

Q1453. Why might learned embeddings outperform manually engineered features for categorical variables?

- A. Manual features are always superior
- B. Embeddings learn dense representations that capture latent semantic relationships between categories in the context of the specific task, beyond what manual encoding schemes can express
- C. Embeddings only work for text
- D. Manual features capture more relationships

Answer: B

Q1454. How does automated feature engineering with deep feature synthesis work?

- A. It uses deep learning only
- B. DFS automatically generates features by stacking aggregation and transformation primitives across relational tables, systematically exploring the feature space
- C. It only creates simple averages
- D. It requires no data

Answer: B

Q1455. What is the theoretical basis for why polynomial feature expansion can lead to overfitting?

- A. Polynomials never cause overfitting
- B. The VC dimension grows rapidly with polynomial degree, increasing model capacity exponentially and requiring exponentially more data to avoid overfitting
- C. Polynomials always underfit
- D. Overfitting is unrelated to feature expansion

Answer: B

Q1456. How does entity embedding of categorical variables from neural networks transfer to tree-based models?

- A. It cannot be transferred
- B. Entity embeddings learned by neural networks can be extracted and used as numerical features in tree-based models, combining the representation learning of deep learning with the strengths of gradient boosting
- C. Tree models always outperform embeddings
- D. Neural embeddings are incompatible with trees

Answer: B

Q1457. What is the stability selection method and how does it improve feature selection reliability?

- A. It selects the most stable features across time
- B. Stability selection repeatedly subsamples data and applies feature selection, keeping features that are consistently selected across many subsamples, controlling false discovery rate
- C. It uses the same selection every time
- D. It only works for regression

Answer: B

Q1458. Why is the order of feature engineering steps important in a preprocessing pipeline?

- A. Order never matters
- B. Transformations are not commutative: scaling before encoding produces different results than encoding before scaling, and incorrect ordering can introduce data leakage or distort distributions
- C. All orderings produce identical results
- D. Only the last step matters

Answer: B

Q1459. How do feature interactions affect the performance gap between linear and nonlinear models?

- A. Interactions are irrelevant to model choice
- B. When important feature interactions exist, linear models require explicit interaction features while tree-based models discover them automatically, explaining their performance advantage
- C. Linear models automatically detect interactions
- D. Nonlinear models cannot use interactions

Answer: B

Q1460. What is the concept of leakage-free feature engineering in production ML systems?

- A. It is impossible to prevent leakage
- B. Ensuring all features are computed using only information available at prediction time and that no future or target information contaminates features during training or serving
- C. Leakage only occurs in training
- D. Production systems never have leakage

Answer: B

Q1461. What is the connection between information theory and deep learning through the Information Bottleneck principle?

- A. Information theory is irrelevant to deep learning
- B. Each layer progressively compresses input information while retaining information relevant to the output, balancing compression and prediction
- C. Information always increases through layers
- D. The Information Bottleneck only applies to autoencoders

Answer: B

Q1462. What is the double descent phenomenon in deep learning?

- A. Models always improve with more parameters
- B. Test error shows a classical U-shaped curve then decreases again in the over-parameterized regime, challenging classical bias-variance intuitions
- C. It only occurs with small datasets
- D. Double descent means training the model twice

Answer: B

Q1463. How do implicit biases of gradient descent favor certain solutions in over-parameterized networks?

- A. Gradient descent always finds the worst solution
- B. SGD implicitly regularizes toward flat minima with small norm weights, which tend to generalize better, even without explicit regularization
- C. Gradient descent has no bias
- D. Implicit bias only affects small networks

Answer: B

Q1464. What is the Sharpness-Aware Minimization technique and why does it improve generalization?

- A. It sharpens image inputs
- B. SAM seeks parameters in regions where the loss landscape is uniformly low rather than just at a single low point, finding flatter minima that generalize better
- C. It makes the loss function sharper
- D. SAM only works for classification

Answer: B

Q1465. What is the role of the Hessian spectrum in understanding neural network optimization?

- A. The Hessian is not relevant to optimization
- B. The eigenvalue distribution of the Hessian reveals the loss landscape geometry, with a few large eigenvalues and many near-zero values indicating low effective dimensionality
- C. The Hessian spectrum is always flat
- D. It only applies to linear models

Answer: B

Q1466. How does neural network pruning maintain accuracy while reducing model size?

- A. Pruning always degrades accuracy significantly
- B. Over-parameterized networks contain redundant connections, and structured or unstructured pruning can remove a large fraction of weights with minimal accuracy loss, especially with fine-tuning
- C. Pruning only removes input features
- D. It only works for convolutional networks

Answer: B

Q1467. What is the difference between post-training quantization and quantization-aware training?

- A. They are identical approaches
- B. Post-training quantization converts a trained model's weights to lower precision, while quantization-aware training simulates quantization during training for better accuracy retention
- C. Post-training is always better
- D. Quantization-aware training uses higher precision

Answer: B

Q1468. Why does batch normalization enable higher learning rates?

- A. It has no effect on learning rates
- B. By normalizing layer inputs, batch norm reduces internal covariate shift and smooths the loss landscape, allowing larger learning rate steps without divergence
- C. It only affects the first layer
- D. Batch norm always reduces learning rates

Answer: B

Q1469. What is the grokking phenomenon in deep learning?

- A. A standard training pattern
- B. A delayed generalization where a model memorizes training data first and only much later suddenly learns to generalize, well past the point of overfitting
- C. Grokking means fast convergence
- D. It only occurs in shallow networks

Answer: B

Q1470. How does the loss of plasticity problem affect continually trained neural networks?

- A. Plasticity always increases with training
- B. Networks gradually lose the ability to learn new tasks as training progresses, due to dead neurons, saturated weights, and reduced effective learning capacity
- C. Loss of plasticity only affects small models
- D. It is the same as catastrophic forgetting

Answer: B

Q1471. How does the Transformer's multi-head attention improve upon single-head attention?

- A. It uses more memory without benefit
- B. Multiple heads attend to different representation subspaces at different positions simultaneously, capturing diverse relational patterns that a single head cannot
- C. More heads always reduce performance
- D. Multi-head is computationally cheaper

Answer: B

Q1472. What is the theoretical basis behind why deeper networks can represent functions more efficiently than wider shallow ones?

- A. Depth has no theoretical advantage
- B. Depth enables compositional representation where complex functions are built hierarchically from simpler sub-functions, achieving exponential efficiency over shallow networks for certain function classes
- C. Wider networks are always more efficient
- D. Deep and shallow networks have equal capacity

Answer: B

Q1473. How does the score-based diffusion model formulation connect to stochastic differential equations?

- A. They are unrelated
- B. Diffusion models can be formulated as continuous-time SDEs where the forward process adds noise and the reverse process removes it, with the score function learned by the neural network guiding denoising
- C. SDEs are only for financial modeling
- D. Diffusion models use deterministic equations only

Answer: B

Q1474. What is the fundamental trade-off in designing efficient Transformer architectures?

- A. There is no trade-off
- B. Reducing the quadratic $O(n^2)$ attention complexity to linear or sub-quadratic while preserving the model's ability to capture long-range dependencies and global context
- C. Efficiency always improves accuracy
- D. Transformer efficiency only depends on model size

Answer: B

Q1475. How does LoRA enable efficient fine-tuning of large pre-trained models?

- A. It retrains all parameters
- B. LoRA freezes pre-trained weights and injects trainable low-rank decomposition matrices into attention layers, dramatically reducing trainable parameters while maintaining performance
- C. It deletes unused parameters
- D. LoRA only works for small models

Answer: B

Q1476. What is the role of the exponential moving average of model weights in improving generalization?

- A. It speeds up training only
- B. EMA maintains a smoothed version of weights that averages over the training trajectory, effectively ensembling models from different training stages and finding flatter minima
- C. EMA always reduces accuracy
- D. It is only used during data preprocessing

Answer: B

Q1477. Why do Vision Transformers partition images into patches rather than processing individual pixels?

- A. Patches are visually more appealing
- B. Processing individual pixels would make the self-attention sequence length equal to image resolution squared, making computation prohibitively expensive, while patches create manageable sequences
- C. Patches always produce better accuracy
- D. Pixel-level processing is computationally cheaper

Answer: B

Q1478. How does the Mixture of Experts architecture achieve conditional computation?

- A. All experts process every input
- B. A gating network routes each input to only a small subset of expert networks, so total computation scales sub-linearly with model size
- C. Experts are selected randomly without any gating
- D. MoE reduces model capacity

Answer: B

Q1479. What is the relationship between contrastive learning and the InfoNCE loss?

- A. They are unrelated
- B. InfoNCE is the standard loss function for contrastive learning that maximizes a lower bound on mutual information between positive pairs while treating other samples as negatives
- C. InfoNCE is only for supervised learning
- D. Contrastive learning does not use loss functions

Answer: B

Q1480. How does the concept of neural scaling laws inform model and dataset design decisions?

- A. Scaling laws are not useful for planning
- B. Empirical power-law relationships between model size, dataset size, compute budget, and loss enable predicting performance of larger models from smaller experiments, guiding resource allocation
- C. Performance scales linearly with compute
- D. Scaling laws only apply to language models

Answer: B

Q1481. What is the fundamental challenge of evaluating open-ended text generation with automatic metrics?

- A. Automatic metrics are always perfect
- B. There are many valid ways to express the same meaning, so reference-based metrics like BLEU penalize valid but differently-worded outputs, failing to capture semantic equivalence
- C. All generated text can be evaluated by word matching
- D. Open-ended generation does not need evaluation

Answer: B

Q1482. How does the rotary position embedding encode positional information in Transformers?

- A. It adds a fixed vector to all positions
- B. RoPE rotates the query and key vectors by position-dependent angles, encoding relative positions through the dot product's natural decay with distance
- C. It uses learned position embeddings
- D. RoPE removes all position information

Answer: B

Q1483. What is the constitutional AI approach to aligning language models?

- A. Writing an AI constitution document
- B. Training the model to critique and revise its own outputs according to a set of principles, reducing reliance on human feedback for each specific case
- C. It is a type of government regulation
- D. Constitutional AI replaces all human oversight

Answer: B

Q1484. What is the difference between sparse and dense retrieval in retrieval-augmented generation?

- A. They produce identical results
- B. Sparse retrieval uses term matching like BM25 while dense retrieval uses learned semantic embeddings, with dense capturing meaning beyond exact word overlap
- C. Sparse retrieval is always better
- D. Dense retrieval only matches exact words

Answer: B

Q1485. How does Direct Preference Optimization simplify the RLHF pipeline?

- A. It adds more complexity than RLHF
- B. DPO directly optimizes the language model policy using preference pairs without needing a separate reward model or reinforcement learning, simplifying the training pipeline
- C. DPO requires a separate reward model
- D. It is identical to standard RLHF

Answer: B

Q1486. What is the mixture of softmaxes approach and how does it address the softmax bottleneck?

- A. It mixes multiple models together
- B. It uses a mixture of multiple softmax distributions to increase the rank of the output distribution matrix, addressing the limitation that a single softmax cannot model complex distributions
- C. Single softmax has no limitations
- D. It reduces model capacity

Answer: B

Q1487. Why is tokenizer choice critically important for multilingual NLP models?

- A. Tokenizers are identical across languages
- B. The tokenizer determines how efficiently each language is represented: languages with poor tokenizer coverage get fragmented into more tokens, increasing compute cost and reducing effective context length
- C. Tokenizers only affect English text
- D. All languages require the same number of tokens

Answer: B

Q1488. What is the in-context learning phenomenon in large language models?

- A. Learning that occurs during pre-training
- B. LLMs can learn to perform new tasks from examples provided in the prompt without any parameter updates, implicitly implementing learning algorithms within their forward pass
- C. It requires fine-tuning the model
- D. In-context learning only works for classification

Answer: B

Q1489. How does the KV-cache optimization speed up autoregressive Transformer inference?

- A. It caches the entire model on disk
- B. It stores previously computed key and value tensors to avoid redundant recomputation at each generation step, reducing cumulative attention computation
- C. KV-cache only works during training
- D. It caches output probabilities

Answer: B

Q1490. What is the challenge of faithfulness in abstractive summarization systems?

- A. Faithfulness is not important for summaries
- B. Abstractive models can generate fluent but factually incorrect summaries that contradict or fabricate information not present in the source document
- C. Extractive summarization has worse faithfulness
- D. All neural summaries are perfectly faithful

Answer: B

Q1491. How do deformable convolutions improve upon standard convolutions for object detection?

- A. They use fixed rectangular receptive fields
- B. Deformable convolutions learn input-dependent offsets for sampling positions, enabling adaptive receptive fields that conform to object shapes
- C. They always produce worse results
- D. Deformable convolutions are faster than standard convolutions

Answer: B

Q1492. What is the concept of equivariance versus invariance in CNN design?

- A. They are the same property
- B. Equivariance means the output transforms predictably with input transformations while invariance means the output remains unchanged
- C. CNNs are always invariant to all transformations
- D. Equivariance prevents learning

Answer: B

Q1493. Why is the anchor-free approach gaining popularity over anchor-based methods in object detection?

- A. Anchors are always better
- B. Anchor-free methods directly predict object centers and sizes without predefined anchor boxes, reducing hyperparameters and simplifying the pipeline
- C. Anchor-free methods cannot detect objects
- D. All modern detectors use anchors

Answer: B

Q1494. What is the fundamental challenge of training object detection models with highly imbalanced foreground-background ratios?

- A. Imbalance does not affect detection
- B. The vast majority of anchor locations are background, which dominate the loss and overwhelm the gradient signal from sparse foreground objects
- C. All locations are equally important
- D. Background samples improve foreground detection

Answer: B

Q1495. How does the CLIP model enable zero-shot image classification?

- A. It memorizes all possible images
- B. CLIP learns a shared embedding space for images and text through contrastive pre-training, enabling classification by comparing image embeddings with text descriptions of classes
- C. Zero-shot classification is impossible
- D. CLIP only works for text

Answer: B

Q1496. What is the role of attention mechanisms in modern image segmentation architectures?

- A. Attention is not used in segmentation
- B. Attention modules capture long-range dependencies between distant pixels, enabling global context aggregation critical for understanding semantic relationships
- C. Attention only works for small images
- D. It replaces all convolutional layers

Answer: B

Q1497. Why is test-time augmentation effective for improving computer vision model predictions?

- A. It has no effect on predictions
- B. TTA applies multiple augmentations to the test image, runs inference on each, and aggregates predictions, effectively ensembling predictions from different perspectives
- C. It only works during training
- D. TTA always reduces accuracy

Answer: B

Q1498. How do neural radiance fields represent 3D scenes from 2D images?

- A. They stack 2D images together
- B. NeRF trains a neural network to map 3D coordinates and viewing direction to color and density, then uses volume rendering to synthesize novel views
- C. They require 3D scanners
- D. NeRF only works for simple objects

Answer: B

Q1499. What is the challenge of domain gap in transferring CV models across different visual domains?

- A. Domain gap does not exist
- B. Models trained on one visual domain perform poorly on another due to differences in appearance, lighting, style, and sensor characteristics, requiring domain adaptation techniques
- C. All visual domains look identical to CNNs
- D. Domain gap only affects text models

Answer: B

Q1500. What are the tradeoffs in the CAP theorem and how do modern distributed databases address them?

A. There are no tradeoffs in distributed systems
B. The CAP theorem states only two of consistency, availability, and partition tolerance can be fully guaranteed. Modern systems like CockroachDB and Spanner make nuanced tradeoffs using techniques like consensus protocols
C. All three guarantees are always achievable
D. CAP theorem is no longer relevant

Answer: B

Q1501. How does Apache Flink achieve exactly-once processing semantics in stream processing?

A. It simply processes each message once
B. Flink uses distributed snapshots (Chandy-Lamport algorithm) for consistent checkpointing combined with two-phase commit for external sinks to guarantee exactly-once end-to-end processing
C. Exactly-once is impossible in streaming
D. It buffers all data in memory

Answer: B

Q1502. What is the challenge of data skew in distributed joins and how can it be mitigated?

A. Data skew has no impact on joins
B. Skewed key distributions cause some partitions to receive disproportionately more data, creating bottleneck tasks. Mitigations include salted keys, broadcast joins for small tables, and adaptive query engines
C. All data is always evenly distributed
D. Skew only affects small datasets

Answer: B

Q1503. How does the Delta Lake transaction log provide ACID guarantees on top of object storage?

A. It uses a traditional relational database
B. Delta Lake maintains a transaction log of ordered JSON files recording every change, using optimistic concurrency control to provide atomicity, consistency, isolation, and durability on cloud object stores
C. ACID is impossible on object storage
D. It modifies the underlying storage system

Answer: B

Q1504. What is the theoretical basis for why bloom filters are space-efficient for approximate membership testing?

A. Bloom filters store exact data
B. Bloom filters use multiple hash functions mapping elements to a bit array, allowing false positives but no false negatives with space usage sub-linear in the number of elements
C. They always produce exact results
D. Bloom filters require more space than hash tables

Answer: B

Q1505. Why is backpressure management critical in streaming data systems?

A. Backpressure is not a concern in streaming
B. When downstream consumers cannot keep up with upstream producers, backpressure prevents system overload by propagating flow control signals to slow producers or buffer data appropriately
C. Streaming systems never have speed mismatches
D. Backpressure always causes data loss

Answer: B

Q1506. How does the concept of materialized views improve query performance in analytical databases?

- A. They have no performance benefit
- B. Materialized views pre-compute and store query results, trading storage space for query speed by avoiding re-computation of expensive aggregations and joins
- C. They only work for simple queries
- D. Materialized views always have stale data

Answer: B

Q1507. What is the log-structured merge tree and why is it used in write-heavy big data systems?

- A. It is a type of logging framework
- B. LSM trees buffer writes in memory, flush sorted segments to disk, and merge them periodically, optimizing write throughput at the cost of read amplification
- C. It only works for read-heavy workloads
- D. LSM trees are slower than B-trees for all operations

Answer: B

Q1508. How does federated query processing enable analytics across heterogeneous data sources?

- A. It requires all data to be in one database
- B. Federated query engines translate a single query into sub-queries for different data sources, execute them in parallel, and combine results, enabling analytics without data migration
- C. It only works with SQL databases
- D. Federated queries are always slower than centralized

Answer: B

Q1509. What is the challenge of maintaining data quality at scale in big data systems?

- A. Data quality is automatic at scale
- B. Scale amplifies data quality issues: schema drift across sources, late-arriving data, duplicate events, and silent failures require automated validation, schema registries, dead letter queues, and data observability platforms
- C. Big data is always high quality
- D. Data quality only matters for small datasets

Answer: B

Q1510. What is the training-serving skew problem and how does a feature store address it?

- A. There is no such problem
- B. Training-serving skew occurs when feature computation differs between training and serving, causing silent performance degradation. Feature stores provide a single source of truth for feature definitions used in both contexts
- C. Feature stores cause more skew
- D. Skew only affects small models

Answer: B

Q1511. How does continuous training differ from traditional periodic retraining in production ML?

- A. They are identical approaches
- B. Continuous training automatically triggers model retraining based on data changes, performance degradation, or schedules, with automated validation and deployment, creating a self-maintaining ML system
- C. Continuous training never stops training
- D. Traditional retraining is always better

Answer: B

Q1512. What are the unique challenges of testing ML systems compared to traditional software?

- A. ML systems are easier to test
- B. ML systems have non-deterministic outputs, depend on training data quality, require testing of data pipelines and model behavior, and lack clear oracles for correct outputs in many cases
- C. Testing is identical for ML and software
- D. ML systems do not need testing

Answer: B

Q1513. How does the concept of ML technical debt manifest in production systems?

- A. ML systems have no technical debt
- B. ML systems accumulate hidden debt through entangled features, undeclared data dependencies, feedback loops, configuration complexity, and the difficulty of monitoring and testing non-deterministic systems
- C. Technical debt is only a software concern
- D. ML debt is always visible

Answer: B

Q1514. What is the role of a serving graph in complex ML deployment architectures?

- A. A graph showing model accuracy
- B. A serving graph orchestrates multiple models and processing steps into a single prediction pipeline, handling feature transformation, model ensemble, post-processing, and fallback logic
- C. It only serves single models
- D. Serving graphs are only theoretical

Answer: B

Q1515. How does data-centric AI differ from the traditional model-centric approach to improving ML systems?

- A. They are the same approach
- B. Data-centric AI focuses on systematically improving the quality and quantity of training data rather than iterating on model architectures, recognizing data quality as the primary performance bottleneck
- C. Data-centric ignores data quality
- D. Model-centric is always superior

Answer: B

Q1516. What is the challenge of implementing real-time feature computation for online model serving?

- A. Real-time features are simple to implement
- B. Real-time features require low-latency computation from streaming data, consistent logic with training features, handling of late-arriving data, and maintaining state across events
- C. Real-time features are identical to batch features
- D. Latency does not matter for serving

Answer: B

Q1517. How does ML governance address compliance and regulatory requirements for AI systems?

- A. Governance is unnecessary for ML
- B. ML governance establishes policies, processes, and tools for model documentation, approval workflows, bias auditing, access control, and audit trails to meet regulatory requirements
- C. Governance only applies to financial models
- D. All ML models are automatically compliant

Answer: B

Q1518. What is the multi-armed bandit approach to model selection in production?

- A. A gambling strategy
- B. Multi-armed bandit algorithms dynamically allocate traffic between model variants based on their observed performance, balancing exploration and exploitation more efficiently than static A/B tests
- C. It always selects the same model
- D. It requires manual selection

Answer: B

Q1519. Why is observability more important than traditional monitoring for production ML systems?

- A. Observability and monitoring are identical
- B. Observability enables investigation of novel, unexpected failure modes using high-cardinality telemetry data, while monitoring only checks predefined metrics against thresholds and cannot diagnose unknown unknowns
- C. Monitoring covers all failure modes
- D. Observability is only for infrastructure

Answer: B

Q1520. Why is the impossibility theorem for fairness significant for AI system design?

- A. The theorem has been disproven
- B. The impossibility theorem shows that several intuitive fairness criteria cannot all be satisfied simultaneously, requiring explicit tradeoff decisions in AI system design
- C. All fairness criteria can be satisfied
- D. The theorem only applies to binary classification

Answer: B

Q1521. How does machine unlearning attempt to address the right to be forgotten in trained AI models?

- A. It deletes the model entirely
- B. Machine unlearning develops algorithms to efficiently remove the influence of specific training data points from a trained model without full retraining, satisfying data deletion requests
- C. Unlearning is impossible for neural networks
- D. Simply deleting data files is sufficient

Answer: B

Q1522. What is the concept of counterfactual fairness in AI decision-making?

- A. Making decisions based on fictional scenarios
- B. A decision is counterfactually fair if it would remain the same in a hypothetical world where the individual's protected attribute had been different, ensuring the attribute had no causal influence
- C. It means treating all groups identically
- D. Counterfactual fairness ignores protected attributes

Answer: B

Q1523. How does the concept of differential privacy provide mathematical guarantees for data protection?

- A. It encrypts all data permanently
- B. Differential privacy ensures that the output of an analysis is statistically indistinguishable whether or not any single individual's data is included, with the privacy budget epsilon quantifying the guarantee
- C. Differential privacy eliminates all utility
- D. It only works for numerical data

Answer: B

Q1524. What is the challenge of ensuring AI transparency across the full ML pipeline?

- A. Transparency is only needed for the model itself
- B. Full transparency requires documenting data collection and consent, preprocessing decisions, model selection rationale, training procedures, evaluation methodology, deployment context, and ongoing monitoring, which is complex and costly
- C. Transparency means making models open source
- D. Full transparency is easily achieved

Answer: B

Q1525. How do adversarial attacks on AI systems raise unique ethical concerns?

- A. Adversarial attacks have no ethical implications
- B. Adversarial vulnerabilities mean AI systems can be deliberately manipulated to make dangerous errors, raising concerns about safety-critical deployment, security, and the responsibility gap when AI is deceived
- C. Adversarial attacks only affect image models
- D. All AI systems are robust to adversarial attacks

Answer: B

Q1526. What is the tension between model performance optimization and fairness in high-stakes AI applications?

- A. There is no tension between performance and fairness
- B. Optimizing for overall accuracy may come at the cost of disparate performance across demographic groups, requiring explicit fairness constraints that may reduce overall accuracy but improve equity
- C. Fairness always improves accuracy
- D. Performance and fairness are identical objectives

Answer: B

Q1527. How does the EU AI Act's risk-based framework categorize AI systems?

- A. All AI systems are treated equally
- B. The Act categorizes AI into unacceptable risk (banned), high risk (strict requirements including conformity assessment), limited risk (transparency obligations), and minimal risk (no specific obligations)
- C. It only regulates social media AI
- D. The Act has no enforcement mechanism

Answer: B

Q1528. What is the concept of AI sovereignty and why is it geopolitically significant?

- A. AI that is self-aware
- B. National control over AI infrastructure, data, research capabilities, and regulatory frameworks, which has become a strategic priority affecting global power dynamics and technology independence
- C. It only matters for military AI
- D. AI sovereignty is not a real concept

Answer: B

Q1529. How does the concept of meaningful human control apply to autonomous AI systems?

- A. Humans must manually approve every AI decision
- B. Meaningful human control requires that humans can understand, intervene in, and override AI decisions at appropriate levels while maintaining effective oversight even as AI systems become more autonomous
- C. Human control is unnecessary for advanced AI
- D. It means AI should never make decisions

Answer: B

Q1530. How does the Segment Anything Model generalize to unseen segmentation tasks?

- A. It memorizes all possible object shapes
- B. SAM is trained on a massive diverse dataset with a promptable architecture that accepts points, boxes, or text to segment any object without task-specific fine-tuning
- C. It uses a fixed set of object categories
- D. It only works for previously seen objects

Answer: B

SUBJECT

Computer Networks & Cloud Computing

900 Multiple Choice Questions

Computer Networks & Cloud Computing — Easy

300 questions

Q1. What is the primary purpose of data communication?

- A. To exchange data between two devices
- B. To design hardware circuit boards
- C. To store data permanently on disks
- D. To compile programs into binaries

Answer: A

Q2. Which of the following is NOT a component of a data communication system?

- A. Compiler
- B. Receiver
- C. Sender
- D. Protocol

Answer: A

Q3. In simplex mode of data transmission, data flows in:

- A. One direction only
- B. Both directions at once
- C. Both ways but alternating
- D. Multiple directions

Answer: A

Q4. Which transmission mode allows data to flow in both directions simultaneously?

- A. Half-duplex
- B. Multiplex
- C. Simplex
- D. Full-duplex

Answer: D

Q5. What does bandwidth refer to in data communication?

- A. The count of devices on the network
- B. The physical width of the cable used
- C. The range of frequencies a channel passes
- D. The length of the transmission medium

Answer: C

Q6. Walkie-talkie is an example of which transmission mode?

- A. Full-duplex
- B. Half-duplex
- C. Simplex
- D. Automatic

Answer: B

Q7. Which of the following is a guided transmission medium?

- A. Infrared
- B. Microwaves
- C. Radio waves
- D. Twisted pair

Answer: D

Q8. What is the unit of data transmission speed?

- A. Amperes (A)
- B. Hertz (Hz)
- C. Watts (W)
- D. Bits per second (bps)

Answer: D

Q9. In data communication, what is a protocol?

- A. A hardware device on a NIC
- B. A software application suite
- C. A set of rules for data exchange
- D. A type of transmission cable

Answer: C

Q10. Which signal type has a continuous varying pattern?

- A. Analog signal
- B. Digital signal
- C. Binary signal
- D. Discrete signal

Answer: A

Q11. What type of network covers a small geographical area like a single building?

- A. LAN
- B. PAN
- C. MAN
- D. WAN

Answer: A

Q12. Which topology connects all devices to a single central cable?

- A. Star
- B. Ring
- C. Mesh
- D. Bus

Answer: D

Q13. How many layers does the OSI model have?

- A. 6
- B. 4
- C. 7
- D. 5

Answer: C

Q14. Which device operates at the physical layer to regenerate signals?

- A. Switch
- B. Repeater
- C. Gateway
- D. Router

Answer: B

Q15. Which network topology provides the highest redundancy?

- A. Full mesh topology
- B. Bus topology
- C. Star topology
- D. Ring topology

Answer: A

Q16. Which layer of the OSI model is responsible for routing?

- A. Data link layer
- B. Network layer
- C. Physical layer
- D. Transport layer

Answer: B

Q17. What does WAN stand for?

- A. Wide Access Node
- B. Wireless Area Network
- C. Web Access Network
- D. Wide Area Network

Answer: D

Q18. In which switching technique is a dedicated path established before data transfer?

- A. Circuit switching
- B. Cell switching
- C. Message switching
- D. Packet switching

Answer: A

Q19. Which networking device connects two different networks and operates at the network layer?

- A. Hub
- B. Router
- C. Repeater
- D. Switch

Answer: B

Q20. What is the TCP/IP model also known as?

- A. Internet model
- B. IEEE model
- C. OSI model
- D. ISO model

Answer: A

Q21. What is the main function of the data link layer?

- A. Node-to-node delivery of frames
- B. Routing packets between networks
- C. Physical signal encoding
- D. End-to-end message delivery

Answer: A

Q22. What is a MAC address?

- A. A port number for applications
- B. A physical address burned into a NIC
- C. An IP address used for routing
- D. A logical address assigned by software

Answer: B

Q23. How many bits are in a MAC address?

- A. 32 bits
- B. 48 bits
- C. 128 bits
- D. 64 bits

Answer: B

Q24. Which error detection method appends a calculated value at the end of the data?

- A. All of the above
- B. Parity check only
- C. Hamming code only
- D. CRC method only

Answer: A

Q25. What does CSMA/CD stand for?

- A. Carrier Sense Media Access with Collision Detection
- B. Central Station Multiple Access with Collision Detection
- C. Carrier Sense Multiple Access with Collision Detection
- D. Carrier Signal Multiple Access with Collision Division

Answer: C

Q26. Which sublayer of the data link layer controls access to the shared medium?

- A. Physical sublayer
- B. Network sublayer
- C. LLC sublayer
- D. MAC sublayer

Answer: D

Q27. What is framing in the data link layer?

- A. Assigning IP addresses to each host
- B. Dividing data into units called frames
- C. Converting digital signals to analog form
- D. Encrypting data for secure transfer

Answer: B

Q28. Which protocol is used to resolve an IP address to a MAC address?

- A. DNS
- B. RARP
- C. DHCP
- D. ARP

Answer: D

Q29. Simple parity check can detect which type of errors?

- A. Only even number of bit errors
- B. Only odd number of bit errors
- C. All types of errors in data
- D. Only burst errors in frames

Answer: B

Q30. In Ethernet, what happens when a collision is detected?

- A. Data is retransmitted right away immediately
- B. The entire network shuts down completely
- C. Devices wait a random time then retransmit
- D. The frame is discarded permanently by all

Answer: C

Q31. What is the size of an IPv4 address?

- A. 64 bits
- B. 128 bits
- C. 16 bits
- D. 32 bits

Answer: D

Q32. What is the main function of a router?

- A. Storing web pages on servers
- B. Forwarding packets between networks
- C. Providing wireless network access
- D. Connecting devices within a LAN

Answer: B

Q33. Which class of IPv4 addresses is used for multicasting?

- A. Class B
- B. Class D
- C. Class C
- D. Class A

Answer: B

Q34. What does the TTL field in an IP header stand for?

- A. Transfer Type Label
- B. Transmission Time Limit
- C. Total Transfer Length
- D. Time to Live

Answer: D

Q35. Which protocol is used to send error messages at the network layer?

- A. UDP
- B. TCP
- C. ARP
- D. ICMP

Answer: D

Q36. What is the default subnet mask for a Class C IPv4 address?

- A. 255.255.0.0
- B. 255.255.255.255
- C. 255.0.0.0
- D. 255.255.255.0

Answer: D

Q37. Which IP address range is reserved for loopback testing?

- A. 169.254.0.0/16
- B. 127.0.0.0/8
- C. 10.0.0.0/8
- D. 192.168.0.0/16

Answer: B

Q38. What tool uses ICMP echo requests to test connectivity?

- A. traceroute
- B. nslookup
- C. ping
- D. netstat

Answer: C

Q39. Which of the following is a private IPv4 address?

- A. 11.0.0.1
- B. 172.32.0.1
- C. 8.8.8.8
- D. 192.168.1.1

Answer: D

Q40. What is the purpose of NAT (Network Address Translation)?

- A. To translate private IPs to public IPs
- B. To compress packets for transfer
- C. To assign domain names to hosts
- D. To encrypt data during transit

Answer: A

Q41. Which two protocols operate at the transport layer?

- A. ARP and RARP
- B. IP and ICMP
- C. HTTP and FTP
- D. TCP and UDP

Answer: D

Q42. Which transport layer protocol provides reliable, connection-oriented service?

- A. IP
- B. ICMP
- C. UDP
- D. TCP

Answer: D

Q43. What is the range of port numbers?

- A. 0 to 49151
- B. 0 to 1023
- C. 0 to 65535
- D. 1 to 65535

Answer: C

Q44. Which well-known port is used by HTTP?

- A. 25
- B. 80
- C. 21
- D. 443

Answer: B

Q45. What type of service does UDP provide?

- A. Unreliable, connectionless
- B. Reliable, connectionless
- C. Reliable, connection-oriented
- D. Unreliable, connection-oriented

Answer: A

Q46. How many steps are in the TCP three-way handshake?

- A. 5
- B. 4
- C. 2
- D. 3

Answer: D

Q47. What is a socket in networking?

- A. A routing protocol name
- B. A type of cable connector
- C. A physical port on a device
- D. An IP address plus port number

Answer: D

Q48. Which application typically uses UDP?

- A. DNS queries
- B. Web browsing
- C. Email sending
- D. File transfer

Answer: A

Q49. What does the checksum field in the UDP header provide?

- A. Error detection
- B. Flow control
- C. Congestion control
- D. Packet sequencing

Answer: A

Q50. What is the purpose of the transport layer?

- A. Network-to-network routing
- B. Node-to-node frame delivery
- C. Signal encoding on media
- D. Process-to-process delivery

Answer: D

Q51. What does DNS stand for?

- A. Digital Network Standard
- B. Data Network Service
- C. Domain Name System
- D. Dynamic Name Server

Answer: C

Q52. Which protocol is used to send emails?

- A. FTP
- B. SMTP
- C. IMAP
- D. POP3

Answer: B

Q53. What port does FTP use for control connection?

- A. 23
- B. 21
- C. 22
- D. 20

Answer: B

Q54. What is HTTP?

- A. A routing protocol for networks
- B. An email delivery protocol
- C. A file transfer protocol only
- D. A protocol for transferring web pages

Answer: D

Q55. Which protocol is used to retrieve emails from a mail server while keeping them on the server?

- A. SMTP
- B. IMAP
- C. POP3
- D. FTP

Answer: B

Q56. Which port does SSH use?

- A. 25
- B. 21
- C. 22
- D. 23

Answer: C

Q57. What is Telnet used for?

- A. Remote login
- B. Email sending
- C. Web browsing
- D. File transfer

Answer: A

Q58. What does URL stand for?

- A. Universal Reference Locator
- B. Universal Resource Locator
- C. Uniform Resource Locator
- D. Unified Resource Link

Answer: C

Q59. Which protocol does a web browser use to access websites?

- A. FTP protocol
- B. SMTP protocol
- C. Telnet protocol
- D. HTTP/HTTPS

Answer: D

Q60. What is the role of a DNS resolver?

- A. Encrypting data for transfer
- B. Sending emails between servers
- C. Hosting websites on servers
- D. Translating domain names to IPs

Answer: D

Q61. What does Wi-Fi stand for?

- A. Wireless Fidelity standard
- B. Wide Field technology
- C. It is a brand name only
- D. Wired Fixture protocol

Answer: C

Q62. Which IEEE standard defines Wi-Fi?

- A. IEEE 802.16
- B. IEEE 802.11
- C. IEEE 802.15
- D. IEEE 802.3

Answer: B

Q63. What frequency band does Bluetooth primarily operate on?

- A. 900 MHz
- B. 60 GHz
- C. 5 GHz
- D. 2.4 GHz

Answer: D

Q64. What is an Access Point (AP) in a wireless network?

- A. A type of directional antenna used
- B. A wireless client device on the network
- C. A device bridging wireless to wired networks
- D. A wireless router with no other function

Answer: C

Q65. What is the typical range of Bluetooth?

- A. Up to 100 meters
- B. Up to 10 meters
- C. Up to 1 kilometer
- D. Up to 1 meter

Answer: B

Q66. What does SSID stand for in wireless networking?

- A. Secure Service Identification
- B. System Signal ID
- C. Service Set Identifier
- D. Standard Service ID

Answer: C

Q67. Which generation of cellular networks introduced mobile broadband data?

- A. 3G
- B. 1G
- C. 4G
- D. 2G

Answer: A

Q68. What is the hidden station problem in wireless networks?

- A. A station that is invisible on the network map
- B. A station with no assigned IP address
- C. A station using strong encryption protocol
- D. Two stations that cannot hear each other cause collisions

Answer: D

Q69. Which wireless technology is designed for very short-range communication used in contactless payments?

- A. WiMAX standard
- B. Wi-Fi standard
- C. Bluetooth Low Energy
- D. NFC technology

Answer: D

Q70. What type of network does ZigBee create?

- A. Satellite network for remote areas
- B. Wide Area Network for cities
- C. Cellular network for mobiles
- D. Low-power mesh network for IoT

Answer: D

Q71. What is cloud computing?

- A. A type of weather forecasting software tool
- B. On-demand delivery of resources via Internet
- C. Local server-based computing on premises
- D. Computing done using aerial drones

Answer: B

Q72. What does IaaS stand for?

- A. Infrastructure as a Service
- B. Information as a Service
- C. Integration as a Service
- D. Internet as a Service

Answer: A

Q73. Which of the following is an example of SaaS?

- A. Microsoft Azure VMs
- B. Google Gmail
- C. Docker containers
- D. Amazon EC2

Answer: B

Q74. What is a public cloud?

- A. A cloud hosted within company premises
- B. A cloud available to the general public
- C. A cloud with no security measures
- D. A cloud available to government only

Answer: B

Q75. What is the main benefit of cloud computing's pay-as-you-go model?

- A. Fixed monthly cost regardless of usage
- B. Paying only for resources you use
- C. No need for any internet access
- D. Unlimited free resources for everyone

Answer: B

Q76. What does PaaS provide?

- A. A platform for developing and deploying apps
- B. Network cables and hardware equipment
- C. Physical servers for data centers
- D. Pre-built software applications for users

Answer: A

Q77. What is a private cloud?

- A. Cloud infrastructure for a single organization
- B. The same thing as a public cloud
- C. A cloud with no internet connection
- D. A cloud with no active users at all

Answer: A

Q78. Which of the following is a major public cloud provider?

- A. Linux operating system
- B. Amazon Web Services (AWS)
- C. Cisco Routers hardware
- D. Microsoft Office suite

Answer: B

Q79. What is virtualization in cloud computing?

- A. Creating virtual versions of resources
- B. Creating virtual reality environments
- C. A type of encryption algorithm
- D. Making data invisible to users

Answer: A

Q80. What is scalability in cloud computing?

- A. The ability to adjust resources on demand
- B. The ability to weigh resources precisely
- C. The speed of the network connection
- D. The physical size of a data center

Answer: A

Q81. What is a firewall?

- A. A hardware component inside a router
- B. A type of computer virus or malware
- C. A device that monitors and controls traffic
- D. A wall constructed from fire materials

Answer: C

Q82. What does encryption do?

- A. Speeds up data transmission rate
- B. Deletes data from the storage
- C. Converts plaintext into ciphertext
- D. Compresses data to save space

Answer: C

Q83. What is a VPN?

- A. A virtual public network for browsing
- B. A very private network with no encryption
- C. A secure encrypted tunnel over public networks
- D. A virtual protocol network for routing

Answer: C

Q84. What is a DDoS attack?

- A. A type of computer virus infection
- B. A data decryption technique for ciphertext
- C. A routing protocol for network traffic
- D. Overwhelming a target with traffic from many sources

Answer: D

Q85. What is the purpose of an IDS (Intrusion Detection System)?

- A. To detect and alert on suspicious activity
- B. To assign IP addresses to devices
- C. To block all incoming network traffic
- D. To encrypt data during transmission

Answer: A

Q86. What is phishing?

- A. A social engineering attack to steal data
- B. A feature built into firewalls
- C. A network communication protocol
- D. A type of recreational fishing activity

Answer: A

Q87. What does SSL/TLS provide?

- A. Secure encrypted communication
- B. Faster network speeds overall
- C. IP address assignment to hosts
- D. File compression for storage

Answer: A

Q88. What is malware?

- A. A type of firewall software
- B. Malicious software designed to harm
- C. Defective hardware components
- D. A network topology design

Answer: B

Q89. What is the difference between symmetric and asymmetric encryption?

- A. Symmetric encryption is always much stronger
- B. They are the same encryption method
- C. Asymmetric uses a single shared key only
- D. Symmetric uses one key; asymmetric uses a key pair

Answer: D

Q90. What is a digital certificate?

- A. An electronic document binding key to identity
- B. A type of password for user accounts
- C. A printed diploma for graduates
- D. An IP address assigned to a server

Answer: A

Q91. What does IoT stand for?

- A. Internal Operating Terminal
- B. Integration of Technology
- C. Internet of Transactions
- D. Internet of Things

Answer: D

Q92. What is 5G?

- A. Fifth generation of mobile networks
- B. A type of Wi-Fi standard
- C. A network security protocol
- D. A cloud storage service

Answer: A

Q93. What is IPv6 primarily designed to address?

- A. Wireless communication needs
- B. The exhaustion of IPv4 addresses
- C. Better security features only
- D. Faster packet routing speeds

Answer: B

Q94. What is a smart city?

- A. A city with many computer stores
- B. A city with a very high population
- C. A virtual city in a video game
- D. A city using IoT to enhance services

Answer: D

Q95. What is the main advantage of fiber optic networks over copper networks?

- A. Better physical flexibility for bending
- B. Easier installation and maintenance work
- C. Much higher bandwidth and longer distance
- D. Lower cost of installation and materials

Answer: C

Q96. What does SDN stand for?

- A. Standard Digital Network
- B. System Data Node
- C. Secure Data Network
- D. Software-Defined Networking

Answer: D

Q97. What is a sensor network?

- A. A network of human sensor operators
- B. A type of local area network design
- C. Distributed sensor nodes monitoring conditions
- D. A social media network platform

Answer: C

Q98. What is Voice over IP (VoIP)?

- A. A routing protocol for packets
- B. Voice communications delivered over IP networks
- C. A type of network hardware device
- D. A video streaming service platform

Answer: B

Q99. What is network convergence?

- A. Networks colliding with each other
- B. The end of all networking technology
- C. Integrating voice, video, data on one network
- D. A type of router configuration mode

Answer: C

Q100. What is Li-Fi?

- A. A type of local area network
- B. A battery charging technology
- C. A fiber optic cable standard
- D. A technology using light for data transfer

Answer: D

Q101. What does the term 'throughput' measure in data communication?

- A. The maximum cable length allowed
- B. The voltage level of a channel
- C. The actual data transfer rate achieved
- D. The number of errors in a signal

Answer: C

Q102. Which type of cable uses light to transmit data signals?

- A. Coaxial cable
- B. Fiber optic
- C. Twisted pair
- D. Ethernet cable

Answer: B

Q103. What is the function of a modem in data communication?

- A. It compresses files before sending
- B. It connects two LANs wirelessly
- C. It modulates and demodulates signals
- D. It filters spam from email traffic

Answer: C

Q104. Which of the following is an unguided transmission medium?

- A. Twisted pair
- B. Fiber optic
- C. Radio waves
- D. Coaxial cable

Answer: C

Q105. What does 'latency' mean in the context of data communication?

- A. The speed of signal modulation
- B. The delay before data transfer begins
- C. The frequency range of a channel
- D. The size of the data payload sent

Answer: B

Q106. In data communication, what does 'noise' refer to?

- A. A type of data encryption method
- B. Unwanted signals that distort data
- C. The volume of audio transmissions
- D. The speed of packet forwarding

Answer: B

Q107. Which topology connects every device to a central hub?

- A. Ring topology
- B. Mesh topology
- C. Bus topology
- D. Star topology

Answer: D

Q108. What does 'attenuation' mean in signal transmission?

- A. Reflection of signal at endpoints
- B. Loss of signal strength over distance
- C. Increase in signal frequency
- D. Conversion of analog to digital

Answer: B

Q109. What is the role of a repeater in communication?

- A. It encrypts data for security
- B. It assigns IP addresses to hosts
- C. It routes packets between networks
- D. It amplifies and regenerates signals

Answer: D

Q110. Which encoding scheme represents data using two voltage levels?

- A. Unipolar encoding
- B. Analog encoding
- C. Multilevel encoding
- D. Bipolar encoding

Answer: A

Q111. What does LAN stand for in computer networking?

- A. Local Area Network
- B. Linear Array Network
- C. Linked Access Node
- D. Large Area Network

Answer: A

Q112. Which device connects multiple networks and routes packets?

- A. Router
- B. Repeater
- C. Bridge
- D. Hub

Answer: A

Q113. What is the primary function of a network switch?

- A. It forwards frames based on MAC address
- B. It assigns IP addresses to devices
- C. It provides wireless connectivity only
- D. It converts digital signals to analog

Answer: A

Q114. How many layers are in the OSI reference model?

- A. Six
- B. Four
- C. Seven
- D. Five

Answer: C

Q115. Which layer of the OSI model handles end-to-end delivery?

- A. Application layer
- B. Transport layer
- C. Network layer
- D. Session layer

Answer: B

Q116. What type of network covers a large geographic area like a country?

- A. WAN
- B. LAN
- C. PAN
- D. MAN

Answer: A

Q117. Which cable type is commonly used in Ethernet LANs?

- A. Coaxial cable only
- B. Satellite links
- C. Fiber optic only
- D. Twisted pair copper

Answer: D

Q118. What is a MAC address used for in networking?

- A. Identifying a device on the internet globally
- B. Encrypting data during wireless transmission
- C. Assigning a domain name to a web server
- D. Uniquely identifying a device on a local network

Answer: D

Q119. Which OSI layer converts data into a format suitable for transmission?

- A. Transport layer
- B. Application layer
- C. Session layer
- D. Presentation layer

Answer: D

Q120. What does a hub do when it receives a data frame?

- A. It broadcasts the frame to all connected ports
- B. It drops the frame if the destination is unknown
- C. It sends the frame only to the destination port
- D. It stores and forwards based on IP addresses

Answer: A

Q121. What is the primary function of the data link layer?

- A. End-to-end process communication
- B. Formatting data for display
- C. Routing packets across networks
- D. Reliable node-to-node data transfer

Answer: D

Q122. What is a frame in the context of the data link layer?

- A. A network routing table entry
- B. A data unit at the data link layer
- C. A type of physical cable connector
- D. A wireless signal modulation type

Answer: B

Q123. Which sublayer of the data link layer controls access to the shared medium?

- A. Media Access Control
- B. Physical Control
- C. Network Control
- D. Logical Link Control

Answer: A

Q124. What is the purpose of a parity bit in error detection?

- A. It compresses the data payload
- B. It encrypts the frame contents
- C. It determines the routing path used
- D. It detects single-bit transmission errors

Answer: D

Q125. What does CRC stand for in error detection?

- A. Central Routing Control
- B. Carrier Rate Conversion
- C. Cyclic Redundancy Check
- D. Circuit Relay Connection

Answer: C

Q126. Which framing method uses a special flag byte to mark frame boundaries?

- A. Clock encoding
- B. Bit stuffing
- C. Character count
- D. Byte stuffing

Answer: D

Q127. What does the LLC sublayer provide in the data link layer?

- A. Physical signal encoding methods
- B. Flow control and error management
- C. IP address assignment to devices
- D. Wireless frequency band selection

Answer: B

Q128. In Ethernet, what is the standard maximum frame size?

- A. 512 bytes
- B. 2048 bytes
- C. 4096 bytes
- D. 1518 bytes

Answer: D

Q129. What happens when a frame fails an error check at the receiver?

- A. The frame is stored indefinitely for later use
- B. The frame is forwarded to the next router
- C. The frame is typically discarded by receiver
- D. The frame is automatically corrected in place

Answer: C

Q130. Which protocol is a well-known data link layer protocol for point-to-point connections?

- A. FTP
- B. PPP
- C. DNS
- D. TCP

Answer: B

Q131. What is the main function of the network layer?

- A. Routing packets from source to destination
- B. Encrypting data for secure transfer
- C. Establishing physical connections
- D. Managing user application interfaces

Answer: A

Q132. Which protocol is primarily used for logical addressing?

- A. TCP
- B. IP
- C. ARP
- D. UDP

Answer: B

Q133. How many bits are in an IPv4 address?

- A. 48 bits
- B. 32 bits
- C. 16 bits
- D. 64 bits

Answer: B

Q134. What does a subnet mask define in IP networking?

- A. The maximum speed of data transfer
- B. The encryption key for data packets
- C. The network and host portions of an IP
- D. The physical address of a network card

Answer: C

Q135. What is the purpose of a default gateway?

- A. It blocks all unauthorized network traffic
- B. It stores all website data on the network
- C. It assigns IP addresses to local devices
- D. It forwards packets to networks outside the LAN

Answer: D

Q136. Which device operates primarily at the network layer?

- A. Router
- B. Hub
- C. Repeater
- D. Switch

Answer: A

Q137. What does ICMP stand for?

- A. Internet Connected Media Protocol
- B. Internet Control Message Protocol
- C. Internal Cache Management Protocol
- D. Integrated Circuit Module Process

Answer: B

Q138. What is the purpose of the TTL field in an IP packet?

- A. To limit the packet's lifetime in network
- B. To encrypt the packet data payload
- C. To identify the transport layer protocol
- D. To store the destination port number

Answer: A

Q139. What command is commonly used to test network connectivity?

- A. ping
- B. tracert
- C. netstat
- D. ipconfig

Answer: A

Q140. What is the loopback address in IPv4?

- A. 192.168.0.1
- B. 255.255.255.0
- C. 10.0.0.1
- D. 127.0.0.1

Answer: D

Q141. What is the main function of the transport layer?

- A. Routing across networks
- B. Physical signal transmission
- C. Data link error detection
- D. End-to-end communication

Answer: D

Q142. Which transport protocol provides reliable data delivery?

- A. UDP
- B. TCP
- C. ARP
- D. ICMP

Answer: B

Q143. What does UDP stand for in networking?

- A. Universal Direct Process
- B. Unverified Data Pathway
- C. Unified Data Protocol
- D. User Datagram Protocol

Answer: D

Q144. What is a port number used for in the transport layer?

- A. Specifying the data encoding format used
- B. Determining the routing path for packets
- C. Identifying a specific application or process
- D. Identifying the physical cable connection used

Answer: C

Q145. Which well-known port number is used by HTTP?

- A. Port 80
- B. Port 53
- C. Port 21
- D. Port 25

Answer: A

Q146. What is a socket in networking terminology?

- A. A physical cable connector on a device
- B. A combination of IP address and port number
- C. A type of wireless antenna for signals
- D. A protocol for error detection in frames

Answer: B

Q147. What type of service does UDP provide?

- A. Connectionless unreliable delivery
- B. Connection-oriented reliable delivery
- C. Connectionless reliable delivery
- D. Connection-oriented unreliable delivery

Answer: A

Q148. What is segmentation in the transport layer?

- A. Routing packets through different network paths
- B. Dividing data into smaller manageable units
- C. Encrypting data before network transmission
- D. Combining multiple packets into one large block

Answer: B

Q149. Which port number is used by HTTPS?

- A. Port 443
- B. Port 110
- C. Port 8080
- D. Port 143

Answer: A

Q150. Which protocol is preferred for real-time video streaming?

- A. TCP
- B. UDP
- C. FTP
- D. SMTP

Answer: B

Q151. What does DNS stand for in networking?

- A. Direct Network Signal
- B. Digital Node Switch
- C. Data Network Service
- D. Domain Name System

Answer: D

Q152. Which protocol is used to send email?

- A. DNS
- B. HTTP
- C. SMTP
- D. FTP

Answer: C

Q153. What is the primary function of FTP?

- A. Sending email messages
- B. Transferring files between hosts
- C. Browsing web pages online
- D. Resolving domain names

Answer: B

Q154. Which protocol is used to retrieve email from a server?

- A. POP3
- B. SNMP
- C. HTTP
- D. SMTP

Answer: A

Q155. What does HTTP stand for?

- A. HyperText Transfer Protocol
- B. Hyper Terminal Transfer Process
- C. High Throughput Test Platform
- D. Host Transfer Text Protocol

Answer: A

Q156. What is the purpose of the DHCP protocol?

- A. Resolving domain name queries
- B. Transferring files to a server
- C. Assigning IP addresses automatically
- D. Encrypting web page content

Answer: C

Q157. Which application layer protocol is used for remote login?

- A. Telnet
- B. DNS
- C. SMTP
- D. FTP

Answer: A

Q158. What is a URL in the context of the web?

- A. A protocol for securing email traffic
- B. A type of network cable connector
- C. The address used to locate web resources
- D. A device that manages web server load

Answer: C

Q159. What does SNMP stand for?

- A. Standard Network Module Platform
- B. Simple Network Management Protocol
- C. System Network Message Protocol
- D. Secure Node Monitoring Process

Answer: B

Q160. Which protocol provides secure web communication?

- A. HTTP
- B. SMTP
- C. HTTPS
- D. FTP

Answer: C

Q161. What does WiFi stand for?

- A. Wave Finder
- B. Wired Filter
- C. Wide Frequency
- D. Wireless Fidelity

Answer: D

Q162. Which IEEE standard defines wireless LAN technology?

- A. 802.16
- B. 802.11
- C. 802.3
- D. 802.15

Answer: B

Q163. What device provides wireless network access to clients?

- A. Modem
- B. Repeater
- C. Router
- D. Access Point

Answer: D

Q164. Which frequency band is commonly used by WiFi networks?

- A. 900 MHz
- B. 5 MHz
- C. 2.4 GHz
- D. 100 GHz

Answer: C

Q165. What is Bluetooth primarily used for?

- A. Short-range device communication
- B. Long-range internet connectivity
- C. Underwater data transmission
- D. Satellite television broadcast

Answer: A

Q166. What does SSID represent in a wireless network?

- A. The channel frequency used by the router
- B. The MAC address of the access point device
- C. The name that identifies a wireless network
- D. The encryption key for network access

Answer: C

Q167. Which generation of mobile networks introduced mobile broadband?

- A. 3G
- B. 4G
- C. 1G
- D. 2G

Answer: A

Q168. What is the main advantage of wireless networks over wired?

- A. Greater mobility for users
- B. Higher data transfer speeds
- C. Stronger signal over distances
- D. Lower electromagnetic interference

Answer: A

Q169. What type of wireless network has the shortest range?

- A. LAN
- B. WAN
- C. PAN
- D. MAN

Answer: C

Q170. What does NFC stand for in wireless technology?

- A. New Frequency Channel
- B. Network File Control
- C. Near Field Communication
- D. Node Forwarding Cache

Answer: C

Q171. What is cloud computing in simple terms?

- A. Connecting computers using only Bluetooth signals
- B. Building physical servers in an office building
- C. Delivering computing services over the internet
- D. Storing data on local hard drives only

Answer: C

Q172. What does IaaS stand for in cloud computing?

- A. Infrastructure as a Service
- B. Integration as a System
- C. Internet as a Service
- D. Information as a Standard

Answer: A

Q173. Which cloud service model provides ready-to-use software applications?

- A. DaaS
- B. SaaS
- C. PaaS
- D. IaaS

Answer: B

Q174. What is a public cloud?

- A. A cloud that uses only open source software tools
- B. A cloud available only to a single organization
- C. A cloud that stores only public government records
- D. A cloud accessible by any user over the internet

Answer: D

Q175. What is a private cloud?

- A. A social media platform with privacy settings
- B. A cloud that is not connected to any network
- C. A cloud dedicated to a single organization's use
- D. A personal computer connected to the internet

Answer: C

Q176. What is a virtual machine in cloud computing?

- A. A software emulation of a computer system
- B. A networking cable connecting two servers
- C. A type of database management software tool
- D. A physical server in a data center building

Answer: A

Q177. Which company provides the AWS cloud platform?

- A. Microsoft
- B. Amazon
- C. Apple
- D. Google

Answer: B

Q178. What is scalability in cloud computing?

- A. The ability to encrypt all stored data securely
- B. The ability to increase or decrease resources as needed
- C. The process of backing up data to external storage
- D. The speed at which data travels between cloud servers

Answer: B

Q179. What does PaaS provide to developers?

- A. Complete ready-to-use software for end users
- B. Network monitoring and security scanning tools
- C. A platform for building and deploying applications
- D. Only physical hardware servers and cables

Answer: C

Q180. What is a hybrid cloud deployment model?

- A. A cloud that stores data in multiple file formats
- B. A combination of public and private cloud services
- C. A cloud that only runs on mobile phone devices
- D. A cloud that uses only renewable energy sources

Answer: B

Q181. What is the primary purpose of a firewall?

- A. Filtering traffic between networks
- B. Assigning IP addresses to hosts
- C. Compressing data for faster transfer
- D. Increasing network bandwidth speed

Answer: A

Q182. What does encryption do to data?

- A. It increases data transfer speed on network
- B. It routes data through the fastest path
- C. It compresses data to save storage space
- D. It converts data into an unreadable format

Answer: D

Q183. What is a denial-of-service (DoS) attack?

- A. Stealing passwords from a database
- B. Modifying web pages without permission
- C. Flooding a system to make it unavailable
- D. Intercepting emails during transmission

Answer: C

Q184. What is the purpose of an antivirus program?

- A. It manages user login credentials securely
- B. It configures firewall rules automatically
- C. It detects and removes malicious software
- D. It speeds up network data transfers

Answer: C

Q185. What is phishing in cybersecurity?

- A. Installing updates on network equipment
- B. Scanning networks to find open ports
- C. Fraudulent attempts to steal sensitive info
- D. Monitoring bandwidth usage on a network

Answer: C

Q186. What does VPN stand for?

- A. Variable Protocol Number
- B. Visual Processing Node
- C. Virtual Private Network
- D. Verified Public Network

Answer: C

Q187. What is malware?

- A. A type of network cable connector
- B. A protocol for secure email transfer
- C. A tool for monitoring network traffic
- D. Software designed to cause system harm

Answer: D

Q188. What is two-factor authentication (2FA)?

- A. Using two different web browsers for access
- B. Requiring two forms of identity verification
- C. Creating two copies of encrypted file data
- D. Connecting through two separate network paths

Answer: B

Q189. What is the role of a security patch?

- A. It changes the user interface design layout
- B. It increases the software processing speed
- C. It adds new features to the software
- D. It fixes known security vulnerabilities

Answer: D

Q190. What does HTTPS protect against?

- A. Eavesdropping on web communication traffic
- B. Physical theft of computer hardware devices
- C. Power surges that damage server components
- D. Hard drive failures that cause data loss

Answer: A

Q191. What does IoT stand for in modern networking?

- A. Internet of Things
- B. Integrated Online Technology
- C. Internal Operating Terminal
- D. Internet of Transfers

Answer: A

Q192. What is 5G in terms of mobile networks?

- A. A satellite communication protocol for space
- B. The fifth version of WiFi technology
- C. A new type of fiber optic cable standard
- D. The fifth generation of mobile network technology

Answer: D

Q193. What is the purpose of smart home technology?

- A. To connect homes directly to satellite networks
- B. To replace all furniture with electronic versions
- C. To automate and remotely control home devices
- D. To eliminate the need for electricity in homes

Answer: C

Q194. What is a sensor network?

- A. A cable network connecting television channels
- B. A network exclusively for video surveillance only
- C. A network of high-speed gaming computers
- D. A group of sensors that collect and share data

Answer: D

Q195. What is artificial intelligence used for in networks?

- A. Designing the physical layout of data center rooms
- B. Automating network management and threat detection
- C. Manufacturing network hardware in smart factories
- D. Replacing all network cables with wireless signals

Answer: B

Q196. What is a smart city concept?

- A. A city that uses only renewable energy exclusively
- B. A city with no physical stores or shopping malls
- C. A city without any public transportation systems
- D. A city using technology and data to improve services

Answer: D

Q197. What does the term 'bandwidth on demand' mean?

- A. Restricting bandwidth during peak usage time periods
- B. Adjusting bandwidth allocation based on real-time needs
- C. Free unlimited internet for all users forever
- D. Permanently increasing bandwidth for all connections

Answer: B

Q198. What is the role of IPv6 in next-generation networks?

- A. It eliminates the need for routers in all networks
- B. It provides a vastly larger address space for devices
- C. It converts all networks to use satellite connections
- D. It replaces WiFi with a faster wireless technology

Answer: B

Q199. What is telemedicine in the context of networks?

- A. A type of network monitoring software for hospitals
- B. A network used exclusively by telephone companies
- C. Remote healthcare delivery using network technology
- D. A method of diagnosing network connection problems

Answer: C

Q200. What is a digital twin in next-generation technology?

- A. A backup copy of a website on a different server
- B. A virtual replica of a physical object or system
- C. A duplicate email account for security purposes
- D. A second network interface card on a computer

Answer: B

Q201. Which type of data transmission sends data in only one direction at a time?

- A. Multiplex transmission
- B. Half-duplex transmission
- C. Simplex transmission
- D. Full-duplex transmission

Answer: B

Q202. What is the smallest unit of data in digital communication?

- A. A single bit representing 0 or 1
- B. A word consisting of 16 bits
- C. A nibble consisting of 4 bits
- D. A byte consisting of 8 bits

Answer: A

Q203. Which device converts digital signals to analog signals for transmission over telephone lines?

- A. A repeater amplifying weak signals
- B. A router directing network traffic
- C. A modem performing signal conversion
- D. A network switch forwarding frames

Answer: C

Q204. What does the term 'baud rate' measure in data communication?

- A. The total bandwidth of a communication link
- B. The number of bytes transferred per minute
- C. The error rate during data transmission
- D. The number of signal changes per second

Answer: D

Q205. Which transmission medium offers the highest bandwidth capacity?

- A. Coaxial copper cable medium
- B. Twisted pair copper cable
- C. Fiber optic cable medium
- D. Unshielded telephone wire

Answer: C

Q206. What is the purpose of a communication protocol?

- A. To define rules governing data exchange
- B. To manufacture networking equipment
- C. To increase the voltage of signals
- D. To physically connect cables to devices

Answer: A

Q207. What is crosstalk in the context of data communication?

- A. Planned interference for signal testing
- B. A method of increasing data capacity
- C. The process of encoding digital data
- D. Unwanted signal coupling between wires

Answer: D

Q208. In serial data transmission, how are bits sent over the communication channel?

- A. Bits are split across multiple channels
- B. Bits are sent one after another sequentially
- C. Multiple bits are sent simultaneously
- D. Bits are grouped in fixed-size blocks

Answer: B

Q209. What does the term 'propagation delay' refer to in networking?

- A. Time to encrypt and decrypt data packets
- B. Time to convert data between formats
- C. Time to process packets inside a router
- D. Time for a signal to travel from source to destination

Answer: D

Q210. Which signal type uses discrete voltage levels to represent data?

- A. Acoustic signal through air medium
- B. Analog signal with continuous waves
- C. Digital signal with distinct levels
- D. Modulated carrier frequency signal

Answer: C

Q211. What is the primary function of a network interface card (NIC)?

- A. To connect a device to a computer network
- B. To store data temporarily during transfers
- C. To route packets between different networks
- D. To provide wireless power to portable devices

Answer: A

Q212. Which network topology arranges all devices in a circular chain?

- A. Mesh topology with full connections
- B. Bus topology with shared backbone
- C. Star topology with central device
- D. Ring topology with circular connection

Answer: D

Q213. What is a Metropolitan Area Network (MAN)?

- A. A network connecting countries across continents
- B. A network spanning a single room or building floor
- C. A personal network within a few meters range
- D. A network covering a city or large campus area

Answer: D

Q214. What does a gateway do in computer networking?

- A. It stores web pages for faster user access
- B. It connects networks using different protocols
- C. It monitors bandwidth usage on local links
- D. It amplifies signals on a single network

Answer: B

Q215. Which layer of the OSI model is responsible for data formatting and encryption?

- A. The session layer managing connections
- B. The transport layer handling segments
- C. The presentation layer handling format
- D. The application layer for user access

Answer: C

Q216. What is a peer-to-peer (P2P) network model?

- A. A model where each device can act as client and server
- B. A model requiring a central database for all data
- C. A model where one server manages all resources
- D. A model using only wireless communication links

Answer: A

Q217. What does bandwidth refer to in computer networking?

- A. The total storage available on a file server
- B. The maximum data transfer capacity of a link
- C. The number of devices connected to a network
- D. The physical length of a network cable

Answer: B

Q218. What is the client-server model in networking?

- A. A model where data is stored exclusively on client machines
- B. A model where devices share resources equally without hierarchy
- C. A model where clients request services from a dedicated server
- D. A model where only one device is connected at any given time

Answer: C

Q219. What does the physical layer of the OSI model deal with?

- A. Logical addressing and routing decisions
- B. Session establishment between applications
- C. Error detection within transmitted frames
- D. Raw bit transmission over physical media

Answer: D

Q220. What is a network protocol stack?

- A. A layered set of protocols working together
- B. A single protocol handling all network tasks
- C. A hardware component inside every router
- D. A physical rack holding network equipment

Answer: A

Q221. What is the purpose of flow control at the data link layer?

- A. To encrypt data before physical transmission
- B. To prevent the sender from overwhelming the receiver
- C. To route frames between different networks
- D. To assign IP addresses to connected devices

Answer: B

Q222. What is the stop-and-wait protocol in data link communication?

- A. The sender transmits all frames at once before waiting
- B. The receiver stops all incoming traffic when busy
- C. The sender waits for a timer before each transmission
- D. The sender sends one frame and waits for acknowledgment

Answer: D

Q223. What does ARQ stand for in data link layer protocols?

- A. Automatic Repeat reQuest for error recovery
- B. Asynchronous Relay Queue for data buffering
- C. Address Resolution Queue for frame delivery
- D. Adaptive Rate Quantization for bandwidth use

Answer: A

Q224. What is the broadcast MAC address used to reach all devices on a network?

- A. 00:00:00:00:00:00 representing zero address
- B. 01:00:00:00:00:00 representing multicast
- C. FF:FF:FF:FF:FF:FF representing all ones
- D. FE:FE:FE:FE:FE:FE representing link local

Answer: C

Q225. What type of error occurs when a single bit in a frame is changed during transmission?

- A. A burst error affecting multiple bits
- B. A framing error at the boundary mark
- C. A single-bit error in the data frame
- D. A sequence error in the frame order

Answer: C

Q226. What is the purpose of the destination address field in an Ethernet frame?

- A. To identify the intended recipient device
- B. To indicate the protocol type being used
- C. To contain error-checking information data
- D. To specify the sender's network location

Answer: A

Q227. What is the role of acknowledgment frames in reliable data delivery?

- A. They establish the initial connection setup
- B. They confirm successful receipt of data frames
- C. They carry user data from sender to receiver
- D. They request retransmission of all sent frames

Answer: B

Q228. What is a collision in the context of Ethernet networks?

- A. When a cable is physically disconnected from a port
- B. When a device sends data to an unknown MAC address
- C. When a frame exceeds the maximum allowed frame size
- D. When two devices transmit data simultaneously on shared media

Answer: D

Q229. What is the function of the type/length field in an Ethernet frame?

- A. It contains the frame check sequence value
- B. It indicates the protocol or payload length
- C. It stores the MAC address of the next hop
- D. It specifies the transmission speed to use

Answer: B

Q230. What does the term 'full-duplex Ethernet' mean?

- A. Data can only flow in one direction at a time
- B. Data is transmitted at double the normal speed
- C. Data can flow in both directions simultaneously
- D. Data is duplicated across two separate cables

Answer: C

Q231. What is the purpose of an IP address in networking?

- A. To specify the application protocol being used
- B. To logically identify a device on a network
- C. To determine the cable type connecting devices
- D. To identify the physical hardware of a device

Answer: B

Q232. What is the difference between a public and a private IP address?

- A. Public addresses are assigned by DHCP while private addresses are manually configured
- B. Private addresses use IPv6 while public addresses always use IPv4 format
- C. Private addresses are used within local networks while public addresses are on the internet
- D. Public addresses are used on local networks while private addresses are on the internet

Answer: C

Q233. What is the default subnet mask for a Class A IPv4 address?

- A. 255.255.255.0 with 24 network bits
- B. 255.255.255.128 with 25 network bits
- C. 255.0.0.0 with 8 network bits only
- D. 255.255.0.0 with 16 network bits

Answer: C

Q234. What is a routing table used for in a router?

- A. To determine the best path for forwarding packets
- B. To filter web content based on user preferences
- C. To encrypt all outgoing traffic at the gateway
- D. To store the physical locations of all devices

Answer: A

Q235. Which class of IPv4 addresses has a first octet range of 128 to 191?

- A. Class D addresses for multicast groups
- B. Class C addresses for small networks
- C. Class A addresses for large networks
- D. Class B addresses for medium networks

Answer: D

Q236. What is static routing in a network?

- A. Routes change dynamically based on current traffic load
- B. Routes are manually configured by a network administrator
- C. Routes are computed using the shortest path algorithm
- D. Routes are automatically learned from neighboring routers

Answer: B

Q237. What does the traceroute (tracert) command show?

- A. The MAC addresses of all devices on a local network
- B. The list of open ports on a destination host
- C. The bandwidth available on each network interface
- D. The path packets take through routers to a destination

Answer: D

Q238. What is an IPv4 header's minimum size in bytes?

- A. 20 bytes without any option fields
- B. 32 bytes without any option fields
- C. 10 bytes without any option fields
- D. 40 bytes without any option fields

Answer: A

Q239. What type of routing protocol learns routes automatically from other routers?

- A. Direct routing for locally connected subnets
- B. Default routing using a single gateway path
- C. Dynamic routing using protocol exchanges
- D. Static routing configured manually by admin

Answer: C

Q240. What is the purpose of the network layer's fragmentation mechanism?

- A. To combine small packets into larger ones for efficiency
- B. To split packets that exceed a link's maximum transfer unit
- C. To compress large packets for faster transmission speed
- D. To encrypt packet contents for secure delivery to hosts

Answer: B

Q241. What is the primary difference between TCP and UDP?

- A. TCP provides reliable delivery while UDP is connectionless
- B. TCP works only on wired networks while UDP works wirelessly
- C. TCP is faster while UDP provides guaranteed delivery
- D. UDP uses larger packets while TCP uses smaller packets

Answer: A

Q242. What is a well-known port number?

- A. A port used exclusively by operating system processes
- B. A dynamically assigned port used for temporary connections
- C. A port in the range 0-1023 reserved for standard services
- D. Any port above 49151 used by client applications

Answer: C

Q243. Which port number does DNS typically use?

- A. Port 23 for domain name queries
- B. Port 53 for domain name queries
- C. Port 67 for domain name queries
- D. Port 110 for domain name queries

Answer: B

Q244. What is multiplexing at the transport layer?

- A. Converting data formats between different application types
- B. Splitting a single application across multiple networks
- C. Encrypting traffic from a single application for security
- D. Combining data from multiple applications for transmission

Answer: D

Q245. What flag is used in the TCP handshake to initiate a connection?

- A. The ACK flag for acknowledgment setup
- B. The SYN flag for synchronization request
- C. The FIN flag for connection finalization
- D. The RST flag for connection reset action

Answer: B

Q246. What application layer service typically uses TCP port 25?

- A. Email sending using SMTP protocol
- B. Secure shell remote login to servers
- C. Web browsing using HTTP protocol
- D. File transfer using FTP data channel

Answer: A

Q247. What is the purpose of sequence numbers in TCP?

- A. To measure the round-trip time for each packet
- B. To encrypt the data payload of each segment
- C. To identify the source application on the host
- D. To track the order of transmitted data bytes

Answer: D

Q248. Why is UDP preferred for live video streaming applications?

- A. UDP guarantees all video frames arrive in order
- B. UDP provides better encryption for video content
- C. UDP has lower overhead and tolerates some data loss
- D. UDP supports larger maximum packet sizes than TCP

Answer: C

Q249. What is demultiplexing at the transport layer?

- A. Converting between different transport protocols dynamically
- B. Combining data from many applications into one stream
- C. Delivering received data to the correct application via port numbers
- D. Splitting data across multiple network paths for speed

Answer: C

Q250. What is an ephemeral port in networking?

- A. A port used exclusively for encrypted communication
- B. A permanently assigned port for a specific service
- C. A reserved port used only by system administrators
- D. A temporary port assigned to a client for a session

Answer: D

Q251. What is a web browser used for in networking?

- A. To monitor bandwidth usage across all links
- B. To access and display web pages from servers
- C. To manage database records on the network
- D. To configure router settings and interfaces

Answer: B

Q252. What port does POP3 use for email retrieval?

- A. Port 143 for email message retrieval
- B. Port 993 for email message retrieval
- C. Port 25 for email message retrieval
- D. Port 110 for email message retrieval

Answer: D

Q253. What is the purpose of TFTP (Trivial File Transfer Protocol)?

- A. To securely transfer files using encrypted communication channels
- B. To transfer large files with error correction and authentication
- C. To provide simple file transfers without authentication features
- D. To synchronize file versions across distributed storage systems

Answer: C

Q254. What does a DNS A record store?

- A. The IPv4 address mapped to a domain name
- B. The alias or canonical name for a domain
- C. The mail server address for the domain
- D. The authoritative nameserver for the domain

Answer: A

Q255. What is the role of an email client application?

- A. It routes email traffic between different mail servers
- B. It converts email protocols from SMTP to HTTP format
- C. It filters spam at the server before delivery to users
- D. It allows users to compose, send, and read email messages

Answer: D

Q256. What does the HTTP GET method request from a server?

- A. It requests to create a new resource on the server
- B. It requests to retrieve a resource from the server
- C. It requests to update an existing resource on server
- D. It requests to delete a resource from the server

Answer: B

Q257. What is IMAP used for in email communication?

- A. Accessing and managing email directly on the server
- B. Sending email messages between mail servers
- C. Filtering incoming email for spam and viruses
- D. Encrypting email content during transmission

Answer: A

Q258. What type of DNS record is used for email routing?

- A. A PTR record for reverse DNS mail lookups
- B. An A record pointing to the mail server IP
- C. A CNAME record aliasing the mail server name
- D. An MX record specifying the mail exchange server

Answer: D

Q259. What does NTP (Network Time Protocol) do?

- A. It transfers files between networked computers
- B. It translates network addresses between formats
- C. It synchronizes clocks across networked devices
- D. It monitors network traffic for security threats

Answer: C

Q260. What is the function of a proxy server?

- A. It assigns MAC addresses to new network devices
- B. It acts as an intermediary between clients and servers
- C. It generates encryption keys for secure connections
- D. It directly connects two LANs using physical cables

Answer: B

Q261. What is the purpose of a wireless router in a home network?

- A. To replace the ISP's connection to the internet
- B. To provide wireless and wired connectivity to the internet
- C. To boost cellular signal strength in buildings
- D. To connect only wired devices to the internet

Answer: B

Q262. What is the 5 GHz frequency band used for in WiFi?

- A. It provides higher data speeds with shorter range
- B. It is used only for outdoor wireless deployments
- C. It is reserved exclusively for Bluetooth devices
- D. It provides longer range but lower data speeds

Answer: A

Q263. What does WPA stand for in wireless security?

- A. Wide Protocol Authentication mechanism
- B. Wireless Protected Architecture standard
- C. Wireless Packet Authorization framework
- D. Wi-Fi Protected Access security protocol

Answer: D

Q264. What is a wireless hotspot?

- A. An area where wireless signal is intentionally blocked
- B. A tool for testing wireless network security
- C. A location providing public wireless internet access
- D. A device that measures wireless signal strength

Answer: C

Q265. What is the main purpose of WEP encryption in wireless networks?

- A. To assign IP addresses to wireless client devices
- B. To increase the speed of wireless data transfer
- C. To manage handoff between wireless access points
- D. To provide basic encryption for wireless communication

Answer: D

Q266. Which generation of mobile networks is primarily known for voice calls?

- A. 4G networks with LTE high-speed service
- B. 2G networks with digital voice capability
- C. 5G networks with ultra-low latency service
- D. 3G networks with mobile broadband data

Answer: B

Q267. What is a wireless repeater (range extender) used for?

- A. To extend the coverage area of an existing wireless network
- B. To secure the wireless network against unauthorized access
- C. To create a new wireless network with different settings
- D. To reduce interference between neighboring wireless networks

Answer: A

Q268. What does the term 'wireless channel' refer to in WiFi?

- A. The cable connecting an access point to a switch
- B. The physical antenna used to transmit signals
- C. A specific frequency range for wireless communication
- D. The security encryption method used by the router

Answer: C

Q269. What type of wireless network does IoT typically use for low-power devices?

- A. Standard WiFi 6 enterprise access points
- B. High-speed 5G millimeter wave connections
- C. Low-power wide-area networks like LoRaWAN
- D. Satellite-based global broadband connections

Answer: C

Q270. What is the role of an antenna in wireless communication?

- A. It encrypts the wireless signal for security
- B. It stores data while waiting for transmission slots
- C. It assigns channels to different wireless clients
- D. It converts electrical signals to radio waves and vice versa

Answer: D

Q271. What is a cloud data center?

- A. A personal computer storing data at home
- B. A facility housing cloud computing resources and servers
- C. A local hard drive connected to a network
- D. A mobile application for storing contacts

Answer: B

Q272. What is an example of Infrastructure as a Service (IaaS)?

- A. Amazon EC2 providing virtual server instances
- B. Google Docs for document editing online
- C. Salesforce for customer relationship management
- D. Microsoft Office 365 for productivity suite

Answer: A

Q273. What is cloud storage?

- A. Data backed up on a local external hard drive
- B. Data stored on remote servers accessed via the internet
- C. Data stored on physical USB drives in an office
- D. Data cached temporarily in a web browser's memory

Answer: B

Q274. What does the term 'on-demand' mean in cloud computing?

- A. Resources must be ordered weeks before they are needed
- B. Resources are shared equally among all users at all times
- C. Resources are available immediately when users need them
- D. Resources are allocated permanently to a single customer

Answer: C

Q275. What is Microsoft Azure?

- A. A cloud computing platform providing various cloud services
- B. A desktop operating system for personal computers
- C. A programming language for web application development
- D. A database management software for local installations

Answer: A

Q276. What is the main advantage of elasticity in cloud computing?

- A. It provides permanent storage that never changes size
- B. It eliminates the need for any internet connectivity
- C. It restricts user access to fixed resource allocations
- D. It automatically adjusts resources based on current demand

Answer: D

Q277. What is Google Cloud Platform (GCP)?

- A. A web browser developed by Google for internet access
- B. A cloud computing service suite offered by Google
- C. A mobile operating system developed by Google team
- D. A social media platform operated by Google company

Answer: B

Q278. What is multi-cloud computing?

- A. Using cloud services from two or more cloud providers
- B. Storing multiple copies of data in one cloud storage bucket
- C. Running multiple applications on one cloud virtual machine
- D. Using multiple data centers from a single cloud provider

Answer: A

Q279. What is a community cloud deployment model?

- A. A cloud shared by the general public without restrictions
- B. A cloud owned and operated by a single private company
- C. A cloud used exclusively by government defense agencies
- D. A cloud shared by organizations with common requirements

Answer: D

Q280. What is a cloud API?

- A. A monitoring tool for tracking cloud resource utilization
- B. A security certificate used to access cloud dashboards
- C. An interface allowing applications to interact with cloud services
- D. A physical connector between cloud data center servers

Answer: C

Q281. What is a strong password policy?

- A. Rules requiring complex, lengthy, and unique passwords
- B. Sharing passwords among trusted team members
- C. Using short passwords for easy memorization
- D. Changing passwords only when a breach is detected

Answer: A

Q282. What is social engineering in the context of cybersecurity?

- A. Building social media platforms for secure communication
- B. Manipulating people into revealing confidential information
- C. Designing network architectures for social media companies
- D. Programming chatbots for automated customer support

Answer: B

Q283. What is the purpose of a security audit?

- A. To install new software on all company computers
- B. To permanently block all external network connections
- C. To increase the bandwidth of the network infrastructure
- D. To evaluate and assess the security posture of a system

Answer: D

Q284. What does the term 'ransomware' refer to?

- A. A tool for managing network access control lists
- B. Software that speeds up your computer's performance
- C. Malware that encrypts data and demands payment for the key
- D. A backup solution that stores copies of important files

Answer: C

Q285. What is the function of an access control list (ACL) on a network device?

- A. It manages the allocation of IP addresses on the local network
- B. It lists all users who have physical access to the server room
- C. It records all login attempts to the network administration panel
- D. It defines rules to permit or deny network traffic on interfaces

Answer: D

Q286. What is a brute force attack?

- A. Systematically trying all possible passwords until finding the correct one
- B. Physically breaking into a data center facility
- C. Flooding a server with traffic to cause denial of service
- D. Intercepting network traffic between two communicating devices

Answer: A

Q287. What is the purpose of network monitoring in security?

- A. To replace firewalls with a more modern security solution
- B. To detect suspicious activities and potential security threats
- C. To increase the speed of all network connections
- D. To permanently store all data passing through the network

Answer: B

Q288. What is a security vulnerability?

- A. A planned feature in software for administrator access
- B. A strength in the system's defense against cyber threats
- C. A weakness that can be exploited to compromise security
- D. A configuration setting that improves network performance

Answer: C

Q289. What is the role of a certificate authority (CA)?

- A. To manage physical access to data center facilities
- B. To issue and verify digital certificates for entities
- C. To monitor network traffic for security threats
- D. To create firewall rules for network perimeter defense

Answer: B

Q290. What does the term 'authentication' mean in network security?

- A. The process of verifying the identity of a user or device
- B. The process of compressing data for faster transmission
- C. The process of routing packets through the fastest path
- D. The process of encrypting data before sending it

Answer: A

Q291. What is an autonomous vehicle in the context of next-generation networks?

- A. A server that operates independently without network connectivity
- B. A self-driving vehicle that relies on network communication for operation
- C. A drone used to deploy networking cables in remote locations
- D. A robotic system that manufactures network equipment automatically

Answer: B

Q292. What is augmented reality (AR) in the context of network applications?

- A. A technique that replaces the real world with a virtual environment
- B. A server architecture used for rendering cloud-based video games
- C. A network protocol designed for transferring 3D model data files
- D. A technology that overlays digital information onto the real world view

Answer: D

Q293. What is the concept of smart grid in next-generation networking?

- A. An intelligent electrical grid that uses network communication for management
- B. A high-performance computing grid for scientific research
- C. A display technology for showing network monitoring dashboards
- D. A network topology where all devices are arranged in a grid pattern

Answer: A

Q294. What is the Internet of Everything (IoE)?

- A. A faster version of the internet with unlimited bandwidth
- B. An extension of IoT connecting people, processes, data, and things
- C. A government initiative to provide internet access everywhere
- D. A protocol for connecting every device using IPv6 addresses

Answer: B

Q295. What is the purpose of wearable technology in next-generation networks?

- A. To replace desktop computers for all computing tasks
- B. To install network equipment without using any tools
- C. To protect network cables from environmental damage
- D. To provide personal health, fitness, and connectivity on the body

Answer: D

Q296. What is Industry 4.0 in the context of next-generation networks?

- A. The fourth major update to enterprise networking equipment
- B. The fourth version of the internet protocol standard
- C. The fourth industrial revolution driven by smart manufacturing and IoT
- D. The fourth generation of mobile wireless communication systems

Answer: C

Q297. What is a drone network in next-generation communication?

- A. A mesh network connecting satellites in geostationary orbit
- B. Unmanned aerial vehicles providing aerial communication coverage
- C. A network of underground cables for secure data transmission
- D. A cluster of servers processing drone surveillance video footage

Answer: B

Q298. What is the concept of connected healthcare?

- A. A social media platform for healthcare professionals to connect
- B. A VPN service designed specifically for medical record access
- C. Using networked devices and systems to improve healthcare delivery
- D. Connecting hospital buildings with high-speed fiber optic cables

Answer: C

Q299. What is a satellite internet constellation?

- A. A group of ground-based antennas pointed at a single satellite
- B. A star topology network named after celestial patterns
- C. A server cluster designed to process astronomical research data
- D. A large fleet of satellites providing global internet coverage

Answer: D

Q300. What is the role of 5G in enabling the Internet of Things?

- A. 5G provides high capacity, low latency, and massive device connectivity for IoT
- B. 5G eliminates the need for any IoT gateway or edge devices
- C. 5G is used only for smartphone communications not IoT devices
- D. 5G replaces WiFi entirely for all IoT device connections

Answer: A

Computer Networks & Cloud Computing — Medium

300 questions

Q301. According to the Nyquist theorem, the maximum data rate of a noiseless channel with bandwidth B and L signal levels is:

- A. $2B \log_2(L)$
- B. $2B / \log_2(L)$
- C. $B \log_2(L)$
- D. $B^2 \log_2(L)$

Answer: A

Q302. What is the Shannon capacity formula for a noisy channel?

- A. $C = B * \text{SNR}$
- B. $C = B \log_2(1 + \text{SNR})$
- C. $C = 2B \log_2(\text{SNR})$
- D. $C = B / \log_2(\text{SNR})$

Answer: B

Q303. Which multiplexing technique divides the bandwidth of a link into frequency bands?

- A. CDM
- B. FDM
- C. TDM
- D. WDM

Answer: B

Q304. In PCM (Pulse Code Modulation), what are the three steps in order?

- A. Encoding, Sampling, Quantization
- B. Quantization, Sampling, Encoding
- C. Sampling, Quantization, Encoding
- D. Sampling, Encoding, Quantization

Answer: C

Q305. What type of encoding uses the transition at the middle of a bit interval to represent data?

- A. Unipolar encoding
- B. Manchester encoding
- C. NRZ-I encoding
- D. NRZ-L encoding

Answer: B

Q306. If a signal-to-noise ratio (SNR) is 1000, what is the SNR in decibels?

- A. 20 dB
- B. 30 dB
- C. 10 dB
- D. 40 dB

Answer: B

Q307. Which type of transmission sends bits one at a time over a single channel?

- A. Serial transmission
- B. Parallel transmission
- C. Synchronous transmission
- D. Isochronous transmission

Answer: A

Q308. In asynchronous serial transmission, what is the purpose of start and stop bits?

- A. Compressing the data being sent
- B. Detecting transmission errors in data
- C. Encrypting the transmitted payload
- D. Synchronizing sender and receiver clocks

Answer: D

Q309. Wavelength Division Multiplexing (WDM) is conceptually similar to which multiplexing technique?

- A. TDM
- B. FDM
- C. CDM
- D. STDM

Answer: B

Q310. What is throughput in data communication?

- A. The delay in delivering packets
- B. Number of bits in each signal
- C. Maximum capacity of a channel
- D. Actual data rate through a point

Answer: D

Q311. In a full mesh topology with n devices, how many links are required?

- A. n
- B. n-1
- C. $n(n-1)/2$
- D. n^2

Answer: C

Q312. Which layer of the OSI model handles data encryption and compression?

- A. Presentation layer
- B. Application layer
- C. Session layer
- D. Transport layer

Answer: A

Q313. What is the main advantage of packet switching over circuit switching?

- A. It has simpler implementation
- B. It uses network resources efficiently
- C. It guarantees fixed bandwidth
- D. It provides much lower latency

Answer: B

Q314. Which TCP/IP layer corresponds to the OSI network layer?

- A. Network access layer
- B. Internet layer
- C. Application layer
- D. Transport layer

Answer: B

Q315. What type of switching does ATM (Asynchronous Transfer Mode) use?

- A. Message switching
- B. Cell switching
- C. Packet switching
- D. Circuit switching

Answer: B

Q316. Which organization is responsible for defining the Ethernet standard (IEEE 802.3)?

- A. W3C
- B. IETF
- C. IEEE
- D. ITU

Answer: C

Q317. What is the function of the session layer in the OSI model?

- A. Detecting errors in frames
- B. Encrypting transmitted data
- C. Managing dialog between systems
- D. Routing packets across networks

Answer: C

Q318. In packet switching, what is the difference between datagram and virtual circuit approaches?

- A. Datagrams always use fixed-size data packets
- B. Virtual circuits do not include packet headers
- C. Virtual circuits set up a path first; datagrams do not
- D. Datagrams offer lower latency than virtual circuits

Answer: C

Q319. A hub operates at which layer of the OSI model?

- A. Network layer
- B. Transport layer
- C. Data link layer
- D. Physical layer

Answer: D

Q320. What is the difference between a bridge and a switch?

- A. Bridges forward packets faster than switches do
- B. Bridges work at Layer 3; switches work at Layer 2
- C. There is no functional difference between them
- D. Switches handle multiple frames at once; bridges process one

Answer: D

Q321. What is the purpose of bit stuffing in HDLC?

- A. To prevent the flag pattern in data
- B. To compress the transmitted data
- C. To encrypt the frame contents
- D. To increase available bandwidth

Answer: A

Q322. In the Go-Back-N ARQ protocol, what happens when an error is detected in frame n?

- A. The entire connection is terminated
- B. Only frame n is retransmitted alone
- C. All frames from n onward are resent
- D. Frame n is skipped and ignored

Answer: C

Q323. What is the window size at the receiver in the Go-Back-N protocol?

- A. Same as sender window size
- B. Equal to N minus 1
- C. Always exactly 1 only
- D. Equal to $2^{(m-1)}$

Answer: C

Q324. In Selective Repeat ARQ, what is the maximum window size if the sequence number has m bits?

- A. $2^{(m-1)}$
- B. 2^m
- C. $2^m - 1$
- D. $2^m + 1$

Answer: A

Q325. What is the purpose of the LLC sublayer?

- A. Collision detection on the wire
- B. Signal encoding and modulation
- C. Physical addressing of devices
- D. Flow and error control for frames

Answer: D

Q326. What is the minimum frame size in standard Ethernet (IEEE 802.3)?

- A. 64 bytes
- B. 46 bytes
- C. 32 bytes
- D. 128 bytes

Answer: A

Q327. Which access method does IEEE 802.11 (Wi-Fi) use?

- A. CSMA/CD method
- B. Token passing
- C. CSMA/CA method
- D. FDMA technique

Answer: C

Q328. In the Hamming code, how many redundancy bits are needed for a 4-bit data word?

- A. 3
- B. 2
- C. 5
- D. 4

Answer: A

Q329. What is a VLAN (Virtual LAN)?

- A. A logical group of devices on a network
- B. A wireless LAN using access points
- C. A LAN without any switch devices
- D. A LAN that uses virtual machine hosts

Answer: A

Q330. What is the exponential backoff algorithm used for in Ethernet?

- A. Increasing the link bandwidth
- B. Performing error correction tasks
- C. Determining wait time after collision
- D. Assigning MAC addresses to NICs

Answer: C

Q331. In CIDR notation, what does /28 mean?

- A. 28 hosts are available on subnet
- B. 28 bits are for network addressing
- C. 28 bits are for host addressing
- D. 28 subnets exist in the network

Answer: B

Q332. How many usable host addresses are available in a /26 subnet?

- A. 64
- B. 30
- C. 62
- D. 126

Answer: C

Q333. Which routing protocol uses the Bellman-Ford algorithm?

- A. RIP
- B. OSPF
- C. BGP
- D. IS-IS

Answer: A

Q334. What is the maximum hop count in RIP?

- A. 255
- B. 16
- C. 15
- D. 8

Answer: C

Q335. OSPF is classified as which type of routing protocol?

- A. Path vector
- B. Link-state
- C. Distance vector
- D. Hybrid

Answer: B

Q336. What is the size of an IPv6 address?

- A. 64 bits
- B. 32 bits
- C. 128 bits
- D. 96 bits

Answer: C

Q337. What is the count-to-infinity problem associated with?

- A. Link-state protocols
- B. Distance vector protocols
- C. Static routing tables
- D. Path vector protocols

Answer: B

Q338. Which field in the IPv4 header is used for fragmentation reassembly?

- A. Header checksum field
- B. Identification field
- C. Protocol field
- D. TTL field

Answer: B

Q339. What does DHCP provide?

- A. Automatic IP address assignment
- B. Email delivery management
- C. Routing information exchange
- D. Domain name resolution services

Answer: A

Q340. Which protocol is used between autonomous systems on the Internet?

- A. RIP
- B. EIGRP
- C. BGP
- D. OSPF

Answer: C

Q341. In TCP, what is the purpose of the window size field?

- A. Controlling flow by advertising buffer space
- B. Specifying the maximum segment size used
- C. Detecting errors in transmitted segments
- D. Identifying the application layer protocol

Answer: A

Q342. What is the TCP mechanism for handling out-of-order segments?

- A. Buffering them indefinitely in memory
- B. Requesting retransmission of all data
- C. Discarding all out-of-order segments
- D. Using sequence numbers to reorder them

Answer: D

Q343. What is the purpose of the TCP FIN flag?

- A. To synchronize the sequence numbers
- B. To establish a new connection
- C. To reset an existing connection
- D. To gracefully terminate a connection

Answer: D

Q344. How many steps are involved in graceful TCP connection termination?

- A. 5
- B. 2
- C. 4
- D. 3

Answer: C

Q345. What is Nagle's algorithm in TCP?

- A. A routing algorithm for IP packets
- B. A congestion control mechanism for TCP
- C. An algorithm that buffers small segments
- D. An error correction algorithm for data

Answer: C

Q346. What is the difference between TCP's flow control and congestion control?

- A. Congestion control is faster than flow control always
- B. Flow control is used in UDP; congestion control is for TCP
- C. Flow control protects receiver; congestion control protects network
- D. They are the same mechanism in TCP

Answer: C

Q347. What port number does HTTPS use?

- A. 143
- B. 443
- C. 80
- D. 993

Answer: B

Q348. What is the maximum size of a TCP segment's data payload if the MSS is 1460 bytes?

- A. 65535 bytes
- B. 1500 bytes
- C. 536 bytes
- D. 1460 bytes

Answer: D

Q349. What is the size of the UDP header?

- A. 12 bytes
- B. 4 bytes
- C. 20 bytes
- D. 8 bytes

Answer: D

Q350. In TCP, what does a RST (Reset) flag indicate?

- A. Acknowledgment of data received
- B. Normal graceful termination
- C. An abrupt connection abort
- D. A request to synchronize

Answer: C

Q351. What is the difference between HTTP/1.1 persistent and non-persistent connections?

- A. Persistent connections use UDP instead of TCP
- B. Persistent keeps TCP open for multiple requests; non-persistent does not
- C. There is no difference between them at all
- D. Non-persistent connections are faster than persistent ones

Answer: B

Q352. What type of DNS record maps a domain name to an IPv4 address?

- A. AAAA record
- B. A record
- C. MX record
- D. CNAME record

Answer: B

Q353. What is the difference between FTP active and passive modes?

- A. Active mode uses TCP; passive mode uses UDP for transfer
- B. Active mode is encrypted; passive mode is not encrypted
- C. There is no functional difference between the two modes
- D. In active mode server initiates data; in passive mode client does

Answer: D

Q354. What HTTP status code indicates 'Not Found'?

- A. 403
- B. 200
- C. 301
- D. 404

Answer: D

Q355. What is the purpose of DNS caching?

- A. To create backup DNS server entries
- B. To block malicious domain name requests
- C. To encrypt DNS queries for privacy
- D. To store resolved names for faster lookups

Answer: D

Q356. Which protocol is used for secure file transfer?

- A. Telnet
- B. HTTP
- C. SFTP
- D. FTP

Answer: C

Q357. What is SNMP used for?

- A. Transferring files
- B. Browsing web pages
- C. Sending emails to users
- D. Managing network devices

Answer: D

Q358. In HTTP, what is the difference between GET and POST methods?

- A. GET is more secure than POST method
- B. GET retrieves data via URL; POST sends data in body
- C. POST is always faster than GET requests
- D. They are identical in every respect

Answer: B

Q359. What is a MIME type used for in email?

- A. Filtering out spam messages
- B. Routing emails between servers
- C. Defining the content type format
- D. Encrypting email messages

Answer: C

Q360. What port does SMTP use?

- A. 110
- B. 21
- C. 25
- D. 80

Answer: C

Q361. Why does Wi-Fi use CSMA/CA instead of CSMA/CD?

- A. Collision detection is impractical in wireless
- B. CSMA/CA uses significantly less power
- C. CSMA/CD is not defined for any wireless
- D. CSMA/CA provides faster throughput overall

Answer: A

Q362. What is the purpose of RTS/CTS mechanism in IEEE 802.11?

- A. To increase the available bandwidth
- B. To solve the hidden station problem
- C. To assign IP addresses to clients
- D. To encrypt the transmitted data frames

Answer: B

Q363. What frequency bands does IEEE 802.11n (Wi-Fi 4) support?

- A. 60 GHz band only
- B. Both 2.4 and 5 GHz
- C. 2.4 GHz band only
- D. 5 GHz band only

Answer: B

Q364. What technology does 4G LTE use for its air interface?

- A. FDMA
- B. OFDMA
- C. TDMA
- D. CDMA

Answer: B

Q365. What is MIMO in wireless networking?

- A. A type of signal modulation used
- B. Multiple antennas to boost throughput
- C. A wireless security protocol standard
- D. A specific radio frequency band

Answer: B

Q366. What is the difference between infrastructure mode and ad-hoc mode in Wi-Fi?

- A. They are the same wireless mode
- B. Ad-hoc mode requires an access point to function
- C. Infrastructure mode is slower than ad-hoc mode
- D. Infrastructure uses an AP; ad-hoc is direct device-to-device

Answer: D

Q367. What is handoff (handover) in cellular networks?

- A. Changing the phone number on a SIM
- B. Turning off the mobile phone device
- C. Transferring a call between cell towers
- D. Connecting a device to a Wi-Fi network

Answer: C

Q368. What is the maximum theoretical speed of IEEE 802.11ac (Wi-Fi 5)?

- A. 9.6 Gbps
- B. 1.3 Gbps
- C. 600 Mbps
- D. 6.9 Gbps

Answer: D

Q369. What does WPA3 improve over WPA2?

- A. Faster wireless data speeds overall
- B. Stronger encryption and resistance to dictionary attacks
- C. Greater wireless signal range coverage
- D. Improved battery life for all devices

Answer: B

Q370. What is frequency reuse in cellular networks?

- A. Assigning same frequencies to non-adjacent cells
- B. Changing frequencies every few seconds
- C. Using only one frequency for all traffic
- D. Using the same frequency in every cell

Answer: A

Q371. What is a hybrid cloud?

- A. A cloud running both Windows and Linux
- B. A mix of public and private clouds together
- C. A cloud without any virtualization layer
- D. A cloud powered by solar energy only

Answer: B

Q372. What is the difference between horizontal and vertical scaling?

- A. Vertical scaling is always the better choice
- B. They are the same scaling approach
- C. Horizontal adds machines; vertical adds power
- D. Horizontal scaling is slower than vertical

Answer: C

Q373. What is a hypervisor?

- A. A network communication protocol
- B. Software that creates and runs VMs
- C. A type of network firewall
- D. A cloud storage service platform

Answer: B

Q374. What is serverless computing?

- A. A model where providers manage servers for you
- B. A type of dedicated hardware platform
- C. Computing without any servers at all
- D. Computing without needing Internet access

Answer: A

Q375. What is the difference between a container and a virtual machine?

- A. Containers share the OS kernel; VMs include full OS
- B. Containers require a hypervisor to operate
- C. They are the same technology in practice
- D. VMs are lighter weight than containers are

Answer: A

Q376. What is multi-tenancy in cloud computing?

- A. Running multiple operating systems at once
- B. Using multiple different cloud providers
- C. Multiple customers sharing the same resources
- D. Having multiple data center locations

Answer: C

Q377. What is an SLA in cloud computing?

- A. A server location address identifier
- B. A software license agreement for vendors
- C. A system load analyzer monitoring tool
- D. A contract defining service quality guarantees

Answer: D

Q378. What is cloud migration?

- A. Moving workloads from on-premises to cloud
- B. Deleting cloud resources permanently
- C. Moving birds through cloud formations
- D. Changing between different cloud providers

Answer: A

Q379. What is object storage in cloud computing?

- A. Storing physical objects in warehouses
- B. A traditional relational database type
- C. Managing data as objects with metadata
- D. RAM-based temporary data storage

Answer: C

Q380. What is auto-scaling in cloud computing?

- A. Permanently reducing allocated resources
- B. Manual scaling done by administrators
- C. Automatically adjusting resources with demand
- D. Scaling only during business hours

Answer: C

Q381. What is the difference between IDS and IPS?

- A. IPS only detects without blocking
- B. IDS actively blocks all traffic
- C. They are identical in function
- D. IDS only alerts; IPS also blocks threats

Answer: D

Q382. What is a man-in-the-middle (MITM) attack?

- A. A brute force password cracking attack
- B. A type of network firewall configuration
- C. An attacker intercepting communication between parties
- D. A physical attack on network cables

Answer: C

Q383. What is ARP spoofing?

- A. Repairing corrupted ARP table entries
- B. A DNS poisoning attack on resolvers
- C. A routing protocol attack on BGP routers
- D. Sending fake ARPs to link attacker MAC to valid IP

Answer: D

Q384. What is the purpose of a DMZ (Demilitarized Zone) in network security?

- A. A restricted military zone on base
- B. A backup network for disaster recovery
- C. A type of encryption algorithm used
- D. A segment separating internal from external networks

Answer: D

Q385. What is port scanning?

- A. A type of data encryption technique
- B. A method to permanently close all ports
- C. Physically inspecting network port hardware
- D. Discovering open ports and services on a target

Answer: D

Q386. What is the purpose of IPSec?

- A. To speed up IP packet routing
- B. To replace TCP with a new protocol
- C. To assign IP addresses to devices
- D. To provide security services at IP layer

Answer: D

Q387. What is the difference between stateful and stateless firewalls?

- A. Stateless firewalls are more secure overall
- B. Stateful tracks connections; stateless filters packets individually
- C. They are the same type of firewall
- D. Stateful firewalls cannot track any connections

Answer: B

Q388. What is a honeypot in network security?

- A. A sweet treat left for hackers to find
- B. A type of firewall with extra logging
- C. An encryption algorithm for secure data
- D. A decoy system to attract and study attackers

Answer: D

Q389. What is the function of a WAF (Web Application Firewall)?

- A. Filtering email messages for spam
- B. Protecting web apps by filtering HTTP traffic
- C. Encrypting web pages for secure storage
- D. Blocking all web traffic on the network

Answer: B

Q390. What is two-factor authentication (2FA)?

- A. Logging into a system twice in succession
- B. Using two different usernames for access
- C. Using two different passwords to log in
- D. Requiring two identification factors to verify identity

Answer: D

Q391. What are the three main use cases of 5G defined by ITU?

- A. Voice, video, and data transfer
- B. Speed, security, and storage capacity
- C. eMBB, URLLC, and mMTC
- D. LAN, WAN, and MAN types

Answer: C

Q392. What is the role of MQTT protocol in IoT?

- A. A lightweight messaging protocol for IoT
- B. Transferring large files between devices
- C. Sending emails from IoT sensors
- D. Browsing websites on IoT devices

Answer: A

Q393. What is OpenFlow in the context of SDN?

- A. A distance vector routing protocol
- B. A type of network firewall
- C. A cloud computing service platform
- D. A protocol between SDN controller and switches

Answer: D

Q394. What is fog computing?

- A. Computing in foggy weather conditions
- B. Computing without any electricity source
- C. A type of cloud storage for backups
- D. Extending cloud computing to the network edge

Answer: D

Q395. What is the main purpose of LPWAN (Low-Power Wide-Area Network)?

- A. Replacing Wi-Fi in home networks
- B. Long-range low-power IoT communication
- C. Streaming high-definition video content
- D. High-speed online gaming connections

Answer: B

Q396. What is an IMS (IP Multimedia Subsystem)?

- A. An internet monitoring system tool
- B. A type of network router device
- C. A framework for multimedia over IP networks
- D. An instant messaging service for users

Answer: C

Q397. What is digital twin technology in the context of networks?

- A. Two identical physical networks side by side
- B. A virtual replica for simulation and monitoring
- C. A backup network for disaster recovery
- D. A VPN technology for secure tunnels

Answer: B

Q398. What is the concept of network automation?

- A. A type of network switch hardware
- B. Networks that never require any maintenance
- C. Using software to auto-configure network devices
- D. Manual network configuration by engineers

Answer: C

Q399. What is Content-Centric Networking (CCN) / Named Data Networking (NDN)?

- A. A paradigm focusing on content not host location
- B. A social networking platform for users
- C. A caching technology for web browsers only
- D. A type of domain name resolution system

Answer: A

Q400. What is Multi-access Edge Computing (MEC)?

- A. Cloud capabilities at the mobile network edge
- B. A type of cloud storage for backups
- C. A Wi-Fi standard for edge networks
- D. Computing on multiple user devices at once

Answer: A

Q401. In Manchester encoding, how is a binary 1 typically represented?

- A. A low-to-high voltage transition at mid-bit
- B. A constant high voltage for the full bit
- C. A high-to-low voltage transition at mid-bit
- D. A constant low voltage for the full bit

Answer: A

Q402. What is the primary advantage of differential Manchester encoding over standard Manchester?

- A. It uses less bandwidth overall
- B. It eliminates the need for any synchronization
- C. It offers better noise immunity via transitions
- D. It provides higher data transfer rate

Answer: C

Q403. Which multiplexing technique assigns unique codes to each channel?

- A. Code Division Multiplexing
- B. Time Division Multiplexing
- C. Wavelength Division Multiplexing
- D. Frequency Division Multiplexing

Answer: A

Q404. What is the main difference between synchronous and asynchronous transmission?

- A. Synchronous uses analog; asynchronous uses digital signals
- B. Synchronous is slower; asynchronous is always faster
- C. Synchronous uses wireless; asynchronous uses wired media
- D. Synchronous uses frames; asynchronous sends individual bytes

Answer: D

Q405. A channel has a bandwidth of 4 kHz and SNR of 31. What is the Shannon capacity?

- A. 24,000 bps
- B. 16,000 bps
- C. 32,000 bps
- D. 20,000 bps

Answer: D

Q406. What causes intersymbol interference (ISI) in digital communication?

- A. Excessive signal amplification at receiver end
- B. Too many repeaters placed along the channel
- C. Spreading of pulses that overlap adjacent symbols
- D. Use of fiber optic cables instead of copper

Answer: C

Q407. In TDM, what happens when a time slot's source has no data to send?

- A. The slot is transmitted empty as a placeholder
- B. The slot is given to another active source
- C. The slot is filled with error correction bits
- D. The entire frame transmission is paused

Answer: A

Q408. What is the purpose of a constellation diagram in digital modulation?

- A. It maps network topology connections visually
- B. It shows signal amplitude and phase combinations
- C. It illustrates the physical layout of cables
- D. It displays the frequency spectrum of a signal

Answer: B

Q409. What is the key advantage of statistical TDM over synchronous TDM?

- A. It dynamically allocates slots to active sources
- B. It provides lower latency for all sources
- C. It removes the need for framing information
- D. It uses wider bandwidth for faster transfers

Answer: A

Q410. What does QAM combine to create multiple signal variations?

- A. Compression and encryption simultaneously
- B. Digital and analog encoding simultaneously
- C. Amplitude and phase modulation together
- D. Frequency and wavelength modulation together

Answer: C

Q411. What is the key difference between a bridge and a switch?

- A. A bridge uses IP addresses; a switch uses MAC addresses
- B. A bridge is wireless; a switch requires wired connections
- C. A bridge works in software; a switch uses hardware ASICs
- D. A bridge operates at Layer 3; a switch at Layer 2 only

Answer: C

Q412. In the TCP/IP model, which layer corresponds to the OSI session and presentation layers?

- A. Network interface layer
- B. Transport layer
- C. Application layer
- D. Internet layer

Answer: C

Q413. What is the purpose of ARP in computer networking?

- A. It assigns domain names to IP addresses
- B. It resolves IP addresses to MAC addresses
- C. It encrypts data at the network layer
- D. It routes packets between different subnets

Answer: B

Q414. What is a collision domain in Ethernet networking?

- A. A network area where firewalls block traffic
- B. A group of devices sharing a broadcast address
- C. A subnet with identical IP address configurations
- D. A segment where simultaneous transmissions can collide

Answer: D

Q415. What is the primary advantage of packet switching over circuit switching?

- A. It offers more efficient use of network bandwidth
- B. It provides guaranteed constant bandwidth to users
- C. It ensures zero packet loss during transmission
- D. It eliminates the need for routing protocols entirely

Answer: A

Q416. What does CSMA/CD stand for and where is it used?

- A. Circuit Switched Multiple Access / Channel Division in MANs
- B. Coded Signal Modulation Access / Carrier Demodulation in PANs
- C. Central Station Multiple Access / Conflict Detection in WANs
- D. Carrier Sense Multiple Access / Collision Detection in LANs

Answer: D

Q417. What is the function of the network layer in the OSI model?

- A. It formats data for presentation to applications
- B. It handles end-to-end error control between processes
- C. It manages logical addressing and packet routing
- D. It establishes and maintains sessions between hosts

Answer: C

Q418. What is the purpose of subnetting in IP networking?

- A. To increase the total number of IP addresses
- B. To divide a network into smaller logical groups
- C. To encrypt data between network segments
- D. To connect wireless and wired networks together

Answer: B

Q419. How does a VLAN differ from a physical LAN?

- A. A VLAN requires separate cables for each group
- B. A VLAN only works with wireless network connections
- C. A VLAN uses IP routing; a LAN uses MAC switching
- D. A VLAN groups devices logically regardless of location

Answer: D

Q420. What is the difference between a broadcast domain and a collision domain?

- A. Collision domains are larger; broadcast domains are per port
- B. Both terms describe exactly the same network concept
- C. Broadcast domains are larger; collision domains are per port
- D. Broadcast domains exist only in wireless network setups

Answer: C

Q421. How does the Go-Back-N ARQ protocol handle frame errors?

- A. It retransmits the error frame and all subsequent frames
- B. It requests the receiver to correct the frame locally
- C. It retransmits only the specific erroneous frame alone
- D. It ignores the error and continues with new frame data

Answer: A

Q422. What is the advantage of Selective Repeat ARQ over Go-Back-N?

- A. It eliminates the need for acknowledgment frames
- B. It only retransmits the specific frames with errors
- C. It provides stronger encryption for transmitted frames
- D. It uses simpler receiver buffer management logic

Answer: B

Q423. What is the purpose of bit stuffing in HDLC framing?

- A. To increase the data transmission speed on the link
- B. To add error correction bits throughout the payload
- C. To prevent data patterns from mimicking frame delimiter
- D. To compress the data payload for efficient transfer

Answer: C

Q424. In CSMA/CA, what does the device do before transmitting?

- A. It requests a dedicated channel from central controller
- B. It sends data immediately without any delay checks
- C. It broadcasts a test packet to check for collisions
- D. It senses the channel and waits if it is already busy

Answer: D

Q425. What is the sliding window protocol's main benefit?

- A. It encrypts frames during transit across network
- B. It converts analog signals to digital frame format
- C. It allows multiple frames in transit simultaneously
- D. It eliminates the need for error detection codes

Answer: C

Q426. What is the binary exponential backoff algorithm used for?

- A. Calculating the optimal frame size for a network
- B. Encoding binary data into Ethernet frame formats
- C. Selecting the fastest route between two endpoints
- D. Determining retransmission wait time after collisions

Answer: D

Q427. How does a learning bridge build its forwarding table?

- A. By examining the source MAC address of incoming frames
- B. By querying a central DHCP server for address mappings
- C. By reading a pre-configured static file at startup
- D. By broadcasting requests to all devices on the network

Answer: A

Q428. What is the purpose of the preamble in an Ethernet frame?

- A. It synchronizes the receiver's clock with the sender
- B. It stores the VLAN tag for traffic prioritization
- C. It carries the destination MAC address information
- D. It contains error correction data for the payload

Answer: A

Q429. In token ring networks, how is channel access managed?

- A. A central controller assigns fixed time slots to nodes
- B. Devices transmit whenever the channel appears free
- C. A special token frame circulates granting send rights
- D. Devices use frequency division to avoid any collisions

Answer: C

Q430. What is the role of the FCS field in a data link frame?

- A. It stores the fragmentation offset for large payloads
- B. It specifies the priority level for quality of service
- C. It identifies the network layer protocol in the payload
- D. It contains the frame check sequence for error detection

Answer: D

Q431. What is the difference between classful and classless IP addressing?

- A. Classful uses fixed boundaries; classless uses variable masks
- B. Classful uses variable masks; classless uses fixed boundaries
- C. Classful supports IPv6; classless only supports IPv4 format
- D. Classful is used in WANs; classless is used in LANs only

Answer: A

Q432. What does CIDR notation /24 mean for an IP address?

- A. The subnet can hold exactly 24 host IP addresses
- B. The first 24 bits are the host portion of address
- C. The address has a total length of 24 bits only
- D. The first 24 bits are the network portion of address

Answer: D

Q433. How does NAT help conserve IPv4 addresses?

- A. It maps multiple private IPs to fewer public IP addresses
- B. It compresses IP headers to use fewer address bits
- C. It assigns temporary addresses that expire after one use
- D. It converts IPv4 addresses directly to IPv6 format

Answer: A

Q434. What is the key difference between distance-vector and link-state routing?

- A. Distance-vector shares routing tables; link-state shares LSA info
- B. Distance-vector shares full topology; link-state shares hop counts
- C. Distance-vector is faster to converge than link-state protocols
- D. Distance-vector uses flooding; link-state uses periodic updates

Answer: A

Q435. What routing protocol uses the Dijkstra algorithm for path computation?

- A. BGP
- B. OSPF
- C. EIGRP
- D. RIP

Answer: B

Q436. What is the purpose of IP fragmentation?

- A. To split large packets to fit the link's maximum unit
- B. To encrypt packets for secure transmission over links
- C. To compress packet headers for faster router processing
- D. To prioritize certain packets over others in the queue

Answer: A

Q437. What is the count-to-infinity problem in distance-vector routing?

- A. Routers keep incrementing hop counts for unreachable routes
- B. Routers count packets infinitely without forwarding any data
- C. Routers exceed their maximum memory for routing table size
- D. The metric value overflows the integer storage in routers

Answer: A

Q438. How does split horizon prevent routing loops?

- A. It divides the routing table into two separate equal halves
- B. It splits the network into isolated routing zones entirely
- C. It uses two different routing protocols simultaneously always
- D. It never advertises a route back to the interface it learned from

Answer: D

Q439. What is the main advantage of IPv6 over IPv4?

- A. IPv6 only works with wireless network connections
- B. IPv6 uses shorter 64-bit addresses for simplicity
- C. IPv6 eliminates the need for any routing protocols
- D. IPv6 provides a vastly larger 128-bit address space

Answer: D

Q440. What is the function of the DHCP protocol?

- A. It dynamically assigns IP addresses to network hosts
- B. It translates MAC addresses into IP address format
- C. It monitors network bandwidth usage across subnets
- D. It encrypts DNS queries for privacy and security

Answer: A

Q441. What is the TCP three-way handshake used for?

- A. To assign port numbers to communicating processes
- B. To encrypt data before transmission starts
- C. To fragment packets for network layer processing
- D. To establish a reliable connection between hosts

Answer: D

Q442. What does the TCP window size control?

- A. The encryption key length for secure connections
- B. The maximum number of TCP connections allowed
- C. The amount of data sent before acknowledgment needed
- D. The physical size of network interface buffers

Answer: C

Q443. How does TCP handle out-of-order segment delivery?

- A. It buffers segments and reorders using sequence numbers
- B. It discards all out-of-order segments immediately
- C. It forwards segments to the application without reordering
- D. It requests the sender to retransmit the entire data stream

Answer: A

Q444. What is the purpose of the TCP checksum field?

- A. It determines the optimal window size for the flow
- B. It identifies the priority level of the data segment
- C. It counts the total number of segments transmitted
- D. It detects errors in the TCP segment during transfer

Answer: D

Q445. What triggers TCP fast retransmit?

- A. Detection of a CRC error in the segment received
- B. A request from the application layer to resend data
- C. Expiration of the retransmission timeout timer
- D. Receipt of three duplicate acknowledgment messages

Answer: D

Q446. What is the difference between flow control and congestion control?

- A. Flow control prevents routing loops; congestion prevents collisions
- B. Flow control is for UDP only; congestion control is TCP only
- C. Flow control manages sender speed; congestion manages network load
- D. Flow control operates at Layer 3; congestion at Layer 4 only

Answer: C

Q447. What is Nagle's algorithm designed to address?

- A. Increasing the throughput for large file transfers
- B. Optimizing the routing of packets across subnets
- C. Preventing unauthorized access to network services
- D. Reducing the overhead from many small TCP segments

Answer: D

Q448. What does the URG flag indicate in a TCP segment?

- A. The receiver should increase its window size
- B. The segment is a retransmission of lost data
- C. The segment contains high-priority urgent data
- D. The connection is being reset by the sender

Answer: C

Q449. Why is UDP preferred over TCP for DNS queries?

- A. UDP provides more reliable delivery for DNS resolution
- B. DNS requires encryption that only UDP provides
- C. UDP's low overhead is ideal for small DNS query packets
- D. DNS queries require the connection setup that UDP offers

Answer: C

Q450. What is a TCP RST segment used for?

- A. To advertise a larger receive window to sender
- B. To initiate a new TCP connection with a server
- C. To abruptly terminate an existing TCP connection
- D. To request retransmission of lost data segments

Answer: C

Q451. What is the difference between POP3 and IMAP for email?

- A. POP3 downloads and deletes; IMAP keeps mail on the server
- B. POP3 uses encryption; IMAP sends messages in plain text
- C. POP3 works offline only; IMAP requires constant connection
- D. POP3 is for sending; IMAP is for receiving email messages

Answer: A

Q452. How does DNS resolution typically work in a recursive query?

- A. The DNS resolver handles the entire lookup on behalf of client
- B. The root server returns the final IP address immediately
- C. The client broadcasts a request to all DNS servers at once
- D. The client queries each DNS server in the hierarchy directly

Answer: A

Q453. What is the purpose of DNS caching?

- A. To speed up repeated lookups by storing recent resolutions
- B. To permanently store all DNS records on every client
- C. To balance the load across multiple DNS server instances
- D. To encrypt DNS queries to prevent eavesdropping attacks

Answer: A

Q454. What is the role of MIME in email communication?

- A. It encrypts the email body for confidential communication
- B. It compresses email attachments to reduce storage space used
- C. It extends email to support non-text attachments and formats
- D. It routes email messages between different mail servers

Answer: C

Q455. What HTTP method is used to submit form data to a server?

- A. HEAD
- B. DELETE
- C. POST
- D. GET

Answer: C

Q456. What is the difference between active and passive FTP modes?

- A. Active is faster; passive provides more reliable file transfer
- B. In active mode the server connects back; in passive the client connects
- C. Active encrypts data; passive transfers data in plain text
- D. Active uses one port; passive uses two ports for data

Answer: B

Q457. What does an HTTP status code 404 indicate?

- A. The client needs to authenticate before accessing resource
- B. The server encountered an internal processing error
- C. The client request was successfully processed by server
- D. The requested resource was not found on the server

Answer: D

Q458. What is the purpose of a reverse proxy server?

- A. It converts HTTP requests into FTP transfer commands
- B. It sits before servers to handle requests on their behalf
- C. It caches web pages for faster client-side browsing
- D. It translates domain names into IP addresses for clients

Answer: B

Q459. How does a CDN improve web content delivery?

- A. It encrypts web traffic to prevent data interception
- B. It compresses all web content to reduce file sizes
- C. It blocks malicious requests before they reach servers
- D. It distributes content to edge servers closer to users

Answer: D

Q460. What is the function of the SSH protocol?

- A. It transfers files using unencrypted connections only
- B. It sends email messages between different mail servers
- C. It provides secure encrypted remote access and commands
- D. It resolves domain names to their IP addresses securely

Answer: C

Q461. What is the key difference between WiFi 5 (802.11ac) and WiFi 6 (802.11ax)?

- A. WiFi 5 uses 2.4 GHz only; WiFi 6 uses 5 GHz band only
- B. WiFi 5 supports MU-MIMO; WiFi 6 does not support it
- C. WiFi 6 introduces OFDMA for better multi-device efficiency
- D. WiFi 6 has shorter range but much higher power consumption

Answer: C

Q462. What security protocol replaced WEP in wireless networks?

- A. SSL
- B. SSH
- C. WPA
- D. IPSec

Answer: C

Q463. What is the purpose of channel bonding in WiFi?

- A. It combines adjacent channels to increase bandwidth
- B. It encrypts data across multiple wireless channels
- C. It reduces interference by separating signal channels
- D. It connects multiple access points into a single unit

Answer: A

Q464. What is MU-MIMO and what problem does it solve?

- A. It modulates signals to reduce power consumption overall
- B. It allows simultaneous transmission to multiple devices
- C. It manages user authentication for wireless network access
- D. It monitors unauthorized intrusions on the wireless network

Answer: B

Q465. How does 5G differ from 4G in terms of architecture?

- A. 5G uses circuit switching; 4G uses packet switching
- B. 5G only works with stationary devices, not mobile ones
- C. 5G uses smaller cells and millimeter wave frequencies
- D. 5G eliminates base stations entirely for all connections

Answer: C

Q466. What is a wireless mesh network?

- A. A network where every device connects to a central hub
- B. A wired backbone connecting multiple wireless access points
- C. A network where nodes relay data for each other wirelessly
- D. A network using satellite links for all communication paths

Answer: C

Q467. What is the role of a wireless controller in enterprise networks?

- A. It encrypts all wired traffic before wireless transmission
- B. It replaces all access points with a single device
- C. It centrally manages and configures multiple access points
- D. It provides internet connectivity to the entire building

Answer: C

Q468. What is the purpose of beamforming in wireless communication?

- A. It focuses the wireless signal toward a specific device
- B. It converts analog signals into digital format wirelessly
- C. It compresses data packets for faster wireless transfer
- D. It spreads the signal equally in all directions around

Answer: A

Q469. What is a handoff (handover) in mobile wireless networks?

- A. The encryption key exchange between device and tower
- B. The initial registration of a device on a mobile network
- C. The process of powering off a mobile device gracefully
- D. The transfer of a connection from one cell to another

Answer: D

Q470. What is the ISM band and why is it significant for WiFi?

- A. A satellite communication band for international broadcasts
- B. An unlicensed band that WiFi and Bluetooth can use freely
- C. A licensed frequency band reserved for military usage
- D. A frequency band used exclusively for medical equipment

Answer: B

Q471. What is the difference between vertical and horizontal scaling?

- A. Vertical is for public cloud; horizontal is for private cloud
- B. Vertical adds more machines; horizontal upgrades existing machines
- C. Vertical upgrades one machine; horizontal adds more machines
- D. Vertical scales storage; horizontal scales computing power only

Answer: C

Q472. What is a container and how does it differ from a virtual machine?

- A. Containers include a full OS; VMs share the host OS kernel
- B. Containers are physical machines; VMs are software abstractions
- C. Containers share the host OS kernel; VMs include a full OS each
- D. Containers only run on private clouds; VMs run on public clouds

Answer: C

Q473. What is the purpose of a hypervisor in cloud computing?

- A. It manages network routing between cloud data centers
- B. It balances internet traffic across multiple web servers
- C. It creates and manages virtual machines on physical hosts
- D. It encrypts data stored in cloud storage services only

Answer: C

Q474. What does serverless computing mean?

- A. Applications run without any servers or infrastructure at all
- B. Applications run entirely on client devices with no backend
- C. The cloud provider manages servers; developers deploy code only
- D. Servers are physically removed from data center facilities

Answer: C

Q475. What is the role of a load balancer in cloud architecture?

- A. It distributes incoming traffic across multiple server instances
- B. It converts database queries into API calls for applications
- C. It monitors the temperature of physical servers in data centers
- D. It compresses data for faster storage on cloud disks

Answer: A

Q476. What is multi-tenancy in cloud computing?

- A. Multiple data centers hosting copies of the same application
- B. Multiple customers sharing the same physical infrastructure
- C. Multiple users accessing a cloud service through one account
- D. Multiple backups of data stored across different regions

Answer: B

Q477. What is the purpose of auto-scaling in cloud environments?

- A. It automatically encrypts data when storage is nearly full
- B. It automatically migrates data between different cloud providers
- C. It automatically adjusts resources based on current demand
- D. It automatically updates software to the latest version only

Answer: C

Q478. What is object storage and how does it differ from block storage?

- A. Object storage is faster; block storage provides better data durability
- B. Object storage is only for images; block storage is for all file types
- C. Object storage uses fixed blocks; block storage uses flexible objects
- D. Object storage stores data as objects with metadata; block uses raw blocks

Answer: D

Q479. What is a cloud region and availability zone?

- A. A region is a city; an availability zone is a building in it
- B. A region is a continent; an availability zone is a country within it
- C. A region is a geographic area; zones are isolated data centers within
- D. A region is a network; zones are different subnets within the network

Answer: C

Q480. What is the shared responsibility model in cloud security?

- A. Security duties are split between the provider and customer
- B. The customer handles all security without provider involvement
- C. A third-party security firm manages all cloud security tasks
- D. The cloud provider handles all security aspects completely

Answer: A

Q481. What is the difference between symmetric and asymmetric encryption?

- A. Symmetric uses one shared key; asymmetric uses a key pair
- B. Symmetric is slower; asymmetric is always faster in practice
- C. Symmetric uses two keys; asymmetric uses a single shared key
- D. Symmetric is for files; asymmetric is only for email messages

Answer: A

Q482. What is a man-in-the-middle (MITM) attack?

- A. An attacker intercepts communication between two parties secretly
- B. An attacker deletes data from the middle of a large database
- C. An attacker blocks only the middle portion of a data stream
- D. An attacker physically sits between two network servers

Answer: A

Q483. What is the purpose of an Intrusion Detection System (IDS)?

- A. It manages user passwords across the entire network
- B. It encrypts all data stored on network servers securely
- C. It monitors network traffic for suspicious activities
- D. It blocks all incoming network traffic automatically

Answer: C

Q484. How does a digital certificate work?

- A. It stores all user passwords in an encrypted database file
- B. It binds a public key to an entity verified by a trusted CA
- C. It routes encrypted traffic through the shortest network path
- D. It compresses encrypted data for faster transfer speeds

Answer: B

Q485. What is the purpose of network segmentation for security?

- A. To compress network traffic for more efficient routing
- B. To increase the overall network bandwidth capacity
- C. To merge multiple networks into one unified network
- D. To isolate network zones and limit breach impact spread

Answer: D

Q486. What is a DDoS attack and how does it differ from DoS?

- A. DDoS targets databases; DoS targets web servers exclusively
- B. DDoS originates from multiple sources; DoS from a single source
- C. DDoS is legal; DoS is the only illegal form of this attack
- D. DDoS uses encryption; DoS uses plain text for attacks

Answer: B

Q487. What is the principle of least privilege in security?

- A. Users should have only the minimum access rights they need
- B. Users should have all permissions to avoid workflow delays
- C. Users should share credentials to simplify access management
- D. Users should change permissions weekly for security rotation

Answer: A

Q488. What is the purpose of a DMZ in network architecture?

- A. It connects two private networks using encrypted tunnels
- B. It manages wireless access point configuration remotely
- C. It compresses all outgoing traffic to save internet bandwidth
- D. It hosts public-facing services isolated from the internal net

Answer: D

Q489. What is port scanning used for in network security?

- A. It assigns port numbers to newly installed server applications
- B. It discovers open ports and services on a target system
- C. It physically scans network cables for damage or breaks
- D. It increases the number of available network ports on switch

Answer: B

Q490. What is the role of TLS in network security?

- A. It manages domain name registration for web applications
- B. It monitors disk space usage on networked file servers
- C. It provides encryption for data in transit between endpoints
- D. It assigns IP addresses to devices on the local network

Answer: C

Q491. How does network automation improve operations?

- A. It eliminates all security vulnerabilities automatically
- B. It reduces manual tasks and human errors in configuration
- C. It replaces networking protocols with machine learning only
- D. It removes the need for any network hardware devices

Answer: B

Q492. What is the concept of intent-based networking?

- A. Networks that intentionally drop packets for load testing
- B. Networks that translate business intent into automated policies
- C. Networks that only connect devices with matching firmware
- D. Networks that administrators configure line by line manually

Answer: B

Q493. What is the role of blockchain in next-generation networks?

- A. It compresses network traffic to reduce bandwidth consumption
- B. It increases WiFi signal range through distributed antennas
- C. It replaces TCP/IP as the fundamental networking protocol
- D. It provides decentralized trust and secure transaction records

Answer: D

Q494. What is multi-access edge computing (MEC)?

- A. A technique for connecting multiple keyboard and mouse devices
- B. A method of storing data on multiple hard drives for backup
- C. Computing at the edge of the mobile network for low latency
- D. A protocol for accessing email from multiple client devices

Answer: C

Q495. How do digital twins benefit network management?

- A. They create physical copies of network equipment for backup
- B. They provide two-factor authentication for network access
- C. They double the network bandwidth by mirroring all traffic
- D. They simulate network behavior for testing without real impact

Answer: D

Q496. What is the significance of LPWAN for IoT deployments?

- A. It provides high-bandwidth connections for video streaming
- B. It replaces all existing WiFi standards with a new protocol
- C. It offers ultra-fast speeds for gaming and virtual reality
- D. It enables long-range, low-power connectivity for IoT sensors

Answer: D

Q497. What is the purpose of network telemetry?

- A. It collects real-time operational data from network devices
- B. It controls satellite communications for space networks
- C. It encrypts all management traffic on the network devices
- D. It provides telephone services over computer network lines

Answer: A

Q498. How does AIOps enhance network operations?

- A. It replaces all network devices with AI-powered hardware
- B. It uses AI and ML to automate and improve IT operations
- C. It generates artificial network traffic for testing purposes
- D. It provides artificial internet connections for offline work

Answer: B

Q499. What is the concept of network as a service (NaaS)?

- A. A free network service provided by government to all citizens
- B. Consuming networking capabilities as a cloud-delivered service
- C. A network that only provides services during business hours
- D. A physical network installed by the service provider on site

Answer: B

Q500. What is the role of quantum computing in future networks?

- A. It eliminates the need for any physical network connections
- B. It replaces all existing encryption methods with weaker ones
- C. It reduces the cost of network hardware to nearly zero price
- D. It promises faster processing and new cryptographic methods

Answer: D

Q501. What is the relationship between bit rate and baud rate in QPSK modulation?

- A. Bit rate is twice the baud rate value
- B. Bit rate equals the baud rate exactly
- C. Bit rate is four times the baud rate
- D. Bit rate is half the baud rate value

Answer: A

Q502. What is the primary purpose of line coding in digital communication?

- A. To multiplex several analog voice channels
- B. To convert digital data into digital signals
- C. To compress data before wireless transmission
- D. To encrypt data for secure communication

Answer: B

Q503. In Frequency Division Multiplexing, what prevents interference between channels?

- A. Time slots assigned to each channel
- B. Unique code sequences per channel
- C. Guard bands separating frequency ranges
- D. Priority levels for each data stream

Answer: C

Q504. What is the advantage of NRZ-I encoding over NRZ-L encoding?

- A. NRZ-I doubles the effective data rate
- B. NRZ-I uses less bandwidth per signal
- C. NRZ-I eliminates all baseline wander
- D. NRZ-I provides better clock synchronization

Answer: D

Q505. A channel has a bandwidth of 6 kHz. Using the Nyquist formula with 4 signal levels, what is the maximum data rate?

- A. 24 kbps with 4 levels encoding
- B. 48 kbps with 4 levels encoding
- C. 36 kbps with 4 levels encoding
- D. 12 kbps with 4 levels encoding

Answer: A

Q506. What is the primary difference between baseband and broadband transmission?

- A. Broadband uses digital signals on the entire bandwidth while baseband divides it
- B. Baseband uses digital signals on the entire bandwidth while broadband divides it
- C. Baseband supports higher data rates than broadband transmission systems
- D. Broadband can only carry a single channel at any given time

Answer: B

Q507. What problem does scrambling solve in data communication systems?

- A. It compresses data for faster throughput
- B. It increases signal power for long cables
- C. It reduces latency in fiber optic networks
- D. It avoids long sequences of constant bits

Answer: D

Q508. In Phase Shift Keying (PSK), what property of the carrier signal is modified?

- A. The phase of the carrier wave signal
- B. The frequency of the carrier wave
- C. The wavelength of the carrier wave
- D. The amplitude of the carrier wave

Answer: A

Q509. What is the role of an equalizer in a data communication receiver?

- A. To amplify signals for longer distances
- B. To compensate for signal distortion caused by the channel
- C. To encrypt incoming data for secure storage
- D. To convert analog signals into digital format

Answer: B

Q510. In Amplitude Shift Keying (ASK), how are binary values represented?

- A. By changing the frequency between two values
- B. By shifting the phase at each bit boundary
- C. By varying the amplitude of the carrier wave
- D. By combining frequency and phase variations

Answer: C

Q511. What is the primary difference between a Layer 2 switch and a Layer 3 switch?

- A. Layer 2 switches support wireless clients while Layer 3 cannot
- B. Layer 3 switches can perform routing based on IP addresses
- C. Layer 2 switches are faster but less reliable than Layer 3 switches
- D. Layer 3 switches only operate in star topology networks

Answer: B

Q512. What is the purpose of the encapsulation process in the OSI model?

- A. To remove headers as data moves down the layers
- B. To compress data at each layer for faster transmission
- C. To encrypt data at every layer for maximum security
- D. To add protocol information as data descends through layers

Answer: D

Q513. What is the function of the Address Resolution Protocol (ARP)?

- A. It assigns dynamic IP addresses to new network devices
- B. It resolves port numbers to specific running applications
- C. It maps IP addresses to their corresponding MAC addresses
- D. It maps domain names to their corresponding IP addresses

Answer: C

Q514. How does a switch learn which MAC addresses are on which ports?

- A. By examining the source MAC address of incoming frames
- B. By reading configuration files loaded at startup
- C. By querying a central MAC address database server
- D. By broadcasting ARP requests to all connected devices

Answer: A

Q515. What is the difference between unicast, multicast, and broadcast transmissions?

- A. Broadcast sends to one, unicast to a group, multicast to all
- B. Multicast sends to one, broadcast to a group, unicast to all
- C. Unicast sends to all, broadcast to a group, multicast to one
- D. Unicast sends to one, multicast to a group, broadcast to all

Answer: D

Q516. What does the data link layer PDU (Protocol Data Unit) consist of?

- A. Frames with MAC headers and trailer fields
- B. Packets with IP headers and payload data
- C. Bits transmitted as raw signals on the wire
- D. Segments with port numbers and sequence info

Answer: A

Q517. What is store-and-forward switching in a network switch?

- A. The switch only forwards broadcast frames to all ports
- B. The switch receives the entire frame before forwarding it
- C. The switch forwards frames before fully receiving them
- D. The switch stores frames indefinitely until requested

Answer: B

Q518. What is the purpose of RARP (Reverse Address Resolution Protocol)?

- A. It translates private IP addresses to public ones
- B. It resolves hostnames into IP address records
- C. It converts MAC addresses back to IP addresses
- D. It maps port numbers to application layer services

Answer: C

Q519. In a network with multiple paths, what advantage does load balancing provide?

- A. It assigns static IP addresses to all network devices
- B. It blocks redundant paths to simplify routing tables
- C. It distributes traffic across paths to optimize utilization
- D. It encrypts all traffic across different network paths

Answer: C

Q520. What is the role of the session layer in the OSI model?

- A. It converts data formats between different system types
- B. It handles physical transmission of bits across the medium
- C. It performs routing and logical addressing for packets
- D. It manages dialog control and synchronization between hosts

Answer: D

Q521. In the Go-Back-N protocol, what does the sender do when a timer expires for an unacknowledged frame?

- A. It waits for the receiver to send a negative acknowledgment
- B. It resends only the specific frame that timed out
- C. It closes the connection and initiates a new session
- D. It resends the timed-out frame and all subsequent sent frames

Answer: D

Q522. What is the difference between error detection and error correction?

- A. Detection finds errors and requests retransmission while correction fixes them directly
- B. Detection finds and fixes errors while correction only finds them
- C. Detection uses parity bits while correction uses checksum only
- D. Detection works only for burst errors while correction handles single-bit errors

Answer: A

Q523. What is the purpose of the interframe gap in Ethernet transmission?

- A. To synchronize clocks of all devices on the network
- B. To encrypt the gap between successive data frames
- C. To carry management data between consecutive frames
- D. To provide minimum idle time between frame transmissions

Answer: D

Q524. How does a two-dimensional parity check improve over simple parity?

- A. It only works with even numbers of data bytes in each block
- B. It can detect and correct single-bit errors in the data block
- C. It uses half the number of check bits compared to simple parity
- D. It eliminates the need for any retransmission of faulty frames

Answer: B

Q525. What is the piggybacking technique in data link layer protocols?

- A. Attaching acknowledgments to outgoing data frames to save bandwidth
- B. Compressing frame headers to reduce per-frame transmission overhead
- C. Duplicating frames across two links for redundancy and reliability
- D. Sending multiple data frames inside a single larger container frame

Answer: A

Q526. In CSMA/CD, what is the minimum frame size related to?

- A. The maximum number of devices on the network segment
- B. The total available bandwidth on the shared cable medium
- C. The round-trip propagation delay of the collision domain
- D. The processor speed of the transmitting network device

Answer: C

Q527. What is the spanning tree port state 'blocking' used for?

- A. The port is disabled and cannot receive any traffic
- B. The port is actively forwarding user data frames
- C. The port does not forward frames to prevent network loops
- D. The port is learning MAC addresses from incoming frames

Answer: C

Q528. What is the purpose of the jam signal in CSMA/CD?

- A. To request priority access to the shared transmission medium
- B. To notify all stations that a collision has been detected
- C. To test the cable for faults before data is transmitted
- D. To synchronize the clocks of all stations on the network

Answer: B

Q529. What is the role of the Logical Link Control (LLC) sublayer in frame handling?

- A. It provides an interface between the network layer and MAC sublayer
- B. It manages collision detection on shared Ethernet segments
- C. It handles physical signal encoding on the transmission medium
- D. It assigns physical addresses to each networked device

Answer: A

Q530. What is a transparent bridge and how does it operate?

- A. It requires manual configuration of all forwarding table entries
- B. It forwards frames based on IP addresses and routing tables
- C. It operates only between networks using different data link protocols
- D. It learns and forwards frames based on MAC addresses invisibly to hosts

Answer: D

Q531. What is the administrative distance used for in routing?

- A. It measures the physical distance between two routers in meters
- B. It ranks the trustworthiness of routing information sources
- C. It calculates the total bandwidth available on a route path
- D. It determines the maximum number of hops allowed on a route

Answer: B

Q532. In OSPF, what is a Link State Advertisement (LSA)?

- A. A message containing information about a router's links and their states
- B. A request sent to the designated router for database synchronization
- C. A packet used to discover neighboring OSPF routers on a link
- D. An acknowledgment confirming receipt of routing table updates

Answer: A

Q533. What is the purpose of the Don't Fragment (DF) flag in the IPv4 header?

- A. It marks the packet as the final fragment in a sequence
- B. It indicates the packet has already been fragmented once
- C. It instructs routers not to fragment the packet along the path
- D. It requests priority handling to avoid fragmentation delays

Answer: C

Q534. How does proxy ARP work in a network?

- A. A switch duplicates ARP requests and sends them to all VLANs
- B. A host caches ARP entries for all devices on the entire internet
- C. A firewall blocks ARP requests from unauthorized network segments
- D. A router answers ARP requests on behalf of hosts on another network

Answer: D

Q535. What is the difference between RIPv1 and RIPv2?

- A. RIPv1 supports VLSM and authentication while RIPv2 does not
- B. RIPv2 supports VLSM and multicast updates while RIPv1 does not
- C. RIPv2 has unlimited hop count while RIPv1 is limited to 15 hops
- D. RIPv1 uses link-state algorithm while RIPv2 uses distance-vector

Answer: B

Q536. What is the purpose of the OSPF hello protocol?

- A. To discover neighbors and maintain adjacency relationships
- B. To redistribute routes from other routing protocols into OSPF
- C. To calculate the shortest path using Dijkstra algorithm
- D. To exchange full routing tables between all OSPF routers

Answer: A

Q537. What is the purpose of the Identification field in the IPv4 header?

- A. It determines the priority level of the packet for QoS processing
- B. It specifies the version of the IP protocol being used
- C. It stores the checksum value for header integrity verification
- D. It uniquely identifies fragments belonging to the same original packet

Answer: D

Q538. What is the longest prefix match rule in IP routing?

- A. The route with the most hops to the destination is preferred
- B. The route learned from the oldest routing protocol is chosen
- C. The route with the most specific (longest) subnet mask is selected
- D. The route with the lowest administrative distance is preferred

Answer: C

Q539. What is the multicast address range in IPv4?

- A. 192.168.0.0 to 192.168.255.255 range
- B. 224.0.0.0 to 239.255.255.255 range
- C. 172.16.0.0 to 172.31.255.255 range
- D. 240.0.0.0 to 255.255.255.254 range

Answer: B

Q540. What is the purpose of ICMP redirect messages?

- A. To request a DNS resolver to redirect queries elsewhere
- B. To redirect all network traffic through a backup gateway
- C. To inform a host of a better route for a specific destination
- D. To notify the sender that a packet has been fragmented

Answer: C

Q541. How does TCP's sliding window mechanism improve throughput?

- A. It compresses data within each segment to reduce packet size
- B. It allows multiple segments to be sent before requiring acknowledgment
- C. It chooses the fastest network path for each individual segment
- D. It eliminates the need for sequence numbers in the TCP header

Answer: B

Q542. What is the purpose of the TCP Maximum Segment Size (MSS) option?

- A. It specifies the largest data payload a segment can carry
- B. It sets the timeout value for retransmission of segments
- C. It limits the total number of segments in a connection
- D. It determines the maximum window size for flow control

Answer: A

Q543. What happens during the TCP connection termination process?

- A. Both sides exchange FIN and ACK segments in a four-step process
- B. A single FIN segment from either side terminates the connection
- C. The server unilaterally closes the connection without client input
- D. One side sends a RST and the connection immediately closes

Answer: A

Q544. What is the TCP retransmission timeout (RTO) based on?

- A. The number of segments currently in the send buffer queue
- B. Estimated round-trip time and its variation for the connection
- C. A fixed interval of exactly 3 seconds for all connections
- D. The bandwidth of the network interface used for transmission

Answer: B

Q545. What is the difference between TCP's cumulative and selective acknowledgment?

- A. Cumulative ACKs work only with UDP while selective ACK works only with TCP
- B. Selective ACK confirms all data up to a point while cumulative identifies specific ranges
- C. Selective ACK requires more bandwidth but cumulative ACK causes more retransmissions
- D. Cumulative ACKs confirm all data up to a point while SACK identifies specific received ranges

Answer: D

Q546. What is the purpose of the TCP congestion window (cwnd)?

- A. It limits the total number of connections a server can accept
- B. It determines the priority of segments during network congestion
- C. It controls how much data the sender can transmit before receiving ACKs
- D. It specifies the maximum receive buffer size at the destination

Answer: C

Q547. What is delayed acknowledgment in TCP?

- A. The sender delays retransmission until multiple timeouts have occurred
- B. The sender delays sending data until it accumulates a full-size segment
- C. The receiver discards ACKs for duplicate segments to save bandwidth
- D. The receiver waits briefly before sending an ACK hoping to piggyback it with data

Answer: D

Q548. How does UDP handle data that arrives out of order?

- A. UDP buffers packets until all segments arrive then delivers them
- B. UDP discards out-of-order packets and requests retransmission
- C. UDP delivers data to the application in whatever order it arrives
- D. UDP reorders segments using sequence numbers before delivery

Answer: C

Q549. What does the TCP window scale option allow?

- A. It enables window sizes larger than the 16-bit field allows
- B. It dynamically adjusts the MSS during data transfer
- C. It reduces the header size for low-bandwidth connections
- D. It compresses window update notifications between peers

Answer: A

Q550. What is the purpose of the TCP keepalive mechanism?

- A. To periodically refresh the routing path for the TCP connection
- B. To detect if the other end of an idle connection is still reachable
- C. To prevent the congestion window from shrinking during idle time
- D. To maintain data throughput at a constant rate during transfers

Answer: B

Q551. What is the difference between iterative and recursive DNS queries?

- A. In iterative, the server resolves the query completely; in recursive, it returns a referral
- B. In recursive, the server resolves the query completely; in iterative, it returns a referral
- C. Recursive queries use TCP while iterative queries always use UDP protocol
- D. Iterative queries bypass caching while recursive queries always check cache first

Answer: B

Q552. What HTTP status code range indicates a client error?

- A. 300-399 redirection response code range
- B. 100-199 informational response codes
- C. 400-499 client error response code range
- D. 200-299 successful response code range

Answer: C

Q553. What is the purpose of HTTP cookies?

- A. To compress web page content for faster loading times
- B. To encrypt communication between browser and web server
- C. To validate the digital certificate of a web server
- D. To store state information for the stateless HTTP protocol

Answer: D

Q554. What is the DNS TTL (Time-to-Live) field used for?

- A. To specify how long a DNS record can be cached before refresh
- B. To set the maximum size of a DNS response message packet
- C. To limit the number of hops a DNS query can traverse
- D. To determine the priority of DNS servers for a domain

Answer: A

Q555. What is the purpose of the HTTP PUT method?

- A. To delete the resource located at the specified URI
- B. To retrieve a resource from the specified server URI
- C. To replace the entire resource at the specified URI
- D. To append data to an existing resource on the server

Answer: C

Q556. What is the CNAME record used for in DNS?

- A. To create an alias that points to another domain name
- B. To specify the text-based description of a DNS zone
- C. To map a domain name to an IPv6 address for hosts
- D. To define the service location for specific protocols

Answer: A

Q557. What is the difference between SFTP and FTPS?

- A. FTPS is faster than SFTP because it has less encryption overhead
- B. SFTP runs over SSH while FTPS adds TLS/SSL to the FTP protocol
- C. FTPS runs over SSH while SFTP adds TLS/SSL to the FTP protocol
- D. SFTP uses port 990 while FTPS uses port 22 for connections

Answer: B

Q558. What is the purpose of the HTTP Content-Type header?

- A. It defines the maximum size of the request or response payload
- B. It specifies the encoding used to compress the response body
- C. It authenticates the client making the HTTP request to the server
- D. It indicates the media type of the resource being sent or received

Answer: D

Q559. What is the role of a DNS root server in the name resolution hierarchy?

- A. It assigns new domain names to registrars for public sale
- B. It caches resolved DNS entries for all ISP customers worldwide
- C. It directs queries to the appropriate top-level domain servers
- D. It stores all DNS records for every domain on the internet

Answer: C

Q560. What is the HTTP 301 status code used to indicate?

- A. The requested resource was temporarily unavailable for maintenance
- B. The requested resource has been permanently moved to a new URL
- C. The server encountered an internal error processing the request
- D. The client must authenticate before accessing the resource

Answer: B

Q561. What is the primary benefit of WiFi 6E over WiFi 6?

- A. WiFi 6E removes backward compatibility to improve overall performance
- B. WiFi 6E adds the 6 GHz band providing more channels and less congestion
- C. WiFi 6E doubles the range of WiFi 6 using advanced antenna designs
- D. WiFi 6E uses a completely new modulation scheme for higher speeds

Answer: B

Q562. What is the purpose of the Distributed Coordination Function (DCF) in WiFi?

- A. It provides contention-based medium access using CSMA/CA for all stations
- B. It coordinates frequency hopping between channels to avoid interference
- C. It reserves dedicated time slots for each connected wireless device
- D. It centrally manages all wireless transmissions through the access point

Answer: A

Q563. What is roaming in enterprise wireless networks?

- A. The process of manually switching between different WiFi networks
- B. The technique of increasing transmission power for broader coverage
- C. The method of assigning static IP addresses to mobile wireless devices
- D. The automatic transition of a client between access points within a network

Answer: D

Q564. What is the difference between omnidirectional and directional antennas?

- A. Directional radiates in all horizontal directions while omnidirectional focuses signal in one
- B. Omnidirectional radiates in all horizontal directions while directional focuses signal in one
- C. Directional antennas have lower gain than omnidirectional antennas in all cases
- D. Omnidirectional works only indoors while directional is designed for outdoor use only

Answer: B

Q565. What does the term 'cell' refer to in cellular network architecture?

- A. A geographic area covered by a single base station
- B. A group of frequencies reserved for emergency calls
- C. The encryption module inside a cellular modem chip
- D. The battery unit that powers each mobile device

Answer: A

Q566. What is the purpose of the association process in WiFi?

- A. It assigns a fixed wireless channel to the connecting client device
- B. It discovers all available WiFi networks in the surrounding area
- C. It establishes a logical connection between a client and an access point
- D. It encrypts the data channel between client and access point

Answer: C

Q567. What is the dual-band capability in modern WiFi routers?

- A. The ability to connect to two internet service providers at the same time
- B. The ability to broadcast two separate SSIDs on the same frequency band
- C. The ability to operate on both 2.4 GHz and 5 GHz frequency bands simultaneously
- D. The ability to support both WPA2 and WPA3 security protocols concurrently

Answer: C

Q568. What is the function of the NAV (Network Allocation Vector) in WiFi?

- A. It acts as a virtual carrier sensing timer to reserve the channel
- B. It allocates IP addresses to wireless network participants
- C. It stores the routing table for wireless mesh networks
- D. It manages authentication keys for wireless client devices

Answer: A

Q569. What is the difference between BSS (Basic Service Set) and ESS (Extended Service Set)?

- A. ESS consists of one AP and its clients while BSS connects multiple ESSs together
- B. BSS consists of one AP and its clients while ESS connects multiple BSSs together
- C. ESS supports only wired connections while BSS supports only wireless connections
- D. BSS uses WPA3 encryption while ESS only supports the older WPA2 standard

Answer: B

Q570. What technology does 5G use to handle many connected IoT devices simultaneously?

- A. Frequency Division Multiple Access with guard bands
- B. Single-carrier modulation with higher power output
- C. Circuit switching with dedicated device channels
- D. Massive MIMO with hundreds of antenna elements

Answer: D

Q571. What is the difference between a Type 1 and Type 2 hypervisor?

- A. Type 1 runs directly on hardware while Type 2 runs on a host operating system
- B. Type 2 provides better performance than Type 1 in all cloud deployments
- C. Type 1 supports only Linux guests while Type 2 supports any guest operating system
- D. Type 2 runs directly on hardware while Type 1 runs on a host operating system

Answer: A

Q572. What is the purpose of a container orchestration platform?

- A. It compiles application source code into container images
- B. It converts virtual machines into lightweight container instances
- C. It provides a graphical interface for designing container architectures
- D. It automates deployment, scaling, and management of containerized applications

Answer: D

Q573. What is the concept of cloud-native microservices architecture?

- A. A single large database shared by all application components in the cloud
- B. A method of migrating legacy mainframe applications directly to cloud servers
- C. An approach where applications are built as small, independent, deployable services
- D. A monolithic application deployed entirely within a single cloud virtual machine

Answer: C

Q574. What is the purpose of a content delivery network (CDN) in cloud architecture?

- A. It manages database replication across multiple cloud data regions
- B. It caches and delivers content from edge locations closest to users
- C. It provides computing resources for running backend application logic
- D. It handles user authentication for cloud-hosted web applications

Answer: B

Q575. What is the difference between block storage and file storage in cloud?

- A. File storage manages raw storage volumes while block storage provides a file system hierarchy
- B. File storage offers higher performance than block storage in all scenarios
- C. Block storage manages raw storage volumes while file storage provides a file system hierarchy
- D. Block storage is only for databases while file storage is only for media content

Answer: C

Q576. What is a cloud function (Function as a Service)?

- A. A dedicated server leased monthly for running specific application workloads
- B. A small piece of code executed in response to events without managing servers
- C. A database query engine hosted in the cloud for analytical processing
- D. A virtual machine running continuously to handle all incoming requests

Answer: B

Q577. What is the purpose of cloud resource tagging?

- A. To label resources with metadata for organization, billing, and management
- B. To encrypt resources with unique security tokens for access control
- C. To assign priority levels to cloud resources during peak usage times
- D. To physically mark server hardware for identification in data centers

Answer: A

Q578. What is the concept of cloud data sovereignty?

- A. The ability of cloud users to access data from any global location
- B. The encryption standard used to protect data stored in cloud servers
- C. The cloud provider's guarantee of permanent data availability worldwide
- D. The legal requirement that data is subject to laws of the country where it resides

Answer: D

Q579. What is a VPC (Virtual Private Cloud)?

- A. An isolated virtual network within a public cloud for a customer's resources
- B. A VPN tunnel connecting an office to the internet through cloud proxy
- C. A physical data center dedicated exclusively to one customer
- D. A virtual machine with enhanced security features enabled by default

Answer: A

Q580. What is the difference between reserved and spot instances in cloud computing?

- A. Reserved instances are temporary while spot instances provide guaranteed availability
- B. Spot instances require long-term contracts while reserved instances are pay-per-use
- C. Spot instances are more expensive but offer better performance than reserved ones
- D. Reserved instances are cheaper with commitment while spot instances offer discounts on spare capacity

Answer: D

Q581. What is the difference between a vulnerability scan and a penetration test?

- A. A penetration test uses automated tools while a vulnerability scan requires manual testing
- B. A vulnerability scan actively exploits weaknesses while a penetration test only identifies them
- C. A vulnerability scan identifies weaknesses automatically while a penetration test actively exploits them
- D. Both vulnerability scans and penetration tests perform identical security assessment functions

Answer: C

Q582. What is the purpose of Multi-Factor Authentication (MFA)?

- A. To allow users to access systems with any single form of identification
- B. To distribute authentication requests across multiple authentication servers
- C. To encrypt user credentials using multiple encryption algorithms simultaneously
- D. To require multiple verification methods from different authentication categories

Answer: D

Q583. What is a SQL injection attack?

- A. An attack that intercepts SQL queries between application and database
- B. An attack that inserts malicious SQL code through application input fields
- C. An attack that floods a database server with excessive connection requests
- D. An attack that physically damages database storage hardware components

Answer: B

Q584. What is the purpose of network access control (NAC)?

- A. To accelerate network traffic through quality of service policies
- B. To manage the physical cabling infrastructure of the enterprise network
- C. To enforce security policies on devices before granting network access
- D. To compress network traffic headers for bandwidth conservation

Answer: C

Q585. What is a Cross-Site Scripting (XSS) attack?

- A. An attack that injects malicious scripts into web pages viewed by other users
- B. An attack that blocks access to a website through distributed denial requests
- C. An attack that redirects website visitors to a cloned phishing site
- D. An attack that exploits expired SSL certificates on insecure web servers

Answer: A

Q586. What is the purpose of RADIUS (Remote Authentication Dial-In User Service)?

- A. To encrypt all data transmitted between wireless clients and access points
- B. To manage DHCP address assignments for remote VPN client connections
- C. To detect and prevent unauthorized wireless access points on the network
- D. To provide centralized authentication, authorization, and accounting for network access

Answer: D

Q587. What is the concept of defense in depth in network security?

- A. Using the deepest encryption algorithm available for all data types
- B. Implementing multiple layers of security controls throughout a system
- C. Deploying security measures only at the network perimeter boundary
- D. Relying on a single strong firewall to protect the entire network

Answer: B

Q588. What is the purpose of a digital signature?

- A. To verify the sender's identity and ensure message integrity
- B. To route signed messages through prioritized network paths
- C. To encrypt the entire message content for confidentiality
- D. To compress messages for faster network transmission speed

Answer: A

Q589. What is a VLAN hopping attack?

- A. An attack that rapidly changes VLAN assignments to cause network instability
- B. An attack that disables VLAN trunking between switches to isolate network segments
- C. An attack where an unauthorized user gains access to a wireless VLAN
- D. An attack that allows an attacker to access traffic on other VLANs bypassing segmentation

Answer: D

Q590. What is the role of hashing in password storage?

- A. It encrypts passwords so they can be decrypted and verified during login
- B. It duplicates passwords across multiple servers for redundancy and availability
- C. It converts passwords into fixed-length values that cannot be reversed to reveal the original
- D. It compresses passwords to reduce the storage space needed in databases

Answer: C

Q591. What is the purpose of a digital twin in network operations?

- A. To duplicate network traffic for load balancing across twin paths
- B. To pair every physical device with an identical standby device
- C. To build a virtual replica of a network for simulation and planning
- D. To create an identical backup of all network hardware equipment

Answer: C

Q592. What is the concept of network function chaining in NFV?

- A. Linking network monitoring tools to create a unified dashboard view
- B. Physically connecting network appliances in a serial chain with cables
- C. Chaining multiple physical firewalls together for redundant protection
- D. Directing traffic through an ordered sequence of virtual network functions

Answer: D

Q593. What is the role of YANG data modeling language in network automation?

- A. It generates graphical network topology diagrams automatically
- B. It models the structure of network device configuration and state data
- C. It defines the programming logic for automated network scripts
- D. It encrypts network automation commands for secure transmission

Answer: B

Q594. What is the concept of network slicing in 5G architecture?

- A. Creating multiple virtual networks on shared physical infrastructure for different needs
- B. Slicing large data packets into smaller fragments for transmission
- C. Dividing the radio spectrum into equal frequency slices for all users
- D. Physically dividing network cables into thinner sections for more connections

Answer: A

Q595. What is the NETCONF protocol used for in network management?

- A. It provides a mechanism to configure and manage network devices programmatically
- B. It encrypts network traffic between data center and branch offices
- C. It monitors network performance through passive packet capture
- D. It distributes network routing updates between autonomous systems

Answer: A

Q596. What is the concept of Vehicle-to-Everything (V2X) communication?

- A. A system for transmitting vehicle diagnostic data to repair shops
- B. A cellular network dedicated to providing entertainment in vehicles
- C. Communication between vehicles and infrastructure, pedestrians, and other vehicles
- D. A satellite navigation system designed exclusively for commercial vehicles

Answer: C

Q597. What is the purpose of an IoT gateway device?

- A. It generates power for battery-free IoT sensor devices
- B. It provides cloud computing resources directly to IoT sensors
- C. It physically protects IoT devices from environmental damage
- D. It bridges communication between IoT devices and cloud or enterprise networks

Answer: D

Q598. What is the concept of haptic communication in next-generation networks?

- A. Communication limited to text-based messaging protocols only
- B. Communication that transmits touch and force feedback over a network
- C. Communication using visible light instead of radio frequencies
- D. Communication between devices using acoustic sound wave signals

Answer: B

Q599. What is the significance of Open RAN (O-RAN) for mobile network operators?

- A. It requires operators to use a single vendor for all radio access equipment
- B. It enables interoperable multi-vendor radio access network components
- C. It eliminates the need for physical cell towers in mobile network design
- D. It restricts RAN deployment to only open-source software implementations

Answer: B

Q600. What is the purpose of a Software-Defined WAN (SD-WAN)?

- A. To eliminate the need for WAN connectivity between branch offices
- B. To replace all WAN links with a single high-speed fiber connection
- C. To physically define WAN boundaries using software-generated maps
- D. To intelligently manage WAN connections using software-based controls

Answer: D

Computer Networks & Cloud Computing — Hard

300 questions

Q601. A channel has a bandwidth of 4 KHz and SNR of 63. What is the maximum achievable data rate according to Shannon's theorem?

- A. 24 Kbps
- B. 12 Kbps
- C. 48 Kbps
- D. 36 Kbps

Answer: A

Q602. In CDMA (Code Division Multiple Access), if a station's chip sequence is $(-1, -1, -1, +1, +1, -1, +1, +1)$ and the received data is $(-1, -1, -3, +1, -1, -3, +1, +1)$, what is the station's data bit?

- A. Indeterminate
- B. Bit is -1
- C. Bit is +1
- D. Bit is 0

Answer: C

Q603. What is the main disadvantage of Differential Manchester encoding compared to standard Manchester encoding?

- A. It is more complex to implement
- B. It requires more bandwidth
- C. It lacks self-clocking ability
- D. It has higher error rates

Answer: A

Q604. In a QAM-16 constellation diagram, how many bits does each signal element represent?

- A. 2 bits
- B. 3 bits
- C. 8 bits
- D. 4 bits

Answer: D

Q605. A signal has passed through 3 cascaded amplifiers, each with a gain of 10 dB. If the original signal power is 5 mW, what is the output power?

- A. 5000 mW
- B. 50000 mW
- C. 500 mW
- D. 50 mW

Answer: A

Q606. In Statistical Time Division Multiplexing (STDM), what is the key advantage over synchronous TDM?

- A. It allocates time slots based on demand
- B. It guarantees lower latency always
- C. It offers simpler implementation
- D. It provides higher total bandwidth

Answer: A

Q607. A digitized voice channel is made by sampling at 8000 samples/sec with 8 bits per sample. What is the required data rate?

- A. 32 Kbps
- B. 128 Kbps
- C. 64 Kbps
- D. 48 Kbps

Answer: C

Q608. Which scrambling technique replaces 8 consecutive zeros with a specific violation pattern to maintain synchronization?

- A. AMI
- B. MLT-3
- C. B8ZS
- D. HDB3

Answer: C

Q609. What is the relationship between bit rate and baud rate when using 8-PSK modulation?

- A. Bit rate = 8 * baud rate
- B. Bit rate = 3 * baud rate
- C. Bit rate = 2 * baud rate
- D. Bit rate = baud rate

Answer: B

Q610. In spread spectrum communication, what happens to the signal bandwidth compared to the original data bandwidth?

- A. It is halved exactly
- B. It decreases significantly
- C. It remains the same
- D. It increases significantly

Answer: D

Q611. In the OSI model, at which layer does the PDU (Protocol Data Unit) change from 'segment' to 'packet'?

- A. From Session to Transport layer
- B. From Transport to Network layer
- C. From Data Link to Physical layer
- D. From Network to Data Link layer

Answer: B

Q612. Which RFC defines the Internet Protocol version 4 (IPv4)?

- A. RFC 793
- B. RFC 791
- C. RFC 2460
- D. RFC 768

Answer: B

Q613. What problem does the Spanning Tree Protocol (STP) solve in switched networks?

- A. Domain name resolution failures
- B. Broadcast storms from switching loops
- C. Automatic IP address assignment
- D. Bandwidth allocation conflicts

Answer: B

Q614. In a hierarchical network design, what is the role of the distribution layer?

- A. Policy-based routing between access and core
- B. High-speed backbone packet forwarding
- C. Connecting end-user devices directly
- D. Providing wireless access to clients

Answer: A

Q615. What is the key concept behind Software-Defined Networking (SDN)?

- A. Hosting all networking functions in the cloud
- B. Eliminating all physical network device hardware
- C. Using only software-based routers for all traffic
- D. Separating the control plane from the data plane

Answer: D

Q616. What is the purpose of MPLS (Multi-Protocol Label Switching) in networking?

- A. To encrypt all network traffic end-to-end
- B. To provide wireless connectivity to devices
- C. To completely replace IP addressing schemes
- D. To speed up forwarding using labels not IPs

Answer: D

Q617. In network performance, what is jitter?

- A. Maximum throughput of a link
- B. Total time to deliver a message
- C. Loss of packets during transit
- D. Variation in packet arrival delay

Answer: D

Q618. What distinguishes a Layer 3 switch from a traditional router?

- A. It performs routing in hardware via ASICs
- B. It operates at a much higher OSI layer
- C. It cannot route between different VLANs
- D. It only works with wireless networks

Answer: A

Q619. In the context of Quality of Service (QoS), what does DiffServ provide?

- A. Guaranteed zero latency for all flows
- B. Per-flow resource reservation for traffic
- C. Automatic bandwidth doubling on demand
- D. Traffic classification using DSCP values

Answer: D

Q620. What is the main purpose of Network Function Virtualization (NFV)?

- A. To create virtual LAN segments on switches
- B. To create backup networks for failover
- C. To virtualize and manage IP address pools
- D. To replace hardware appliances with software

Answer: D

Q621. If the generator polynomial for CRC is $x^3 + 1$, what is the binary divisor?

- A. 1001
- B. 1101
- C. 1111
- D. 1011

Answer: A

Q622. In the HDLC protocol, what type of frame is used for flow and error control without carrying any data?

- A. U-frame
- B. P-frame
- C. I-frame
- D. S-frame

Answer: D

Q623. In a sliding window protocol with a window size of 7 and sequence numbers 0-7, what is the maximum sender window size in Go-Back-N?

- A. 7
- B. 8
- C. 3
- D. 4

Answer: A

Q624. What is the purpose of the IEEE 802.1Q tag in Ethernet frames?

- A. Quality of Service marking
- B. VLAN identification tagging
- C. Frame payload encryption
- D. Error correction coding

Answer: B

Q625. What is the maximum number of VLANs supported by IEEE 802.1Q?

- A. 65535
- B. 1024
- C. 4094
- D. 256

Answer: C

Q626. What is the Hamming distance needed to correct t errors?

- A. $2t$
- B. $t + 1$
- C. $2t + 1$
- D. t

Answer: C

Q627. In Ethernet, why is there a minimum frame size requirement of 64 bytes?

- A. To reduce processing overhead on devices
- B. To support VLAN tagging in the header
- C. To maintain backward compatibility with hubs
- D. To ensure collision detection within the network

Answer: D

Q628. What is the difference between persistent and non-persistent CSMA?

- A. There is no functional difference between the two modes
- B. Non-persistent is significantly faster than persistent always
- C. Persistent transmits when idle; non-persistent waits randomly before sensing
- D. Persistent CSMA relies on a token passing mechanism

Answer: C

Q629. In the PPP (Point-to-Point Protocol) frame, what is the role of the LCP (Link Control Protocol)?

- A. Establishing and maintaining the link
- B. Resolving domain names for hosts
- C. Assigning IP addresses to endpoints
- D. Routing packets across the network

Answer: A

Q630. CRC-32 used in Ethernet can detect all burst errors of length up to how many bits?

- A. 32
- B. 64
- C. 24
- D. 16

Answer: A

Q631. Given the IP address 172.16.45.14/21, what is the network address?

- A. 172.16.32.0
- B. 172.16.40.0
- C. 172.16.45.0
- D. 172.16.44.0

Answer: B

Q632. In OSPF, what is an Area Border Router (ABR)?

- A. A router that runs the BGP protocol
- B. A router operating within a single OSPF area
- C. A router connecting two or more OSPF areas
- D. A router connecting to external networks only

Answer: C

Q633. What technique does split horizon use to solve the count-to-infinity problem?

- A. Not advertising routes back through learned interfaces
- B. Using TTL values to limit the route propagation
- C. Sending updates to directly connected routers only
- D. Encrypting all routing update messages sent

Answer: A

Q634. In IPv6, what replaces broadcast?

- A. No replacement exists
- B. Multicast and anycast
- C. Unicast addressing
- D. Simplex mode

Answer: B

Q635. What is the purpose of the OSPF Designated Router (DR)?

- A. To connect to external autonomous systems
- B. To serve as the default gateway for hosts
- C. To handle DHCP requests from all clients
- D. To reduce adjacencies on multi-access networks

Answer: D

Q636. A company needs to create 6 subnets from a /24 network. What is the most efficient subnet mask to use?

- A. /29
- B. /26
- C. /28
- D. /27

Answer: D

Q637. What is the function of the Protocol field in the IPv4 header?

- A. To indicate the total length of the packet
- B. To identify the upper-layer protocol used
- C. To store the header checksum for errors
- D. To specify the version of IP being used

Answer: B

Q638. In BGP, what is the AS_PATH attribute used for?

- A. Listing traversed ASes to prevent loops
- B. Measuring the link bandwidth capacity
- C. Encrypting routing updates over links
- D. Assigning IP addresses to BGP peers

Answer: A

Q639. What is the main advantage of VLSM (Variable Length Subnet Masking)?

- A. It enables faster packet routing overall
- B. It simplifies network configuration greatly
- C. It provides better security for subnets
- D. It allows subnets of different sizes for efficiency

Answer: D

Q640. What is route aggregation (supernetting) and why is it important?

- A. Combining network prefixes to reduce routing tables
- B. Assigning multiple IP addresses to a single host
- C. Creating backup routes for network redundancy
- D. Splitting a network into smaller subnets for isolation

Answer: A

Q641. In TCP congestion control, what triggers the slow start algorithm to switch to congestion avoidance?

- A. When all sent data has been acknowledged
- B. When the receiver window fills up completely
- C. After a retransmission timeout event occurs
- D. When congestion window reaches the ssthresh

Answer: D

Q642. What happens in TCP fast retransmit?

- A. The sender waits for a timeout before resending
- B. All unacknowledged segments are retransmitted
- C. The receiver requests immediate retransmission
- D. The sender resends after three duplicate ACKs

Answer: D

Q643. In the TCP state diagram, what state does a client enter after sending a SYN?

- A. SYN_RECEIVED
- B. ESTABLISHED
- C. SYN_SENT
- D. LISTEN

Answer: C

Q644. What is the TIME_WAIT state in TCP used for?

- A. Waiting for lost data to be retransmitted
- B. Waiting for the server to respond to requests
- C. Performing the initial connection setup phase
- D. Ensuring old duplicate segments expire first

Answer: D

Q645. What is TCP SACK (Selective Acknowledgment)?

- A. A compression method for TCP segments
- B. An extension to report non-contiguous data blocks
- C. A method to acknowledge every single byte received
- D. A security feature for TCP authentication

Answer: B

Q646. What is the Silly Window Syndrome in TCP?

- A. A routing loop caused by TCP misconfiguration
- B. A security vulnerability in older TCP stacks
- C. A DNS resolution error affecting TCP connections
- D. Inefficient exchange of tiny segments due to small windows

Answer: D

Q647. In TCP Reno, what happens when three duplicate ACKs are received?

- A. The cwnd is reset to 1 MSS and slow start begins
- B. The TCP connection is terminated by the sender
- C. The cwnd is doubled and transmission continues
- D. The ssthresh is halved and fast recovery is entered

Answer: D

Q648. What is the purpose of the TCP PSH (Push) flag?

- A. To indicate urgent data in the stream
- B. To establish a new TCP connection
- C. To deliver buffered data to the app now
- D. To reset the existing TCP connection

Answer: C

Q649. What is a SYN flood attack?

- A. Sending many SYNs without completing handshakes
- B. Flooding the network with excessive UDP packets
- C. Sending too much data at once to overflow buffers
- D. Sending many malformed DNS query messages

Answer: A

Q650. What distinguishes QUIC from TCP?

- A. QUIC is only used for email protocols
- B. QUIC does not support stream multiplexing
- C. QUIC is slower but more reliable than TCP
- D. QUIC runs over UDP with built-in encryption

Answer: D

Q651. In DNS, what is a recursive query?

- A. A query where the server fully resolves the name
- B. A query that only checks local cache entries
- C. A query that returns only partial results
- D. A query that bypasses all DNS caches

Answer: A

Q652. What is the key improvement in HTTP/2 over HTTP/1.1?

- A. It uses UDP instead of TCP for speed
- B. It removes all encryption and TLS support
- C. It uses a completely different port number
- D. It multiplexes requests over a single connection

Answer: D

Q653. What is DNS amplification attack?

- A. Modifying DNS responses to redirect users
- B. Using open resolvers to flood a target with traffic
- C. Blocking all DNS queries on the network
- D. Stealing DNS records from name servers

Answer: B

Q654. What is the purpose of DNSSEC?

- A. To speed up DNS query resolution time
- B. To authenticate and verify DNS responses
- C. To encrypt DNS queries for user privacy
- D. To replace DNS with a new modern protocol

Answer: B

Q655. What is a RESTful API?

- A. An API that only supports GET requests
- B. An API using HTTP methods for stateless access
- C. An API that only works during off-peak hours
- D. An API built on SOAP protocol framework

Answer: B

Q656. What is DNS-over-HTTPS (DoH)?

- A. A new DNS record type for websites
- B. A faster version of DNS using caching
- C. A protocol encapsulating DNS within HTTPS
- D. A way to host DNS on web servers only

Answer: C

Q657. In HTTP, what is the purpose of the ETag header?

- A. Session management cookies
- B. Resource versioning for caching
- C. User authentication tokens
- D. Content encoding indicators

Answer: B

Q658. What is the function of a reverse proxy?

- A. To block all incoming network traffic
- B. To hide the client identity from servers
- C. To forward requests to backend servers
- D. To replace DNS resolution entirely

Answer: C

Q659. What is the purpose of a CDN (Content Delivery Network)?

- A. To encrypt emails during transit
- B. To replace web servers entirely
- C. To provide DNS resolution services
- D. To distribute content globally for speed

Answer: D

Q660. What is WebSocket and how does it differ from HTTP?

- A. WebSocket is only designed for file transfers
- B. WebSocket provides full-duplex persistent communication
- C. WebSocket uses UDP instead of TCP for transport
- D. WebSocket is a faster version of the HTTP protocol

Answer: B

Q661. What is beamforming in 5G and Wi-Fi 6?

- A. Using a single omnidirectional antenna only
- B. Focusing signal energy toward specific devices
- C. Broadcasting signals in all directions equally
- D. Reducing signal power to save energy

Answer: B

Q662. What is the key technology difference between 5G mmWave and sub-6 GHz 5G?

- A. Sub-6 GHz provides faster speeds than mmWave
- B. They are identical in performance and range
- C. mmWave offers extreme speed but limited range; sub-6 gives wider coverage
- D. mmWave has better penetration through buildings

Answer: C

Q663. What is MU-MIMO and how does it differ from SU-MIMO?

- A. They are the same MIMO technology in practice
- B. MU-MIMO uses fewer antennas than SU-MIMO does
- C. SU-MIMO is newer and more advanced than MU-MIMO
- D. MU-MIMO serves multiple users at once; SU-MIMO serves one

Answer: D

Q664. What is OFDMA and how does it improve Wi-Fi 6?

- A. A wireless security protocol for encryption
- B. A technique dividing channels into sub-channels for users
- C. A type of data encryption for Wi-Fi
- D. A new dedicated radio frequency band

Answer: B

Q665. What is network slicing in 5G?

- A. Creating virtual networks on shared infrastructure
- B. Reducing overall network speed for stability
- C. Dividing frequency bands into smaller chunks
- D. Physically cutting and splicing network cables

Answer: A

Q666. What is the exposed station problem in wireless networks?

- A. A station without any encryption configured on it
- B. A security vulnerability allowing eavesdropping attacks
- C. A station that unnecessarily defers its own transmission
- D. A station that is visible to all users on the network

Answer: C

Q667. In Wi-Fi 6 (802.11ax), what is BSS Coloring?

- A. Marking frames to differentiate overlapping BSSs
- B. Changing the LED color on the router hardware
- C. A security feature for frame encryption
- D. A QoS mechanism for traffic prioritization

Answer: A

Q668. What is the role of the Mobility Management Entity (MME) in LTE architecture?

- A. Radio signal transmission to devices
- B. Control-plane functions like auth and handover
- C. Forwarding user data between nodes
- D. Content delivery to mobile users

Answer: B

Q669. What is the KRACK (Key Reinstallation Attack) vulnerability?

- A. A DNS spoofing attack on wireless networks
- B. An attack exploiting the WPA2 four-way handshake
- C. A brute force password cracking technique
- D. A physical attack on router hardware devices

Answer: B

Q670. What is the purpose of Target Wake Time (TWT) in Wi-Fi 6?

- A. To increase the transmit power of antennas
- B. To keep all devices active and always on
- C. To change wireless channels automatically
- D. To schedule wake times to save device power

Answer: D

Q671. What is the CAP theorem in distributed cloud systems?

- A. A theorem about total network capacity limits
- B. A theorem about reducing cloud service costs
- C. A theorem about choosing two of three guarantees
- D. A theorem about CPU performance scaling

Answer: C

Q672. What is cloud-native architecture?

- A. Cloud storage solutions without compute services
- B. Applications designed for on-premises servers only
- C. Native mobile apps that run on phones only
- D. Apps built with microservices, containers, and CI/CD

Answer: D

Q673. What is the shared responsibility model in cloud security?

- A. Provider secures infrastructure; customer secures data
- B. The cloud provider handles all security alone
- C. No one is responsible for cloud security
- D. The customer handles all security responsibilities

Answer: A

Q674. What is Kubernetes?

- A. A container orchestration platform
- B. A programming language for cloud
- C. A cloud provider platform service
- D. A type of virtual machine image

Answer: A

Q675. What is edge computing and how does it relate to cloud computing?

- A. A complete replacement for cloud computing
- B. Processing data near the source to cut latency
- C. Computing at the edge of a physical desk
- D. Computing only on mobile phone devices

Answer: B

Q676. What is Infrastructure as Code (IaC)?

- A. Managing infrastructure via configuration files
- B. Coding without using a computer at all
- C. A type of cloud storage service
- D. Writing code directly on physical servers

Answer: A

Q677. What is vendor lock-in in cloud computing and how can it be mitigated?

- A. A security feature that protects cloud accounts
- B. Being locked out of a vendor's office building
- C. Dependency on one provider, mitigated by open standards
- D. A pricing model for reserved cloud instances

Answer: C

Q678. What is a cloud service mesh?

- A. A type of cloud object storage bucket
- B. An infrastructure layer for service communication
- C. Physical network cables in a cloud data center
- D. A wireless network within cloud data centers

Answer: B

Q679. What is the difference between disaster recovery (DR) and high availability (HA) in cloud?

- A. HA prevents downtime; DR recovers from disasters
- B. DR is always faster than HA to implement
- C. HA is only applicable to database services
- D. They are the same concept with different names

Answer: A

Q680. What is FinOps in the context of cloud computing?

- A. A cloud provider offering compute services
- B. A practice for managing cloud spending effectively
- C. Financial operations for stock market trading
- D. A type of encryption for financial data

Answer: B

Q681. What is the difference between IPSec tunnel mode and transport mode?

- A. Tunnel encrypts entire packet; transport encrypts payload only
- B. Transport mode encrypts the entire IP packet
- C. They are identical in operation and security
- D. Tunnel mode provides less security than transport

Answer: A

Q682. What is a zero-day vulnerability?

- A. A vulnerability that takes zero days to fix
- B. A vulnerability with zero security impact
- C. An unknown flaw exploited before a patch exists
- D. A vulnerability found on the software release day

Answer: C

Q683. What is the TLS 1.3 handshake improvement over TLS 1.2?

- A. It reduces handshake to 1-RTT with 0-RTT option
- B. It uses more round trips for security
- C. It removes all encryption for faster speed
- D. It uses significantly longer encryption keys

Answer: A

Q684. What is DNS cache poisoning?

- A. Physically destroying DNS server hardware
- B. Clearing the DNS cache on all resolvers
- C. A DNS backup and recovery method
- D. Injecting false DNS entries to redirect users

Answer: D

Q685. What is network segmentation and why is it important for security?

- A. Creating one large flat network for simplicity
- B. Dividing networks into segments to limit attacks
- C. A type of data encryption for networks
- D. Disconnecting all network cables physically

Answer: B

Q686. What is the principle of zero trust security?

- A. Having zero security measures on the network
- B. Trusting only traffic from external sources
- C. Trusting all internal network traffic by default
- D. Never trust, always verify every access request

Answer: D

Q687. What is the BGP hijacking attack?

- A. A DNS spoofing attack on name servers
- B. A wireless deauthentication attack
- C. Stealing physical router hardware devices
- D. Announcing unauthorized IP prefixes via BGP

Answer: D

Q688. What is the purpose of SIEM (Security Information and Event Management)?

- A. Optimizing network speed and performance
- B. Managing email systems for organizations
- C. Securing file transfers between servers
- D. Collecting and analyzing security events for threats

Answer: D

Q689. What is perfect forward secrecy (PFS) in cryptographic protocols?

- A. A type of network firewall technology
- B. Ensuring past sessions stay safe if keys leak
- C. Using the same encryption key forever
- D. A backup encryption method for storage

Answer: B

Q690. What is the difference between signature-based and anomaly-based intrusion detection?

- A. They are the same detection method
- B. Signature-based can detect brand new attacks
- C. Anomaly-based detection is less effective overall
- D. Signature matches known patterns; anomaly detects deviations

Answer: D

Q691. What is the concept of intent-based networking (IBN)?

- A. A fully manual network configuration approach
- B. Admins define outcomes and the system auto-configures
- C. Networking based only on user browsing intentions
- D. A type of routing protocol for WANs

Answer: B

Q692. What is Segment Routing (SR) and why is it important for next-gen networks?

- A. DNS-based routing for domain resolution
- B. A type of firewall filtering technique
- C. Routing through physical cable segments only
- D. Source routing using segment lists to simplify forwarding

Answer: D

Q693. What is the role of AI/ML in next-generation network management?

- A. Only powering customer service chatbot systems
- B. Creating virtual network topologies from scratch
- C. Replacing all network hardware with software
- D. Enabling predictive analytics and self-healing networks

Answer: D

Q694. What is the concept of network disaggregation?

- A. Breaking network cables into smaller pieces
- B. Separating network hardware from software layers
- C. A type of encryption for network traffic
- D. Combining all network functions into one box

Answer: B

Q695. What is the P4 programming language used for in networking?

- A. Web development and frontend design
- B. Database management and SQL queries
- C. Cloud computing resource provisioning
- D. Programming how devices process packets

Answer: D

Q696. What is Quantum Key Distribution (QKD) in the context of next-gen network security?

- A. Distributing physical keys by courier
- B. A cloud service for key management
- C. A type of VPN tunnel technology
- D. Using quantum mechanics for provably secure keys

Answer: D

Q697. What is Time-Sensitive Networking (TSN) and which industries benefit from it?

- A. A wireless protocol for mobile devices only
- B. IEEE standards for deterministic Ethernet communication
- C. A faster version of standard Ethernet protocol
- D. A cloud service for time synchronization

Answer: B

Q698. What is the concept of network as a service (NaaS)?

- A. Free networking for all users
- B. Networking infrastructure provided on-demand as a service
- C. A programming language for networks
- D. A type of physical network router

Answer: B

Q699. What is the significance of Open RAN (O-RAN) in 5G networks?

- A. Disaggregating RAN with open interfaces for interop
- B. A security protocol for radio access
- C. A new radio frequency band for 5G
- D. A type of antenna for signal boosting

Answer: A

Q700. What is the concept of self-organizing networks (SON) in mobile communications?

- A. A manual configuration tool for base stations
- B. A type of routing protocol for mobile data
- C. Networks that physically rearrange themselves
- D. Networks that auto-configure, optimize, and self-heal

Answer: D

Q701. A noiseless channel has 5 kHz bandwidth and uses 16 signal levels. What is the maximum data rate per Nyquist?

- A. 80,000 bps
- B. 40,000 bps
- C. 20,000 bps
- D. 60,000 bps

Answer: B

Q702. In 64-QAM, how many bits are represented by each signal element?

- A. 6 bits per element
- B. 8 bits per element
- C. 4 bits per element
- D. 5 bits per element

Answer: A

Q703. What is the primary limitation of increasing signal levels to boost data rate?

- A. It increases the physical size of transmitters
- B. It makes the receiver more susceptible to noise
- C. It reduces the frequency of carrier signals
- D. It requires longer transmission cables overall

Answer: B

Q704. A signal has a power of 50 mW and noise power of 5 uW. What is the SNR in decibels?

- A. 30 dB
- B. 20 dB
- C. 40 dB
- D. 50 dB

Answer: C

Q705. In spread spectrum communication, what is the purpose of the chipping code?

- A. It compresses data to reduce bandwidth usage
- B. It encrypts data using symmetric key algorithms
- C. It corrects errors in received data automatically
- D. It spreads the signal across a wider frequency band

Answer: D

Q706. Why does frequency hopping spread spectrum (FHSS) change frequencies during transmission?

- A. To increase the overall data throughput rate
- B. To reduce power consumption at the transmitter
- C. To synchronize timing between sender and receiver
- D. To avoid prolonged interference on any one band

Answer: D

Q707. What is the impact of jitter on real-time data communication applications?

- A. It causes variation in packet arrival times
- B. It reduces the overall power consumption
- C. It increases the bandwidth of the channel
- D. It permanently corrupts all transmitted data

Answer: A

Q708. In PCM, what determines the number of quantization levels?

- A. The carrier frequency of the channel used
- B. The sampling rate of the analog signal
- C. The number of bits per sample allocated
- D. The bandwidth of the transmission medium

Answer: C

Q709. What is quantization noise and when does it become significant?

- A. Distortion from using excessively long fiber optic cables
- B. Interference caused by adjacent frequency band signals
- C. Noise from power supply fluctuations in the transmitter
- D. Error introduced by rounding analog values to discrete levels

Answer: D

Q710. How does OFDM improve data transmission in broadband systems?

- A. By amplifying signal power at regular interval points
- B. By splitting data across many orthogonal subcarriers
- C. By compressing data before modulation on the channel
- D. By using a single high-frequency carrier for all data

Answer: B

Q711. In a hierarchical network design, what is the function of the distribution layer?

- A. It enforces policies, filtering, and inter-VLAN routing
- B. It provides high-speed backbone connectivity between sites
- C. It manages wireless access points and mobile devices
- D. It connects end-user devices to the network directly

Answer: A

Q712. What problem does the Spanning Tree Protocol (STP) solve?

- A. It prevents IP address conflicts on the network
- B. It encrypts data flowing between network switches
- C. It eliminates broadcast storms from Layer 2 loops
- D. It balances traffic load across multiple gateways

Answer: C

Q713. How does software-defined networking (SDN) differ from traditional networking?

- A. SDN eliminates the need for any network protocols
- B. SDN separates the control plane from the data plane
- C. SDN replaces switches with direct cable connections
- D. SDN uses only wireless media for all communication

Answer: B

Q714. What is the purpose of MPLS in modern networks?

- A. To assign MAC addresses to network interfaces dynamically
- B. To encrypt all traffic between two endpoint routers
- C. To translate domain names into corresponding IP addresses
- D. To label packets for fast forwarding along predefined paths

Answer: D

Q715. In the context of QoS, what is traffic shaping?

- A. Controlling the rate of outgoing traffic for compliance
- B. Compressing packets to reduce overall bandwidth usage
- C. Physically rerouting cables for optimal signal paths
- D. Encrypting traffic to prevent unauthorized access

Answer: A

Q716. What is the role of an autonomous system (AS) in internet routing?

- A. A device that converts between IPv4 and IPv6 addresses
- B. A protocol for encrypting data across network boundaries
- C. A collection of routers under a single administrative policy
- D. A single router that manages all global internet traffic

Answer: C

Q717. How does network function virtualization (NFV) reduce costs?

- A. By removing the need for IP addressing in modern networks
- B. By replacing dedicated hardware with software on commodity servers
- C. By eliminating all network cables and using wireless only
- D. By using a single protocol for all communication types

Answer: B

Q718. What is the convergence time of a routing protocol?

- A. The interval between consecutive routing table backups
- B. The time taken to encrypt all routing table entries
- C. The delay between sending and receiving a single packet
- D. The time for all routers to agree on the network topology

Answer: D

Q719. What is the primary advantage of intent-based networking (IBN)?

- A. It allows network configuration through business intent policies
- B. It eliminates security threats through hardware isolation
- C. It provides unlimited bandwidth to all connected devices
- D. It removes the need for switches and routers entirely

Answer: A

Q720. How does segment routing simplify traffic engineering in networks?

- A. By physically segmenting the network into isolated zones
- B. By compressing routing tables to reduce memory usage
- C. By encrypting each packet segment individually for security
- D. By encoding the forwarding path in the packet header itself

Answer: D

Q721. In Go-Back-N with a window size of 7, what is the minimum number of sequence number bits needed?

- A. 3 bits
- B. 5 bits
- C. 4 bits
- D. 2 bits

Answer: A

Q722. Why can't the window size in Go-Back-N equal the sequence number space?

- A. The sender buffer would overflow causing system failures
- B. The receiver could confuse new frames with retransmissions
- C. It would exceed the maximum transmission unit size limit
- D. It would require more bandwidth than the link can handle

Answer: B

Q723. What is the maximum efficiency of pure ALOHA protocol?

- A. 12.5 percent
- B. 36.8 percent
- C. 18.4 percent
- D. 50.0 percent

Answer: C

Q724. How does slotted ALOHA improve over pure ALOHA?

- A. It doubles throughput by restricting starts to time slots
- B. It uses frequency division to separate all transmissions
- C. It eliminates collisions entirely through token passing
- D. It compresses frames to reduce the chance of overlap

Answer: A

Q725. In Selective Repeat, what is the maximum window size for n-bit sequence numbers?

- A. $2^{(n+1)}$
- B. $2^{(n-1)}$
- C. 2^n minus one
- D. 2^n

Answer: B

Q726. What problem does the hidden terminal issue cause in wireless LANs?

- A. The access point fails to detect any transmitted frames
- B. All devices on the network lose their IP configurations
- C. Encrypted frames are decrypted by unauthorized receivers
- D. Two stations that cannot hear each other collide at the AP

Answer: D

Q727. How does the RTS/CTS mechanism address the hidden terminal problem?

- A. By reserving the channel before data frame transmission
- B. By increasing transmission power to reach all stations
- C. By encrypting frames to prevent unauthorized decoding
- D. By assigning fixed frequencies to each wireless device

Answer: A

Q728. What is the minimum frame size requirement in standard Ethernet and why?

- A. 256 bytes, to guarantee sufficient error correction data space
- B. 64 bytes, to ensure collision detection before transmission ends
- C. 32 bytes, to maintain compatibility with token ring networks
- D. 128 bytes, to provide enough space for header and payload

Answer: B

Q729. What is the difference between persistent and non-persistent CSMA?

- A. Persistent CSMA continuously monitors; non-persistent waits randomly
- B. Persistent CSMA encrypts data; non-persistent sends it plainly
- C. Persistent CSMA uses tokens; non-persistent uses frequency slots
- D. Persistent CSMA is wireless only; non-persistent is wired only

Answer: A

Q730. How does 802.1Q VLAN tagging modify an Ethernet frame?

- A. It appends VLAN information after the frame check sequence
- B. It changes the preamble to include VLAN routing information
- C. It replaces the destination MAC address with a VLAN identifier
- D. It inserts a 4-byte tag between source MAC and type field

Answer: D

Q731. Given the IP address 172.16.45.0/26, how many usable host addresses are available?

- A. 64 hosts
- B. 126 hosts
- C. 62 hosts
- D. 30 hosts

Answer: C

Q732. What is the purpose of route summarization (supernetting)?

- A. It assigns additional IP addresses to overloaded subnets
- B. It combines multiple smaller routes into a single entry
- C. It splits a large route into more specific smaller routes
- D. It encrypts routing updates between adjacent routers

Answer: B

Q733. How does BGP differ from OSPF in terms of routing scope?

- A. BGP and OSPF both operate only within a single AS domain
- B. BGP handles inter-AS routing; OSPF is for intra-AS routing
- C. BGP is for LANs only; OSPF is for WANs exclusively used
- D. BGP is for intra-AS routing; OSPF handles inter-AS routing

Answer: B

Q734. What is the role of the Protocol field in the IPv4 header?

- A. It identifies which transport layer protocol is encapsulated
- B. It contains the checksum for the IP header verification
- C. It indicates the total number of fragments for this packet
- D. It specifies the version of IP being used for routing

Answer: A

Q735. What mechanism does OSPF use to handle large networks efficiently?

- A. It uses only static routes for all network destinations
- B. It limits the number of routers to 50 per network
- C. It compresses all LSAs into a single combined broadcast
- D. It divides the network into areas with Area 0 backbone

Answer: D

Q736. In IPv6, what is the purpose of the flow label field?

- A. It identifies packets belonging to the same traffic flow
- B. It stores the source and destination MAC addresses
- C. It contains the fragment offset for reassembly process
- D. It specifies the maximum hop limit for the IP packet

Answer: A

Q737. How does Anycast addressing work in IPv6?

- A. It delivers packets to the nearest device in a group
- B. It assigns unique addresses to every network interface
- C. It sends packets to all devices in the network at once
- D. It translates IPv4 addresses to IPv6 format directly

Answer: A

Q738. What is the tunneling technique used for IPv6 transition?

- A. It encapsulates IPv6 packets inside IPv4 packets for transit
- B. It drops IPv6 packets when they reach an IPv4 only network
- C. It converts IPv6 addresses into compatible IPv4 addresses
- D. It replaces all IPv4 routers with IPv6 routers instantly

Answer: A

Q739. What is the purpose of ECMP (Equal-Cost Multi-Path) routing?

- A. It assigns different costs to each available network path
- B. It encrypts traffic flowing across multiple network paths
- C. It selects the single lowest cost path for all packets
- D. It distributes traffic across multiple equal-cost routes

Answer: D

Q740. How does Path MTU Discovery work in IP networking?

- A. It finds the smallest MTU along the path to avoid fragmentation
- B. It measures the total bandwidth available on the route
- C. It identifies all routers along the path using ICMP requests
- D. It discovers the fastest path between source and receiver

Answer: A

Q741. How does TCP Reno handle packet loss detection through triple duplicate ACKs?

- A. It halves the congestion window and enters fast recovery phase
- B. It resets the congestion window to one segment immediately
- C. It doubles the congestion window to recover lost bandwidth
- D. It closes the connection and re-establishes a new session

Answer: A

Q742. What is the silly window syndrome in TCP?

- A. When tiny amounts of data are exchanged very inefficiently
- B. When large segments are sent through low bandwidth links
- C. When DNS resolution takes excessively long to complete
- D. When the routing table contains incorrect gateway entries

Answer: A

Q743. What is the purpose of TCP SACK (Selective Acknowledgment)?

- A. It encrypts acknowledgments to prevent spoofing attacks
- B. It selects which port number to use for the connection
- C. It allows the receiver to report non-contiguous blocks received
- D. It aggregates multiple acknowledgments into a single message

Answer: C

Q744. How does TCP handle the TIME_WAIT state after connection close?

- A. It sends additional data segments to flush the pipeline
- B. It immediately releases all resources for the connection
- C. It waits for 2x MSL to ensure all packets are delivered
- D. It resets the sequence numbers for the next connection

Answer: C

Q745. What is the TCP slow start threshold (ssthresh) used for?

- A. It sets the minimum segment size for transmission allowed
- B. It defines the timeout period for retransmission of lost segments
- C. It determines when to switch from slow start to congestion avoidance
- D. It limits the maximum number of simultaneous TCP connections

Answer: C

Q746. What problem does TCP head-of-line blocking cause?

- A. It prevents new connections from being established on server
- B. A lost segment delays delivery of all subsequent segments
- C. It blocks DNS resolution until the TCP connection times out
- D. It causes the routing table to become incorrectly populated

Answer: B

Q747. How does QUIC (used in HTTP/3) improve over TCP+TLS?

- A. It replaces IP addressing with a new addressing scheme used
- B. It uses larger packet sizes for higher throughput overall
- C. It combines transport and encryption in fewer round trips
- D. It eliminates the need for any congestion control mechanism

Answer: C

Q748. What is the purpose of ECN (Explicit Congestion Notification)?

- A. To notify applications when their port number is in use
- B. To signal congestion without dropping packets at the router
- C. To extend the congestion window beyond its normal limits
- D. To encrypt traffic flowing through congested network links

Answer: B

Q749. What is the difference between TCP Tahoe and TCP Reno on timeout?

- A. Tahoe uses fast recovery; Reno resets cwnd to one segment
- B. Both reset cwnd to one on timeout, Reno adds fast recovery for dup ACKs
- C. Reno uses slow start only; Tahoe uses congestion avoidance only
- D. Tahoe closes the connection; Reno continues transmitting data

Answer: B

Q750. What is Multipath TCP (MPTCP) designed to achieve?

- A. To reduce the MTU size for faster processing by routers
- B. To encrypt multiple TCP streams using a single shared key
- C. To use multiple network paths for a single TCP connection
- D. To create multiple connections using a single source port

Answer: C

Q751. What is DNS amplification attack and how does it work?

- A. Attacker spoofs the victim's IP and sends small queries to get large replies
- B. Attacker overwhelms DNS servers by sending millions of registration requests
- C. Attacker intercepts DNS responses and replaces them with false information
- D. Attacker modifies DNS cache entries to redirect users to malicious websites

Answer: A

Q752. What is the purpose of DNSSEC (DNS Security Extensions)?

- A. It filters malicious domain names from DNS resolution
- B. It replaces DNS with a more secure naming system entirely
- C. It encrypts all DNS query and response traffic end-to-end
- D. It authenticates DNS responses using digital signatures

Answer: D

Q753. How does HTTP/2 improve performance over HTTP/1.1?

- A. It multiplexes multiple requests over a single connection
- B. It uses a text-based protocol for easier debugging
- C. It increases the maximum URL length for complex queries
- D. It eliminates the need for TLS encryption on all connections

Answer: A

Q754. What is the CAP theorem in distributed application systems?

- A. A system must choose between Caching, Addressing, and Protocol compatibility
- B. A system can have Compression, Authentication, and Performance simultaneously
- C. A system balances Concurrency, Atomicity, and Persistence for transactions
- D. A system can guarantee only two of Consistency, Availability, and Partition tolerance

Answer: D

Q755. What is the purpose of WebSocket protocol?

- A. It provides a new addressing scheme for web-based resources
- B. It enables full-duplex communication over a single TCP connection
- C. It replaces HTTP entirely for all web communications now
- D. It compresses static web content for faster initial page loads

Answer: B

Q756. What is the difference between gRPC and REST APIs?

- A. gRPC works without any network; REST requires internet connectivity
- B. gRPC uses Protocol Buffers and HTTP/2; REST uses JSON and HTTP/1.1 typically
- C. gRPC is only for web browsers; REST is for all client applications
- D. gRPC uses XML encoding; REST uses binary serialization for data

Answer: B

Q757. How does DNS over HTTPS (DoH) improve privacy?

- A. It stores DNS records locally to avoid network queries
- B. It blocks all tracking cookies from websites visited
- C. It encrypts DNS queries within HTTPS to prevent snooping
- D. It anonymizes the user's IP address from DNS servers

Answer: C

Q758. What is the purpose of HSTS (HTTP Strict Transport Security)?

- A. It blocks HTTP traffic from reaching the web server entirely
- B. It compresses HTTPS traffic for faster encrypted page loads
- C. It generates SSL certificates automatically for all domains
- D. It forces browsers to always use HTTPS for a specific domain

Answer: D

Q759. What is GraphQL's main advantage over traditional REST APIs?

- A. It eliminates the need for a database on the server backend
- B. It automatically encrypts all data exchanged between nodes
- C. It provides built-in authentication for all API endpoints
- D. It allows clients to request exactly the data fields they need

Answer: D

Q760. What is the role of service mesh in microservices architecture?

- A. It manages inter-service communication, observability, and security
- B. It provides a physical network topology for container hosting
- C. It stores persistent data shared across all microservice nodes
- D. It replaces all application code with auto-generated services

Answer: A

Q761. What is the KRACK attack on WPA2 wireless networks?

- A. It exploits the four-way handshake to reinstall encryption keys
- B. It jams the wireless frequency to deny service to all users
- C. It brute-forces the WiFi password using parallel computing
- D. It clones the access point to create a fake network nearby

Answer: A

Q762. How does Massive MIMO differ from conventional MIMO systems?

- A. It uses a single antenna instead of multiple antenna arrays
- B. It employs hundreds of antennas at the base station for capacity
- C. It eliminates the need for any digital signal processing steps
- D. It operates exclusively in the sub-1 GHz frequency band range

Answer: B

Q763. What is the purpose of network slicing in 5G networks?

- A. It splits the frequency spectrum into equal fixed portions
- B. It creates virtual networks tailored to specific use cases
- C. It physically divides the antenna into separate segments
- D. It reduces the cell tower height for urban deployments

Answer: B

Q764. What is the near-far problem in wireless CDMA systems?

- A. A strong nearby signal overwhelms weaker signals from far users
- B. Near users experience more interference than far away users
- C. Distant users are unable to connect to any base station
- D. The antenna cannot distinguish between near and far frequencies

Answer: A

Q765. What is the role of the EPC (Evolved Packet Core) in LTE?

- A. It provides power supply management for all base stations
- B. It manages the physical radio transmission between towers
- C. It controls the display settings on mobile user devices
- D. It handles core network functions like authentication and routing

Answer: D

Q766. How does carrier aggregation improve LTE-Advanced performance?

- A. It compresses data before transmission to save bandwidth space
- B. It reduces the number of connected devices to improve speed
- C. It combines multiple frequency bands for greater total bandwidth
- D. It increases the transmission power of the base station antenna

Answer: C

Q767. What is the purpose of the BSS color field in WiFi 6?

- A. It identifies the color scheme for network management GUIs
- B. It determines the physical location of each connected device
- C. It assigns priority levels based on the device manufacturer
- D. It distinguishes overlapping BSSs to reduce co-channel interference

Answer: D

Q768. What is the main challenge of millimeter wave (mmWave) in 5G?

- A. It provides extremely low data rates unsuitable for broadband
- B. It has very limited range and is easily blocked by obstacles
- C. It only works in underwater communication environments now
- D. It requires excessively large antennas that are impractical

Answer: B

Q769. What is Device-to-Device (D2D) communication in LTE/5G?

- A. Communication between two base stations without user devices
- B. Transfer of device configuration settings during factory reset process
- C. Communication between a device and its cloud backup storage server
- D. Direct communication between nearby devices without base station relay

Answer: D

Q770. How does full-duplex wireless communication differ from current systems?

- A. It uses two separate antennas for upload and download traffic
- B. It doubles the frequency band allocation for each user device
- C. It transmits and receives on the same frequency simultaneously
- D. It alternates between transmission and reception at high speed

Answer: C

Q771. What is the purpose of Kubernetes in cloud-native architectures?

- A. It orchestrates deployment, scaling, and management of containers
- B. It provides virtual machine hypervisor management for data centers
- C. It encrypts all network traffic between microservices automatically
- D. It replaces traditional databases with distributed file storage

Answer: A

Q772. What is the concept of immutable infrastructure in cloud?

- A. Infrastructure that is immune to all security vulnerabilities
- B. Hardware that never needs physical maintenance or replacement
- C. Infrastructure that cannot be modified after initial deployment
- D. Servers that automatically update their own operating systems

Answer: C

Q773. What is a service level agreement (SLA) in cloud computing?

- A. A software license for cloud application source code access
- B. A network protocol for secure cloud data transfer operations
- C. A contract defining guaranteed performance and uptime metrics
- D. A hardware specification for physical cloud server components

Answer: C

Q774. How does a cloud-native application differ from a traditional application?

- A. It requires manual scaling and dedicated hardware for deployment
- B. It uses monolithic architecture deployed on single physical servers
- C. It runs only on desktop computers without internet access
- D. It is designed for cloud environments using microservices and containers

Answer: D

Q775. What is the purpose of Infrastructure as Code (IaC)?

- A. Managing and provisioning infrastructure through code and templates
- B. Writing application code that runs on cloud infrastructure
- C. Converting physical servers into virtual machines automatically
- D. Encoding network traffic for secure cloud communication links

Answer: A

Q776. What is cloud bursting and when is it used?

- A. When data is deleted from cloud storage due to inactivity timeout
- B. When multiple cloud providers merge their resources into one pool
- C. When a cloud server physically overheats and shuts itself down
- D. When private cloud overflows workloads to public cloud during peaks

Answer: D

Q777. What is the purpose of a chaos engineering approach in cloud?

- A. To overload systems deliberately to find performance limits
- B. To randomly generate test data for application benchmarking
- C. To create confusing network topologies that test admin skills
- D. To intentionally introduce failures and test system resilience

Answer: D

Q778. What is the difference between cloud-native and cloud-enabled applications?

- A. Cloud-native is less scalable; cloud-enabled offers better performance
- B. Cloud-native is built for cloud; cloud-enabled is migrated legacy apps
- C. Cloud-native uses containers; cloud-enabled uses only virtual machines
- D. Cloud-native requires bare metal; cloud-enabled works on any platform

Answer: B

Q779. What is edge computing and how does it relate to cloud?

- A. It replaces cloud computing entirely with local device processing
- B. It stores all data at the center of the cloud data center only
- C. It processes data at the network edge, reducing latency to cloud
- D. It encrypts data at the boundary of the corporate network only

Answer: C

Q780. What is the FinOps framework in cloud computing?

- A. A financial auditing tool for cloud provider tax compliance
- B. A protocol for encrypting financial transactions in the cloud
- C. A cloud-based accounting software for small business finance
- D. A practice for managing cloud costs through collaboration and data

Answer: D

Q781. What is a zero-day vulnerability?

- A. A security issue that affects zero users in the network
- B. A bug that was fixed exactly zero days after discovery
- C. A vulnerability that only exists for zero days before expiring
- D. A flaw exploited before the vendor has released any patch

Answer: D

Q782. How does IPSec provide security at the network layer?

- A. It replaces IP addresses with encrypted identifiers in headers
- B. It filters packets based on source and destination IP only
- C. It blocks all ICMP traffic to prevent network reconnaissance
- D. It encrypts and authenticates IP packets using AH and ESP

Answer: D

Q783. What is the purpose of a honeypot in network security?

- A. It provides backup power to critical network infrastructure
- B. It distributes network traffic across multiple security zones
- C. It encrypts sensitive data stored on production file servers
- D. It acts as a decoy system to attract and study attackers

Answer: D

Q784. What is the Diffie-Hellman key exchange used for?

- A. It securely establishes a shared secret over an insecure channel
- B. It authenticates users using their biometric data remotely
- C. It compresses public keys for efficient network transmission
- D. It encrypts stored files on shared network drives securely

Answer: A

Q785. What is ARP spoofing and what can it enable?

- A. It intercepts DHCP requests to assign wrong IP addresses
- B. It floods the network with fake ARP requests for DoS attacks
- C. It modifies DNS records to redirect traffic to malicious sites
- D. It sends fake ARP replies to link attacker's MAC to victim's IP

Answer: D

Q786. What is the purpose of certificate pinning in mobile apps?

- A. It creates a backup copy of the server certificate on device
- B. It pins certificate files to email attachments for verification
- C. It stores certificates locally to work offline without network
- D. It restricts which certificates are trusted for a specific domain

Answer: D

Q787. What is the concept of zero trust security architecture?

- A. It never trusts and always verifies every access request made
- B. It trusts all internal network traffic without verification
- C. It grants permanent access after initial identity verification
- D. It removes all firewalls to simplify the network architecture

Answer: A

Q788. How does a SYN flood attack exploit the TCP handshake?

- A. It sends many SYN requests without completing the handshake
- B. It completes thousands of handshakes to exhaust server memory
- C. It modifies the SYN packet to inject malicious code payloads
- D. It intercepts SYN-ACK packets to steal session credentials

Answer: A

Q789. What is the role of SIEM in enterprise security?

- A. It provides physical access control for server room facilities
- B. It encrypts all stored data across the enterprise network
- C. It replaces firewalls with software-defined security boundaries
- D. It aggregates and analyzes security events from multiple sources

Answer: D

Q790. What is perfect forward secrecy (PFS) in TLS?

- A. It ensures past sessions cannot be decrypted if long-term key leaks
- B. It guarantees that no data will ever be lost in transit
- C. It provides unlimited bandwidth for encrypted network connections
- D. It permanently stores encryption keys for future session recovery

Answer: A

Q791. What is the concept of network disaggregation?

- A. Separating network hardware from software for vendor flexibility
- B. Dividing the network into smaller broadcast domains only
- C. Removing all redundant paths from the network topology
- D. Combining all network functions into a single proprietary box

Answer: A

Q792. How does quantum key distribution (QKD) enhance security?

- A. It replaces all digital certificates with quantum identifiers
- B. It generates longer passwords using quantum random number math
- C. It encrypts data using the physical properties of fiber optics
- D. It uses quantum mechanics to create unbreakable encryption keys

Answer: D

Q793. What is deterministic networking and why is it important?

- A. Networks that randomly route traffic for load balancing
- B. Networks that only operate during predetermined time windows
- C. Networks that determine user identity before allowing access
- D. Networks guaranteeing bounded latency and zero packet loss

Answer: D

Q794. What is the purpose of a digital twin network (DTN)?

- A. A second physical network installed for disaster recovery
- B. A virtual model of a network for simulation and optimization
- C. A backup network that activates when the primary fails
- D. A network that only carries duplicated data for redundancy

Answer: B

Q795. What challenges does post-quantum cryptography aim to address?

- A. The high cost of managing digital certificates across networks
- B. The latency introduced by current encryption and decryption steps
- C. The threat of quantum computers breaking current encryption methods
- D. The difficulty of implementing encryption on IoT sensor devices

Answer: C

Q796. What is the concept of programmable data planes in next-gen networks?

- A. Data planes that can be programmed by end users directly
- B. Network connections that adjust their physical medium dynamically
- C. Data storage planes that automatically organize stored data files
- D. Network hardware whose packet processing logic can be defined in code

Answer: D

Q797. How does federated learning benefit network intelligence?

- A. It centralizes all data in one location for faster processing
- B. It federates multiple networks into one large unified network
- C. It trains AI models across distributed nodes without sharing raw data
- D. It distributes stored data evenly across all network node devices

Answer: C

Q798. What is the role of SASE in next-generation enterprise networking?

- A. It provides physical security for enterprise server room access
- B. It replaces all enterprise applications with cloud alternatives
- C. It separates network management from security administration
- D. It converges networking and security into a unified cloud service

Answer: D

Q799. What is the concept of computing continuum from cloud to edge?

- A. A method of continuously compressing data as it crosses network
- B. A single massive computer that handles all processing tasks
- C. A seamless spectrum of computing resources from cloud to device
- D. A continuous network connection that never drops under any load

Answer: C

Q800. How does 6G envision advancing beyond 5G capabilities?

- A. By eliminating all wireless communication in favor of wired only
- B. By reducing network speeds to improve reliability and coverage
- C. By providing terahertz frequencies, AI integration, and sensing abilities
- D. By using only ground-based infrastructure without any satellites

Answer: C

Q801. In a 256-QAM modulation scheme, how many bits does each symbol carry?

- A. 10 bits per symbol in the constellation
- B. 8 bits per symbol in the constellation
- C. 12 bits per symbol in the constellation
- D. 6 bits per symbol in the constellation

Answer: B

Q802. A noisy channel has a bandwidth of 3 kHz and SNR of 127. What is the Shannon capacity?

- A. 35 kbps maximum capacity
- B. 28 kbps maximum capacity
- C. 14 kbps maximum capacity
- D. 21 kbps maximum capacity

Answer: D

Q803. What is the fundamental problem with unipolar NRZ encoding for long data streams?

- A. It consumes excessive bandwidth per symbol
- B. It cannot operate above 1 Mbps data rates
- C. It introduces a DC component causing baseline wander
- D. It requires complex decoding circuitry at receivers

Answer: C

Q804. In Direct Sequence Spread Spectrum (DSSS), what determines the spreading factor?

- A. The ratio of chip rate to data bit rate in the system
- B. The number of time slots available in each frame period
- C. The amplitude variation range of the modulated signal
- D. The number of frequency hops per second in the pattern

Answer: A

Q805. Why is the Shannon limit considered a theoretical upper bound for channel capacity?

- A. It assumes error-free coding with infinite complexity at the encoder
- B. It assumes zero latency in all signal processing hardware stages
- C. It assumes synchronous clocking across all connected network nodes
- D. It assumes ideal filtering with zero guard bands between channels

Answer: A

Q806. What distinguishes Differential Phase Shift Keying (DPSK) from conventional PSK?

- A. DPSK requires a coherent carrier reference at the receiver
- B. DPSK uses amplitude changes instead of phase changes
- C. DPSK operates at double the frequency of conventional PSK
- D. DPSK encodes data as changes between consecutive phase values

Answer: D

Q807. What is the Nyquist bandwidth for a channel carrying a signal with a baud rate of 10,000 symbols/s?

- A. 10 kHz minimum channel bandwidth
- B. 5 kHz minimum channel bandwidth
- C. 20 kHz minimum channel bandwidth
- D. 15 kHz minimum channel bandwidth

Answer: B

Q808. In CDMA, what happens when a receiver multiplies a received composite signal by a station's chip sequence?

- A. It amplifies signals from all stations equally for decoding
- B. It compresses the composite signal to reduce overall bandwidth
- C. It filters out the target station's data by orthogonal separation
- D. It encrypts the signal with the station's unique security key

Answer: C

Q809. What is the main advantage of OFDM over single-carrier modulation in wideband channels?

- A. OFDM eliminates the need for error correction coding entirely
- B. OFDM is resilient to multipath fading by using parallel subcarriers
- C. OFDM doubles the available spectrum through signal compression
- D. OFDM reduces transmitter power consumption significantly

Answer: B

Q810. A PCM system samples a signal at 16,000 samples/s using 12-bit quantization. What is the resulting data rate?

- A. 128 kbps output data stream
- B. 224 kbps output data stream
- C. 192 kbps output data stream
- D. 160 kbps output data stream

Answer: C

Q811. What is the primary purpose of the Rapid Spanning Tree Protocol (RSTP) over STP?

- A. RSTP enables routing capabilities within Layer 2 switch devices
- B. RSTP provides significantly faster convergence after topology changes
- C. RSTP supports more VLANs per switch than standard STP protocol
- D. RSTP eliminates the need for root bridge election in networks

Answer: B

Q812. In the context of SDN, what is the function of the southbound API?

- A. It enables communication between the controller and network devices
- B. It provides network analytics to management applications
- C. It connects external applications to the SDN controller layer
- D. It manages user authentication for the SDN control plane

Answer: A

Q813. What is the purpose of ECMP (Equal-Cost Multi-Path) in network design?

- A. To encrypt multiple paths for secure data delivery
- B. To create backup paths that activate only during failure
- C. To distribute traffic across multiple equal-cost routes
- D. To select the single lowest-cost path and block others

Answer: C

Q814. What is a fabric network architecture in data center design?

- A. A legacy three-tier design using core, distribution, and access layers
- B. A bus topology where all servers share a single backbone cable
- C. A ring-based topology connecting servers in a circular pattern
- D. A flat, scalable topology using spine-and-leaf switch arrangement

Answer: D

Q815. What is control plane policing (CoPP) used for in network infrastructure?

- A. To protect the router's control plane from excessive traffic
- B. To prioritize user data traffic over management traffic
- C. To balance routing protocol traffic across multiple interfaces
- D. To encrypt all control plane communications between routers

Answer: A

Q816. What is the difference between in-band and out-of-band network management?

- A. In-band management works only for wired networks and out-of-band only for wireless
- B. In-band uses production network paths while out-of-band uses dedicated management channels
- C. Out-of-band uses production network paths while in-band uses dedicated management channels
- D. Out-of-band management requires specialized hardware while in-band uses standard NICs

Answer: B

Q817. What is the purpose of BFD (Bidirectional Forwarding Detection) in networks?

- A. To authenticate routing protocol messages between peer routers
- B. To encrypt bidirectional traffic between two network devices
- C. To balance network traffic bidirectionally across redundant paths
- D. To provide rapid failure detection between adjacent forwarding engines

Answer: D

Q818. In MPLS, what is the purpose of the label distribution protocol (LDP)?

- A. It converts IP headers into MPLS labels at network edges
- B. It establishes label-switched paths by distributing label bindings
- C. It encrypts MPLS labels to prevent label spoofing attacks
- D. It monitors label utilization and optimizes path selection

Answer: B

Q819. What is the concept of network telemetry and how does it differ from SNMP polling?

- A. Telemetry pushes data in real-time while SNMP uses periodic polling
- B. Telemetry is limited to local networks while SNMP works globally
- C. Telemetry requires manual data collection while SNMP is automated
- D. Telemetry only monitors hardware while SNMP monitors software too

Answer: A

Q820. What is the role of a route reflector in large-scale BGP deployments?

- A. It converts eBGP routes into OSPF routes for internal distribution
- B. It encrypts BGP update messages between internal BGP speakers
- C. It reduces the need for full mesh iBGP peering by reflecting routes
- D. It filters malicious BGP routes from external autonomous systems

Answer: C

Q821. What is the throughput efficiency of the pure ALOHA protocol expressed as a fraction?

- A. Approximately 18.4% maximum channel utilization
- B. Approximately 36.8% maximum channel utilization
- C. Approximately 50.0% maximum channel utilization
- D. Approximately 25.0% maximum channel utilization

Answer: A

Q822. In the CRC error detection scheme, if the generator polynomial is $x^4 + x + 1$, how many check bits are appended?

- A. 5 check bits are appended to the data
- B. 3 check bits are appended to the data
- C. 6 check bits are appended to the data
- D. 4 check bits are appended to the data

Answer: D

Q823. What is the maximum throughput efficiency of slotted ALOHA compared to pure ALOHA?

- A. Slotted ALOHA triples throughput to about 55.2% by using time slots
- B. Slotted ALOHA quadruples throughput to about 73.6% by time slots
- C. Slotted ALOHA doubles throughput to about 36.8% by using time slots
- D. Slotted ALOHA achieves the same throughput as pure ALOHA protocol

Answer: C

Q824. In a Hamming code with 11-bit data words, how many redundancy bits are required?

- A. 3 redundancy bits for 11 data bits
- B. 4 redundancy bits for 11 data bits
- C. 5 redundancy bits for 11 data bits
- D. 6 redundancy bits for 11 data bits

Answer: B

Q825. What is the difference between 1-persistent and p-persistent CSMA?

- A. p-persistent transmits immediately when idle while 1-persistent transmits with probability p
- B. p-persistent is used only in wireless networks while 1-persistent is for wired only
- C. 1-persistent waits a random time while p-persistent waits a fixed time period
- D. 1-persistent transmits immediately when idle while p-persistent transmits with probability p

Answer: D

Q826. What is the purpose of the BPDU (Bridge Protocol Data Unit) in spanning tree operation?

- A. It exchanges topology information between bridges for loop prevention
- B. It carries user data frames between bridges in the network
- C. It encrypts forwarding table entries for secure bridge operation
- D. It assigns IP addresses to bridge ports during initialization

Answer: A

Q827. In Link Aggregation (IEEE 802.3ad), what is the primary benefit of combining multiple physical links?

- A. It enables different speed links to operate at the same data rate
- B. It increases bandwidth and provides link redundancy as one logical link
- C. It eliminates the need for spanning tree on aggregated link ports
- D. It reduces the MAC address table size on connected switches

Answer: B

Q828. What is the channel capture effect in CSMA/CD networks?

- A. When the broadcast domain captures all unicast frames on the segment
- B. When a switch port captures and mirrors all traffic for monitoring
- C. When one station dominates the channel because its backoff counter resets to zero after success
- D. When a collision causes all stations to simultaneously attempt retransmission

Answer: C

Q829. What is the purpose of the 802.1Q native VLAN concept?

- A. It handles untagged frames on a trunk link by assigning them to a default VLAN
- B. It reserves bandwidth for management traffic between switches
- C. It encrypts inter-VLAN routing traffic on Layer 3 switches
- D. It assigns all unknown unicast traffic to a quarantine VLAN

Answer: A

Q830. How does the token bucket algorithm differ from the leaky bucket algorithm for traffic shaping?

- A. Token bucket discards excess tokens while leaky bucket stores them indefinitely
- B. Leaky bucket allows bursts up to a limit while token bucket enforces a constant rate
- C. Leaky bucket uses tokens for rate control while token bucket uses timers only
- D. Token bucket allows bursts up to a limit while leaky bucket enforces a constant rate

Answer: D

Q831. Given the address 10.45.128.0/17, how many host addresses are available in this subnet?

- A. 32,766 usable host addresses available
- B. 131,070 usable host addresses available
- C. 65,534 usable host addresses available
- D. 16,382 usable host addresses available

Answer: A

Q832. In OSPF, what is the purpose of stub areas?

- A. They serve as the backbone area connecting all other OSPF areas
- B. They handle inter-area routing between non-backbone OSPF areas
- C. They contain all external routes and redistribute them to other areas
- D. They block external route advertisements to reduce routing table size

Answer: D

Q833. What is the purpose of the BGP LOCAL_PREF attribute?

- A. It tracks the sequence of autonomous systems a route has traversed
- B. It specifies the physical distance to the destination network prefix
- C. It indicates the preference for exit paths within an autonomous system
- D. It assigns a community tag for route filtering between providers

Answer: C

Q834. What is the difference between NAT and PAT (Port Address Translation)?

- A. PAT maps each private IP to a unique public IP while NAT uses port-based multiplexing
- B. PAT supports only TCP traffic while NAT supports both TCP and UDP protocols
- C. NAT operates at the data link layer while PAT operates at the network layer only
- D. NAT maps each private IP to a unique public IP while PAT maps many to one using ports

Answer: D

Q835. What is the poison reverse technique in distance-vector routing?

- A. Setting the metric to infinity for routes learned from a neighbor back to that neighbor
- B. Sending routing updates with encrypted metrics to prevent tampering by attackers
- C. Removing failed routes immediately from the routing table without waiting for timers
- D. Blocking all routing updates from a specific neighbor for a fixed timeout period

Answer: A

Q836. In IPv6, what is a link-local address and what prefix does it use?

- A. A multicast group address using the ff00::/8 address prefix
- B. An address valid only on a single link using the fe80::/10 prefix
- C. A loopback address using the ::1/128 single host address
- D. A globally routable address using the 2001::/16 address prefix

Answer: B

Q837. What is the purpose of the OSPF virtual link?

- A. To connect non-backbone areas that are not physically adjacent to Area 0
- B. To redistribute external routes from BGP into the OSPF routing domain
- C. To load balance traffic between two equal-cost paths in the backbone
- D. To provide encrypted tunnels between OSPF routers across the internet

Answer: A

Q838. What is Gratuitous ARP and what is its primary use case?

- A. An ARP message that assigns new IP addresses to devices joining a network
- B. An ARP reply sent in response to a multicast discovery query from a router
- C. An ARP request where the sender asks for its own IP to detect address conflicts
- D. An ARP update used to flush old MAC entries from neighboring device tables

Answer: C

Q839. In BGP route selection, what is the MED (Multi-Exit Discriminator) attribute used for?

- A. To indicate the maximum number of paths available for equal-cost load balancing
- B. To determine the originating router's preference for internal path selection
- C. To specify the exit point for traffic leaving the local autonomous system
- D. To suggest preferred entry points into an AS to neighboring autonomous systems

Answer: D

Q840. What is the purpose of the Router Solicitation and Router Advertisement messages in IPv6?

- A. They establish IPsec security associations between IPv6 endpoints
- B. They enable hosts to discover routers and obtain network configuration
- C. They provide DNS resolution services for IPv6 host name lookups
- D. They are used for multicast group membership management on a link

Answer: B

Q841. In TCP congestion control, what is the additive increase multiplicative decrease (AIMD) principle?

- A. Window stays constant during data transfer and doubles when congestion is cleared
- B. Window increases exponentially during slow start and resets to one on any timeout
- C. Window doubles on each ACK and halves on congestion detection for rapid adaptation
- D. Window increases by one MSS per RTT and halves on congestion for stable convergence

Answer: D

Q842. What is the TCP CUBIC algorithm and when is it commonly used?

- A. A retransmission algorithm that groups lost segments into cubic-sized recovery blocks
- B. A congestion control algorithm using a cubic function for window growth in high-BDP networks
- C. A scheduling algorithm that assigns cubic priority levels to competing TCP flows
- D. A flow control algorithm that triples the window size after each successful round trip

Answer: B

Q843. What is the TCP fast open (TFO) mechanism designed to reduce?

- A. The CPU cost of computing TCP checksums on high-throughput connections
- B. The overhead of maintaining idle TCP connection state on servers
- C. The latency of TCP connection establishment by sending data in the SYN
- D. The bandwidth consumed by TCP keepalive probes on long-lived connections

Answer: C

Q844. What is the fundamental problem that QUIC's connection migration solves?

- A. The inability to maintain a session when a client's IP address changes
- B. The difficulty of encrypting UDP traffic across multiple network paths
- C. The complexity of implementing flow control over unreliable datagram service
- D. The overhead of establishing separate TLS and TCP handshake sequences

Answer: A

Q845. In TCP, what is the purpose of the Karn's algorithm?

- A. It prevents RTT estimation from using retransmitted segments' measurements
- B. It selects the best MSS value based on the path MTU discovery results
- C. It calculates optimal initial window size based on network bandwidth estimates
- D. It determines when to switch from slow start to congestion avoidance phase

Answer: A

Q846. What is the BBR (Bottleneck Bandwidth and Round-trip) congestion control approach?

- A. A delay-based algorithm that reduces rate only when latency exceeds a threshold
- B. A model-based algorithm that estimates bottleneck bandwidth and minimum RTT
- C. A loss-based algorithm that halves the window on any detected packet loss
- D. A hybrid algorithm combining slow start with multiplicative decrease on timeout

Answer: B

Q847. What problem does TCP's Nagle algorithm interaction with delayed ACK cause?

- A. It causes excessive retransmissions due to premature timeouts on the sender
- B. It causes connection resets when both algorithms activate simultaneously
- C. It leads to buffer overflow on the receiver due to rapid small packet arrival
- D. It creates a latency problem where small writes are delayed up to 500ms

Answer: D

Q848. What is the purpose of the SCTP (Stream Control Transmission Protocol)?

- A. To reduce TCP header size for low-bandwidth satellite communication links
- B. To encrypt all transport layer communications without needing TLS overhead
- C. To provide multi-homing and multi-streaming over a single transport association
- D. To replace UDP entirely for all real-time multimedia streaming applications

Answer: C

Q849. How does Explicit Congestion Notification (ECN) reduce packet loss?

- A. It allows routers to signal congestion by marking packets rather than dropping them
- B. It encrypts congestion signals to prevent third parties from dropping packets
- C. It reserves bandwidth for critical flows to prevent congestion from occurring
- D. It duplicates packets across alternate paths to bypass congested network links

Answer: A

Q850. What is the difference between TCP Tahoe and TCP NewReno in handling multiple losses?

- A. NewReno doubles the ssthresh on loss while Tahoe halves it before entering slow start
- B. Tahoe enters fast recovery for each loss while NewReno resets to slow start on every loss
- C. NewReno can handle multiple losses in one window without exiting fast recovery while Tahoe cannot
- D. Tahoe uses selective acknowledgments while NewReno relies only on cumulative acknowledgments

Answer: C

Q851. What is HTTP/3's primary underlying transport protocol and why was it chosen?

- A. QUIC was chosen because it eliminates head-of-line blocking at the transport layer
- B. TCP was chosen because it provides built-in reliable ordered delivery
- C. DCCP was chosen because it provides congestion control without reliability overhead
- D. SCTP was chosen because it supports multi-homing natively across interfaces

Answer: A

Q852. What is the purpose of the DNS SRV record?

- A. It maps domain names to IPv4 addresses for web server access
- B. It creates an alias pointing one domain name to another domain
- C. It stores arbitrary text data associated with a DNS domain name
- D. It specifies the host and port for a specific service within a domain

Answer: D

Q853. What is the key difference between HTTP/2 server push and WebSocket communication?

- A. WebSocket allows the server to proactively send resources while server push provides full-duplex channels
- B. WebSocket requires HTTP/2 as a transport while server push works independently of version
- C. Server push allows the server to proactively send resources while WebSocket provides full-duplex communication
- D. Server push works only with GET requests while WebSocket works with all HTTP methods

Answer: C

Q854. What is the OCSP (Online Certificate Status Protocol) used for in HTTPS?

- A. To negotiate the cipher suite for TLS encrypted connections
- B. To check the revocation status of a digital certificate in real time
- C. To generate and distribute new SSL certificates to web servers
- D. To compress certificate chains for faster TLS handshake completion

Answer: B

Q855. What is the concept of HTTP content negotiation?

- A. The mechanism where client and server agree on the best representation of a resource
- B. The process of establishing a TCP connection before sending HTTP requests
- C. The method of splitting large HTTP responses into smaller chunked transfers
- D. The technique of compressing HTTP headers to reduce request overhead

Answer: A

Q856. What is DNS zone transfer and which record type initiates it?

- A. The encryption of DNS query data using DNSKEY records for secure transfer
- B. The migration of domain registration between different DNS registrars
- C. The process of delegating a subdomain to a different nameserver using NS records
- D. The replication of DNS data between servers initiated by the SOA record

Answer: D

Q857. What is the purpose of HPACK header compression in HTTP/2?

- A. It combines multiple HTTP requests into a single TCP segment for efficiency
- B. It encrypts HTTP headers to prevent man-in-the-middle header inspection
- C. It reduces header overhead by encoding headers using static and dynamic tables
- D. It compresses the HTTP response body for faster content delivery

Answer: C

Q858. What is the purpose of the DANE (DNS-based Authentication of Named Entities) protocol?

- A. To encrypt all DNS queries between recursive resolvers and stub resolvers
- B. To authenticate DNS client requests using challenge-response mechanisms
- C. To provide automatic DNS failover when primary nameservers become unavailable
- D. To bind TLS certificates to domain names using TLSA DNS records

Answer: D

Q859. What is the difference between HTTP caching with ETag and Last-Modified headers?

- A. ETag uses a content hash for validation while Last-Modified uses a timestamp
- B. ETag works only with static content while Last-Modified handles dynamic content
- C. Last-Modified uses a content hash while ETag relies on file size comparisons
- D. Last-Modified is more precise than ETag for detecting content modifications

Answer: A

Q860. What is the Server-Sent Events (SSE) mechanism and how does it compare to WebSocket?

- A. SSE requires a dedicated transport protocol while WebSocket uses standard HTTP
- B. SSE provides unidirectional server-to-client streaming over HTTP while WebSocket is bidirectional
- C. SSE provides bidirectional communication while WebSocket is server-to-client only
- D. SSE supports binary data frames while WebSocket is limited to text-only messages

Answer: B

Q861. What is the purpose of the 802.11r (Fast BSS Transition) standard?

- A. It increases wireless data throughput by doubling channel bandwidth
- B. It reduces roaming latency by pre-authenticating with target access points
- C. It adds support for the 6 GHz frequency band in enterprise deployments
- D. It provides backward compatibility with legacy 802.11b/g wireless clients

Answer: B

Q862. What is the SAE (Simultaneous Authentication of Equals) protocol in WPA3?

- A. A multi-factor authentication system requiring biometric verification for WiFi access
- B. A certificate-based authentication protocol replacing pre-shared keys entirely
- C. A key exchange protocol replacing the WPA2 four-way handshake with stronger security
- D. A centralized authentication protocol that requires a RADIUS server for all clients

Answer: C

Q863. What is the concept of small cells in 5G network densification?

- A. Deploying low-power base stations covering small areas to increase capacity
- B. Compressing cell tower signals into smaller packets for faster processing
- C. Reducing the physical size of 5G base station equipment for portability
- D. Using smaller frequency bands to achieve higher penetration through walls

Answer: A

Q864. What is the purpose of the FILS (Fast Initial Link Setup) protocol in WiFi?

- A. It extends the range of WiFi signals using advanced beamforming techniques
- B. It provides redundant link connectivity between access points and switches
- C. It enables file sharing between WiFi clients without requiring an access point
- D. It accelerates initial WiFi connection setup for IoT and mobile devices

Answer: D

Q865. What is Non-Orthogonal Multiple Access (NOMA) being considered for in future wireless systems?

- A. Eliminating the need for base stations by enabling pure device-to-device networks
- B. Reducing antenna count requirements by using single-element directional antennas
- C. Allowing multiple users to share the same frequency and time resources via power levels
- D. Providing stronger encryption for wireless data transmission security

Answer: C

Q866. What is the purpose of the Opportunistic Wireless Encryption (OWE) in WiFi?

- A. It opportunistically increases bandwidth by combining adjacent channels
- B. It encrypts wireless management frames to prevent deauthentication attacks
- C. It provides authenticated encryption on password-protected WiFi networks
- D. It encrypts open WiFi network traffic without requiring a shared password

Answer: D

Q867. What is the concept of coordinated multi-point (CoMP) in LTE-Advanced?

- A. Multiple base stations jointly processing signals to reduce interference at cell edges
- B. A single base station dynamically adjusting its coverage area based on demand
- C. Coordinating spectrum allocation between licensed and unlicensed frequency bands
- D. Coordinating multiple mobile devices to form a mesh network

Answer: A

Q868. What is the Multi-Link Operation (MLO) feature in WiFi 7 (802.11be)?

- A. It enables a device to connect to multiple SSIDs at the same time
- B. It allows a single device to simultaneously use multiple frequency bands and channels
- C. It provides link-level encryption across multiple concurrent sessions
- D. It connects multiple access points using wireless backhaul links

Answer: B

Q869. What is the Reconfigurable Intelligent Surface (RIS) technology being explored for 6G?

- A. A reconfigurable processor that adapts its architecture for different protocols
- B. A software-defined antenna that changes its frequency band dynamically
- C. A self-healing network fabric that repairs physical cable breaks automatically
- D. A surface with controllable elements that can reflect and shape wireless signals

Answer: D

Q870. What is the Private 5G network concept and what advantage does it offer enterprises?

- A. A consumer 5G plan with guaranteed bandwidth for home users
- B. A virtual private network tunnel running exclusively over public 5G infrastructure
- C. A dedicated 5G network operated by an enterprise for its own use with controlled QoS
- D. A 5G network that operates only in airplane mode for in-flight connectivity

Answer: C

Q871. What is the concept of a service mesh in cloud-native architecture?

- A. A physical network connecting cloud data centers across regions using fiber links
- B. A mesh of virtual machines providing redundant compute capacity for applications
- C. A dedicated infrastructure layer handling service-to-service communication in microservices
- D. A security framework that meshes multiple firewall rules into a single policy

Answer: C

Q872. What is the purpose of a cloud-native GitOps approach?

- A. It converts Git commit messages into automated cloud billing invoices
- B. It uses Git repositories as the single source of truth for declarative infrastructure and deployments
- C. It provides a Git-based code editor hosted in the cloud for collaborative programming
- D. It synchronizes Git repositories across multiple cloud providers for code redundancy

Answer: B

Q873. What is the twelve-factor app methodology in cloud application design?

- A. A pricing model with twelve tiers of service for different cloud workload sizes
- B. A methodology with twelve principles for building scalable, maintainable cloud-native apps
- C. A security framework defining twelve encryption standards for cloud applications
- D. A testing framework requiring twelve automated test suites before cloud deployment

Answer: B

Q874. What is the purpose of a cloud access security broker (CASB)?

- A. It provides visibility, compliance, and threat protection for cloud service usage
- B. It negotiates pricing between cloud customers and cloud service providers
- C. It brokers compute resources between different cloud provider regions
- D. It distributes network traffic across multiple cloud availability zones

Answer: A

Q875. What is the concept of cloud-native observability and its three pillars?

- A. Compute, storage, and networking resources monitored for utilization
- B. Encryption, authentication, and authorization for securing cloud services
- C. Availability, durability, and scalability measured across cloud deployments
- D. Metrics, logs, and distributed traces for understanding system behavior

Answer: D

Q876. What is the purpose of a cloud landing zone?

- A. A disaster recovery site that activates when the primary cloud region fails
- B. A physical location where cloud data center construction begins
- C. A pre-configured, secure, multi-account cloud environment following best practices
- D. A staging environment where applications are tested before going to production

Answer: C

Q877. What is the blast radius concept in cloud architecture?

- A. The coverage area of a cloud provider's content delivery network edges
- B. The processing speed of a distributed cloud computing cluster under load
- C. The scope of impact when a failure or security incident occurs in the system
- D. The maximum geographic distance a cloud data center can serve effectively

Answer: C

Q878. What is the strangler fig pattern in cloud migration?

- A. A pattern that monitors cloud costs and automatically shuts down idle resources
- B. A pattern of incrementally replacing legacy components with new cloud services over time
- C. A pattern where legacy systems are immediately replaced by cloud-native versions
- D. A pattern for distributing database queries across multiple cloud-hosted replicas

Answer: B

Q879. What is the concept of data gravity in cloud computing?

- A. The gravitational pull affecting satellite-based cloud network connections
- B. The force required to transfer large datasets between different clouds
- C. The physical weight of storage hardware in cloud data center facilities
- D. The tendency of data to attract applications and services to its location

Answer: D

Q880. What is the difference between cloud repatriation and cloud migration?

- A. Migration moves workloads to the cloud while repatriation moves them back on-premises
- B. Migration is automated while repatriation requires manual intervention for every workload
- C. Repatriation moves workloads from on-premises to cloud while migration moves within clouds
- D. Repatriation upgrades cloud services while migration downgrades them for cost savings

Answer: A

Q881. What is the difference between IKEv1 and IKEv2 in IPsec VPN establishment?

- A. IKEv2 uses fewer message exchanges and supports MOBIKE for seamless mobility
- B. IKEv2 only supports pre-shared keys while IKEv1 supports certificate authentication
- C. IKEv1 is more efficient with fewer exchanges while IKEv2 requires more overhead
- D. IKEv1 natively supports NAT traversal while IKEv2 requires additional configuration

Answer: A

Q882. What is the purpose of certificate transparency (CT) logs?

- A. To manage the lifecycle of internal certificates within an organization
- B. To store private keys securely for all issued digital certificates
- C. To encrypt certificate revocation lists before distributing them to clients
- D. To provide a public, verifiable record of all issued TLS certificates for a domain

Answer: D

Q883. What is a side-channel attack in cryptographic systems?

- A. An attack that intercepts encrypted data during transmission through network side channels
- B. An attack that exploits physical implementation details like timing, power, or electromagnetic emissions
- C. An attack that uses social engineering to obtain encryption keys from system administrators
- D. An attack that directly breaks the mathematical foundation of an encryption algorithm

Answer: B

Q884. What is the MITRE ATT&CK framework used for in cybersecurity?

- A. It defines compliance requirements for financial services security standards
- B. It offers automated penetration testing tools for discovering network vulnerabilities
- C. It catalogs adversary tactics, techniques, and procedures based on real-world observations
- D. It provides a standardized encryption algorithm for securing network protocols

Answer: C

Q885. What is the purpose of HPKP (HTTP Public Key Pinning) and why was it deprecated?

- A. It pinned port numbers to services but was deprecated due to conflicts with NAT devices
- B. It pinned public keys to prevent MitM attacks but risked permanent site lockout on misconfiguration
- C. It pinned DNS records to prevent cache poisoning but was deprecated due to complexity
- D. It pinned session cookies to prevent hijacking but was deprecated due to privacy concerns

Answer: B

Q886. What is the difference between SOAR and SIEM in security operations?

- A. SIEM replaces human analysts entirely while SOAR assists them with data visualization
- B. SOAR collects and analyzes security data while SIEM automates response and orchestration workflows
- C. SOAR is used only in cloud environments while SIEM works exclusively in on-premises setups
- D. SIEM collects and analyzes security data while SOAR automates response and orchestration workflows

Answer: D

Q887. What is the concept of a supply chain attack in cybersecurity?

- A. An attack that compromises software or hardware through a trusted third-party vendor
- B. An attack targeting the physical delivery of hardware networking equipment
- C. An attack on the logistics systems used for shipping cloud data center components
- D. An attack that disrupts the supply of electricity to critical infrastructure systems

Answer: A

Q888. What is the purpose of RPKI (Resource Public Key Infrastructure) for BGP security?

- A. It provides authentication for BGP peer sessions using digital certificates
- B. It monitors BGP routing tables for anomalies and generates security alerts
- C. It validates the authorization of ASes to originate specific IP prefixes
- D. It encrypts all BGP route advertisements between autonomous systems

Answer: C

Q889. What is the concept of microsegmentation in zero trust network security?

- A. Breaking large security policies into micro-policies for faster processing
- B. Segmenting encryption keys into smaller fragments for distributed storage
- C. Creating fine-grained security zones around individual workloads or applications
- D. Dividing network traffic into microsecond time slots for precise scheduling

Answer: C

Q890. What is the TLS 1.3 0-RTT resumption feature and what security risk does it introduce?

- A. It allows data transfer without any handshake but introduces vulnerability to replay attacks
- B. It pre-negotiates cipher suites but increases susceptibility to downgrade attacks
- C. It caches session keys permanently but makes them vulnerable to key extraction
- D. It eliminates encryption overhead entirely but allows passive eavesdropping on initial data

Answer: A

Q891. What is the concept of Terahertz (THz) communication being explored for 6G?

- A. A technique for compressing data to one-terabyte packets for faster delivery
- B. A method of using terraforming technology to improve signal propagation distance
- C. Communication using frequencies between 0.1 and 10 THz for extremely high data rates
- D. Communication using thermal radiation from electronic components for data encoding

Answer: C

Q892. What is the concept of Ambient IoT in next-generation networks?

- A. Battery-free IoT devices that harvest energy from ambient sources like RF signals
- B. IoT devices that only function in controlled indoor environments
- C. Background IoT processes running on cloud servers during low-demand periods
- D. IoT sensors designed specifically for monitoring air quality and pollution

Answer: A

Q893. What is the purpose of Integrated Sensing and Communication (ISAC) in 6G?

- A. Integrating multiple ISP connections into a single logical communication pipe
- B. Using communication signals simultaneously for both data transmission and radar sensing
- C. Integrating wired and wireless sensors into a single network management platform
- D. Combining sensory input devices like cameras and microphones for video conferencing

Answer: B

Q894. What is the role of the ORAN Alliance's RIC (RAN Intelligent Controller)?

- A. It manages billing and subscription information for mobile network users
- B. It controls the physical orientation of radio antennas at cell sites
- C. It provides intelligent, programmable control and optimization of RAN functions
- D. It encrypts all radio access network traffic using quantum-resistant algorithms

Answer: C

Q895. What is the concept of semantic communication in future networks?

- A. Communication that focuses on transmitting the meaning rather than exact bit sequences
- B. Communication that uses natural language processing to translate between languages
- C. A standardized naming convention for network protocols and services
- D. A technique for embedding metadata into network packet headers

Answer: A

Q896. What is the purpose of Non-Terrestrial Networks (NTN) in 3GPP standards?

- A. Providing networking services exclusively for underground data centers
- B. Creating isolated networks that operate without any connection to the internet
- C. Deploying submarine cables for transoceanic high-capacity data transmission
- D. Integrating satellite and aerial platforms into the cellular network ecosystem

Answer: D

Q897. What is the concept of AI-native networking in 6G vision?

- A. Training AI models exclusively on network traffic data from previous 5G deployments
- B. Using AI only for post-deployment network troubleshooting and incident reports
- C. Replacing all network protocols with AI-generated custom protocols for each connection
- D. Designing networks where AI is a fundamental architectural component from the ground up

Answer: D

Q898. What is the concept of digital twin network (DTN) in next-generation network management?

- A. Connecting two data centers using twin fiber optic cables for maximum bandwidth
- B. Creating a real-time virtual model of the network for prediction, optimization, and verification
- C. Mirroring all network traffic to a second network for security analysis
- D. Deploying two identical physical networks for redundancy and failover purposes

Answer: B

Q899. What is the purpose of the ETSI Multi-access Edge Computing (MEC) standard?

- A. It manages authentication of edge devices connecting to core networks
- B. It standardizes the encryption protocols used at network edge devices
- C. It defines a framework for running applications at the network edge with low latency
- D. It specifies hardware requirements for edge computing server equipment

Answer: C

Q900. What is the concept of zero-touch network and service management (ZSM)?

- A. Fully automated end-to-end network management without human intervention
- B. Networks that operate without any software updates or maintenance
- C. Networks requiring no physical cables or hardware for operation
- D. Touch-free interfaces for configuring network devices using gestures

Answer: A

SUBJECT

Cyber Security

1,440 Multiple Choice Questions

Cyber Security — Easy

480 questions

Q1. What is cyber security?

- A. Assembling and configuring enterprise hardware components
- B. Developing and deploying scalable application software
- C. Architecting and building responsive web applications
- D. Protecting computer systems, networks, and data from digital attacks

Answer: D

Q2. Which of the following is a goal of cyber security?

- A. Improving user interface design
- B. Increasing internet speed
- C. Reducing hardware costs
- D. Protecting confidentiality of data

Answer: D

Q3. What does the CIA triad stand for in cyber security?

- A. Computer, Internet, Application
- B. Confidentiality, Integrity, Availability
- C. Central Intelligence Agency
- D. Code, Integration, Authentication

Answer: B

Q4. Which of the following is an example of a cyber threat?

- A. Phishing email
- B. Power outage
- C. Software update
- D. Hardware upgrade

Answer: A

Q5. What is the primary purpose of a firewall?

- A. Reduce file sizes through compression algorithms
- B. Filter incoming and outgoing network traffic
- C. Maintain redundant copies of critical data
- D. Optimize internet speed and network throughput

Answer: B

Q6. Which of the following is considered personal identifiable information (PII)?

- A. Public news article
- B. National ID card number
- C. Weather forecast
- D. Open-source code

Answer: B

Q7. What is a vulnerability in the context of cyber security?

- A. A security tool for scanning and detecting vulnerabilities
- B. A category of self-replicating malicious software
- C. A standard network communication protocol for data transfer
- D. A weakness in a system that can be exploited

Answer: D

Q8. What is the role of antivirus software?

- A. Detect and remove malicious software
- B. Expand available system memory resources
- C. Organize and catalog system file structures
- D. Schedule and generate automated data backups

Answer: A

Q9. Which organization is responsible for cyber security standards globally?

- A. ISO (International Organization for Standardization)
- B. UNICEF within enterprise security environments
- C. WHO within enterprise security environments
- D. FIFA within enterprise security environments

Answer: A

Q10. What does the term 'threat' mean in cyber security?

- A. A potential cause of an unwanted incident that may harm a system
- B. A routine patch applied to update system software
- C. A high-level language used for development
- D. A hardware component inside wireless networking routers

Answer: A

Q11. What does confidentiality mean in information security?

- A. Data is regularly backed up to remote locations
- B. Data is continuously available to all authorized users
- C. Data is only accessible to authorized individuals
- D. Data remains unmodified and accurately preserved

Answer: C

Q12. What does integrity mean in information security?

- A. Data remains accurate and unaltered during storage and transmission
- B. Data is available 24/7 within enterprise security environments
- C. Data is deleted after use within enterprise security environments
- D. Data is always encrypted regardless of the specific situation or context

Answer: A

Q13. What does availability mean in the CIA triad?

- A. Data can be accessed by anyone in security contexts
- B. Data is encrypted within enterprise security environments
- C. Data is stored in the cloud within enterprise security environments
- D. Authorized users can access data and resources when needed

Answer: D

Q14. Which of the following is a strong password practice?

- A. Using password123 within enterprise security environments
- B. Using your name as password within enterprise security environments
- C. Using a mix of uppercase, lowercase, numbers, and special characters
- D. Using the same password for all accounts

Answer: C

Q15. What is encryption?

- A. Deleting files permanently within enterprise security environments
- B. Converting data into a coded format to prevent unauthorized access
- C. Compressing files to save space within enterprise security environments
- D. Backing up data to cloud within enterprise security environments

Answer: B

Q16. What is the purpose of a backup?

- A. To encrypt files using established cryptographic standards
- B. To speed up computer processing performance speed
- C. To delete unnecessary files within the data management framework
- D. To create a copy of data for recovery in case of data loss

Answer: D

Q17. Which of the following is a type of security control?

- A. Gaming control
- B. Preventive control
- C. Marketing control
- D. Advertising control

Answer: B

Q18. What is authentication?

- A. Verifying the identity of a user or system
- B. Applying cryptographic ciphers to protect data
- C. Compressing files to reduce disk storage consumption
- D. Architecting and building responsive web applications

Answer: A

Q19. What does a password manager do?

- A. Attempts to crack and compromise stored passwords
- B. Broadcasts and shares passwords on public platforms
- C. Permanently removes and erases stored passwords
- D. Securely stores and manages passwords for various accounts

Answer: D

Q20. Which is an example of 'something you know' in authentication?

- A. Smart card
- B. Retina scan
- C. Password
- D. Fingerprint

Answer: C

Q21. What is cryptography?

- A. Study of databases within enterprise security environments
- B. Study of networks within enterprise security environments
- C. Study of computers within enterprise security environments
- D. The practice of securing communication through encoding messages

Answer: D

Q22. What is plaintext in cryptography?

- A. The original readable message before encryption
- B. A cryptographic key applied to data protection workflows
- C. A type of cipher applied to data protection workflows
- D. Encrypted data within enterprise security environments

Answer: A

Q23. What is ciphertext?

- A. A hash function within modern computing environments
- B. A type of key within modern computing environments
- C. A readable message within modern computing environments
- D. The encrypted, unreadable form of a message

Answer: D

Q24. What is a cryptographic key?

- A. A physical metal key for a locked server room
- B. A piece of information used to encrypt and decrypt data
- C. A strain of self-replicating malicious software programs
- D. A physical cable used for network connections

Answer: B

Q25. Which of the following is a symmetric encryption algorithm?

- A. RSA
- B. ECC
- C. Diffie-Hellman
- D. AES

Answer: D

Q26. What is decryption?

- A. Converting plaintext to ciphertext
- B. Converting ciphertext back to plaintext
- C. Compressing data files for reduced storage space consumption
- D. Deleting encrypted data in security contexts

Answer: B

Q27. What is the Caesar cipher?

- A. An asymmetric algorithm within modern computing environments
- B. A hash function within modern computing environments
- C. A modern encryption standard applied to data protection workflows
- D. A simple substitution cipher that shifts letters by a fixed number

Answer: D

Q28. What is the purpose of a digital signature?

- A. To compress files for reducing disk storage consumption
- B. To sign documents physically within enterprise computing infrastructure
- C. To verify the authenticity and integrity of a digital message
- D. To encrypt all data using established cryptographic standards

Answer: C

Q29. Which of the following is an asymmetric encryption algorithm?

- A. Blowfish
- B. DES
- C. RSA
- D. AES

Answer: C

Q30. What is a public key in asymmetric cryptography?

- A. A key that is freely shared and used for encrypting messages
- B. A key that must be kept secret used in enterprise computing environments
- C. A password within the identity management system
- D. A physical key within modern computing environments

Answer: A

Q31. What is a firewall?

- A. A physical fireproof barrier within a building structure
- B. A category of self-replicating malicious software
- C. A high-level language used for development used in enterprise computing environments
- D. A network security device that monitors and filters network traffic

Answer: D

Q32. What does VPN stand for?

- A. Very Personal Network
- B. Verified Public Network
- C. Virtual Private Network
- D. Visual Processing Node

Answer: C

Q33. What is the purpose of a VPN?

- A. To reduce data sizes through compression methods
- B. To increase internet speed for improved operational efficiency
- C. To block all websites within enterprise computing infrastructure
- D. To create a secure, encrypted connection over a public network

Answer: D

Q34. What protocol is used for secure web browsing?

- A. FTP
- B. HTTPS
- C. Telnet
- D. HTTP

Answer: B

Q35. What is a network intrusion?

- A. Adding new devices to a network
- B. Upgrading network hardware
- C. Unauthorized access to a computer network
- D. Installing network software

Answer: C

Q36. Which port number is commonly used for HTTPS?

- A. 21
- B. 80
- C. 25
- D. 443

Answer: D

Q37. What is a DMZ in networking?

- A. A network segment between internal and external networks to host public-facing services
- B. A military zone within enterprise computing environments
- C. A wireless network used in enterprise network infrastructure
- D. A network routing device for managing enterprise traffic flows

Answer: A

Q38. What does IDS stand for in network security?

- A. Integrated Development System
- B. Internal Data Storage
- C. Internet Download Service
- D. Intrusion Detection System

Answer: D

Q39. What is packet filtering?

- A. Encrypting all packets regardless of the specific situation or context
- B. Deleting all packets regardless of the specific situation or context
- C. Compressing network packets within enterprise security environments
- D. Examining packets and allowing or blocking them based on defined rules

Answer: D

Q40. What is the function of a proxy server in security?

- A. To act as an intermediary between users and the internet, providing anonymity and filtering
- B. To directly connect to the internet within enterprise computing infrastructure
- C. To create viruses for enterprise computing environments and management
- D. To store passwords within enterprise computing infrastructure and management

Answer: A

Q41. What is OS hardening?

- A. Making the OS run faster within enterprise security environments
- B. Increasing storage capacity within enterprise security environments
- C. Installing new applications within enterprise security environments
- D. The process of securing an operating system by reducing its attack surface

Answer: D

Q42. Why is it important to keep operating systems updated?

- A. To fix security vulnerabilities and improve system stability
- B. To make the computer slower within enterprise computing infrastructure
- C. To increase file storage within the data management framework
- D. To change the wallpaper within enterprise computing infrastructure

Answer: A

Q43. What is a user account in operating system security?

- A. A software application used for enterprise computing operations
- B. A network connection used in enterprise network infrastructure
- C. An identity used to access and interact with the OS with specific permissions
- D. A social media profile within modern computing environments

Answer: C

Q44. What is the purpose of file permissions in an OS?

- A. To compress files for reducing disk storage consumption
- B. To encrypt all files automatically
- C. To make files larger within the data management framework
- D. To control who can read, write, or execute files

Answer: D

Q45. What is a system log?

- A. A file that records events, errors, and activities within the operating system
- B. A strain of self-replicating malicious software programs
- C. A physical network cable used for connecting infrastructure devices
- D. A physical security measure for protecting building infrastructure

Answer: A

Q46. What is the purpose of disabling unnecessary services on an OS?

- A. To free up disk space only within enterprise computing infrastructure
- B. To reduce the attack surface by eliminating potential entry points
- C. To make the OS look cleaner within enterprise computing infrastructure
- D. To speed up boot time only for improved operational efficiency

Answer: B

Q47. What is a security patch?

- A. A software update that fixes a specific security vulnerability
- B. A decorative element within modern computing environments
- C. A physical patch on hardware installed within computing infrastructure
- D. A category of self-replicating malicious software

Answer: A

Q48. Which is a basic OS security practice?

- A. Disabling the firewall within enterprise security environments
- B. Running all programs as administrator
- C. Using strong passwords and enabling automatic updates
- D. Sharing admin passwords within enterprise security environments

Answer: C

Q49. What is Windows Defender?

- A. A built-in antimalware component of Windows operating system
- B. A file manager for organizing system directories
- C. A standard web browser for accessing internet sites
- D. A first-person video game for entertainment

Answer: A

Q50. What is the root account in Linux?

- A. The superuser account with unrestricted access to all commands and files
- B. A regular user account within the identity management system
- C. A temporary account within the identity management system
- D. A guest account within the identity management system

Answer: A

Q51. What does HTTPS indicate in a website URL?

- A. The website is an older legacy version of the platform
- B. The website loads with higher overall speed performance
- C. The website is free to access without any subscription
- D. The website uses encrypted communication via SSL/TLS

Answer: D

Q52. What is a phishing website?

- A. A government website providing official public services
- B. A secure website verified by a certificate authority
- C. A website about recreational fishing techniques and tips
- D. A fake website designed to trick users into revealing personal information

Answer: D

Q53. What is a web application firewall (WAF)?

- A. A standard web browser used for accessing internet-hosted content
- B. An online platform used for creating and hosting web content
- C. A physical barrier installed to protect building infrastructure
- D. A security solution that filters and monitors HTTP traffic between a web app and the internet

Answer: D

Q54. What is the purpose of CAPTCHA on websites?

- A. To deliberately slow down website page loading speed
- B. To distinguish between human users and automated bots
- C. To compress and optimize images for web page loading
- D. To encrypt data using symmetric cipher algorithms

Answer: B

Q55. What is a cookie in web security?

- A. A category of self-replicating malicious software
- B. A small piece of data stored by a web browser for session management
- C. A baked food snack or treat served at parties
- D. A security tool for scanning and detecting vulnerabilities

Answer: B

Q56. Why check for a padlock icon in the browser address bar?

- A. It shows the site operates without any advertisements
- B. It means the site loads with high-performance speed
- C. It means the site is popular and heavily visited
- D. It indicates the connection is encrypted with SSL/TLS

Answer: D

Q57. What is a secure password practice for web accounts?

- A. Using the same password everywhere
- B. Using single-character passwords
- C. Writing passwords on sticky notes
- D. Using unique, strong passwords for each website

Answer: D

Q58. What is the purpose of two-factor authentication on websites?

- A. To deliberately make the login process much slower
- B. To completely block all users from accessing the site
- C. To add an extra layer of security beyond just a password
- D. To encrypt the entire website using a secret algorithm

Answer: C

Q59. What does 'session' mean in web security?

- A. A formal class held at a school or university within modern computing environments
- B. A software application used for enterprise computing operations
- C. A period of interaction between a user and a web application, tracked by a session identifier
- D. A physical cable used for network connections used in enterprise network infrastructure

Answer: C

Q60. What is input validation in web security?

- A. Checking keyboard hardware functionality and connectivity
- B. Verifying that user-submitted data meets expected criteria before processing
- C. Formatting and styling text content on a web page
- D. Compressing and reducing the size of input data

Answer: B

Q61. What is malware?

- A. A physical hardware component inside computer systems
- B. Software intentionally designed to cause damage to a computer, server, or network
- C. Legitimate helpful software designed for user productivity
- D. An operating system for managing system resources used in enterprise computing environments

Answer: B

Q62. What is a computer virus?

- A. A written record or logbook for tracking daily operational activities
- B. Malicious code that attaches to programs and replicates when the host program runs
- C. A routine system update for improving stability used in enterprise computing environments
- D. A self-replicating strain of malicious software targeting systems

Answer: B

Q63. What is a computer worm?

- A. A physical hardware component used in computing infrastructure
- B. Self-replicating malware that spreads across networks without needing a host program
- C. A software application used for enterprise computing operations
- D. A hardware component inside wireless networking routers

Answer: B

Q64. What is a Trojan horse in computing?

- A. A physical hardware component used in computing infrastructure
- B. Malware disguised as legitimate software to trick users into installing it
- C. An antivirus program designed to detect known threats
- D. A wooden horse artifact from ancient Greek history

Answer: B

Q65. What is ransomware?

- A. Malware that encrypts files and demands payment for the decryption key
- B. Free and open-source software available for download
- C. A backup tool for creating copies of important data
- D. An antivirus update that patches known vulnerabilities

Answer: A

Q66. What is spyware?

- A. Software that secretly monitors and collects user information without consent
- B. A security tool for scanning and detecting vulnerabilities
- C. A telescope used for astronomical observation used in enterprise computing environments
- D. A standard web browser for accessing internet sites

Answer: A

Q67. What is adware?

- A. A security program designed to protect against malware
- B. An advertising company that creates marketing campaigns
- C. Software that automatically displays or downloads unwanted advertisements
- D. A helpful tool for managing and organizing digital files

Answer: C

Q68. What is a phishing attack?

- A. A routine security scan assessing system vulnerabilities used in enterprise computing environments
- B. A standard network communication protocol for data transmission
- C. A recreational fishing technique using specialized gear used in enterprise computing environments
- D. A social engineering attack using deceptive messages to trick victims into revealing sensitive information

Answer: D

Q69. What is a Denial of Service (DoS) attack?

- A. A network upgrade that improves bandwidth and throughput used in enterprise computing environments
- B. Refusing customer service in a retail environment across computing environments
- C. A software application used for enterprise computing operations
- D. An attack that overwhelms a system with traffic to make it unavailable to legitimate users

Answer: D

Q70. What is the best defense against malware?

- A. Completely avoiding and never using a computer at all
- B. Only using free open-source software for all tasks
- C. Permanently disconnecting the system from internet access
- D. Using updated antivirus software and practicing safe computing habits

Answer: D

Q71. What is multi-factor authentication (MFA)?

- A. Using multiple different passwords for one account
- B. Using two or more different verification methods to confirm identity
- C. Displaying multiple login screens for a single service
- D. Using multiple different usernames for one account

Answer: B

Q72. What are the three main factors of authentication?

- A. Name, address, phone within enterprise security environments
- B. Hardware, software, firmware within enterprise security environments
- C. Username, password, email within enterprise security environments
- D. Something you know, something you have, something you are

Answer: D

Q73. What is a biometric authentication method?

- A. Authenticating identity using a traditional password or passphrase
- B. Using physical characteristics like fingerprints or facial recognition to verify identity
- C. Using a smart card hardware token for two-factor authentication
- D. Using a numeric PIN code for verifying user identity at login

Answer: B

Q74. What is role-based access control (RBAC)?

- A. Access permissions assigned based on a user's role within the organization
- B. No formal access control mechanism exists in the system
- C. Everyone has the same unrestricted access to resources
- D. Access is assigned randomly without specific criteria

Answer: A

Q75. What is a PIN (Personal Identification Number)?

- A. A numeric password used to authenticate a user's identity
- B. A symmetric key data encryption algorithm
- C. A network identifier assigned to connected devices
- D. A physical metal pin used for hardware components

Answer: A

Q76. What is single sign-on (SSO)?

- A. An authentication scheme allowing a user to log in once and access multiple systems
- B. A next-generation firewall for filtering network traffic flows
- C. Creating a single user account for system administration
- D. Using one password for everything without variation

Answer: A

Q77. What is the purpose of a username?

- A. To speed up computer processing performance speed
- B. To identify a specific user account in a system
- C. To create automated backup copies of user data
- D. To encrypt data using symmetric cipher algorithms

Answer: B

Q78. What is account lockout?

- A. Normally logging out of an active and authenticated user session
- B. A security mechanism that locks an account after too many failed login attempts
- C. Permanently deleting and removing a user account from the system
- D. Changing a user account password for improved security posture

Answer: B

Q79. What is a one-time password (OTP)?

- A. A shared password used by multiple user accounts
- B. A permanent password that never expires or changes
- C. A password valid for only one login session or transaction
- D. A master password that unlocks all system accounts

Answer: C

Q80. What is authorization?

- A. Permanently deleting files and data from the system storage drives
- B. Determining what resources and actions an authenticated user is allowed to access
- C. Creating and provisioning a new user account entry in the directory
- D. Verifying and confirming who you are as a user during the login process

Answer: B

Q81. What is secure software development?

- A. Writing detailed technical documentation for all application code modules
- B. Only testing for bugs and defects at the very end of the development cycle
- C. Writing application code quickly without any formal quality review process
- D. Integrating security measures throughout the entire software development lifecycle

Answer: D

Q82. What is the SDLC?

- A. A high-level language used for development used in enterprise computing environments
- B. A database management system for storing data records
- C. A security tool for scanning and detecting vulnerabilities
- D. A structured process for planning, creating, testing, and deploying software

Answer: D

Q83. Why is input validation important in secure coding?

- A. To ensure user input is safe and prevent malicious data from being processed
- B. To make the user interface look visually appealing
- C. To speed up the application processing and response time
- D. To reduce the overall file size of application resources

Answer: A

Q84. What is a software vulnerability?

- A. A routine patch applied to update and improve existing system software
- B. A category of self-replicating malicious software designed for harm
- C. A weakness or flaw in software that can be exploited to compromise security
- D. A planned and well-documented software feature built for end users

Answer: C

Q85. What is code review in secure development?

- A. Copying code from the internet without attribution or proper licensing
- B. Examining source code to identify security flaws, bugs, and code quality issues
- C. Permanently deleting old obsolete source code from repository branches
- D. Writing code faster without any verification of quality or correctness

Answer: B

Q86. What does 'security by design' mean?

- A. Adding security measures only after production deployment
- B. Building security into software from the earliest stages of development
- C. Installing antivirus software on the development machine
- D. Using only encryption as the single security measure

Answer: B

Q87. What is the purpose of error handling in secure coding?

- A. To deliberately cause the application to crash on errors
- B. To log all error details publicly for transparency
- C. To ignore all errors and continue processing without alerts
- D. To gracefully manage errors without exposing sensitive information

Answer: D

Q88. What is a security bug?

- A. A physical hardware failure requiring component replacement
- B. A network connectivity issue affecting system availability
- C. A software defect that could be exploited to compromise system security
- D. A planned feature request from the product stakeholders

Answer: C

Q89. What is the purpose of using HTTPS in web applications?

- A. To block unwanted advertisements from displaying
- B. To encrypt data transmitted between client and server
- C. To compress web pages for reduced bandwidth usage
- D. To optimize the website for faster loading performance

Answer: B

Q90. Why should default passwords be changed?

- A. To comply with organizational branding requirements
- B. To improve overall system performance and speed
- C. To make the login process more difficult for users
- D. Because default passwords are publicly known and easily exploited

Answer: D

Q91. What does Wi-Fi stand for?

- A. Wide Firewall
- B. Wired Finder
- C. Wireless Fidelity
- D. Wireless Firewall

Answer: C

Q92. What is the purpose of a Wi-Fi password?

- A. To slow down the connection within enterprise computing infrastructure
- B. To restrict network access to authorized users and encrypt wireless traffic
- C. To identify the network name across the enterprise network infrastructure
- D. To increase internet speed for improved operational efficiency

Answer: B

Q93. What is Bluetooth?

- A. A short-range wireless technology for exchanging data between devices
- B. A dental condition affecting tooth enamel coloring
- C. A wired connection using physical Ethernet cables
- D. A strain of self-replicating malicious software programs

Answer: A

Q94. Why should you avoid unknown public Wi-Fi networks?

- A. They are always blocked by network providers
- B. They always have slow data transfer speeds
- C. They consume too much device battery power
- D. They may be set up by attackers to intercept your data

Answer: D

Q95. What is the purpose of screen lock on a mobile device?

- A. To speed up device processing and app performance
- B. To conserve and extend device battery life span
- C. To prevent unauthorized physical access to the device and its data
- D. To improve the display quality and screen resolution

Answer: C

Q96. What is mobile device encryption?

- A. Converting all data on the device into an unreadable format requiring authentication
- B. Making the phone processor run faster and smoother
- C. A screen saver that activates after inactivity timeout used in enterprise computing environments
- D. Compressing files to reduce disk storage consumption

Answer: A

Q97. What should you do if your mobile device is lost or stolen?

- A. Do nothing at all and wait patiently for the situation to resolve itself
- B. Buy a new replacement device immediately without taking any other action
- C. Use remote wipe and tracking features, change passwords, and report to authorities
- D. Wait patiently for the lost or stolen device to be returned by someone

Answer: C

Q98. What is the risk of outdated mobile operating systems?

- A. Known security vulnerabilities remain unpatched and exploitable
- B. Newer wallpapers and themes are not available to install
- C. Applications appear different with slightly changed layouts
- D. The phone appears visually old and aesthetically outdated

Answer: A

Q99. What is an SSID in wireless networking?

- A. A hardware component inside wireless networking routers
- B. A symmetric key data encryption algorithm
- C. A wireless security protocol for data encryption
- D. The name that identifies a specific wireless network

Answer: D

Q100. Why should you keep Bluetooth off when not in use?

- A. To reduce the attack surface and prevent unauthorized connections
- B. To make phone calls connect and complete faster overall
- C. To customize and change the phone color theme appearance
- D. To improve camera quality and photo capture resolution

Answer: A

Q101. What is cloud computing?

- A. A weather forecasting system using satellite data analysis
- B. Storing files exclusively on your local desktop hard drive
- C. Delivering computing services like servers, storage, and software over the internet
- D. A type of external hard drive for portable data storage

Answer: C

Q102. What are the three main cloud service models?

- A. Small, Medium, Large
- B. RAM, CPU, GPU
- C. IaaS, PaaS, SaaS
- D. HTTP, FTP, SMTP

Answer: C

Q103. What is the shared responsibility model in cloud security?

- A. Both the provider and customer share responsibility for different aspects of security
- B. Only the customer is responsible for all security measures including infrastructure
- C. The cloud provider is solely responsible for everything including customer data
- D. No one is responsible for implementing security controls in cloud environments

Answer: A

Q104. What is a public cloud?

- A. Cloud infrastructure shared among multiple organizations over the internet
- B. A government-operated cloud for public sector agencies
- C. A cloud deployment visible to everyone on the internet
- D. A private physical server in an on-premises data center

Answer: A

Q105. What is a private cloud?

- A. A secret cloud that nobody uses or accesses at all
- B. Cloud infrastructure exclusively used by a single organization
- C. A specific type of weather and atmospheric phenomena
- D. A personal desktop computer used for individual tasks

Answer: B

Q106. Why is data encryption important in cloud computing?

- A. To protect data confidentiality on shared infrastructure and during transmission
- B. To organize files better with structured directory layouts
- C. To make data load faster with improved performance speed
- D. To reduce storage costs through efficient data management

Answer: A

Q107. What is multi-tenancy in cloud computing?

- A. Multiple users connected to one physical computer terminal for shared access
- B. Having multiple user accounts on different cloud services for various purposes
- C. A cloud computing platform designed for hosting distributed enterprise workloads
- D. Multiple customers sharing physical infrastructure while data remains logically isolated

Answer: D

Q108. What is a cloud access security broker (CASB)?

- A. A next-generation firewall for filtering network traffic flows
- B. A cloud provider offering infrastructure as a service
- C. A security tool between users and cloud services to enforce security policies
- D. An automated backup solution for enterprise data recovery needs

Answer: C

Q109. What is the risk of misconfigured cloud storage?

- A. Faster data access speeds and improved system performance
- B. Accidental public exposure of sensitive data due to incorrect access permissions
- C. Better performance throughput for cloud-hosted applications
- D. Lower costs for cloud infrastructure resource consumption

Answer: B

Q110. What is IAM in cloud security?

- A. A framework for managing digital identities and controlling user access to cloud resources
- B. A social media feature for managing online profile settings
- C. A file management system for organizing stored documents used in enterprise computing environments
- D. An email service for sending and receiving electronic messages

Answer: A

Q111. What is digital forensics?

- A. Repairing damaged computers and replacing failed hardware components in infrastructure
- B. A standard software application used for routine enterprise computing operations
- C. Installing and configuring software applications on production computing systems
- D. Collecting, analyzing, and preserving digital evidence from electronic devices for investigation

Answer: D

Q112. What is digital evidence?

- A. A physical security measure for protecting building infrastructure
- B. Physical fingerprints collected at a crime scene within enterprise security environments
- C. A software application used for enterprise computing operations
- D. Any data stored or transmitted using a computer that supports or refutes a theory in a legal case

Answer: D

Q113. Why is chain of custody important in digital forensics?

- A. It reduces storage space required for evidence preservation
- B. It makes the forensic investigation process faster
- C. It ensures evidence integrity and admissibility by documenting who handled it and how
- D. Continuous monitoring is not important for security for modern enterprise security environments

Answer: C

Q114. What is a forensic image?

- A. A digital photograph taken with a camera for documentation
- B. A bit-by-bit exact copy of a storage device for forensic examination
- C. A screenshot captured from a computer display for reference
- D. A photograph of a physical crime scene by investigators

Answer: B

Q115. What types of devices can contain digital evidence?

- A. Only USB drives can contain relevant digital evidence for forensic investigation
- B. Only desktop computers can contain relevant digital evidence for investigation
- C. Only smartphones can contain relevant digital evidence for forensic investigation
- D. Computers, smartphones, tablets, servers, IoT devices, USB drives, and any digital storage media

Answer: D

Q116. What is the first step in a digital forensics investigation?

- A. Deleting suspicious files from the compromised system
- B. Restarting the computer to clear temporary memory contents
- C. Installing antivirus software on the compromised system
- D. Identifying and securing the scene to prevent evidence tampering

Answer: D

Q117. What is file recovery in digital forensics?

- A. Creating new files on the forensic examination workstation
- B. Encrypting files within enterprise security environments
- C. Retrieving deleted, damaged, or lost files from storage media
- D. Compressing files to reduce disk storage consumption

Answer: C

Q118. What is metadata in digital forensics?

- A. A high-level language used for development used in enterprise computing environments
- B. A strain of self-replicating malicious software programs
- C. Data about data, such as file creation date, author, and modification history
- D. A standard network communication protocol for data transfer

Answer: C

Q119. What is a write blocker used for in forensics?

- A. To block network traffic across the enterprise network infrastructure
- B. To block websites within enterprise computing infrastructure
- C. To prevent any write operations to the evidence drive during examination
- D. To block writing on paper within enterprise computing infrastructure

Answer: C

Q120. What is the purpose of hashing in digital forensics?

- A. To create a unique digital fingerprint of evidence to verify its integrity
- B. To reduce data sizes through compression methods and management
- C. To speed up computers for improved operational efficiency
- D. To encrypt files using established cryptographic standards

Answer: A

Q121. What is an incident in cyber security?

- A. A routine patch applied to update system software used in enterprise computing environments
- B. An event that threatens the security, confidentiality, integrity, or availability of information assets
- C. A physical hardware component used in computing infrastructure
- D. A scheduled maintenance within enterprise computing environments

Answer: B

Q122. What is incident response?

- A. An organized approach to addressing and managing a security breach to limit damage
- B. Generating redundant backup copies of system data for disaster recovery
- C. Installing routine software updates on regularly scheduled maintenance days
- D. Ignoring security events and hoping they resolve on their own naturally

Answer: A

Q123. Why is having an incident response plan important?

- A. It ensures quick, organized, and effective response to security incidents
- B. Continuous monitoring is not important for security
- C. It eliminates all threats regardless of the specific situation or context
- D. It replaces security tools within enterprise security environments

Answer: A

Q124. What is the first phase of incident response?

- A. Recovery
- B. Containment
- C. Preparation
- D. Eradication

Answer: C

Q125. What is an incident response team?

- A. A marketing team managing brand awareness campaigns used in enterprise computing environments
- B. Trained personnel responsible for detecting, analyzing, and responding to security incidents
- C. A software application used for enterprise computing operations
- D. A group of malicious hackers targeting an organization used in enterprise computing environments

Answer: B

Q126. What should you do if you suspect a security incident?

- A. Report it immediately to the appropriate security team
- B. Try to fix it yourself without informing anyone
- C. Delete all files to eliminate any potential evidence
- D. Ignore it and continue with normal operations as usual

Answer: A

Q127. What is the purpose of incident documentation?

- A. To blame someone specific for causing the incident
- B. To delete evidence of the incident from all records
- C. To create unnecessary bureaucratic paperwork and overhead
- D. To maintain a detailed record of the incident, actions taken, and findings

Answer: D

Q128. What is the difference between an event and an incident?

- A. Events are worse than incidents in terms of damage across computing environments
- B. They are fundamentally identical concepts with no meaningful differences
- C. Incidents are planned events scheduled by the security team
- D. An event is any observable occurrence; an incident negatively impacts security or violates policy

Answer: D

Q129. What is the purpose of the containment phase?

- A. To ignore the incident and let it resolve on its own
- B. To limit the scope and damage by isolating affected systems
- C. To start a new project unrelated to the current incident
- D. To spread and propagate the incident to other systems

Answer: B

Q130. What is a security alert?

- A. A fire alarm system installed in a physical building
- B. A notification indicating a potential security event requiring investigation
- C. A performance warning about high system resource usage
- D. A software update notification from the operating system

Answer: B

Q131. What is security monitoring?

- A. Monitoring attendance records for employee time tracking
- B. Watching security guards on physical surveillance cameras
- C. Continuous observation and analysis of systems and networks for security threats
- D. A video surveillance system for monitoring building access

Answer: C

Q132. What is a security log?

- A. A record of security-relevant events and activities within a system or network
- B. A wooden log used for building physical fire structures
- C. A login page for entering user authentication credentials
- D. A security guard's written diary of daily activity notes

Answer: A

Q133. What is an alert in security monitoring?

- A. A routine system update for improving stability used in enterprise computing environments
- B. A user account type with specifically assigned access permissions
- C. A notification when a security tool detects suspicious or potentially malicious activity
- D. An alarm clock that wakes people up at a set time used in enterprise computing environments

Answer: C

Q134. What is the purpose of monitoring network traffic?

- A. To count the total number of packets transmitted on network
- B. To slow down overall network performance and throughput
- C. To increase available bandwidth for improved download speeds
- D. To detect unusual patterns, unauthorized access, and potential security threats

Answer: D

Q135. What is a dashboard in security monitoring?

- A. A file manager for organizing system directories used in enterprise computing environments
- B. A visual interface displaying real-time security metrics, alerts, and status information
- C. A web browser application for accessing internet-hosted content
- D. A car dashboard displaying vehicle speed and fuel information

Answer: B

Q136. Why is continuous monitoring important?

- A. Continuous monitoring exists solely to use more electricity
- B. Continuous monitoring is not important for security for modern enterprise security environments
- C. Threats can occur at any time and continuous monitoring ensures quick detection
- D. Continuous monitoring exists only to generate more data

Answer: C

Q137. What is log management?

- A. Collecting, storing, analyzing, and retaining log data from various systems
- B. Managing a personal blog for publishing online content
- C. Managing a lumber yard for processing wood materials
- D. A task management tool for organizing project work items

Answer: A

Q138. What type of events should be logged for security?

- A. Login attempts, access to sensitive data, system changes, and privilege escalations
- B. Only successful events should be logged for monitoring and compliance reporting
- C. Only errors and crashes should be logged for security analysis and remediation
- D. Only network events should be logged for analysis and security monitoring

Answer: A

Q139. What is a false positive in security monitoring?

- A. A true security threat confirmed by analysis investigation
- B. A false identity used for unauthorized system access
- C. A system error causing application crashes or instability
- D. An alert that incorrectly identifies normal activity as a security threat

Answer: D

Q140. What is the purpose of antivirus monitoring?

- A. To continuously watch for and detect malware infections in real-time
- B. To delete all files from the system storage permanently
- C. To sell software licenses for antivirus subscriptions
- D. To slow down the computer for detailed system analysis

Answer: A

Q141. What is cybercrime?

- A. Playing computer games for recreational entertainment and leisure purposes
- B. Criminal activity involving computers, networks, or digital devices as tools or targets
- C. Using computers for normal everyday work productivity and communications
- D. Building websites for establishing personal or commercial online presence

Answer: B

Q142. What is PECA 2016 in Pakistan?

- A. An education policy for governing school curriculum standards
- B. Pakistan's primary legislation addressing cybercrimes, data protection, and electronic fraud
- C. A health regulation for managing public medical facilities used in enterprise computing environments
- D. A traffic law regulating vehicle movement on roads used in enterprise computing environments

Answer: B

Q143. What is intellectual property in the digital context?

- A. A physical hardware component used in computing infrastructure
- B. Creative works, inventions, designs, and software protected by law from unauthorized use
- C. Internet speed and bandwidth available for data transfers
- D. Physical property like land and buildings owned by individuals

Answer: B

Q144. What is data privacy?

- A. Making data public and freely accessible for anyone across computing environments
- B. The right of individuals to control how their personal information is collected, used, and shared
- C. Hiding your computer from others by placing it out of sight
- D. Deleting all data from storage devices permanently across computing environments

Answer: B

Q145. What is unauthorized access in cyber law?

- A. Logging into your own user account with correct credentials
- B. Downloading free software from authorized official sources
- C. Using public Wi-Fi for browsing the internet content
- D. Accessing a computer system or network without permission

Answer: D

Q146. What is online harassment?

- A. Using digital communication to threaten, intimidate, or harm another person
- B. Social media marketing campaigns for brand awareness
- C. Sending friendly emails to colleagues for collaboration
- D. Online gaming with other players for entertainment purposes

Answer: A

Q147. What is software piracy?

- A. Software testing for validating application feature quality
- B. The unauthorized copying, distribution, or use of copyrighted software
- C. Open-source development contributing code to public projects
- D. Sailing ships equipped with computer navigation technology

Answer: B

Q148. What is digital consent?

- A. A symmetric key data encryption algorithm applied to data protection workflows
- B. A computer command executed from the terminal or console used in enterprise computing environments
- C. A digital signature used for verifying message authenticity used in enterprise computing environments
- D. Permission given by an individual for their personal data to be collected, processed, or shared

Answer: D

Q149. What is identity theft in the digital context?

- A. Creating a new email account for personal communication
- B. Changing your username on a social media platform account
- C. Forgetting your password and requesting an account reset
- D. Stealing personal information to impersonate someone for fraud or other crimes

Answer: D

Q150. What is the purpose of acceptable use policies (AUP)?

- A. To deploy and install application software and management
- B. To restrict all computer use for security purposes and management
- C. To increase internet speed for improved operational efficiency
- D. To define rules and guidelines for responsible and secure use of IT resources

Answer: D

Q151. What is AI in cyber security?

- A. A next-generation firewall security appliance deployed across enterprise environments
- B. The use of machine learning and intelligent algorithms to detect, prevent, and respond to threats
- C. A physical security measure for protecting building infrastructure
- D. A software application used for enterprise computing operations

Answer: B

Q152. What is the Internet of Things (IoT)?

- A. A network of physical devices connected to the internet that collect and exchange data
- B. A web browser application for accessing internet-hosted content
- C. A social media platform for sharing posts and updates used in enterprise computing environments
- D. A new internet service provider offering broadband access

Answer: A

Q153. Why are IoT devices a security concern?

- A. They process data too fast for security tools to monitor
- B. They consume too much electricity for normal household budgets
- C. They are too expensive for consumers to purchase and use
- D. Many have weak security, default passwords, and limited update capabilities

Answer: D

Q154. What is blockchain technology?

- A. A next-generation firewall for filtering network traffic flows
- B. A cryptographic algorithm used for protecting sensitive data
- C. A distributed, immutable ledger that records transactions across many computers
- D. A database backup system for creating redundant data copies

Answer: C

Q155. What is ransomware-as-a-service (RaaS)?

- A. A backup service for managing enterprise data resources
- B. A legitimate cloud service for hosting web applications
- C. A data recovery service for restoring encrypted file backups
- D. A criminal model where ransomware developers lease tools to other criminals

Answer: D

Q156. What is biometric authentication technology?

- A. A standard network communication protocol for data transfer
- B. A cryptographic algorithm used for protecting sensitive data
- C. Using unique physical or behavioral characteristics for identity verification
- D. Using only passwords for user identity verification login

Answer: C

Q157. What is cloud-native security?

- A. A physical security measure specifically designed for protecting building infrastructure
- B. Weather-based security monitoring using atmospheric sensors and satellite imaging data
- C. Security approaches designed specifically for cloud environments including containers and serverless
- D. A self-replicating strain of malicious software designed for targeting computing systems

Answer: C

Q158. What is a deepfake?

- A. AI-generated synthetic media realistically depicting people saying or doing things they never did
- B. A software application used for enterprise computing operations
- C. A deep sea photograph taken underwater for marine research used in enterprise computing environments
- D. A symmetric key data encryption algorithm applied to data protection workflows

Answer: A

Q159. What is 5G and why does it matter for security?

- A. A strain of self-replicating malicious software programs
- B. The fifth generation of personal computer hardware architecture
- C. The fifth generation of mobile network technology bringing new security challenges
- D. A security certification for mobile network professionals

Answer: C

Q160. What is cyber resilience?

- A. A next-generation firewall for filtering network traffic flows
- B. Never experiencing any cyber attacks or security incidents
- C. An automated backup solution for enterprise data recovery needs
- D. An organization's ability to continuously deliver outcomes despite adverse cyber events

Answer: D

Q161. What does the term 'attack surface' refer to in cyber security?

- A. The graphical interface of a security application
- B. The total number of employees in an organization
- C. The set of points where an attacker can try to enter
- D. The physical area covered by a wireless network

Answer: C

Q162. Which term describes an individual who attempts to gain unauthorized access to systems?

- A. Developer
- B. Hacker
- C. Analyst
- D. Administrator

Answer: B

Q163. What is the main purpose of security awareness training?

- A. Training staff to configure enterprise-level firewalls
- B. Teaching employees to write secure code for applications
- C. Educating users to recognize and avoid security threats
- D. Certifying professionals for network administration roles

Answer: C

Q164. Which of the following best describes a 'white hat' hacker?

- A. A hacker who creates malware to disrupt critical services
- B. An ethical hacker who tests systems with proper authorization
- C. A malicious hacker who exploits systems for personal gain
- D. A hacker who sells stolen data on underground marketplaces

Answer: B

Q165. What does 'data breach' mean in cyber security?

- A. Authorized sharing of data between partner organizations
- B. Unauthorized access and disclosure of sensitive information
- C. Scheduled backup of organizational data to offsite storage
- D. Routine migration of data between different server systems

Answer: B

Q166. What is the function of encryption in cyber security?

- A. Deleting outdated files from the system permanently
- B. Compressing files to reduce overall storage capacity
- C. Converting data into unreadable format for protection
- D. Scanning networks to detect unauthorized intrusions

Answer: C

Q167. Which type of cyber attack tricks users into revealing sensitive information via fake emails?

- A. Sniffing
- B. Phishing
- C. Brute force
- D. Keylogging

Answer: B

Q168. What is the purpose of a security policy in an organization?

- A. Defining rules and procedures for protecting information assets
- B. Listing hardware specifications for all organizational devices
- C. Documenting employee performance reviews and evaluations
- D. Recording financial transactions and quarterly budget reports

Answer: A

Q169. Which of the following is NOT a pillar of the CIA triad?

- A. Confidentiality
- B. Integrity
- C. Availability
- D. Scalability

Answer: D

Q170. What does the term 'insider threat' refer to?

- A. A natural disaster affecting the data center infrastructure
- B. A hardware malfunction causing temporary system downtime
- C. A security threat originating from within the organization
- D. An external attacker probing the network from the internet

Answer: C

Q171. What is the purpose of access control in information security?

- A. Monitoring hardware temperature for system stability
- B. Compressing files to save disk storage space usage
- C. Restricting who can view or use resources in a system
- D. Speeding up data transfer across network connections

Answer: C

Q172. Which type of security control is a locked door to a server room?

- A. Logical
- B. Technical
- C. Administrative
- D. Physical

Answer: D

Q173. What does 'integrity' mean in the context of information security?

- A. Keeping sensitive information hidden from unauthorized people
- B. Providing fast network speeds for data transfer and access
- C. Making sure systems are always available to authorized users
- D. Ensuring data is accurate and has not been tampered with

Answer: D

Q174. What is a security baseline?

- A. A software tool for scanning networks against vulnerabilities
- B. A type of firewall used at the network perimeter gateway
- C. A minimum set of security standards for systems to meet
- D. The highest possible security level an organization uses

Answer: C

Q175. Which of the following is an example of a detective security control?

- A. Audit log
- B. Encryption
- C. Backup
- D. Firewall

Answer: A

Q176. What is the purpose of a password policy?

- A. Documenting disaster recovery procedures for data centers
- B. Defining rules for creating and managing strong passwords
- C. Specifying network bandwidth allocation per each department
- D. Listing all approved software applications for installation

Answer: B

Q177. What does 'availability' ensure in the CIA triad?

- A. Data is hidden from all external network users at all times
- B. Data cannot be modified by any internal user of the system
- C. Systems and data are accessible when needed by authorized users
- D. All network traffic is encrypted between source and destination

Answer: C

Q178. Which term describes the process of verifying a user's claimed identity?

- A. Encryption
- B. Accounting
- C. Authorization
- D. Authentication

Answer: D

Q179. What is social engineering in the context of security?

- A. Building secure social media platforms for enterprise users
- B. Creating automated social network bots for data collection
- C. Manipulating people into revealing confidential information
- D. Designing user-friendly interfaces for security applications

Answer: C

Q180. What is the primary function of a security patch?

- A. Upgrading hardware components for better system performance
- B. Adding new features and capabilities to existing software
- C. Fixing known security vulnerabilities in software or systems
- D. Migrating data from legacy storage to modern cloud services

Answer: C

Q181. What is plaintext in cryptography?

- A. A hash value generated from a digital signature
- B. Data that has been encrypted using a cipher algorithm
- C. A secret key used for decrypting encoded messages
- D. Original readable data before encryption is applied

Answer: D

Q182. Which type of cryptography uses the same key for encryption and decryption?

- A. Steganography
- B. Asymmetric
- C. Symmetric
- D. Hashing

Answer: C

Q183. What is the output of a hash function called?

- A. Digest
- B. Plaintext
- C. Token
- D. Ciphertext

Answer: A

Q184. Which algorithm is a widely used symmetric encryption standard?

- A. AES
- B. RSA
- C. ECC
- D. DSA

Answer: A

Q185. What is the primary purpose of a digital signature?

- A. Verifying the authenticity and integrity of a message
- B. Generating random encryption keys for secure sessions
- C. Compressing data to reduce network bandwidth consumption
- D. Encrypting the entire message content for confidentiality

Answer: A

Q186. What does RSA stand for in cryptography?

- A. Rivest-Shamir-Adleman
- B. Recursive Signing Architecture
- C. Random Security Algorithm
- D. Rapid Symmetric Authentication

Answer: A

Q187. What is a certificate authority (CA) responsible for?

- A. Monitoring network traffic for suspicious activity and intrusions
- B. Manufacturing hardware security modules for data center operations
- C. Issuing and managing digital certificates for secure communications
- D. Developing encryption algorithms for government security agencies

Answer: C

Q188. Which of the following is an asymmetric encryption algorithm?

- A. DES
- B. RSA
- C. RC4
- D. AES

Answer: B

Q189. What is the purpose of a cryptographic key?

- A. A software license required to install encryption applications
- B. A physical device used to unlock encrypted hardware storage units
- C. A parameter that determines the output of a cryptographic algorithm
- D. A network address assigned to encrypted communication channels

Answer: C

Q190. What property of hash functions means the output cannot be reversed to find the input?

- A. Key derivation
- B. Collision resistance
- C. One-way property
- D. Block chaining

Answer: C

Q191. What is the primary function of a network firewall?

- A. Managing user accounts across the enterprise network
- B. Storing backup copies of important network data files
- C. Increasing the overall speed of network connections
- D. Filtering traffic based on predefined security rules

Answer: D

Q192. What does a VPN provide for network communications?

- A. Load balancing across multiple server instances
- B. Faster download speeds for large file transfers
- C. Encrypted tunnel for secure data transmission
- D. Automatic virus scanning of all network traffic

Answer: C

Q193. What is a DMZ in network security?

- A. A protocol for encrypting database connections between application server nodes
- B. A tool for monitoring bandwidth usage across all organizational network links
- C. A network segment that separates internal networks from untrusted external networks
- D. A type of malware that spreads automatically across connected wireless networks

Answer: C

Q194. What does IDS stand for in network security?

- A. Intrusion Detection System
- B. Integrated Display Server
- C. Internet Data Service
- D. Internal Defense System

Answer: A

Q195. What is the purpose of network segmentation?

- A. Dividing a network into smaller segments to improve security and performance
- B. Removing all firewalls to simplify the overall network architecture
- C. Increasing the physical cable length between network connected devices
- D. Combining multiple small networks into one large unified network system

Answer: A

Q196. Which protocol provides secure web browsing by encrypting HTTP traffic?

- A. HTTPS
- B. SNMP
- C. FTP
- D. SMTP

Answer: A

Q197. What is a proxy server used for in network security?

- A. Acting as an intermediary between users and the internet for security
- B. Directly connecting users to the internet without any filtering
- C. Replacing the router as the main network traffic forwarding device
- D. Generating encryption keys for all network-connected user devices

Answer: A

Q198. What type of attack floods a server with traffic to make it unavailable?

- A. DoS attack
- B. SQL injection
- C. Spoofing
- D. Phishing

Answer: A

Q199. What does NAT stand for in networking?

- A. Network Address Translation
- B. Network Audit Technique
- C. Network Access Terminal
- D. Node Authentication Token

Answer: A

Q200. What is the purpose of an access control list (ACL) on a network device?

- A. Recording the browsing history of all network users for auditing
- B. Listing all devices currently connected to the network for inventory
- C. Managing software licenses for applications installed on the server
- D. Defining rules that permit or deny traffic based on specified criteria

Answer: D

Q201. What is the primary role of user account management in OS security?

- A. Managing the color scheme and desktop wallpaper preferences
- B. Controlling who can access the system and their permission levels
- C. Increasing the processing speed of all running applications
- D. Improving the graphical user interface design for better usability

Answer: B

Q202. What is the purpose of file permissions in an operating system?

- A. Organizing files alphabetically in directory listings
- B. Setting the default font size for text documents
- C. Controlling who can read, write, or execute a file
- D. Determining how fast files can be read from disk

Answer: C

Q203. What does the 'root' account represent in Linux operating systems?

- A. A standard user account for regular daily computing tasks
- B. A disabled account used only for emergency system recovery
- C. The superuser account with full administrative privileges
- D. A guest account with limited temporary access to the system

Answer: C

Q204. What is the purpose of operating system patching?

- A. Adding new entertainment features to the operating system
- B. Upgrading the hardware drivers for better graphics performance
- C. Changing the default language settings of the operating system
- D. Fixing security vulnerabilities and bugs in the OS software

Answer: D

Q205. What is a host-based firewall?

- A. A centralized network device protecting the entire corporate network
- B. A hardware appliance installed at the internet service provider
- C. A firewall installed on individual computers to control their traffic
- D. A cloud service that scans emails for malicious file attachments

Answer: C

Q206. What is the function of an OS audit log?

- A. Compressing old files to free up available disk storage space
- B. Speeding up the boot process by caching startup configurations
- C. Recording system events and user activities for security review
- D. Displaying real-time CPU usage statistics on the desktop screen

Answer: C

Q207. What is the purpose of disabling unnecessary services on a server?

- A. Improving the visual appearance of the server management console
- B. Increasing the amount of available RAM for user applications
- C. Reducing the attack surface by eliminating potential entry points
- D. Enabling faster network communication between all server nodes

Answer: C

Q208. What does UAC stand for in Windows security?

- A. Unified Audit Console
- B. Update Administration Center
- C. User Account Control
- D. Universal Access Code

Answer: C

Q209. What is the purpose of disk encryption on a laptop?

- A. Protecting data if the laptop is lost or stolen by making it unreadable
- B. Increasing the read and write speed of the hard drive significantly
- C. Preventing the laptop from connecting to unsecured wireless networks
- D. Reducing the storage capacity needed for files on the encrypted disk

Answer: A

Q210. What is a security hardening guide for an operating system?

- A. A document providing steps to secure an OS by reducing vulnerabilities
- B. A manual for repairing physical damage to computer hardware components
- C. A guide for increasing internet connection speeds through OS tweaks
- D. A tutorial for installing entertainment software on the operating system

Answer: A

Q211. What is cross-site scripting (XSS)?

- A. A protocol for sharing authentication tokens between different websites
- B. An attack that injects malicious scripts into web pages viewed by users
- C. A method for optimizing website loading speed across browsers
- D. A technique for cross-browser testing of web application functionality

Answer: B

Q212. What is SQL injection?

- A. A method of creating structured database tables and indexes
- B. An attack that inserts malicious SQL code into application queries
- C. A tool for migrating data between different database platforms
- D. A database optimization technique for faster query results

Answer: B

Q213. What does HTTPS provide that HTTP does not?

- A. Faster page loading speeds for all website resources
- B. Larger file upload limits for web form submissions
- C. Encrypted communication between browser and server
- D. Better search engine rankings for website content

Answer: C

Q214. What is a web application firewall (WAF)?

- A. A security tool that filters and monitors HTTP traffic to web applications
- B. A network switch that connects multiple web servers in a data center
- C. A device that increases the bandwidth available to web server connections
- D. A software program that designs web page layouts and visual templates

Answer: A

Q215. What is the purpose of input validation in web applications?

- A. Converting user input into different languages automatically
- B. Improving the visual appearance of input fields on web forms
- C. Ensuring user input meets expected formats before processing it
- D. Speeding up form submission by compressing input data first

Answer: C

Q216. What does CSRF stand for in web security?

- A. Cached Session Recovery Feature
- B. Cross-Site Request Forgery
- C. Central Server Response Filter
- D. Client-Side Resource Failure

Answer: B

Q217. What is the purpose of the same-origin policy in web browsers?

- A. Preventing web pages from making requests to different domains than their origin
- B. Ensuring all web pages load at the same speed regardless of server location
- C. Requiring all websites to use identical CSS stylesheets for consistent design
- D. Blocking users from opening multiple browser tabs at the same time period

Answer: A

Q218. What is a session token used for in web applications?

- A. Measuring the loading time of web pages for performance optimization
- B. Encrypting all data stored in the web application's backend database
- C. Generating random numbers for CAPTCHA verification on login forms
- D. Identifying and maintaining a user's authenticated session with the server

Answer: D

Q219. What is clickjacking?

- A. Tricking users into clicking hidden elements overlaid on legitimate web pages
- B. A technique for optimizing click-through rates on web advertisements
- C. Stealing user passwords by recording their keyboard strokes while typing
- D. A method of bypassing CAPTCHA challenges using automated mouse clicks

Answer: A

Q220. What is the OWASP Top 10?

- A. A ranking of the top ten fastest web servers available for commercial purchase
- B. A list of the ten most critical web application security risks updated regularly
- C. A collection of ten programming languages recommended for web development
- D. A standard defining ten required features for modern web browser software

Answer: B

Q221. What is a computer virus?

- A. A software tool used for optimizing computer system performance
- B. A program that replicates by attaching itself to other programs or files
- C. A hardware component that protects computers from power surges
- D. A network protocol designed for secure file transfer operations

Answer: B

Q222. What is ransomware?

- A. Malware that encrypts files and demands payment for decryption
- B. A tool that helps organize files into categorized folder systems
- C. A program that provides free access to premium software licenses
- D. Software that speeds up internet browsing on older computers

Answer: A

Q223. How does a worm differ from a virus?

- A. A worm only affects hardware while a virus only affects software
- B. A worm requires user interaction while a virus spreads automatically
- C. A worm spreads independently without needing to attach to host files
- D. A worm is harmless while a virus always causes system destruction

Answer: C

Q224. What is a Trojan horse in cyber security?

- A. A security tool that automatically detects and removes all known viruses
- B. Malware disguised as legitimate software to trick users into installing it
- C. An encryption method used to protect files during network transmission
- D. A firewall configuration that blocks all incoming network traffic completely

Answer: B

Q225. What is spyware designed to do?

- A. Improve computer performance by monitoring resource usage patterns
- B. Protect sensitive files by creating encrypted backup copies of data
- C. Secretly collect information about a user's activities without consent
- D. Block malicious websites from loading in the user's web browser

Answer: C

Q226. What is a botnet?

- A. A secure private network used by government agencies for communications
- B. A legitimate network of robots used in manufacturing and automation
- C. A network of compromised computers controlled remotely by an attacker
- D. A high-speed connection linking multiple data centers across regions

Answer: C

Q227. What is a keylogger?

- A. A device that locks the keyboard to prevent unauthorized system access
- B. A feature that creates keyboard shortcuts for frequently used programs
- C. A program that records every keystroke made by the user on a computer
- D. A tool that tests keyboard hardware for mechanical defects and issues

Answer: C

Q228. What is adware?

- A. A service that manages advertising campaigns across multiple platforms
- B. Software that displays unwanted advertisements on the user's device
- C. A program that blocks all advertisements while browsing the internet
- D. A tool used by marketers to create professional advertisement designs

Answer: B

Q229. What is a brute force attack?

- A. Exploiting software bugs to gain unauthorized access to a database
- B. Physically damaging computer hardware to disrupt system operations
- C. Trying every possible combination to guess a password or encryption key
- D. Flooding a network with traffic to deny service to legitimate users

Answer: C

Q230. What is a man-in-the-middle attack?

- A. A strategy for overwhelming servers with simultaneous connection requests
- B. A technique for bypassing firewalls by tunneling through allowed ports
- C. A method of hiding malware inside legitimate software update packages
- D. An attack where the attacker intercepts communication between two parties

Answer: D

Q231. What is multi-factor authentication (MFA)?

- A. Logging into multiple applications with a single set of credentials
- B. Creating multiple user accounts for the same person in a system
- C. Using two or more verification methods to confirm a user's identity
- D. Setting up multiple passwords for a single user account for safety

Answer: C

Q232. What is single sign-on (SSO)?

- A. A technique for creating unique passwords for every individual application
- B. A security measure requiring users to sign in separately to each application
- C. A method of encrypting passwords before storing them in the user database
- D. A system allowing users to access multiple applications with one login credential

Answer: D

Q233. What is a biometric authentication method?

- A. Creating long and complex passwords that are difficult for attackers to guess
- B. Using a physical security token device to generate one-time access codes
- C. Using physical characteristics like fingerprints or facial features to verify identity
- D. Sending a verification code to a registered phone number for confirmation

Answer: C

Q234. What is the purpose of account lockout policies?

- A. Preventing users from accessing their accounts outside of regular business hours
- B. Forcing users to change their passwords after every single successful login event
- C. Temporarily disabling accounts after multiple failed login attempts to stop attacks
- D. Permanently deleting user accounts that have been inactive for over twelve months

Answer: C

Q235. What are the three factors of authentication?

- A. First name, last name, and date of birth for verification
- B. Something you know, something you have, and something you are
- C. Email address, phone number, and physical mailing address
- D. Username, password, and security question answer combination

Answer: B

Q236. What is role-based access control (RBAC)?

- A. Giving every user in the organization identical access rights to all resources
- B. Restricting access based on the time of day the user attempts to log in
- C. Assigning permissions to users based on their organizational role or job function
- D. Allowing users to choose which resources they want to access independently

Answer: C

Q237. What is a one-time password (OTP)?

- A. A password that can be used unlimited times until the user changes it
- B. A default system password assigned to all new employee user accounts
- C. A permanent password set during initial account creation for the user
- D. A password that is valid for only one login session or transaction

Answer: D

Q238. What is the principle of least privilege in access control?

- A. Restricting access to systems only during designated working hours daily
- B. Requiring users to request permission for every individual file they access
- C. Granting users only the minimum access needed for their specific job duties
- D. Providing all users with administrator-level access for operational convenience

Answer: C

Q239. What is a CAPTCHA used for?

- A. Tracking user browsing behavior across different websites for analytics
- B. Measuring the loading speed of web pages for performance optimization
- C. Encrypting user passwords before storing them in the application database
- D. Distinguishing human users from automated bots during online interactions

Answer: D

Q240. What is an access control list (ACL)?

- A. A list defining which users or groups have permission to access specific resources
- B. A catalog of all software applications installed on the organizational network
- C. A record of all users who have ever logged into the system since installation
- D. A directory listing all hardware devices connected to the computer network

Answer: A

Q241. What is secure software development?

- A. Writing code that prioritizes speed over all other quality attributes
- B. Creating software that focuses exclusively on user interface design
- C. Building software with security integrated throughout the development lifecycle
- D. Developing software as quickly as possible without any testing phases

Answer: C

Q242. What is the purpose of code review in secure development?

- A. Measuring the execution speed of individual functions in code
- B. Formatting code to follow consistent visual styling conventions
- C. Translating code from one programming language to another one
- D. Examining code to find security flaws and bugs before deployment

Answer: D

Q243. What does SDLC stand for in software development?

- A. System Data Logging Control
- B. Software Development Life Cycle
- C. Secure Database Login Credentials
- D. Standard Digital Library Catalog

Answer: B

Q244. What is input sanitization in secure coding?

- A. Formatting input data to display it in a more readable user interface
- B. Cleaning user input to remove or escape potentially harmful characters
- C. Converting all input text to uppercase for consistent data storage
- D. Compressing input data to reduce the database storage requirements

Answer: B

Q245. What is the principle of 'secure by default' in software design?

- A. Software should allow all features by default for maximum user flexibility
- B. Software should disable all functionality until security is configured
- C. Software should require users to manually enable all security features
- D. Software should be configured with the most secure settings out of the box

Answer: D

Q246. What is a buffer overflow vulnerability?

- A. When a hard drive runs out of storage space preventing new files from being saved
- B. When a network connection transfers more data than the bandwidth can accommodate
- C. When a database table exceeds its maximum number of allowed records or entries
- D. When a program writes more data to a buffer than it can hold causing memory corruption

Answer: D

Q247. What is static application security testing (SAST)?

- A. Testing a running application by sending malicious requests to endpoints
- B. Monitoring application behavior in production for suspicious activity
- C. Analyzing source code without executing it to find security vulnerabilities
- D. Scanning network traffic for signs of application-level security attacks

Answer: C

Q248. Why should sensitive data like passwords never be stored in plain text?

- A. Because plain text passwords can be read directly if the database is compromised
- B. Because plain text passwords cannot be validated during the login process flow
- C. Because plain text takes up more storage space than encrypted text in databases
- D. Because plain text is incompatible with modern database management systems today

Answer: A

Q249. What is the purpose of error handling in secure coding?

- A. Displaying all error details to help users troubleshoot their problems
- B. Managing errors gracefully without revealing sensitive system information
- C. Ignoring errors so the application continues running without interruption
- D. Logging errors only to the browser console for developer convenience

Answer: B

Q250. What is the purpose of using version control in secure development?

- A. Compiling code into executable format for production server deployment
- B. Tracking code changes enabling rollback and audit trail of modifications
- C. Encrypting all source code files to prevent unauthorized code access
- D. Automatically fixing security vulnerabilities found in application code

Answer: B

Q251. What is the most secure wireless encryption protocol currently available?

- A. WPA
- B. WEP
- C. WPS
- D. WPA3

Answer: D

Q252. What is a rogue access point?

- A. A legitimate access point that has been configured with a strong password
- B. A backup access point that activates when the primary one stops working
- C. A guest network access point provided for visitors in a corporate office
- D. An unauthorized wireless access point connected to a network without approval

Answer: D

Q253. What is the purpose of a mobile device management (MDM) solution?

- A. Providing free mobile applications to all employees for personal usage
- B. Upgrading the hardware components inside company-issued mobile phones
- C. Increasing the battery life of mobile devices through power optimization
- D. Centrally managing and securing mobile devices used within an organization

Answer: D

Q254. What is Bluetooth security's main concern?

- A. Excessive battery drain caused by Bluetooth hardware power consumption
- B. Slow data transfer speeds compared to modern WiFi network connections
- C. Unauthorized pairing and data interception over short-range connections
- D. Incompatibility between different Bluetooth device manufacturers products

Answer: C

Q255. What is the risk of connecting to public WiFi networks?

- A. Attackers can intercept unencrypted data transmitted over the shared network
- B. Public WiFi networks require stronger passwords than private home networks
- C. Public WiFi networks always provide faster speeds than private connections
- D. Public WiFi automatically installs security software on connected devices

Answer: A

Q256. What does BYOD stand for in mobile security?

- A. Backup Your Online Data
- B. Build Your Own Database
- C. Bring Your Own Device
- D. Browse Your Own Domain

Answer: C

Q257. What is the purpose of remote wipe capability on mobile devices?

- A. Cleaning the device screen remotely using an automated cleaning mechanism
- B. Deleting old text messages to free up storage space on the mobile device
- C. Removing unused applications to improve the mobile device's performance
- D. Erasing all data on a lost or stolen device to prevent unauthorized access

Answer: D

Q258. What is an evil twin attack in wireless security?

- A. Connecting two identical routers to the same network causing IP address conflicts
- B. Creating a fake WiFi access point that mimics a legitimate one to steal data
- C. Sending duplicate data packets across a wireless network to cause congestion
- D. Installing two antivirus programs that conflict with each other on a device

Answer: B

Q259. What is the security risk of jailbreaking or rooting a mobile device?

- A. It enables the device to connect to more WiFi networks simultaneously
- B. It removes built-in security controls making the device vulnerable to malware
- C. It increases the device's processing speed by unlocking hidden CPU features
- D. It improves battery life by disabling unnecessary background system processes

Answer: B

Q260. What is the purpose of a VPN on a mobile device?

- A. Increasing mobile data speed by compressing all network traffic packets
- B. Encrypting network traffic to protect data when using untrusted networks
- C. Extending the WiFi range of the device to connect to distant networks
- D. Blocking all incoming phone calls from unknown or unregistered numbers

Answer: B

Q261. What is the shared responsibility model in cloud security?

- A. A model where the cloud provider handles all security with no customer involvement
- B. A model where the customer handles all security without any provider assistance
- C. A model where security responsibilities are split between the cloud provider and customer
- D. A model where third-party vendors manage all security for both provider and customer

Answer: C

Q262. What are the three main cloud service models?

- A. LAN, WAN, and MAN
- B. CPU, RAM, and SSD
- C. IaaS, PaaS, and SaaS
- D. HTTP, FTP, and SMTP

Answer: C

Q263. What is a cloud access security broker (CASB)?

- A. A person who negotiates cloud service contracts on behalf of organizations
- B. A network device that routes traffic between different cloud provider data centers
- C. A certification program for cloud security professionals in the IT industry
- D. A security tool that sits between users and cloud services to enforce policies

Answer: D

Q264. What is the risk of misconfigured cloud storage buckets?

- A. Data transfer speeds decrease significantly making applications unusable
- B. Sensitive data may be exposed publicly if access controls are not properly set
- C. Cloud storage costs increase dramatically due to inefficient configuration
- D. The cloud provider revokes the organization's account for policy violations

Answer: B

Q265. What is multi-tenancy in cloud computing?

- A. A customer maintaining separate accounts for each department within an organization
- B. Multiple customers sharing the same physical infrastructure in a cloud environment
- C. A single customer using multiple cloud providers simultaneously for redundancy
- D. Multiple data centers connected together to form one large computing cluster

Answer: B

Q266. What is the purpose of identity and access management (IAM) in cloud environments?

- A. Backing up all cloud data to local on-premises servers for redundancy
- B. Monitoring the physical security of data centers operated by cloud providers
- C. Optimizing the performance of applications running on cloud infrastructure
- D. Managing who can access cloud resources and what actions they can perform

Answer: D

Q267. What is data sovereignty in cloud computing?

- A. The cloud provider having complete ownership of all customer data stored
- B. Data being subject to the laws of the country where it is physically stored
- C. Users having the ability to access their data from any country worldwide
- D. Data being automatically encrypted by the cloud provider at all times

Answer: B

Q268. What is the purpose of encryption at rest in cloud storage?

- A. Reducing the overall cost of cloud storage by compressing encrypted data
- B. Speeding up data retrieval from cloud storage by compressing encrypted files
- C. Preventing cloud administrators from modifying data stored in the cloud
- D. Protecting stored data from unauthorized access even if storage is compromised

Answer: D

Q269. What is a virtual private cloud (VPC)?

- A. An isolated section of cloud infrastructure dedicated to a single customer
- B. A private data center that does not connect to any public cloud service
- C. A cloud-based VPN service for encrypting all internet traffic globally
- D. A virtual machine running antivirus software in the cloud environment

Answer: A

Q270. What is shadow IT in the context of cloud security?

- A. Employees using unauthorized cloud services without IT department approval
- B. IT staff working remotely from home offices instead of the corporate campus
- C. Cloud services that operate only during nighttime hours for cost savings
- D. Backup cloud infrastructure that remains inactive until primary systems fail

Answer: A

Q271. What is digital forensics?

- A. The technique of building databases for storing criminal record information
- B. The process of collecting and analyzing digital evidence from electronic devices
- C. The method of encrypting court documents for secure legal proceedings
- D. The practice of designing user interfaces for law enforcement applications

Answer: B

Q272. Why is preserving the chain of custody important in digital forensics?

- A. It ensures evidence integrity and admissibility in legal proceedings
- B. It speeds up the analysis process by organizing evidence chronologically
- C. It prevents investigators from accessing evidence until trial begins
- D. It reduces the cost of forensic investigations by limiting evidence scope

Answer: A

Q273. What is a forensic image of a hard drive?

- A. A compressed backup of only the active files on the hard drive
- B. A log file recording all read and write operations on the drive
- C. A bit-for-bit exact copy of the entire drive including deleted data
- D. A photograph taken of the physical hard drive for documentation

Answer: C

Q274. What is a write blocker used for in digital forensics?

- A. Limiting the number of files that can be saved on a USB drive
- B. Stopping malware from writing malicious files to the hard drive
- C. Preventing any modification to evidence media during examination
- D. Blocking users from writing documents on shared network drives

Answer: C

Q275. What type of evidence can be found in computer RAM?

- A. Permanent files that remain stored even after the computer is completely shut down
- B. Hardware serial numbers and manufacturer information for all installed components
- C. Running processes, encryption keys, and temporary data that disappears when powered off
- D. Operating system installation files that are needed only during the initial setup

Answer: C

Q276. What is the purpose of hashing in digital forensics?

- A. Compressing evidence data to save storage space on forensic servers
- B. Verifying that evidence has not been altered by comparing hash values
- C. Converting binary evidence data into human-readable text documents
- D. Encrypting evidence files to prevent unauthorized access by others

Answer: B

Q277. What is volatile data in forensic investigations?

- A. Data that is stored in cloud services and accessible from anywhere
- B. Data that is lost when a device is powered off such as RAM contents
- C. Data that is permanently stored on hard drives and never changes
- D. Data that is encrypted and cannot be accessed without proper keys

Answer: B

Q278. What are log files used for in forensic analysis?

- A. Managing software licenses for forensic tools used during analysis
- B. Storing user passwords in an encrypted format for secure retrieval
- C. Recording system events that help reconstruct activities and timelines
- D. Tracking the physical location of all devices on the network map

Answer: C

Q279. What is metadata in the context of digital evidence?

- A. Encrypted data that requires a special key to read and understand
- B. Data about data, such as file creation dates, authors, and modification times
- C. Backup copies of files stored on a separate external storage device
- D. The main content of a document file like the text or images within it

Answer: B

Q280. What is the first step a forensic investigator should take at a crime scene?

- A. Immediately turning off all computers to preserve the current state
- B. Securing the scene and documenting the state of all devices present
- C. Deleting unnecessary files to focus the investigation on key evidence
- D. Connecting all devices to the internet for remote forensic analysis

Answer: B

Q281. What is an incident response plan?

- A. A training syllabus for teaching new employees about organizational policies at work
- B. A schedule for performing regular system maintenance and software update installations
- C. A financial plan for budgeting cyber security spending across all fiscal year quarters
- D. A documented strategy for detecting, responding to, and recovering from security incidents

Answer: D

Q282. What is the first phase of the incident response lifecycle according to NIST?

- A. Preparation
- B. Eradication
- C. Containment
- D. Recovery

Answer: A

Q283. What is the purpose of incident containment?

- A. Shutting down the entire organization's network permanently until the threat is gone
- B. Deleting all data on affected systems to remove any traces of the attack completely
- C. Limiting the spread and impact of a security incident to prevent further damage
- D. Notifying the media about the security incident for public awareness and transparency

Answer: C

Q284. What is a CSIRT?

- A. Central Server Infrastructure Review Tool
- B. Cyber Security Integration and Reporting Technology
- C. Computer Security Incident Response Team
- D. Cloud Service Infrastructure Recovery Technique

Answer: C

Q285. What is the purpose of a lessons learned review after an incident?

- A. Assigning blame to the individuals who caused or failed to prevent the incident
- B. Deleting all records of the incident to protect the organization's public reputation
- C. Calculating the exact financial bonus for team members who resolved the incident
- D. Identifying improvements to prevent similar incidents and enhance response processes

Answer: D

Q286. What is the difference between an event and an incident in cyber security?

- A. An event is always malicious while an incident is always a harmless system occurrence
- B. An event is any observable occurrence while an incident is an event that harms security
- C. An event requires investigation while an incident is resolved automatically by tools
- D. An event only happens on servers while an incident only happens on user workstations

Answer: B

Q287. What is the role of incident triage?

- A. Automatically resolving all incidents using predefined scripts without human input
- B. Prioritizing incidents based on severity and impact to allocate resources effectively
- C. Forwarding all incidents to external law enforcement for criminal investigation
- D. Permanently closing all incident tickets without any investigation or resolution

Answer: B

Q288. What is eradication in incident response?

- A. Notifying affected users about the incident and providing credit monitoring services
- B. Removing the root cause of the incident and any attacker artifacts from systems
- C. Documenting the incident for future reference and compliance reporting purposes
- D. Restoring systems to normal operation after the threat has been fully eliminated

Answer: B

Q289. What is the purpose of an incident response communication plan?

- A. Scheduling regular team meetings for discussing non-security project deliverables
- B. Creating social media posts about the organization's daily business operations
- C. Marketing the organization's security capabilities to attract potential customers
- D. Defining who to notify, when, and how during a security incident for coordination

Answer: D

Q290. What is the recovery phase in incident response?

- A. Documenting all evidence collected during the investigation for legal review
- B. Identifying the initial attack vector used by the threat actor to gain access
- C. Purchasing new hardware to replace all systems affected by the security breach
- D. Restoring affected systems to normal operation and verifying they are secure

Answer: D

Q291. What is a SIEM system?

- A. A tool that collects and analyzes security event data from across the network
- B. A protocol for encrypting email communications between users securely
- C. A physical device that blocks malware at the network perimeter gateway
- D. A type of firewall specifically designed for protecting cloud environments

Answer: A

Q292. What is the purpose of security log monitoring?

- A. Deleting old log files to free up storage space on the server hard drives
- B. Converting log files into PDF documents for printing and physical archival
- C. Detecting suspicious activities and security incidents by reviewing system logs
- D. Creating backup copies of log files for disaster recovery planning purposes

Answer: C

Q293. What is a security alert?

- A. A notification generated when a monitoring system detects potentially suspicious activity
- B. A marketing email from a security vendor promoting their latest product features
- C. A scheduled reminder for IT staff to change their passwords at regular intervals
- D. A system notification about available software updates that need to be installed

Answer: A

Q294. What is a false positive in security monitoring?

- A. A system error that prevents the monitoring tool from generating any alerts
- B. A genuine security incident that was correctly detected by the monitoring system
- C. An alert triggered by normal activity that is incorrectly identified as malicious
- D. A disabled alert rule that no longer triggers on any type of security event

Answer: C

Q295. What is network traffic analysis used for in security?

- A. Testing network cable quality to ensure reliable physical connections work
- B. Monitoring network communications to identify anomalies and potential threats
- C. Designing new network architectures for improved organizational connectivity
- D. Increasing network bandwidth by optimizing data packet routing efficiency

Answer: B

Q296. What is endpoint detection and response (EDR)?

- A. A process for repairing damaged endpoint hardware components like keyboards and screens
- B. A method for deploying software updates to all endpoint devices simultaneously
- C. A system for tracking the physical location of all laptop devices issued to employees
- D. A security tool that monitors endpoints for threats and provides investigation capabilities

Answer: D

Q297. What is log aggregation?

- A. Distributing log data across multiple servers for redundant backup storage
- B. Converting log files from text format to binary format for faster reading
- C. Deleting duplicate log entries to reduce the overall volume of stored data
- D. Collecting logs from multiple sources into a centralized location for analysis

Answer: D

Q298. What is a security dashboard?

- A. A physical panel of buttons used to arm and disarm building security alarms
- B. A software application for managing employee access badges and ID cards
- C. A visual display showing security metrics, alerts, and system status in real time
- D. A printed report summarizing monthly security statistics for executive review

Answer: C

Q299. What is vulnerability scanning?

- A. Testing network cable connections for signal quality and potential interference issues
- B. Automated scanning of systems to identify known security weaknesses and misconfigurations
- C. Physically inspecting server hardware for signs of tampering or unauthorized modifications
- D. Manually reviewing source code to find programming errors and logic flaws in software

Answer: B

Q300. What is the purpose of a security operations center (SOC)?

- A. Providing technical support to employees experiencing computer problems
- B. Developing new software applications for the organization's customers
- C. Managing the physical security of the building including locks and guards
- D. Centrally monitoring and responding to security events across the organization

Answer: D

Q301. What is the primary purpose of data protection laws?

- A. Setting minimum salary requirements for IT professionals in organizations
- B. Increasing internet speeds for all citizens within a country's borders
- C. Regulating the manufacturing standards for computer hardware products
- D. Protecting individuals' personal data from misuse and unauthorized processing

Answer: D

Q302. What does GDPR stand for?

- A. General Digital Privacy Regulation
- B. Global Digital Privacy Requirement
- C. General Data Protection Regulation
- D. Government Data Processing Rules

Answer: C

Q303. What is intellectual property in the context of cyber law?

- A. Employee skills and knowledge that are gained through professional training programs
- B. Network bandwidth that is allocated to a specific organization by an internet provider
- C. Physical computer hardware that is owned by an organization and stored in data centers
- D. Creations of the mind like inventions, designs, and software that are legally protected

Answer: D

Q304. What is the purpose of a privacy policy on a website?

- A. Displaying advertisements that are relevant to the user's browsing history
- B. Informing users how their personal data is collected, used, and protected
- C. Preventing users from accessing certain pages on the website based on age
- D. Encrypting all data transmitted between the user's browser and the server

Answer: B

Q305. What is ethical hacking?

- A. Hacking into systems without permission to demonstrate security weaknesses publicly
- B. Using social engineering to test employee awareness without organizational approval
- C. Creating malware for research purposes without deploying it in real environments
- D. Authorized testing of systems to find vulnerabilities before malicious hackers exploit them

Answer: D

Q306. What is a non-disclosure agreement (NDA) used for in cyber security?

- A. Allowing employees to share company secrets with competitors for personal benefit
- B. Disclosing all security vulnerabilities publicly to raise awareness about threats
- C. Legally binding parties to keep shared confidential information private and secure
- D. Granting customers full access to the organization's proprietary source code

Answer: C

Q307. What does the term 'compliance' mean in cyber security?

- A. Completing all security training courses within the required time period
- B. Competing with other organizations to achieve the highest security ratings
- C. Following established laws, regulations, and standards for information security
- D. Purchasing the most expensive security tools available on the market today

Answer: C

Q308. What is responsible disclosure in cyber security?

- A. Publishing all discovered vulnerabilities immediately on social media platforms
- B. Ignoring discovered vulnerabilities to avoid any potential legal consequences
- C. Selling vulnerability information to the highest bidder on dark web forums
- D. Reporting vulnerabilities to the vendor privately before making them public

Answer: D

Q309. What is the purpose of cyber crime laws?

- A. Providing free internet access to all citizens as a fundamental digital right
- B. Defining illegal cyber activities and establishing penalties for those who commit them
- C. Mandating that all organizations must use open-source software exclusively
- D. Regulating the prices of cyber security software products sold in the market

Answer: B

Q310. What is digital consent in data privacy?

- A. Allowing any organization to access user data without requiring authorization
- B. Obtaining user permission before collecting and processing their personal data
- C. Storing user data indefinitely regardless of whether consent was ever provided
- D. Automatically collecting user data without informing them about the practice

Answer: B

Q311. What is the Internet of Things (IoT) in cyber security?

- A. A cloud storage service that allows unlimited data storage for connected devices
- B. A security protocol designed to protect web browsers from malicious website content
- C. A new type of internet connection that is faster than traditional broadband service
- D. Network of connected devices that collect and exchange data creating security challenges

Answer: D

Q312. What is artificial intelligence's role in cyber security?

- A. Enhancing threat detection, automating responses, and analyzing large security datasets
- B. Replacing all human security professionals with fully autonomous robot defenders
- C. Creating new types of encryption that are impossible to break by any method
- D. Eliminating all cyber threats permanently through advanced predictive algorithms

Answer: A

Q313. What is blockchain technology's relevance to cyber security?

- A. Eliminating the need for encryption because blockchain is inherently secure
- B. Providing tamper-resistant record-keeping and decentralized trust mechanisms
- C. Replacing all existing databases with faster and cheaper storage solutions
- D. Creating untraceable communications that cannot be monitored by anyone

Answer: B

Q314. What is quantum computing's potential impact on cyber security?

- A. Eliminating all cyber threats because quantum computers cannot be compromised
- B. Making all existing security tools obsolete and completely useless immediately
- C. Threatening current encryption methods by potentially breaking widely used algorithms
- D. Having no impact on cyber security because quantum computing is only theoretical

Answer: C

Q315. What is zero trust architecture?

- A. A network design that eliminates all firewalls to simplify security management
- B. A security model that trusts no one by default and verifies every access request
- C. A framework that trusts all internal users and only verifies external connections
- D. A policy requiring zero passwords for user authentication across all systems

Answer: B

Q316. What is DevSecOps?

- A. Integrating security practices into the DevOps software development pipeline
- B. A cloud service that automatically secures deployed applications and databases
- C. A specialized operating system designed for security-focused development work
- D. A certification for developers who specialize in building secure applications

Answer: A

Q317. What is ransomware-as-a-service (RaaS)?

- A. A legitimate cloud service that helps organizations recover from ransomware attacks
- B. A government program providing free ransomware decryption tools to affected organizations
- C. A business model where ransomware developers sell or lease their tools to other criminals
- D. A cybersecurity insurance product that covers ransomware payment costs for businesses

Answer: C

Q318. What is edge computing and its security implications?

- A. A computing approach that only works with wired connections and cannot use wireless links
- B. A method of strengthening network perimeter defenses at the edge of the corporate network
- C. A technique for storing backup data at the geographical edges of a country's territory
- D. Processing data closer to where it is generated, creating new distributed security challenges

Answer: D

Q319. What is a deepfake?

- A. A backup copy of data that is stored in a deeply encrypted format for maximum security
- B. A social media account that pretends to be a real person for legitimate marketing goals
- C. A type of encryption that creates fake data to confuse attackers during a breach event
- D. AI-generated synthetic media that can convincingly impersonate real people in video or audio

Answer: D

Q320. What is cyber resilience?

- A. The speed at which an internet connection recovers after a temporary service outage
- B. The physical durability of computer hardware against environmental damage threats
- C. A software feature that automatically restarts applications after an unexpected crash
- D. An organization's ability to prepare for, respond to, and recover from cyber incidents

Answer: D

Q321. What does the term 'cyber hygiene' refer to?

- A. Cleaning computer hardware regularly
- B. Routine practices to maintain system health and security
- C. Installing only free software
- D. Using the internet only during business hours

Answer: B

Q322. Which of the following is an example of a cyber attack?

- A. Backing up files to a cloud server
- B. Sending an encrypted email
- C. Deploying ransomware on a network
- D. Installing a security patch

Answer: C

Q323. What is the primary goal of information security?

- A. Making systems run faster
- B. Protecting information from unauthorized access, use, or destruction
- C. Reducing hardware costs
- D. Increasing internet bandwidth

Answer: B

Q324. What does the term 'exploit' mean in cyber security?

- A. A security certification program
- B. A piece of software or technique that takes advantage of a vulnerability
- C. A type of antivirus software
- D. A network monitoring tool

Answer: B

Q325. Which professional is responsible for protecting an organization's computer systems?

- A. Database administrator
- B. Cyber security analyst
- C. Graphic designer
- D. Help desk technician

Answer: B

Q326. What is a 'patch' in the context of cyber security?

- A. A physical repair to damaged hardware
- B. A software update that fixes security vulnerabilities
- C. A type of network cable
- D. A backup storage device

Answer: B

Q327. What does 'authentication' verify in cyber security?

- A. The speed of a network connection
- B. The identity of a user or system
- C. The size of a file
- D. The cost of a software license

Answer: B

Q328. Which of the following is a common indicator of a phishing email?

- A. The email comes from a known colleague
- B. The email has a professional signature
- C. The email contains urgent language and suspicious links
- D. The email is about a scheduled meeting

Answer: C

Q329. What is a 'risk' in cyber security?

- A. Any computer connected to the internet
- B. The potential for loss or damage when a threat exploits a vulnerability
- C. A type of encryption algorithm
- D. A network protocol

Answer: B

Q330. What is the purpose of a security awareness program?

- A. To train employees on recognizing and responding to security threats
- B. To install firewalls on every computer
- C. To replace antivirus software
- D. To encrypt all company emails automatically

Answer: A

Q331. What is a security incident?

- A. A scheduled system update
- B. An event that threatens the confidentiality, integrity, or availability of information
- C. A routine password change
- D. A planned network maintenance window

Answer: B

Q332. What is the purpose of a firewall in network security?

- A. To increase internet speed
- B. To monitor and control incoming and outgoing network traffic based on security rules
- C. To store backup files
- D. To manage email accounts

Answer: B

Q333. What is the difference between a threat and a vulnerability?

- A. They are the same thing
- B. A threat is a potential danger while a vulnerability is a weakness that can be exploited
- C. A vulnerability is more dangerous than a threat
- D. Threats only exist online while vulnerabilities are physical

Answer: B

Q334. What does 'non-repudiation' ensure in information security?

- A. That data is always available
- B. That a user cannot deny having performed a specific action
- C. That passwords are never reused
- D. That firewalls block all traffic

Answer: B

Q335. What is a security policy?

- A. A type of antivirus software
- B. A document that outlines an organization's rules and procedures for protecting information
- C. A hardware device for encrypting files
- D. A programming language for security tools

Answer: B

Q336. Which of the following is a preventive security control?

- A. Security camera
- B. Intrusion detection system
- C. Access control lock on a door
- D. Audit log review

Answer: C

Q337. What is data classification in information security?

- A. Sorting data alphabetically
- B. Categorizing data based on its sensitivity level to apply appropriate protection
- C. Deleting old data from servers
- D. Compressing data to save storage space

Answer: B

Q338. What does the term 'least privilege' mean in security?

- A. Giving all users administrator access
- B. Granting users only the minimum access necessary to perform their job functions
- C. Removing all access permissions from users
- D. Using the cheapest security software available

Answer: B

Q339. What is two-factor authentication (2FA)?

- A. Using two different passwords
- B. A security method requiring two different types of verification to prove identity
- C. Logging in from two different devices
- D. Having two user accounts

Answer: B

Q340. What is the purpose of an acceptable use policy (AUP)?

- A. To define how employees may use organizational IT resources
- B. To list all software installed on company computers
- C. To track employee internet browsing history
- D. To set pricing for IT services

Answer: A

Q341. What is the difference between encryption and hashing?

- A. They are the same process
- B. Encryption is reversible with a key while hashing is a one-way process
- C. Hashing is always faster than encryption
- D. Encryption does not use keys

Answer: B

Q342. What is a private key in asymmetric cryptography?

- A. A key shared with everyone publicly
- B. A key kept secret by the owner and used for decryption or signing
- C. A key used only for hashing
- D. A password used to log into websites

Answer: B

Q343. What is the purpose of SSL/TLS in web communications?

- A. To speed up website loading times
- B. To encrypt communications between a web browser and web server
- C. To compress images on websites
- D. To block advertisements on web pages

Answer: B

Q344. Which of the following is an example of a hash algorithm?

- A. AES
- B. RSA
- C. SHA-256
- D. Diffie-Hellman

Answer: C

Q345. What does 'key length' refer to in cryptography?

- A. The physical length of a USB security key
- B. The number of bits in a cryptographic key, which affects encryption strength
- C. The number of characters in a password
- D. The duration a key remains valid

Answer: B

Q346. What is a digital certificate used for?

- A. To store encrypted files permanently
- B. To verify the identity of a website or entity and bind a public key to that identity
- C. To block malware from downloading
- D. To increase internet connection speed

Answer: B

Q347. What does symmetric encryption mean?

- A. Different keys are used for encryption and decryption
- B. The same key is used for both encryption and decryption
- C. No key is needed for encryption
- D. Only the sender needs a key

Answer: B

Q348. What is a cryptographic protocol?

- A. A type of computer virus
- B. A set of rules that governs how cryptographic operations are performed in communications
- C. A hardware device for encryption
- D. A programming language for writing security software

Answer: B

Q349. Why should encryption keys be kept secret?

- A. Because keys are expensive to generate
- B. Because anyone with the key can decrypt the protected data
- C. Because keys expire after one use
- D. Because keys take up storage space

Answer: B

Q350. What is end-to-end encryption?

- A. Encrypting only the first and last packets of a transmission
- B. Encrypting data at the sender's end so that only the intended recipient can decrypt it
- C. Encrypting data only when stored on a server
- D. Using two firewalls at each end of a network

Answer: B

Q351. What is an IP address used for in networking?

- A. To encrypt data on a network
- B. To uniquely identify a device on a network
- C. To increase network speed
- D. To store files on a server

Answer: B

Q352. What is a router's role in network security?

- A. It only provides wireless internet access
- B. It directs network traffic between networks and can enforce security policies through access control lists
- C. It stores all network passwords
- D. It replaces the need for firewalls

Answer: B

Q353. What is the purpose of network address translation (NAT) in security?

- A. To speed up internet connections significantly
- B. To hide internal IP addresses from external networks
- C. To encrypt all outgoing network packets
- D. To replace the need for antivirus software

Answer: B

Q354. What does the term 'bandwidth' refer to in networking?

- A. The physical width of a network cable
- B. The maximum amount of data that can be transmitted over a network connection in a given time
- C. The number of computers on a network
- D. The strength of a WiFi signal

Answer: B

Q355. What is the purpose of a network switch in a local area network?

- A. To connect a network to the internet
- B. To connect multiple devices within a network and forward data to the correct destination device
- C. To encrypt all network traffic
- D. To scan for viruses on the network

Answer: B

Q356. What is network monitoring used for?

- A. To increase the speed of all network connections
- B. To observe and analyze network traffic for performance issues and security threats
- C. To replace all security controls
- D. To generate encryption keys

Answer: B

Q357. What does DNS stand for and what is its primary function?

- A. Data Network Storage; stores files on a network
- B. Domain Name System; translates domain names into IP addresses
- C. Digital Network Security; encrypts network traffic
- D. Direct Network Service; manages network connections

Answer: B

Q358. What is a network protocol?

- A. A hardware device used to connect networks
- B. A set of rules governing how data is transmitted and received over a network
- C. A type of network cable
- D. A software program for browsing the internet

Answer: B

Q359. What is the function of an Intrusion Prevention System (IPS)?

- A. To provide wireless internet access
- B. To detect and actively block malicious network traffic in real time
- C. To create backups of network data
- D. To assign IP addresses to devices

Answer: B

Q360. What is a VLAN used for in network security?

- A. To increase the physical range of a WiFi network
- B. To logically segment a network into separate broadcast domains for improved security and management
- C. To encrypt all data on the network
- D. To replace physical firewalls

Answer: B

Q361. What is a one-time pad in cryptography?

- A. A cipher that uses a random key as long as the message and is used only once
- B. A reusable encryption key for multiple messages
- C. A type of hash function
- D. A digital certificate format

Answer: A

Q362. Which type of encryption is faster for large amounts of data: symmetric or asymmetric?

- A. Asymmetric encryption
- B. Both are equally fast
- C. Neither can handle large data
- D. Symmetric encryption

Answer: D

Q363. What does the term 'key pair' refer to in asymmetric cryptography?

- A. Two identical keys used for encryption
- B. A public key and a corresponding private key
- C. Two symmetric keys shared between parties
- D. A key and its backup copy

Answer: B

Q364. What is the purpose of Base64 encoding?

- A. To encrypt data securely
- B. To compress data for storage
- C. To represent binary data in an ASCII string format
- D. To generate cryptographic hashes

Answer: C

Q365. What does the acronym AES stand for?

- A. Asymmetric Encryption System
- B. Advanced Encryption Standard
- C. Automated Encoding Scheme
- D. Applied Encryption Software

Answer: B

Q366. What is a cryptographic hash collision?

- A. When two different inputs produce the same hash output
- B. When a hash function fails to produce output
- C. When encryption and decryption keys match
- D. When two users share the same password

Answer: A

Q367. What is the role of a random number generator in cryptography?

- A. To create user passwords automatically
- B. To generate unpredictable values for keys, nonces, and IVs
- C. To speed up encryption algorithms
- D. To compress data before encryption

Answer: B

Q368. What is the difference between encoding and encryption?

- A. They are the same thing
- B. Encoding transforms data for usability while encryption transforms data for confidentiality
- C. Encoding is more secure than encryption
- D. Encryption is faster than encoding

Answer: B

Q369. Which of the following is NOT a symmetric encryption algorithm?

- A. AES
- B. DES
- C. Blowfish
- D. RSA

Answer: D

Q370. What is a digital envelope in cryptography?

- A. A physical container for cryptographic keys
- B. A method combining symmetric and asymmetric encryption to securely transmit data
- C. A type of email encryption only
- D. A digital signature format

Answer: B

Q371. What is the purpose of a firewall built into an operating system?

- A. To speed up internet connections
- B. To filter incoming and outgoing network traffic based on security rules
- C. To defragment the hard drive
- D. To manage user passwords

Answer: B

Q372. What is the principle of least privilege in operating system security?

- A. Giving all users administrator access
- B. Users and programs should only have the minimum permissions needed to perform their tasks
- C. Disabling all user accounts except root
- D. Removing all security software

Answer: B

Q373. What is a software update in the context of OS security?

- A. A new version of a video game
- B. A patch or fix released to address security vulnerabilities or bugs in the OS
- C. A change to the desktop wallpaper
- D. An increase in storage capacity

Answer: B

Q374. What does antivirus software do on an operating system?

- A. Increases internet speed
- B. Detects, prevents, and removes malicious software
- C. Manages file storage
- D. Creates system backups

Answer: B

Q375. Why should users avoid running programs as administrator unnecessarily?

- A. It makes the computer slower
- B. Programs running with admin privileges can make system-wide changes that could be harmful if compromised
- C. Admin accounts have fewer features
- D. There is no reason to avoid it

Answer: B

Q376. What is a strong password policy for OS user accounts?

- A. Using the word 'password' for all accounts
- B. Requiring passwords with minimum length, complexity, and regular changes
- C. Allowing blank passwords for convenience
- D. Using the same password for all accounts

Answer: B

Q377. What is a boot password in OS security?

- A. A password required before the operating system loads
- B. The default password for all applications
- C. A password used only for email
- D. A password for internet browsing

Answer: A

Q378. What happens when automatic updates are disabled on an OS?

- A. The system becomes faster
- B. The system may miss critical security patches and remain vulnerable to known threats
- C. Nothing changes
- D. The system automatically becomes more secure

Answer: B

Q379. What is a guest account on an operating system?

- A. The main administrator account
- B. A limited-privilege account for temporary users with restricted access
- C. An account with full system permissions
- D. A hidden account used by hackers

Answer: B

Q380. Why is it important to lock your computer screen when stepping away?

- A. To save electricity only
- B. To prevent unauthorized physical access to your active session and data
- C. It is not important
- D. To improve display quality

Answer: B

Q381. What is a web browser's address bar padlock icon used to indicate?

- A. The website is free of malware
- B. The connection to the website is encrypted using HTTPS
- C. The website content is verified as accurate
- D. The website loads faster than others

Answer: B

Q382. What is the purpose of the robots.txt file on a website?

- A. To block all hackers from accessing the site
- B. To instruct search engine crawlers which parts of the site to avoid indexing
- C. To encrypt website data
- D. To manage user passwords

Answer: B

Q383. What is a URL redirect vulnerability?

- A. When a website loads too slowly
- B. When a website redirects users to a malicious site using unvalidated redirect parameters
- C. When a website's URL is too long
- D. When a website changes its domain name

Answer: B

Q384. Why should sensitive forms on websites use POST requests instead of GET?

- A. POST requests are faster
- B. GET requests include data in the URL which can be cached, logged, and visible in browser history
- C. POST requests do not require a server
- D. GET requests cannot carry data

Answer: B

Q385. What is the purpose of rate limiting on a web application?

- A. To make the website faster
- B. To restrict the number of requests a user can make in a given time period to prevent abuse
- C. To increase server storage
- D. To improve website appearance

Answer: B

Q386. What is a Content Delivery Network (CDN) and does it improve security?

- A. A type of encryption algorithm
- B. A distributed network of servers that can improve performance and provide basic DDoS protection
- C. A database management system
- D. A web programming language

Answer: B

Q387. What is an HTTP cookie used for in web applications?

- A. To encrypt web pages
- B. To store small pieces of data on the client's browser for session management, preferences, and tracking
- C. To speed up DNS resolution
- D. To compress web traffic

Answer: B

Q388. What is cross-origin resource sharing (CORS)?

- A. A type of web attack
- B. A mechanism that allows web pages to request resources from a different domain than the one that served the page
- C. A database query language
- D. A server-side scripting language

Answer: B

Q389. Why is it important to use prepared statements when interacting with databases in web applications?

- A. They make queries run faster
- B. They prevent SQL injection by separating SQL code from user data
- C. They reduce database storage usage
- D. They are required by all programming languages

Answer: B

Q390. What does the HttpOnly flag on a cookie do?

- A. Makes the cookie work only over HTTP, not HTTPS
- B. Prevents the cookie from being accessed by client-side JavaScript
- C. Encrypts the cookie value
- D. Deletes the cookie after the browser closes

Answer: B

Q391. What is a password manager and why should you use one?

- A. A tool that remembers only one password
- B. A tool that generates and securely stores unique complex passwords for all your accounts
- C. A tool that removes password requirements
- D. A tool that shares passwords with others

Answer: B

Q392. What does 'something you know' refer to as an authentication factor?

- A. A fingerprint scan
- B. A smart card
- C. Knowledge-based credentials such as a password, PIN, or security question answer
- D. A one-time code from an app

Answer: C

Q393. What is the difference between a PIN and a password?

- A. There is no difference
- B. A PIN is typically a short numeric code while a password is usually longer and can contain letters, numbers, and symbols
- C. A PIN is more secure than a password
- D. A password can only contain numbers

Answer: B

Q394. What is session timeout and why is it important for security?

- A. When a website permanently deletes your account
- B. Automatically ending a user's session after a period of inactivity to prevent unauthorized access
- C. When the server crashes from too many users
- D. When your internet connection drops

Answer: B

Q395. What is two-step verification?

- A. Creating two separate accounts
- B. An authentication process that requires two separate verification steps, typically a password and a code sent to your device
- C. Logging in twice with the same password
- D. Using two different browsers

Answer: B

Q396. What is the purpose of a security question during account recovery?

- A. To make the login process slower
- B. To verify identity by asking a question that only the legitimate account owner should be able to answer
- C. To test the user's general knowledge
- D. To encrypt the user's data

Answer: B

Q397. What is the difference between logging in and logging out?

- A. They are the same action
- B. Logging in authenticates a user and creates a session; logging out ends the session and revokes access
- C. Logging out gives more access than logging in
- D. Only administrators can log out

Answer: B

Q398. Why should you not share your passwords with others?

- A. Shared passwords use more server resources
- B. Sharing passwords means others can access your accounts, and you lose control over who does what with your credentials
- C. Passwords stop working when shared
- D. It is only a recommendation, not important

Answer: B

Q399. What is a brute-force attack on a login page?

- A. A physical attack on the server
- B. An automated attack that systematically tries all possible password combinations until the correct one is found
- C. Refreshing the login page repeatedly
- D. Typing the password very quickly

Answer: B

Q400. What is an authenticator app used for?

- A. To browse the web securely
- B. To generate time-based one-time passwords (TOTP) for two-factor authentication
- C. To manage email accounts
- D. To encrypt files on your phone

Answer: B

Q401. What is a code vulnerability?

- A. A feature that makes code run faster
- B. A weakness or flaw in software code that could be exploited by an attacker
- C. A type of programming language
- D. A software license agreement

Answer: B

Q402. Why should developers avoid hardcoding secrets like API keys in source code?

- A. Hardcoded keys make the code too long
- B. Hardcoded secrets can be exposed in version control systems, making them accessible to anyone with code access
- C. Hardcoded keys improve performance
- D. There is no reason to avoid it

Answer: B

Q403. What is the purpose of logging in application development?

- A. To slow down the application
- B. To record events, errors, and user activities for debugging, monitoring, and security auditing
- C. To encrypt application data
- D. To manage database connections

Answer: B

Q404. What is dependency management in software development?

- A. Managing the development team
- B. Tracking and managing the external libraries and frameworks a project relies on to ensure they are secure and up-to-date
- C. Managing database dependencies
- D. Organizing code files into folders

Answer: B

Q405. What is a security requirement in software development?

- A. A request for more developers
- B. A specification that defines the security features and controls a software application must implement
- C. A requirement to use a specific programming language
- D. A deadline for project completion

Answer: B

Q406. Why is it important to handle errors properly in secure coding?

- A. Error handling makes code longer
- B. Proper error handling prevents sensitive information from being exposed in error messages and maintains application stability
- C. Error handling is only for debugging
- D. Error handling slows down the application

Answer: B

Q407. What is the purpose of unit testing in secure development?

- A. To test the entire application at once
- B. To test individual components or functions in isolation to verify they behave correctly, including security-related behavior
- C. To test network performance
- D. To test user interface design

Answer: B

Q408. What is the concept of 'defense in depth' in application development?

- A. Using one very strong security control
- B. Implementing multiple layers of security controls so that if one layer fails, others still provide protection
- C. Defending only the database
- D. Developing applications quickly without testing

Answer: B

Q409. What is the difference between authentication and input validation in secure coding?

- A. They are the same concept
- B. Authentication verifies user identity while input validation ensures data meets expected format and constraints before processing
- C. Input validation replaces the need for authentication
- D. Authentication validates all input automatically

Answer: B

Q410. Why should developers use parameterized queries instead of string concatenation for database queries?

- A. Parameterized queries are shorter to write
- B. Parameterized queries separate SQL code from data, preventing SQL injection attacks that string concatenation enables
- C. String concatenation is faster
- D. There is no practical difference

Answer: B

Q411. What is airplane mode on a mobile device?

- A. A mode that makes the phone look like an airplane
- B. A setting that disables all wireless communications including cellular, WiFi, and Bluetooth
- C. A mode that increases battery life by 10 times
- D. A special mode for use during flights only

Answer: B

Q412. Why should you verify WiFi network names before connecting?

- A. To ensure the fastest connection
- B. Attackers can create fake networks with similar names to legitimate ones to intercept your data
- C. WiFi names affect download speeds
- D. Network names determine the security level automatically

Answer: B

Q413. What is the purpose of app permissions on mobile devices?

- A. To make apps run faster
- B. To control which device resources and data an app can access, such as camera, location, and contacts
- C. To manage app storage size
- D. To determine which users can install the app

Answer: B

Q414. What is a hotspot on a mobile device?

- A. A device that gets physically hot
- B. A feature that shares the device's cellular data connection with other devices via WiFi
- C. A location with the strongest cell signal
- D. A type of malware

Answer: B

Q415. Why should you download apps only from official app stores?

- A. Official apps are always free
- B. Official app stores review apps for malware and security issues, reducing the risk of installing malicious software
- C. Unofficial apps have better features
- D. It makes no difference where you download apps

Answer: B

Q416. What is NFC on a mobile device?

- A. New Firewall Configuration
- B. Near Field Communication, a short-range wireless technology used for contactless payments and data transfer
- C. Network Frequency Controller
- D. Next-gen File Compression

Answer: B

Q417. What does encryption of mobile device storage protect against?

- A. App crashes
- B. Unauthorized access to data on the device if it is lost, stolen, or physically accessed by an attacker
- C. Slow internet speeds
- D. Battery drain

Answer: B

Q418. What is a wireless access point?

- A. A wired ethernet port
- B. A hardware device that allows wireless devices to connect to a wired network using WiFi
- C. A type of mobile application
- D. A satellite internet receiver

Answer: B

Q419. Why is it important to keep your mobile device's operating system updated?

- A. Updates only add new emojis
- B. Updates include security patches that fix known vulnerabilities that attackers could exploit
- C. Updates always make the device slower
- D. Updates are optional and unimportant

Answer: B

Q420. What is a SIM card in a mobile phone?

- A. A security camera
- B. A small chip that identifies the subscriber and allows the phone to connect to the cellular network
- C. A type of mobile storage for photos
- D. A screen protector

Answer: B

Q421. What is a cloud storage service?

- A. A physical warehouse for storing computers
- B. An online service that allows users to store, access, and share files over the internet
- C. A type of local hard drive
- D. A mobile application for photography

Answer: B

Q422. What is the difference between a public cloud and a private cloud?

- A. There is no difference
- B. A public cloud shares infrastructure among multiple organizations while a private cloud is dedicated to a single organization
- C. Public clouds are free while private clouds are paid
- D. Private clouds do not use the internet

Answer: B

Q423. What is a hybrid cloud environment?

- A. A cloud that only works during certain hours
- B. A combination of public and private cloud infrastructure that allows data and applications to be shared between them
- C. A cloud service that does not use encryption
- D. A backup-only cloud solution

Answer: B

Q424. Why is it important to use strong passwords for cloud accounts?

- A. Cloud services do not support strong passwords
- B. Weak passwords can be easily guessed or cracked, giving attackers access to all your cloud-stored data and services
- C. Strong passwords are only needed for local accounts
- D. Password strength does not affect cloud security

Answer: B

Q425. What is a cloud service level agreement (SLA)?

- A. A type of cloud encryption
- B. A contract between the cloud provider and customer that defines the expected level of service, availability, and security responsibilities
- C. A programming language for cloud apps
- D. A cloud storage format

Answer: B

Q426. What is data backup in cloud computing?

- A. Deleting old data from the cloud
- B. Creating copies of data stored in the cloud to protect against data loss from accidental deletion, corruption, or attacks
- C. Moving all data to a local hard drive
- D. Encrypting data at rest

Answer: B

Q427. What is the role of encryption in transit for cloud services?

- A. It makes data transfer faster
- B. It protects data while it is being transmitted between your device and the cloud server, preventing eavesdropping
- C. It compresses data for storage
- D. It is only used for email

Answer: B

Q428. What is a cloud region in cloud computing?

- A. A marketing term with no technical meaning
- B. A geographic area containing data centers where cloud resources are hosted, affecting latency and data residency compliance
- C. A type of cloud pricing model
- D. The cloud provider's headquarters location

Answer: B

Q429. What is multi-factor authentication's importance for cloud services?

- A. MFA is not useful for cloud services
- B. MFA adds extra security layers beyond passwords for cloud accounts, significantly reducing the risk of unauthorized access even if passwords are compromised
- C. MFA slows down cloud performance
- D. MFA is only for on-premises systems

Answer: B

Q430. What is a cloud compliance requirement?

- A. A requirement to use only one cloud provider
- B. Regulatory standards and laws that organizations must follow when storing and processing data in the cloud, such as data protection and privacy regulations
- C. A requirement to back up data daily
- D. A cloud provider's internal policy only

Answer: B

Q431. What is the purpose of taking photographs at a digital crime scene?

- A. For social media posts
- B. To document the state of evidence including screen displays, cable connections, and device locations before any changes are made
- C. To test the camera equipment
- D. For marketing materials

Answer: B

Q432. What is a forensic copy different from a regular file copy?

- A. They are exactly the same thing
- B. A forensic copy is a bit-for-bit exact duplicate of an entire storage device including deleted files and empty space, while a regular copy only copies visible files
- C. A forensic copy is smaller in size
- D. A regular copy includes more data than a forensic copy

Answer: B

Q433. What role does a hash value play in digital forensics?

- A. It encrypts the evidence
- B. It creates a unique digital fingerprint of evidence that can be used to verify the evidence has not been altered
- C. It compresses evidence files
- D. It translates evidence into different languages

Answer: B

Q434. What is a forensic workstation?

- A. Any regular office computer
- B. A specialized computer configured for analyzing digital evidence with forensic tools, write-blocking capabilities, and proper security controls
- C. A computer used only for word processing
- D. A server that stores websites

Answer: B

Q435. Why is it important to document every step during a forensic investigation?

- A. For billing purposes only
- B. To maintain a detailed record that supports the integrity of evidence and can be presented in legal proceedings
- C. Documentation is optional in forensics
- D. To create a user manual for the software

Answer: B

Q436. What are the common sources of digital evidence in a computer investigation?

- A. Only the hard drive
- B. Hard drives, RAM, network logs, email, browser history, USB devices, cloud accounts, and mobile devices
- C. Only printed documents
- D. Only internet browsing history

Answer: B

Q437. What is the difference between logical and physical acquisition in digital forensics?

- A. They are the same method
- B. Logical acquisition copies files and folders accessible by the OS, while physical acquisition copies the entire storage media bit-by-bit including deleted data
- C. Logical acquisition is always better
- D. Physical acquisition only copies visible files

Answer: B

Q438. What is an evidence bag used for in digital forensics?

- A. To carry personal items
- B. An anti-static, tamper-evident container used to store and transport physical digital evidence while maintaining chain of custody
- C. A regular plastic bag
- D. A bag for storing paper documents

Answer: B

Q439. What is deleted file recovery in digital forensics?

- A. Permanently destroying files
- B. The process of recovering files that have been deleted by users, as the data may still exist on the storage device until overwritten
- C. Downloading files from the internet
- D. Creating backup copies of files

Answer: B

Q440. What is the role of an expert witness in digital forensics?

- A. To hack into the suspect's computer
- B. A qualified professional who presents and explains digital evidence and forensic findings in court in a way that judges and juries can understand
- C. To destroy evidence after the trial
- D. To represent the defendant

Answer: B

Q441. What is the goal of incident response in cybersecurity?

- A. To prevent all incidents from ever occurring
- B. To systematically handle security incidents to minimize damage, reduce recovery time, and learn from the experience
- C. To blame individuals for security failures
- D. To shut down all systems permanently after an attack

Answer: B

Q442. Who should be notified when a security incident is detected?

- A. Only the IT department
- B. The incident response team, management, and potentially affected parties as defined in the incident response plan
- C. No one, to avoid causing panic
- D. Only external media

Answer: B

Q443. What is the purpose of an incident response plan?

- A. To create work for the IT team
- B. To provide a structured, pre-defined approach for handling security incidents so the team can respond quickly and effectively
- C. To document past projects
- D. To track employee attendance

Answer: B

Q444. What should an employee do if they receive a suspicious email?

- A. Forward it to all colleagues to warn them
- B. Report it to the IT security team and do not click any links or open attachments
- C. Delete it and forget about it
- D. Reply to the sender asking if it is legitimate

Answer: B

Q445. What is incident containment?

- A. Putting all computers in a sealed room
- B. Taking immediate actions to limit the scope and damage of a security incident, such as isolating affected systems from the network
- C. Ignoring the incident until it resolves itself
- D. Deleting all data from affected systems

Answer: B

Q446. What is the recovery phase of incident response?

- A. Recovering the cost of the incident from insurance
- B. Restoring affected systems and services to normal operation after the threat has been eliminated
- C. Recovering deleted emails
- D. Recovering employee passwords

Answer: B

Q447. Why should an organization conduct post-incident reviews?

- A. To assign blame to individuals
- B. To analyze what happened, identify what worked well and what needs improvement, and update procedures to better handle future incidents
- C. Post-incident reviews are a waste of time
- D. To create a public news report about the incident

Answer: B

Q448. What is the role of backup systems in incident response?

- A. Backups are not related to incident response
- B. Backups provide clean copies of data and systems that can be used to restore operations after an incident like ransomware or data corruption
- C. Backups only store old, unnecessary data
- D. Backups are used only for testing purposes

Answer: B

Q449. What is a security incident log?

- A. A wooden log used for physical security
- B. A record that documents the details of a security incident including timeline, actions taken, evidence collected, and people involved
- C. A log of employee login times
- D. A record of software purchases

Answer: B

Q450. What is the difference between a security event and a security incident?

- A. They mean exactly the same thing
- B. A security event is any observable occurrence in a system; a security incident is an event that actually threatens or breaches security policies
- C. An event is more serious than an incident
- D. An incident always involves data loss

Answer: B

Q451. What is the purpose of an intrusion detection system (IDS)?

- A. To speed up network performance
- B. To monitor network traffic or system activity for malicious behavior or policy violations and generate alerts
- C. To encrypt all network traffic
- D. To manage user passwords

Answer: B

Q452. What is a security baseline in monitoring?

- A. The lowest level of security an organization can have
- B. A documented standard of normal system behavior and configuration that deviations from can indicate security issues
- C. The first security tool purchased
- D. A baseline measurement of network speed

Answer: B

Q453. What is the purpose of monitoring failed login attempts?

- A. To count how many employees forget their passwords
- B. To detect potential brute-force attacks, unauthorized access attempts, or compromised accounts
- C. Failed logins do not need to be monitored
- D. To measure system performance

Answer: B

Q454. What is a security information and event management (SIEM) alert?

- A. A marketing email from a security vendor
- B. A notification generated when the SIEM detects activity matching predefined rules or patterns that may indicate a security threat
- C. A reminder to update passwords
- D. An error message in the SIEM software

Answer: B

Q455. Why is it important to monitor both inbound and outbound network traffic?

- A. Only inbound traffic needs monitoring
- B. Inbound monitoring detects incoming attacks while outbound monitoring detects data exfiltration, malware communication, and compromised systems beaconing to attacker infrastructure
- C. Outbound traffic is always safe
- D. Monitoring both directions is too expensive to be worthwhile

Answer: B

Q456. What is a false negative in security monitoring?

- A. A false alarm about a security threat
- B. A real security threat that the monitoring system fails to detect, allowing the attack to proceed without alerting the security team
- C. A negative review of security software
- D. A test alert used during security drills

Answer: B

Q457. What is the purpose of monitoring privileged account activity?

- A. Privileged accounts do not need monitoring
- B. Privileged accounts have elevated access that can cause significant damage if misused; monitoring their activity detects unauthorized use, insider threats, and compromised admin accounts
- C. Monitoring privileged accounts violates privacy
- D. Only guest accounts need monitoring

Answer: B

Q458. What is real-time monitoring versus periodic monitoring?

- A. They are the same approach with different names
- B. Real-time monitoring continuously analyzes events as they occur for immediate detection, while periodic monitoring reviews data at scheduled intervals
- C. Real-time monitoring is less effective
- D. Periodic monitoring catches more threats

Answer: B

Q459. What is the role of antivirus alerts in security monitoring?

- A. Antivirus alerts are always false positives
- B. Antivirus alerts notify security teams when malware is detected on endpoints, enabling investigation and response to potential infections
- C. Antivirus alerts only matter for personal computers
- D. Antivirus alerts replace the need for other monitoring

Answer: B

Q460. What is unauthorized access to a computer system?

- A. Using your own computer at home
- B. Accessing a computer system or network without permission from the owner, which is illegal in most jurisdictions
- C. Installing software with admin approval
- D. Browsing a public website

Answer: B

Q461. What is the purpose of a terms of service agreement for online services?

- A. To make registration more complicated
- B. To define the rules, rights, and responsibilities for using an online service, forming a legal agreement between the user and the provider
- C. To collect user email addresses
- D. To slow down the sign-up process

Answer: B

Q462. What is cyber harassment?

- A. Asking someone a question online
- B. Using digital communications to threaten, intimidate, or repeatedly harass another person, which is a punishable offense under cyber laws
- C. Sending a friend request on social media
- D. Posting a restaurant review online

Answer: B

Q463. What is the purpose of a privacy notice on a website?

- A. To decorate the website footer
- B. To inform users about what personal data is collected, how it is used, stored, and shared, and what rights users have regarding their data
- C. To increase website traffic
- D. To prevent users from leaving the website

Answer: B

Q464. What is digital copyright?

- A. The right to copy any digital content freely
- B. Legal protection that gives creators exclusive rights to their original digital works such as software, music, videos, and written content
- C. A license to hack into copyrighted systems
- D. A type of Creative Commons license only

Answer: B

Q465. Why is it important to obtain consent before collecting personal data?

- A. Consent is only a recommendation, not a requirement
- B. Data protection laws require organizations to obtain informed consent before collecting and processing personal data to respect individual privacy rights
- C. Consent is only needed for government organizations
- D. Collecting data without consent is always allowed

Answer: B

Q466. What is an acceptable use policy (AUP)?

- A. A marketing document for software products
- B. A set of rules that define how an organization's IT resources may be used, covering permitted activities, prohibited behaviors, and consequences for violations
- C. A software installation guide
- D. A network speed guarantee

Answer: B

Q467. What is the legal consequence of sending spam emails in many jurisdictions?

- A. There are no legal consequences for spam
- B. Sending unsolicited commercial emails can result in fines and penalties under anti-spam legislation such as CAN-SPAM Act, GDPR, or PECA 2016
- C. Spam is only illegal if it contains malware
- D. Only governments can send bulk emails

Answer: B

Q468. What does it mean to report a cybercrime?

- A. Posting about the crime on social media
- B. Formally notifying law enforcement authorities (like FIA in Pakistan) about a cyber crime so they can investigate and take legal action
- C. Telling your friends about the incident
- D. Writing a blog post about the experience

Answer: B

Q469. What is a smart device in the context of cybersecurity?

- A. Any device that uses electricity
- B. A device connected to the internet that can collect, send, and receive data, such as smart speakers, thermostats, and cameras
- C. A device that is physically large
- D. A computer without internet access

Answer: B

Q470. What is the concept of zero trust in simple terms?

- A. Trusting no technology at all and using only paper
- B. A security model based on the principle of 'never trust, always verify' where every access request is authenticated and authorized regardless of location
- C. Trusting all internal users but not external ones
- D. A model where security is not needed

Answer: B

Q471. What is a supply chain attack in cybersecurity?

- A. An attack on physical delivery trucks
- B. An attack that targets a less-secure element in the supply chain (like a software vendor or service provider) to gain access to a larger target organization
- C. An attack that only affects manufacturers
- D. Stealing products from a warehouse

Answer: B

Q472. What is automation in cybersecurity?

- A. Replacing all security staff with robots
- B. Using technology to automatically perform security tasks like scanning for vulnerabilities, responding to alerts, and updating security rules without manual intervention
- C. Making all passwords automatic
- D. Automating the creation of cyber attacks

Answer: B

Q473. What is a digital twin in technology?

- A. A twin sibling who works in technology
- B. A virtual replica of a physical system, process, or device that can be used for simulation, testing, and monitoring
- C. A backup copy of a digital file
- D. A second computer monitor

Answer: B

Q474. What is the role of machine learning in detecting cyber threats?

- A. Machine learning creates cyber threats
- B. Machine learning can analyze large volumes of data to identify patterns and anomalies that may indicate cyber threats, helping detect attacks that traditional rule-based systems might miss
- C. Machine learning replaces all cybersecurity professionals
- D. Machine learning only works for email filtering

Answer: B

Q475. What is the dark web?

- A. The internet when your screen brightness is low
- B. A part of the internet that requires special software (like Tor) to access and is not indexed by standard search engines, often associated with illegal activities
- C. Any website with a dark color scheme
- D. A private corporate network

Answer: B

Q476. What is multi-cloud strategy?

- A. Using many computers in one office
- B. Using services from multiple cloud providers to avoid vendor lock-in, improve resilience, and leverage different providers' strengths
- C. Running the same application multiple times
- D. Having multiple internet connections

Answer: B

Q477. What is an intrusion prevention system (IPS) capable of doing beyond detection?

- A. Only generating reports
- B. Automatically blocking or preventing detected malicious traffic in real-time before it reaches the target
- C. Encrypting all network traffic
- D. Managing user accounts

Answer: B

Q478. What is online fraud in the context of cyber law?

- A. Selling products online at a discount
- B. Using deceptive practices on the internet to illegally obtain money, data, or other benefits from victims
- C. Creating a website for a legitimate business
- D. Posting product reviews online

Answer: B

Q479. What is wearable technology and what security concerns does it raise?

- A. Clothing made from special fabrics
- B. Internet-connected devices worn on the body like smartwatches and fitness trackers that collect personal health and location data, raising privacy and data security concerns
- C. Any technology that is portable
- D. Protective equipment for construction workers

Answer: B

Q480. What is the significance of 5G technology for cybersecurity?

- A. 5G only improves download speeds with no security impact
- B. 5G introduces new attack surfaces due to network slicing, massive IoT connectivity, and edge computing, but also provides improved security features over 4G
- C. 5G eliminates all cybersecurity concerns
- D. 5G technology is identical to WiFi

Answer: B

Cyber Security — Medium

480 questions

Q481. Which of the following best describes 'defense in depth'?

- A. Using only encryption as the single security method
- B. Layering multiple security controls throughout a system
- C. Depending solely on firewalls for perimeter protection
- D. Relying on only one strong security control measure

Answer: B

Q482. What is the difference between a threat and a risk in cyber security?

- A. A risk is a type of malware designed for compromising system integrity
- B. They are the same thing with no meaningful distinction between them
- C. A threat only applies to hardware targeting enterprise network infrastructure
- D. A threat is a potential danger while risk is the likelihood and impact of that threat

Answer: D

Q483. Which of the following is an example of a physical security control?

- A. Intrusion detection system
- B. Data-at-rest symmetric encryption
- C. Biometric access to server room
- D. Network firewall rule configuration

Answer: C

Q484. What is the purpose of a security policy in an organization?

- A. To define rules and guidelines for protecting information assets
- B. To minimize infrastructure and hardware expenditures
- C. To expand network bandwidth allocation and throughput
- D. To design and build user-facing application interfaces

Answer: A

Q485. Which cyber security principle states that users should only have the minimum access necessary?

- A. Principle of least privilege
- B. Defense in depth
- C. Security by obscurity
- D. Separation of duties

Answer: A

Q486. What is social engineering in cyber security?

- A. Manipulating people into revealing confidential information
- B. Engineering software for social platforms
- C. Building social media applications
- D. Designing social networks within enterprise security environments

Answer: A

Q487. Which of the following is a component of the NIST Cybersecurity Framework?

- A. Identify, Protect, Detect, Respond, Recover
- B. Design, Build, Test, Deploy lifecycle
- C. Input, Process, Output data flow model
- D. Plan, Do, Check, Act continuous cycle

Answer: A

Q488. What is the purpose of a risk assessment in cyber security?

- A. To design new websites for enterprise computing environments
- B. To identify, analyze, and evaluate security risks
- C. To install antivirus software and management
- D. To write application code and management

Answer: B

Q489. What does 'non-repudiation' mean in information security?

- A. Ensuring a party cannot deny having performed an action
- B. Denying access to all users regardless of the specific situation or context
- C. Blocking network traffic within enterprise security environments
- D. Encrypting all messages regardless of the specific situation or context

Answer: A

Q490. Which of the following best describes an 'attack surface'?

- A. The total sum of vulnerabilities accessible to an attacker
- B. A strain of self-replicating malicious software programs
- C. A security certification for mobile network professionals
- D. The physical area of a computer across computing environments

Answer: A

Q491. What is the difference between symmetric and asymmetric encryption?

- A. Symmetric uses two keys, asymmetric uses one across computing environments
- B. There is no difference in any deployment scenario or context
- C. Symmetric is always stronger regardless of the specific situation or context
- D. Symmetric uses one key for both encryption and decryption, asymmetric uses a key pair

Answer: D

Q492. What is a hash function used for in security?

- A. Applying symmetric ciphers for secure data transfer
- B. Generating a fixed-size output from input data for integrity verification
- C. Generating redundant backup copies of system data
- D. Compressing files to reduce disk storage consumption

Answer: B

Q493. Which is an example of a detective security control?

- A. Intrusion Detection System (IDS)
- B. Firewall within enterprise security environments
- C. Data-at-rest symmetric encryption
- D. Access control list in security contexts

Answer: A

Q494. What is the purpose of an access control list (ACL)?

- A. To define which users or systems are granted access to specific resources
- B. To encrypt data using symmetric cipher algorithms
- C. To compress files for reducing disk storage consumption
- D. To generate comprehensive network topology diagrams

Answer: A

Q495. What is the difference between identification and authentication?

- A. Neither is used in security within enterprise security environments
- B. They are fundamentally the same thing in security
- C. Identification claims an identity, authentication proves it
- D. Authentication comes before identification

Answer: C

Q496. What is a security token in authentication?

- A. A backup device for managing enterprise data resources
- B. A physical or digital device that generates one-time passwords for authentication
- C. A network firewall rule for filtering inbound traffic
- D. A software application used for enterprise computing operations

Answer: B

Q497. What is the purpose of security auditing?

- A. To deploy and install application software across enterprise computing systems
- B. To architect and design new information systems and management
- C. To increase available bandwidth for improved download speeds
- D. To systematically evaluate the security of a system by testing against established criteria

Answer: D

Q498. Which security principle involves dividing critical tasks among multiple people?

- A. Need to know
- B. Least privilege
- C. Defense in depth
- D. Separation of duties

Answer: D

Q499. What is a security baseline?

- A. A next-generation firewall security appliance
- B. A backup schedule for creating regular recovery snapshots
- C. The fastest network speed within enterprise security environments
- D. A minimum set of security controls required for a system

Answer: D

Q500. What type of attack targets the availability component of the CIA triad?

- A. Network eavesdropping attack
- B. Denial of Service (DoS)
- C. Spear phishing email attack
- D. Credential data theft attack

Answer: B

Q501. What is the key size of AES-256?

- A. 256 bits
- B. 192 bits
- C. 512 bits
- D. 128 bits

Answer: A

Q502. What is a block cipher?

- A. A cipher that encrypts one bit at a time
- B. A cryptographic algorithm used for protecting sensitive data
- C. A cipher that encrypts fixed-size blocks of data at a time
- D. A cipher that uses no key applied to data protection workflows

Answer: C

Q503. What is a stream cipher?

- A. A cipher that encrypts entire files at once
- B. A cipher that encrypts data one bit or byte at a time
- C. A cipher that uses block processing
- D. A cipher that only works with video

Answer: B

Q504. What is the Diffie-Hellman key exchange used for?

- A. Hashing passwords within enterprise security environments
- B. Compressing data files for reduced storage space consumption
- C. Encrypting data directly within enterprise security environments
- D. Securely exchanging cryptographic keys over an insecure channel

Answer: D

Q505. Which hash algorithm produces a 256-bit output?

- A. SHA-256
- B. SHA-512
- C. MD5
- D. SHA-1

Answer: A

Q506. What is a digital certificate?

- A. A cryptographic algorithm used for protecting sensitive data
- B. An electronic document that uses a digital signature to bind a public key with an identity
- C. A software application used for enterprise computing operations
- D. A physical security measure for protecting building infrastructure

Answer: B

Q507. What is the role of a Certificate Authority (CA)?

- A. To design websites for enterprise computing environments
- B. To develop software across enterprise computing systems
- C. To design and distribute malicious software programs
- D. To issue, manage, and revoke digital certificates

Answer: D

Q508. Why is MD5 considered insecure for cryptographic purposes?

- A. It requires too much memory within enterprise security environments
- B. It is too slow within enterprise security environments
- C. It is vulnerable to collision attacks where two different inputs produce the same hash
- D. It produces too long an output within enterprise security environments

Answer: C

Q509. What is the purpose of a nonce in cryptography?

- A. A number used only once to prevent replay attacks
- B. To compress files for reducing disk storage consumption
- C. To create automated backup copies of user data
- D. To encrypt data using symmetric cipher algorithms

Answer: A

Q510. What is PKI (Public Key Infrastructure)?

- A. A next-generation firewall for filtering network traffic flows
- B. A framework for managing digital certificates and public-key encryption
- C. A category of self-replicating malicious software
- D. A standard network communication protocol for data transfer

Answer: B

Q511. What is the difference between an IDS and an IPS?

- A. IDS only detects and alerts while IPS can also block malicious traffic
- B. They are the same thing with no meaningful distinction between them
- C. IPS is slower than IDS within enterprise security environments
- D. IDS blocks traffic while IPS only detects

Answer: A

Q512. What is network segmentation?

- A. Deleting network components within enterprise security environments
- B. A type of network attack used in enterprise network infrastructure
- C. Dividing a network into smaller segments to improve security and performance
- D. Combining all networks into one regardless of the specific situation or context

Answer: C

Q513. What is ARP spoofing?

- A. An attack where fake ARP messages link an attacker's MAC address with a legitimate IP
- B. A symmetric key data encryption algorithm applied to data protection workflows
- C. A firewall configuration deployed across enterprise environments
- D. A legitimate network technique used in enterprise network infrastructure

Answer: A

Q514. What is the purpose of network access control (NAC)?

- A. To speed up the network across the enterprise network infrastructure
- B. To restrict access to a network based on device compliance and user identity
- C. To install network cables across the enterprise network infrastructure
- D. To generate comprehensive network topology diagrams

Answer: B

Q515. What is a honeypot in network security?

- A. A decoy system designed to attract attackers and study their methods
- B. A network switch used in enterprise network infrastructure
- C. A strain of self-replicating malicious software programs
- D. A type of firewall for filtering wireless network traffic

Answer: A

Q516. What is SSL/TLS used for?

- A. Optimal network routing and efficient packet forwarding across systems
- B. Persistent data storage and streamlined data retrieval management
- C. Providing encrypted communication between a client and server over a network
- D. Lossless data file compression techniques for long-term storage optimization

Answer: C

Q517. What is a VLAN and how does it improve security?

- A. A Virtual LAN that logically segments a network, isolating traffic between groups
- B. A virtual private network encrypted tunnel connection
- C. A next-generation firewall for filtering network traffic flows
- D. A software application used for enterprise computing operations

Answer: A

Q518. What is port scanning?

- A. Sending packets to specific ports to discover which services are running
- B. A cryptographic algorithm used for protecting sensitive data
- C. A physical inspection of ports within modern computing environments
- D. A backup method for managing enterprise data resources

Answer: A

Q519. What is a stateful firewall?

- A. A firewall without rules deployed across enterprise environments
- B. A firewall that tracks the state of active connections and makes decisions based on context
- C. A software-only firewall deployed across enterprise environments
- D. A firewall that only checks packet headers deployed across enterprise environments

Answer: B

Q520. What is DNS security and why is it important?

- A. Continuous monitoring is not important for security for modern enterprise security environments
- B. It is only about encrypting DNS without any additional considerations needed
- C. Protecting DNS infrastructure from attacks like DNS spoofing to ensure correct domain name resolution
- D. DNS security only affects speed without any additional considerations needed

Answer: C

Q521. What is mandatory access control (MAC) in operating systems?

- A. The OS enforces access rules based on security labels that users cannot override
- B. Users set their own permissions within enterprise security environments
- C. All users have equal access with no meaningful distinction between them
- D. No access control is needed in any deployment scenario or context

Answer: A

Q522. What is Address Space Layout Randomization (ASLR)?

- A. A standard network communication protocol for data transmission
- B. A file management technique for organizing directories used in enterprise computing environments
- C. A symmetric key data encryption algorithm applied to data protection workflows
- D. A security technique that randomly arranges memory address space of processes to prevent exploitation

Answer: D

Q523. What is the purpose of SELinux?

- A. To provide mandatory access control security policies in Linux
- B. To install and manage software packages efficiently
- C. To optimize and speed up the Linux kernel performance
- D. To create and manage virtual machine environments

Answer: A

Q524. What is a chroot jail in Linux security?

- A. An operation that changes the root directory for a process, isolating it from the rest of the file system
- B. A physical jail or correctional facility structure within modern computing environments
- C. A network configuration for firewall rule settings deployed across enterprise environments
- D. A specific type of user account with limited access within the identity management system

Answer: A

Q525. What is the purpose of Windows Group Policy?

- A. To centrally manage and configure OS settings, user permissions, and security policies
- B. To create social media groups for enterprise communication
- C. To logically group related files within directories and management
- D. To organize and arrange desktop icons into folders and management

Answer: A

Q526. What is Data Execution Prevention (DEP)?

- A. An automated feature for scheduling data backups for managing enterprise data resources
- B. A security feature that marks memory regions as non-executable to prevent code execution from those areas
- C. A built-in file encryption tool for data protection applied to data protection workflows
- D. A network traffic monitoring feature for analysis used in enterprise network infrastructure

Answer: B

Q527. What is the Windows Registry and why is it important for security?

- A. An automated backup solution for enterprise data recovery needs
- B. A user account type with specifically assigned access permissions
- C. A hierarchical database storing OS and application settings that if compromised can affect system security
- D. A next-generation firewall security appliance deployed across enterprise environments

Answer: C

Q528. What is process isolation in operating systems?

- A. Keeping each process's memory space separate to prevent unauthorized access between processes
- B. A specific type of preemptive multitasking method used in enterprise computing environments
- C. Running all system processes in shared memory space across computing environments
- D. Immediately stopping and terminating all processes across computing environments

Answer: A

Q529. What is the purpose of audit logging in an OS?

- A. To record security-relevant events for monitoring, investigation, and compliance
- B. To identify and remove obsolete system files and management
- C. To add computational overhead slowing the system and management
- D. To reduce data sizes through compression methods and management

Answer: A

Q530. What is sandboxing in OS security?

- A. Playing in a physical sandbox recreation area across computing environments
- B. A type of data encryption applied to data protection workflows
- C. Running applications in an isolated environment to limit their access to system resources
- D. A next-generation firewall for filtering network traffic flows

Answer: C

Q531. What is Cross-Site Scripting (XSS)?

- A. A responsive web design technique used for structuring and organizing page layouts
- B. A front-end JavaScript framework commonly used for building interactive web applications
- C. A vulnerability where malicious scripts are injected into trusted websites and executed in browsers
- D. A written record or logbook used by teams for tracking daily operational activities

Answer: C

Q532. What is SQL injection?

- A. An attack that inserts malicious SQL code through user input to manipulate the database
- B. A database management system for storing structured records
- C. A software application used for enterprise computing operations
- D. An optimization technique for improving SQL query speed used in enterprise computing environments

Answer: A

Q533. What is Cross-Site Request Forgery (CSRF)?

- A. An attack that tricks a browser into making unintended requests to a site where the user is authenticated
- B. A form validation technique for checking user inputs within the identity management system
- C. A standard type of user login authentication system within the identity management system
- D. A type of HTTP cookie used for session tracking within modern computing environments

Answer: A

Q534. What is the Same-Origin Policy?

- A. A web browser application for accessing internet-hosted content
- B. A security mechanism restricting scripts from one origin from accessing resources of another origin
- C. An intellectual property copyright protection law within modern computing environments
- D. A common website design pattern for page layout within modern computing environments

Answer: B

Q535. What is Content Security Policy (CSP)?

- A. An HTTP header specifying which content sources are allowed to be loaded by a web page
- B. A cascading style sheets framework for responsive page layouts
- C. A front-end JavaScript library for dynamic web development
- D. A content management system for organizing digital resources

Answer: A

Q536. What is the purpose of HTTP security headers?

- A. To format and style web content for visual presentation and management
- B. To compress and optimize images for web page loading and management
- C. To instruct browsers to enable security features protecting against common web attacks
- D. To make websites load faster with better performance and management

Answer: C

Q537. What is session hijacking?

- A. Properly logging out of an active user session across computing environments
- B. Creating a new user session for web interaction across computing environments
- C. An attack where an attacker takes over a valid user session by stealing the session ID
- D. An automated session timeout after inactivity period used in enterprise computing environments

Answer: C

Q538. What is the OWASP Top 10?

- A. A ranked list of the most widely used web browsers
- B. A programming framework for building web app front-ends
- C. A curated list of the top ten most popular websites
- D. A regularly updated list of the most critical web application security risks

Answer: D

Q539. What is clickjacking?

- A. A software application used for enterprise computing operations
- B. A gaming technique for improving response speed used in enterprise computing environments
- C. An attack tricking users into clicking hidden elements by layering invisible frames
- D. A user account type with specifically assigned access permissions

Answer: C

Q540. What is the purpose of parameterized queries?

- A. To separate SQL code from user data, ensuring input is treated as data not code
- B. To compress and reduce the size of database stored data
- C. To improve and optimize query execution performance speed
- D. To format and style the output of database query results

Answer: A

Q541. What is a keylogger?

- A. A tool for managing and organizing cryptographic keys
- B. Software or hardware that records keystrokes to capture sensitive information
- C. A software application used for enterprise computing operations
- D. A cryptographic algorithm used for protecting sensitive data

Answer: B

Q542. What is a botnet?

- A. A helpful automated robot network for task assistance used in enterprise computing environments
- B. A chat bot service providing automated customer support used in enterprise computing environments
- C. A network monitoring tool for tracking bandwidth utilization
- D. A network of compromised computers controlled remotely for coordinated malicious activities

Answer: D

Q543. How does a DDoS attack differ from DoS?

- A. DDoS attacks are generally considered less dangerous than standard single-source DoS attacks
- B. DDoS uses multiple compromised systems to flood the target, making it harder to mitigate
- C. They are functionally identical attack frameworks with no meaningful differences between them
- D. DoS attacks typically use significantly more compromised systems than DDoS attacks

Answer: B

Q544. What is a rootkit?

- A. A disk utility for managing storage partitions used in enterprise computing environments
- B. Malware designed to gain unauthorized root-level access and hide its presence
- C. A legitimate system administration tool for configuration
- D. A standard operating system boot loader for initialization

Answer: B

Q545. What is social engineering?

- A. A software application suite used for enterprise computing operations and communications
- B. Building and developing social media networking platforms for public user interaction
- C. Psychological manipulation to trick people into making security mistakes or revealing information
- D. Engineering and designing interactive social interaction systems for enterprise organizations

Answer: C

Q546. What is a drive-by download attack?

- A. A software application used for enterprise computing operations
- B. A standard file transfer method between connected systems used in enterprise computing environments
- C. Downloading files while driving in a moving vehicle
- D. Automatic download of malware when visiting a compromised website without user interaction

Answer: D

Q547. What is a watering hole attack?

- A. A water cooling attack damaging hardware cooling systems
- B. Attacking municipal water supply infrastructure systems
- C. Compromising websites frequently visited by a targeted group to infect their computers
- D. A network flooding attack using volumetric traffic used in enterprise computing environments

Answer: C

Q548. What is fileless malware?

- A. A clean program verified by security code audit analysis
- B. Malware that operates without any source code at all
- C. Malware that operates entirely in memory without writing files to disk
- D. A type of antivirus program for real-time scanning

Answer: C

Q549. What is a man-in-the-middle (MITM) attack?

- A. A professional mediator in business negotiations within modern computing environments
- B. A software application used for enterprise computing operations
- C. An attack where the attacker secretly intercepts and possibly alters communication between two parties
- D. A proxy server that forwards requests between systems used in enterprise computing environments

Answer: C

Q550. What are indicators of compromise (IOCs)?

- A. Performance metrics for measuring system throughput
- B. Signs indicating system improvement and optimization
- C. Software licensing keys for verifying installations
- D. Observable artifacts or evidence suggesting a system has been breached

Answer: D

Q551. What is the difference between authentication and authorization?

- A. Neither process is used in modern security frameworks
- B. Authentication verifies identity, authorization determines permissions
- C. Authorization always comes before the authentication phase
- D. They are functionally identical frameworks with no differences

Answer: B

Q552. What is LDAP used for in access control?

- A. A file format used for structured data storage for managing enterprise data resources
- B. A high-level language used for development within modern computing environments
- C. A software application used for enterprise computing operations
- D. A protocol for accessing distributed directory information services for centralized authentication

Answer: D

Q553. What is attribute-based access control (ABAC)?

- A. Access control decisions based solely on user age and demographics
- B. Access control evaluating attributes of users, resources, actions, and environment
- C. Randomized access control without any defined criteria or policies
- D. Access control based on file system attributes exclusively without context

Answer: B

Q554. What is a federated identity?

- A. A physical security measure for protecting building infrastructure
- B. A fake identity created for fraudulent purposes used in enterprise computing environments
- C. A temporary identity assigned for a single session used in enterprise computing environments
- D. A system allowing users to use the same identity credentials across multiple organizations

Answer: D

Q555. What is the purpose of a RADIUS server?

- A. To provide centralized authentication, authorization, and accounting for network access
- B. To measure and calculate distances between network nodes and management
- C. To increase and optimize overall network throughput speed
- D. To create and deploy enterprise wireless network infrastructure

Answer: A

Q556. What is privilege escalation?

- A. Installing new software applications on the system
- B. Gaining elevated access rights beyond what was initially authorized
- C. Receiving a job promotion or advancement opportunity
- D. Creating a new user account with default permissions

Answer: B

Q557. What is the purpose of access control matrices?

- A. To define access rights of subjects to objects in a structured format
- B. To encrypt data using symmetric cipher algorithms
- C. To compress files for reducing disk storage consumption
- D. To create and format data management spreadsheets

Answer: A

Q558. What is SAML (Security Assertion Markup Language)?

- A. A high-level language used for development within modern computing environments
- B. A software application used for enterprise computing operations
- C. An XML-based standard for exchanging authentication and authorization data between providers
- D. A database management system for storing structured records

Answer: C

Q559. What is the purpose of session management in authentication?

- A. To schedule tasks within enterprise computing infrastructure
- B. To securely maintain a user's authenticated state throughout their interaction
- C. To manage network sessions across the enterprise network infrastructure
- D. To manage computer memory within enterprise computing infrastructure

Answer: B

Q560. What is adaptive authentication?

- A. A biometric method within enterprise computing environments
- B. An approach adjusting security requirements based on risk factors like location and behavior
- C. A credential management mechanism for user authentication flows
- D. Fixed authentication rules within enterprise security environments

Answer: B

Q561. What is the OWASP Secure Coding Practices guide?

- A. A high-level language used for development within modern computing environments
- B. A database design guide for managing enterprise data resources
- C. A comprehensive checklist of secure coding practices organized by vulnerability category
- D. A software application used for enterprise computing operations

Answer: C

Q562. What is threat modeling in software development?

- A. Testing application performance under heavy load conditions
- B. Creating threat actors for penetration testing exercises
- C. Running antivirus scans on development machine environments
- D. A structured approach to identifying and addressing security threats during design

Answer: D

Q563. What is static application security testing (SAST)?

- A. Testing network performance under heavy traffic loads
- B. Testing user interfaces for usability and accessibility
- C. Analyzing source code without executing it to find security vulnerabilities
- D. Testing the application while it is actively running

Answer: C

Q564. What is dynamic application security testing (DAST)?

- A. Testing database query performance under heavy loads
- B. Testing a running application by simulating attacks to find vulnerabilities
- C. Testing source code without running it across computing environments
- D. Testing code compilation for syntax errors and warnings

Answer: B

Q565. What is the principle of fail-safe defaults?

- A. Logging all defaults regardless of the specific situation or context
- B. Ignoring security failures within enterprise security environments
- C. Always granting access by default in security contexts
- D. Denying access by default and requiring explicit permission grants

Answer: D

Q566. What is dependency management in secure development?

- A. Managing and coordinating team members across projects
- B. Tracking, updating, and securing third-party libraries and components
- C. Managing and configuring enterprise server infrastructure
- D. Managing database servers and optimizing query performance

Answer: B

Q567. What is the purpose of security testing in the SDLC?

- A. To deliberately delay the project delivery timeline
- B. To create comprehensive technical documentation materials
- C. To identify and fix security vulnerabilities before deployment to production
- D. To test and validate hardware components for compatibility

Answer: C

Q568. What is 'shift left' in security?

- A. Integrating security practices earlier in the development lifecycle
- B. Moving project deadlines earlier in the release timeline
- C. Moving physical servers to the left of the data center
- D. Changing code alignment and formatting in source files

Answer: A

Q569. What is secure session management?

- A. Properly creating, maintaining, and destroying sessions to prevent hijacking and fixation
- B. Sharing sessions between multiple users simultaneously
- C. Deleting all cookies immediately after each page request
- D. Never creating sessions for any user interactions at all regardless of the deployment context or scenario

Answer: A

Q570. What is the purpose of output encoding?

- A. Making output look visually pretty and well-formatted
- B. Compressing output data for reduced bandwidth consumption
- C. Encrypting output using symmetric cipher key algorithms
- D. Converting output to a safe format to prevent injection attacks like XSS

Answer: D

Q571. What is WPA3?

- A. A web application framework for building online services
- B. A mobile application for managing device Wi-Fi settings
- C. The latest Wi-Fi security protocol with stronger encryption and forward secrecy
- D. A Bluetooth version for short-range wireless communication

Answer: C

Q572. What is an evil twin attack?

- A. A physical hardware component used in computing infrastructure
- B. A mirrored display for screen sharing purposes used in enterprise computing environments
- C. A twin SIM card used for dual number support used in enterprise computing environments
- D. A rogue access point mimicking a legitimate Wi-Fi network to intercept traffic

Answer: D

Q573. What is Mobile Device Management (MDM)?

- A. A physical hardware component used within enterprise computing infrastructure systems
- B. A standard software application used for routine enterprise computing operations
- C. A mobile device application designed for portable computing operations and tasks
- D. Software allowing organizations to remotely manage, secure, and enforce policies on mobile devices

Answer: D

Q574. What are the risks of jailbreaking or rooting a mobile device?

- A. It improves battery life and extends usage time
- B. It removes built-in security protections and exposes the device to malware
- C. It makes the phone processor run significantly faster
- D. There are no security risks at all from this practice

Answer: B

Q575. What is a rogue access point?

- A. An officially authorized network infrastructure device
- B. A Bluetooth device used for short-range data exchange
- C. An unauthorized wireless access point installed without administrator knowledge
- D. A mobile device application for portable computing operations

Answer: C

Q576. What is WEP and why is it insecure?

- A. An outdated wireless encryption protocol with known vulnerabilities crackable in minutes
- B. A physical hardware component used in computing infrastructure
- C. A cryptographic algorithm used for protecting sensitive data
- D. A next-generation firewall for filtering network traffic flows

Answer: A

Q577. What is app sandboxing on mobile devices?

- A. Playing games in a sandbox within enterprise security environments
- B. A testing method for validating application functionality
- C. Isolating each app in its own environment to prevent unauthorized data access
- D. A game development technique for creating mobile games

Answer: C

Q578. What is SIM swapping?

- A. Upgrading SIM cards to a newer generation of mobile network technology standard
- B. Sharing SIM cards between multiple mobile phone devices for network access
- C. An attack convincing a carrier to transfer a victim's number to a new SIM for account takeover
- D. Changing phone cases for a different physical appearance and improved protection

Answer: C

Q579. What is the purpose of a VPN on a mobile device?

- A. To encrypt internet traffic when using public or untrusted networks
- B. To make phone calls cheaper through VoIP routing
- C. To speed up mobile data transfer download bandwidth
- D. To improve camera quality and photo capture resolution

Answer: A

Q580. What is BYOD and its security implications?

- A. A software application used for enterprise computing operations
- B. A policy allowing personal devices for work, creating data separation and security challenges
- C. A physical hardware component used in computing infrastructure
- D. A music streaming service platform for audio content used in enterprise computing environments

Answer: B

Q581. What is the difference between IaaS, PaaS, and SaaS security responsibilities?

- A. Customer always manages everything in every cloud model
- B. Security is not needed in any cloud computing environment for modern enterprise security environments
- C. In IaaS customer manages most; in PaaS provider manages more; in SaaS provider manages most
- D. Provider manages everything in all cloud service models

Answer: C

Q582. What is cloud workload protection?

- A. A backup solution for creating redundant data recovery copies
- B. A load balancing technique for distributing network traffic used in enterprise computing environments
- C. Security solutions protecting workloads including VMs, containers, and serverless functions
- D. Physical security of data center server hardware infrastructure

Answer: C

Q583. What is DLP in cloud environments?

- A. Encrypting all data using symmetric cipher algorithm keys
- B. Compressing data files for reduced storage space consumption
- C. Creating redundant copies of data for backup recovery
- D. Technologies preventing sensitive data from being inappropriately shared or exposed

Answer: D

Q584. What is cloud security posture management (CSPM)?

- A. A standard software application used for routine enterprise computing operations
- B. A cloud infrastructure provider offering a range of core computing service options
- C. An automated solution monitoring cloud infrastructure for misconfigurations and compliance violations
- D. A cloud storage manager designed for organizing and managing hosted data resources

Answer: C

Q585. What is server-side encryption in cloud storage?

- A. The cloud provider encrypts data at rest on the server before storing it
- B. Encrypting data on the client device before transmission
- C. A next-generation firewall for filtering network traffic flows
- D. A compression method for reducing cloud storage data sizes

Answer: A

Q586. What is the purpose of cloud audit logging?

- A. To record all activities and API calls for security monitoring and compliance
- B. To manage and administer user accounts and access rights
- C. To add computational overhead slowing the system and management
- D. To increase storage usage for expanded data capacity

Answer: A

Q587. What is a virtual private cloud (VPC)?

- A. A logically isolated section of public cloud where resources launch in a customer-defined virtual network
- B. A physical hardware component used in computing infrastructure
- C. A virtual private network encrypted tunnel connection used in enterprise computing environments
- D. A physical security measure for protecting building infrastructure

Answer: A

Q588. What is container security in cloud environments?

- A. An automated backup solution designed for enterprise data recovery needs and processes
- B. Protecting containerized apps through image scanning, runtime protection, and secure orchestration
- C. Securing physical shipping containers during transportation across global supply chains
- D. A physical security measure designed for protecting building infrastructure and premises

Answer: B

Q589. What are cloud security groups?

- A. User groups for managing organizational access permissions
- B. Social media groups for enterprise team communication
- C. A type of cloud storage for archiving large data volumes
- D. Virtual firewalls controlling inbound and outbound traffic to cloud resources

Answer: D

Q590. What is least privilege in cloud IAM?

- A. Using shared credentials for all users in the environment
- B. Granting only the minimum permissions necessary for specific tasks
- C. Completely blocking all access to cloud resources
- D. Giving everyone administrator access to all resources

Answer: B

Q591. What is volatile evidence in digital forensics?

- A. Evidence that is explosive within enterprise security environments
- B. Evidence that never changes in any deployment scenario or context
- C. Physical evidence within enterprise security environments
- D. Data in temporary storage like RAM that is lost when the system is powered off

Answer: D

Q592. What is the order of volatility in evidence collection?

- A. Collect backup data first within enterprise security environments
- B. The order does not matter in any deployment scenario or context
- C. Collect most volatile data first: registers, cache, RAM, disk, then remote logs
- D. Collect hard drive data first within enterprise security environments

Answer: C

Q593. What is steganography analysis in digital forensics?

- A. Detecting and extracting hidden data concealed within files like images or audio
- B. Network analysis within enterprise security environments
- C. Analyzing dinosaur fossils within enterprise security environments
- D. Photo editing within enterprise security environments

Answer: A

Q594. What is network forensics?

- A. Designing networks within enterprise security environments
- B. Capture, recording, and analysis of network traffic to investigate security incidents
- C. Installing network cables within enterprise security environments
- D. Configuring routers within enterprise security environments

Answer: B

Q595. What is a forensic toolkit?

- A. A physical toolbox used in enterprise computing environments
- B. A development environment within modern computing environments
- C. A collection of software tools for digital evidence collection, analysis, and reporting
- D. A repair kit within enterprise computing environments

Answer: C

Q596. What is timeline analysis in digital forensics?

- A. A project management technique within modern computing environments
- B. Creating project schedules within enterprise security environments
- C. Reconstructing a chronological sequence of events using timestamps from multiple sources
- D. A time management tool used in enterprise computing environments

Answer: C

Q597. What is email forensics?

- A. Writing and composing emails within enterprise security environments for communication
- B. Examining email headers, content, attachments, and metadata to trace origin and authenticity
- C. Reading and reviewing emails within enterprise security environments for monitoring
- D. Deleting and archiving emails within enterprise security environments for compliance

Answer: B

Q598. What is the role of a forensic examiner's report?

- A. A detailed document presenting methodology, findings, and conclusions for legal proceedings
- B. A blog post published online for broad public readership review and engagement
- C. A news article published within enterprise computing environments for staff awareness
- D. A personal diary used for recording daily reflections and activities for reference

Answer: A

Q599. What is memory forensics?

- A. Testing memory speed and latency within enterprise security environments for capacity
- B. Analysis of RAM contents to find evidence of malware, connections, and encryption keys
- C. A memory stress test within enterprise computing environments to evaluate stability
- D. Adding more RAM modules within enterprise security environments for performance

Answer: B

Q600. What is anti-forensics?

- A. A type of forensic analysis technique within enterprise computing environments
- B. A standard software application used for enterprise computing operations and tasks
- C. Techniques used to prevent or obstruct forensic investigation, such as data wiping and timestamp manipulation
- D. Good forensic practices applied within enterprise security environments for compliance

Answer: C

Q601. What are the six phases of the NIST incident response lifecycle?

- A. Design, Develop, Test, Release, Maintain, Retire across computing environments
- B. Preparation, Detection and Analysis, Containment, Eradication, Recovery, Post-Incident Activity
- C. None of these in any deployment scenario or context across computing environments
- D. Plan, Build, Test, Deploy, Monitor, Retire within enterprise security environments

Answer: B

Q602. What is triage in incident response?

- A. A network tool used in enterprise network infrastructure
- B. Prioritizing incidents based on severity, impact, and urgency to allocate resources
- C. A strain of self-replicating malicious software programs
- D. Medical treatment within enterprise security environments

Answer: B

Q603. What is the eradication phase?

- A. Deleting the incident report within enterprise security environments
- B. Removing the root cause including malware, compromised accounts, and exploited vulnerabilities
- C. Erasing all data regardless of the specific situation or context
- D. Shutting down all systems regardless of the specific situation or context

Answer: B

Q604. What is the recovery phase?

- A. Restoring affected systems to normal operation and monitoring for recurrence
- B. Recovering deleted emails within enterprise security environments
- C. A data backup process for managing enterprise data resources
- D. Finding lost items within enterprise security environments

Answer: A

Q605. What is a lessons learned review?

- A. A post-incident analysis of what happened, what worked, what failed, and how to improve
- B. A security audit deployed across enterprise environments used in enterprise computing environments
- C. A training course within enterprise computing environments
- D. A school exam within enterprise computing environments

Answer: A

Q606. What is the role of a CISO during a major incident?

- A. To write code within enterprise computing infrastructure and management
- B. To provide strategic leadership, coordinate response, and make critical containment decisions
- C. To fix computers within enterprise computing infrastructure and management
- D. To design websites for enterprise computing environments and management

Answer: B

Q607. What is a playbook in incident response?

- A. A game manual within modern computing environments
- B. A software manual used in enterprise computing environments
- C. A children's book within modern computing environments
- D. Predefined procedures for responding to specific types of security incidents

Answer: D

Q608. What is the purpose of incident classification?

- A. To categorize incidents by type and severity for appropriate response and resource allocation
- B. To file paperwork within the data management framework and management
- C. To rank team members within enterprise computing infrastructure and management
- D. To organize files within the data management framework and management

Answer: A

Q609. What is threat hunting?

- A. A recreational entertainment application for consumer devices
- B. Hunting animals within enterprise security environments
- C. Proactively searching through networks to detect threats that evaded existing controls
- D. Installing antivirus software on the development machine

Answer: C

Q610. What is the importance of communication during incident response?

- A. Communication should wait until after resolution
- B. Clear, timely communication with stakeholders is essential for effective coordination
- C. Only technical communication matters without any additional considerations needed
- D. Communication is not needed in any deployment scenario or context

Answer: B

Q611. What is a SIEM system?

- A. A physical hardware component used in computing infrastructure
- B. Security Information and Event Management: aggregates, correlates, and analyzes security data
- C. A next-generation firewall security appliance deployed across enterprise environments
- D. A physical cable used for network connections used in enterprise network infrastructure

Answer: B

Q612. What is the difference between signature-based and anomaly-based detection?

- A. Anomaly-based detection only works effectively in offline analysis and forensic modes
- B. Signature-based detection is always considered better than anomaly-based detection methods
- C. They are functionally identical detection frameworks with no meaningful differences at all
- D. Signature-based matches known patterns; anomaly-based detects deviations from normal behavior

Answer: D

Q613. What is a Security Operations Center (SOC)?

- A. A software company that develops security products deployed across enterprise environments
- B. A physical hardware component used in computing infrastructure
- C. A centralized facility where a team continuously monitors, detects, and responds to security incidents
- D. A cloud computing platform for hosting enterprise workloads

Answer: C

Q614. What is log correlation?

- A. Deleting old logs to free up storage space on servers
- B. Sorting logs alphabetically for easier navigation across computing environments
- C. Comparing log file sizes for storage capacity planning
- D. Analyzing and connecting related events across multiple sources to identify threat patterns

Answer: D

Q615. What is UEBA?

- A. Security analytics using machine learning to establish baselines and detect anomalous user behavior
- B. A social media management tool used for content scheduling tasks and engagement tracking
- C. A user interface design tool used for creating web application page layouts and mockups
- D. An email analytics tool used for tracking message delivery metrics and campaign performance

Answer: A

Q616. What is the purpose of vulnerability scanning?

- A. To automatically identify known security weaknesses in systems and configurations
- B. To compress files for reducing disk storage consumption
- C. To create new vulnerabilities in systems for testing and management
- D. To scan paper documents and convert them to digital format

Answer: A

Q617. What is threat detection in security monitoring?

- A. A recreational entertainment application for consumer devices
- B. Identifying potential security threats through analysis of security data, patterns, and behaviors
- C. A weather detection system for forecasting atmospheric events used in enterprise computing environments
- D. Creating new threats to test system resilience capabilities

Answer: B

Q618. What is the role of a SOC analyst?

- A. Writing application source code for building new software features and functionality
- B. Monitoring alerts, investigating potential incidents, performing triage, and escalating threats
- C. Managing database servers and optimizing query performance for application workloads
- D. Architecting and building responsive enterprise web applications and user interfaces

Answer: B

Q619. What is a honeypot in security monitoring?

- A. A VPN service for encrypting internet traffic connections used in enterprise computing environments
- B. A decoy system designed to attract attackers and provide intelligence about attack methods
- C. A next-generation firewall for filtering network traffic flows
- D. A jar of honey used for food and cooking purposes used in enterprise computing environments

Answer: B

Q620. What is the importance of log retention policies?

- A. Logs should be deleted immediately after they are created
- B. Logs should never be kept or retained for any period regardless of the deployment context or scenario
- C. Ensuring historical data availability for investigation, compliance, and trend analysis
- D. Retention is only about managing physical storage capacity

Answer: C

Q621. What are the main offenses covered under PECA 2016?

- A. Only hacking is covered as a cybercrime offense under PECA
- B. Unauthorized access, data damage, electronic fraud, cyber terrorism, spamming, and spoofing
- C. Only online shopping fraud is covered under PECA across computing environments
- D. Only spam emails are covered as offenses under the PECA act

Answer: B

Q622. What is FIA's role in cybercrime in Pakistan?

- A. FIA only handles immigration cases and border control
- B. The designated agency for investigating and prosecuting cybercrimes under PECA 2016
- C. FIA only handles customs and trade regulation enforcement
- D. FIA has no role in cybercrime investigation or prosecution

Answer: B

Q623. What is the GDPR?

- A. A Pakistani law governing domestic trade and commerce
- B. An EU regulation providing comprehensive data protection rights for EU citizens
- C. A US trade law governing international commerce and tariffs
- D. A UN resolution addressing global peacekeeping and diplomacy

Answer: B

Q624. What is computer forensics evidence admissibility?

- A. Digital evidence is never admissible in any court proceedings regardless of the deployment context or scenario
- B. Evidence must meet standards including authenticity, integrity, chain of custody, and relevance
- C. All digital evidence is automatically admissible without conditions
- D. Only email evidence is admissible while other digital evidence is excluded

Answer: B

Q625. What is the difference between ethical and illegal hacking?

- A. There is no difference between ethical and illegal hacking or notable impact in any deployment scenario
- B. Ethical hacking is authorized security testing with permission; illegal hacking is unauthorized access
- C. Illegal hacking is faster than ethical hacking in all scenarios
- D. Ethical hacking uses better tools than illegal hacking methods

Answer: B

Q626. What is cyber defamation?

- A. Publishing false or damaging statements about a person through digital media
- B. A cryptographic algorithm used for protecting sensitive data
- C. A standard network communication protocol for data transfer
- D. Building a good reputation through positive online interactions

Answer: A

Q627. What is the right to be forgotten?

- A. Forgetting your password and requesting an account reset
- B. A memory improvement technique for better data recall used in enterprise computing environments
- C. An individual's right to request deletion of their personal data under certain conditions
- D. A data backup policy for creating redundant recovery copies

Answer: C

Q628. What are the penalties for cybercrime under PECA 2016?

- A. No penalties exist under PECA for cybercrime offenses regardless of the deployment context or scenario
- B. Only verbal warnings are issued for cybercrime violations
- C. Penalties range from fines to imprisonment up to 14 years depending on offense severity
- D. Only monetary fines are imposed without imprisonment penalties

Answer: C

Q629. What is the concept of data sovereignty?

- A. The principle that data is subject to the laws of the country where it is stored or processed
- B. Data belongs to everyone equally and cannot be controlled by any single entity
- C. Data is always free and not subject to any jurisdictional laws or regulations at all
- D. Data has its own independent government and governance structure for management

Answer: A

Q630. What is responsible disclosure?

- A. Privately reporting vulnerabilities to the vendor and allowing time for patching before public disclosure
- B. Publishing vulnerabilities immediately for public disclosure
- C. Never reporting discovered vulnerabilities to anyone in any deployment scenario or context
- D. Selling discovered vulnerabilities to the highest bidder across computing environments

Answer: A

Q631. How is AI enhancing cyber security defenses?

- A. AI is not used in cybersecurity defenses in any capacity regardless of the deployment context or scenario
- B. AI is only used for creating malicious viruses and malware
- C. AI enhances threat detection through behavioral analytics, automates response, and identifies patterns
- D. AI completely replaces all human security staff permanently

Answer: C

Q632. What are the security challenges of edge computing?

- A. There are no security challenges with this approach in any deployment scenario or context
- B. Edge computing completely eliminates all security threats across computing environments
- C. Edge computing is inherently more secure than cloud computing across computing environments
- D. Distributed processing creates challenges including physical security, limited resources, and managing many nodes

Answer: D

Q633. What is quantum computing's potential impact on cybersecurity?

- A. Quantum computers could break current public-key encryption like RSA and ECC
- B. Quantum computing makes all existing encryption much stronger
- C. Quantum computing has no impact on cybersecurity at all
- D. Quantum computing only affects hardware and not software security

Answer: A

Q634. What is zero trust network access (ZTNA)?

- A. A VPN replacement designed for establishing encrypted network tunnel connections
- B. A model verifying identity and context before granting access to specific applications
- C. A type of firewall designed for filtering wireless network traffic and connections
- D. Trusting all users and devices without any verification or authentication needed

Answer: B

Q635. What is the security concern with smart home devices?

- A. Smart homes are completely secure without any vulnerabilities
- B. Vulnerable devices could be used for surveillance, network entry, or botnet recruitment
- C. Smart homes have no security issues or concerns whatsoever
- D. Only cost is a concern with smart home devices not security

Answer: B

Q636. What is DevSecOps as an emerging trend?

- A. Integration of security into DevOps pipelines, automating security testing throughout development
- B. A high-level language used for development within modern computing environments
- C. A database management system for storing structured records
- D. An operating system for managing system resources used in enterprise computing environments

Answer: A

Q637. What is Security Mesh Architecture?

- A. A chain-link barrier for physical boundary demarcation purposes
- B. A physical perimeter fence for building boundary security
- C. A mesh Wi-Fi network for extending wireless home coverage used in enterprise computing environments
- D. A distributed security approach deploying controls closer to assets for composable, scalable security

Answer: D

Q638. What are the security implications of autonomous vehicles?

- A. There are no implications for this technology whatsoever regardless of the deployment context or scenario
- B. They are completely safe and have no vulnerabilities regardless of the deployment context or scenario
- C. Risks including remote hijacking, sensor manipulation, data privacy, and V2X communication security
- D. Only traffic violations are a concern not cybersecurity regardless of the deployment context or scenario

Answer: C

Q639. What is privacy-enhancing technology (PET)?

- A. A social media privacy setting for managing profile visibility used in enterprise computing environments
- B. A physical hardware component used in computing infrastructure
- C. Technologies protecting privacy while enabling data use, like homomorphic encryption and differential privacy
- D. Pet care technology for monitoring animal health activity across computing environments

Answer: C

Q640. What is the convergence of IT and OT security?

- A. A standard software application used for enterprise computing operations and system management workflows
- B. Integration of information technology and operational technology security as industrial systems connect to IT networks
- C. A physical network cable used for connecting and linking enterprise infrastructure devices together
- D. Merging two companies through a corporate acquisition deal within enterprise security environments

Answer: B

Q641. How does a risk assessment differ from a vulnerability assessment in cyber security?

- A. Risk assessment only examines hardware while vulnerability assessment examines only software components
- B. Risk assessment is performed annually while vulnerability assessment is performed on a daily recurring basis
- C. Risk assessment evaluates likelihood and impact of threats while vulnerability assessment identifies weaknesses
- D. Risk assessment uses automated tools while vulnerability assessment relies solely on manual review processes

Answer: C

Q642. What is the primary difference between a threat and a vulnerability?

- A. Threats are potential dangers while vulnerabilities are exploitable weaknesses in systems
- B. Threats only affect hardware while vulnerabilities only affect software applications
- C. Threats are always external while vulnerabilities are exclusively internal system issues
- D. Threats require physical access while vulnerabilities require remote network connectivity

Answer: A

Q643. Which security model provides access based on sensitivity labels assigned to data and users?

- A. Discretionary Access Control model
- B. Attribute-Based Access Control model
- C. Role-Based Access Control model
- D. Mandatory Access Control model

Answer: D

Q644. What is the concept of 'least privilege' in cyber security?

- A. Granting users maximum access rights for operational efficiency and user convenience
- B. Restricting all users from accessing any systems outside of regular business hours
- C. Allowing temporary administrator rights to all employees during onboarding processes
- D. Providing users only the minimum access necessary to perform their specific job duties

Answer: D

Q645. What is the purpose of a security audit?

- A. Removing outdated hardware from the organization data center facilities
- B. Training employees on new software features and application updates
- C. Installing new security software on all systems throughout the organization
- D. Evaluating how well security policies and controls are being implemented

Answer: D

Q646. Which of the following best describes 'non-repudiation' in information security?

- A. Preventing unauthorized users from accessing any classified information assets
- B. Ensuring the sender cannot deny having sent a particular message or transaction
- C. Verifying the physical identity of users through biometric scanning technology
- D. Guaranteeing that data remains available during scheduled maintenance windows

Answer: B

Q647. What is a 'zero-day vulnerability'?

- A. A flaw known to attackers but unknown to the vendor with no patch
- B. A weakness that requires zero user interaction to fully remediate
- C. A vulnerability that has been patched within zero days of discovery
- D. A vulnerability that only exists in systems less than one day old

Answer: A

Q648. How does 'security through obscurity' differ from proper security measures?

- A. It involves frequent penetration testing together with regular comprehensive code reviews
- B. It uses the latest encryption algorithms combined with multi-layered network segmentation
- C. It enforces strict password policies alongside mandatory multi-factor authentication rules
- D. It relies on keeping system design secret rather than implementing robust security controls

Answer: D

Q649. What is the main objective of a business continuity plan in relation to cyber security?

- A. Recruiting additional security staff for the incident response team
- B. Maximizing profit margins during normal business operating conditions
- C. Ensuring critical operations continue during and after a security incident
- D. Replacing all legacy systems with modern cloud-based infrastructure

Answer: C

Q650. Which framework is commonly used for managing and reducing cyber security risk?

- A. Agile Development Framework
- B. ITIL Service Management Framework
- C. NIST Cybersecurity Framework
- D. PMBOK Project Management Framework

Answer: C

Q651. What is the key difference between identification, authentication, and authorization?

- A. Identification encrypts data, authentication decrypts data, and authorization transfers data
- B. They are three identical processes that serve the same purpose in every security system
- C. Identification grants access, authentication logs activity, and authorization claims identity
- D. Identification claims identity, authentication verifies it, and authorization grants permissions

Answer: D

Q652. Which security principle states that systems should fail in a secure manner?

- A. Separation of duties
- B. Open design
- C. Defense in depth
- D. Fail-secure design

Answer: D

Q653. What is the purpose of separation of duties in security?

- A. Ensuring no single person has complete control over a critical process or transaction
- B. Separating development and production servers into distinct physical data centers
- C. Assigning each employee exactly one security tool to manage and maintain daily
- D. Dividing the network into multiple segments for improved overall performance speed

Answer: A

Q654. How does a hash function contribute to data integrity?

- A. It creates redundant copies of data across multiple geographic locations
- B. It produces a fixed-size digest that changes if the original data is modified
- C. It encrypts data so that only authorized users can access and read the content
- D. It compresses data to reduce storage space requirements on disk drives

Answer: B

Q655. What is the concept of 'need to know' in information security?

- A. Access to information is limited to those who require it for their duties
- B. Users need to know all system administrator passwords for emergencies
- C. All employees should know everything about the organization security posture
- D. Security teams must share all threat intelligence with every staff member

Answer: A

Q656. Which type of security control aims to discourage potential attackers?

- A. Corrective control
- B. Deterrent control
- C. Recovery control
- D. Detective control

Answer: B

Q657. What is the purpose of a security classification scheme?

- A. Ranking employees by their performance scores during annual reviews
- B. Sorting hardware inventory by purchase date for warranty tracking
- C. Categorizing data based on sensitivity to determine protection levels
- D. Organizing network cables by color codes for physical identification

Answer: C

Q658. What does the principle of 'complete mediation' require?

- A. System backups must be created before any configuration change is applied
- B. Security policies must be reviewed by every employee before implementation
- C. All network traffic must be encrypted using the strongest available algorithms
- D. Every access request must be checked against access control policies each time

Answer: D

Q659. What is a compensating control in security?

- A. A redundant backup system that activates during scheduled maintenance
- B. An alternative control used when a primary control cannot be implemented
- C. A monitoring tool that compensates for network latency issues automatically
- D. A financial payment made to users affected by a security breach event

Answer: B

Q660. How does the principle of 'economy of mechanism' improve security?

- A. By limiting the number of employees with access to security dashboards
- B. By automating all security processes to eliminate any human involvement
- C. By keeping security designs simple and small to reduce potential flaws
- D. By reducing costs through outsourcing all security operations and functions

Answer: C

Q661. What is the main advantage of elliptic curve cryptography over traditional RSA?

- A. ECC is immune to all known quantum computing attacks on encryption
- B. ECC provides equivalent security with significantly smaller key sizes than RSA
- C. ECC does not require any mathematical computations for key generation
- D. ECC is easier to implement in software compared to RSA algorithm libraries

Answer: B

Q662. How does a message authentication code (MAC) differ from a digital signature?

- A. MACs use symmetric keys shared between parties while digital signatures use asymmetric keys
- B. MACs can only be used for email while digital signatures work across all protocols
- C. MACs are longer in byte length than digital signatures for all algorithm types
- D. MACs provide confidentiality while digital signatures provide only availability protection

Answer: A

Q663. What is a key exchange protocol used to securely share keys over an insecure channel?

- A. File Transfer Protocol
- B. Diffie-Hellman protocol
- C. Simple Mail Transfer Protocol
- D. Address Resolution Protocol

Answer: B

Q664. What is a rainbow table attack in cryptography?

- A. An attack using precomputed hash values to reverse password hashes quickly
- B. A social engineering attack using colorful phishing emails to attract attention
- C. A method of intercepting encrypted traffic using spectrum analysis techniques
- D. A brute-force attack that tries every possible character combination one at a time

Answer: A

Q665. What is the purpose of adding a salt to password hashing?

- A. To compress the hash output into a shorter string for efficient storage
- B. To increase the computational speed of the hashing algorithm execution
- C. To encrypt the hash output so it cannot be read from the database directly
- D. To add random data preventing identical passwords from producing identical hashes

Answer: D

Q666. Which mode of AES encryption uses an initialization vector and chains blocks together?

- A. Cipher Block Chaining mode
- B. Output Feedback Block mode
- C. Electronic Codebook mode
- D. Counter Frequency mode

Answer: A

Q667. What is the role of a public key infrastructure (PKI)?

- A. Managing the lifecycle of digital certificates including issuance and revocation
- B. Storing all encryption keys in a single centralized public-facing database
- C. Generating one-time passwords for multi-factor authentication processes
- D. Providing public internet access to organizations through dedicated gateways

Answer: A

Q668. What is a known-plaintext attack?

- A. An attack where the attacker guesses the encryption key using random brute force tries
- B. An attack where the attacker has access to both plaintext and its corresponding ciphertext
- C. An attack where the attacker can only observe encrypted data passing over the network
- D. An attack where the attacker modifies plaintext before it reaches the encryption system

Answer: B

Q669. What does perfect forward secrecy guarantee?

- A. That compromise of long-term keys does not compromise past session keys
- B. That all future encryption keys are derived from a single master key
- C. That messages are guaranteed to arrive in the exact order they were sent
- D. That encrypted messages can never be decrypted even by the intended recipient

Answer: A

Q670. What is steganography?

- A. An encryption algorithm that uses steganographic keys for cipher operations
- B. A protocol for securely transmitting keys between two authenticated parties
- C. A technique for breaking encrypted messages using statistical analysis methods
- D. A method of hiding data within other non-secret data like images or audio

Answer: D

Q671. What is the difference between an IDS and an IPS?

- A. IDS operates at the application layer while IPS operates at the physical layer only
- B. IDS is hardware-based while IPS is exclusively deployed as a software solution
- C. IDS detects and alerts on threats while IPS can also actively block malicious traffic
- D. IDS blocks threats automatically while IPS only generates alerts for manual review

Answer: C

Q672. How does a stateful firewall differ from a stateless packet filter?

- A. A stateful firewall tracks active connections and makes decisions based on connection state
- B. A stateless filter provides more security because it examines every packet independently
- C. A stateless filter remembers previous packets while stateful firewalls process in isolation
- D. A stateful firewall only inspects packet headers while stateless filters inspect payloads

Answer: A

Q673. What is ARP spoofing and what risk does it pose?

- A. Sending fake ARP messages to link an attacker's MAC address with a legitimate IP address
- B. Intercepting DNS queries to redirect users to malicious websites for credential theft
- C. Exploiting TCP three-way handshake to create half-open connections exhausting resources
- D. Flooding a switch with random MAC addresses to force it into hub mode for broadcasting

Answer: A

Q674. What is the purpose of DNS Security Extensions (DNSSEC)?

- A. Increasing the speed of DNS resolution through distributed caching mechanisms
- B. Authenticating DNS responses to prevent DNS spoofing and cache poisoning attacks
- C. Encrypting all DNS query traffic between clients and DNS resolver servers
- D. Blocking malicious domain names from being resolved by any DNS resolver server

Answer: B

Q675. What is network address translation's security benefit?

- A. It hides internal IP addresses from external networks making them harder to target
- B. It automatically patches vulnerabilities on all devices connected to the network
- C. It blocks all incoming traffic regardless of any configured firewall rule policies
- D. It encrypts all outgoing traffic automatically before reaching the internet gateway

Answer: A

Q676. What is a honeynet in network security?

- A. A network of honeypots designed to attract and study attacker behavior and techniques
- B. A high-speed encrypted network used exclusively for secure financial data transfers
- C. A mesh network topology that provides redundant paths for critical data transmission
- D. A backup network that automatically activates when the primary network fails entirely

Answer: A

Q677. How does port security help protect a network switch?

- A. By encrypting all traffic passing through individual switch ports automatically
- B. By increasing the bandwidth available on each switch port for better performance
- C. By monitoring port temperatures to prevent overheating and hardware failure events
- D. By limiting the number of MAC addresses allowed on a port to prevent MAC flooding

Answer: D

Q678. What does the term 'east-west traffic' refer to in network security?

- A. Traffic flowing between cloud service providers and on-premises data centers
- B. Traffic flowing between the internal network and external internet destinations
- C. Traffic flowing between the organization headquarters and remote branch offices
- D. Traffic flowing between internal servers and systems within the same network

Answer: D

Q679. What is the role of a network tap in security monitoring?

- A. A device that copies network traffic for monitoring without disrupting data flow
- B. A protocol that encrypts monitoring data before sending it to analysis systems
- C. A software agent installed on endpoints to collect local network usage metrics
- D. A tool that terminates suspicious network connections automatically in real time

Answer: A

Q680. What is the purpose of 802.1X network access control?

- A. Providing port-based authentication to control device access to the network
- B. Encrypting wireless traffic between access points and connected user devices
- C. Managing IP address allocation through dynamic host configuration services
- D. Monitoring network performance metrics like latency and packet loss rates

Answer: A

Q681. How does Address Space Layout Randomization (ASLR) improve OS security?

- A. It encrypts all memory contents to prevent unauthorized data access by programs
- B. It randomizes memory addresses making buffer overflow exploits harder to execute
- C. It allocates more memory to security processes to improve threat detection speed
- D. It defragments memory to prevent crashes caused by memory fragmentation issues

Answer: B

Q682. What is the security purpose of chroot in Linux?

- A. Automatically rotating log files to prevent disk space from being exhausted
- B. Scheduling system updates to run during off-peak hours for minimal impact
- C. Isolating a process by changing its apparent root directory to limit access
- D. Granting root access to all users temporarily during maintenance periods

Answer: C

Q683. What is the purpose of Data Execution Prevention (DEP) in modern operating systems?

- A. Preventing programs from executing code stored in non-executable memory regions
- B. Encrypting executable files on disk to prevent unauthorized code modification
- C. Stopping users from deleting critical system files and configuration settings
- D. Blocking network traffic from reaching applications running on protected ports

Answer: A

Q684. How does mandatory access control (MAC) differ from discretionary access control (DAC) in OS security?

- A. MAC is used only in Windows while DAC is used exclusively in Linux environments
- B. MAC uses system-enforced policies based on labels while DAC lets owners set permissions
- C. MAC is less secure than DAC because it gives users more control over their resources
- D. MAC allows resource owners to set permissions while DAC uses system-enforced policies

Answer: B

Q685. What is the role of SELinux in operating system security?

- A. Scheduling automatic system updates and security patches for Linux distributions
- B. Providing mandatory access control policies for Linux kernel security enforcement
- C. Managing network connections and firewall rules for Linux-based server systems
- D. Creating encrypted virtual machines for running untrusted applications safely

Answer: B

Q686. What is privilege escalation in the context of OS security?

- A. An attacker gaining higher privileges than originally authorized on the system
- B. An automated process that upgrades user accounts during scheduled maintenance
- C. A legitimate administrator granting additional permissions to a user upon request
- D. A security feature that temporarily elevates privileges for approved operations

Answer: A

Q687. What is the purpose of the Windows Security Event Log?

- A. Recording login attempts, policy changes, and security-related system events
- B. Maintaining a list of installed programs and their current version numbers
- C. Tracking hardware performance metrics like CPU temperature and fan speeds
- D. Storing application crash reports and software compatibility error messages

Answer: A

Q688. What is sandboxing in the context of OS security?

- A. Storing sensitive data in encrypted containers on external storage devices
- B. Running applications in an isolated environment to limit their system access
- C. Testing network connections in a simulated environment before deployment
- D. Backing up system files to a separate partition for disaster recovery use

Answer: B

Q689. How does process isolation improve operating system security?

- A. It makes all processes run faster by allocating equal CPU time to each one
- B. It automatically terminates any process that has been running too long
- C. It prevents one process from accessing the memory space of another process
- D. It allows processes to share memory freely for improved data collaboration

Answer: C

Q690. What is the security benefit of using a minimal OS installation?

- A. It increases system performance by using all available disk space for caching data
- B. It automatically configures all security settings without administrator intervention
- C. It provides more user-friendly features for non-technical administrative staff
- D. It reduces the attack surface by installing only necessary components and services

Answer: D

Q691. How does Content Security Policy (CSP) help prevent XSS attacks?

- A. By encrypting all script files before they are sent to the client browser
- B. By requiring users to manually approve each script before it executes
- C. By specifying which content sources the browser should trust and execute
- D. By automatically removing all JavaScript code from web page responses

Answer: C

Q692. What is the difference between stored XSS and reflected XSS?

- A. Stored XSS uses SQL queries while reflected XSS uses JavaScript code exclusively
- B. Stored XSS requires physical access while reflected XSS requires network level access
- C. Stored XSS only affects administrators while reflected XSS targets all regular users
- D. Stored XSS persists on the server while reflected XSS is returned in the immediate response

Answer: D

Q693. What is an insecure direct object reference (IDOR) vulnerability?

- A. When a web server returns a 404 error because the requested resource was not found
- B. When an application exposes internal object references allowing unauthorized data access
- C. When JavaScript code references DOM elements that have not been loaded on the page
- D. When a database query fails because of incorrect table reference in the SQL statement

Answer: B

Q694. How do anti-CSRF tokens protect web applications?

- A. By embedding unique tokens in forms that the server validates on submission
- B. By encrypting all form data before submission to the server for processing
- C. By blocking all cross-origin requests regardless of their source or purpose
- D. By requiring users to re-enter their password before every form submission

Answer: A

Q695. What is the purpose of HTTP security headers like X-Frame-Options?

- A. Defining the caching duration for static resources like images and style sheets
- B. Controlling whether a browser can render a page inside a frame to prevent clickjacking
- C. Setting the default character encoding for all text displayed on the web page
- D. Increasing the maximum file size that can be uploaded through web application forms

Answer: B

Q696. What is a Server-Side Request Forgery (SSRF) attack?

- A. An attack where the client forges server certificates to intercept encrypted communications
- B. An attack where database queries are modified to return unauthorized data to the attacker
- C. An attack where session cookies are stolen through cross-site scripting vulnerabilities
- D. An attack where the server is tricked into making requests to unintended internal resources

Answer: D

Q697. What is the security risk of exposing detailed error messages to users?

- A. Detailed errors increase bandwidth usage because they are larger than generic messages
- B. Detailed errors slow down the web application because they consume more server memory
- C. Detailed errors may reveal system architecture, paths, and technology stack to attackers
- D. Detailed errors cause browsers to crash when they contain too many special characters

Answer: C

Q698. How does parameterized querying prevent SQL injection?

- A. By encrypting all SQL queries before they are sent to the database server
- B. By separating SQL code from user data so input is treated as data not commands
- C. By limiting the number of queries a user can execute within a given time period
- D. By automatically detecting and removing SQL keywords from all user input fields

Answer: B

Q699. What is the purpose of the Strict-Transport-Security (HSTS) header?

- A. Limiting the number of concurrent connections a browser can make to the server
- B. Requiring users to authenticate before accessing any page on the web application
- C. Compressing HTTP responses to reduce bandwidth consumption for mobile devices
- D. Forcing browsers to only communicate with the server using HTTPS connections

Answer: D

Q700. What is DOM-based XSS?

- A. An XSS attack that targets the server-side rendering engine of the web application
- B. An XSS attack that modifies DNS records to redirect users to malicious websites
- C. An XSS attack that only affects web applications using document object models
- D. An XSS attack where the payload is executed through client-side DOM manipulation

Answer: D

Q701. How does polymorphic malware evade antivirus detection?

- A. By changing its code structure with each infection while maintaining its core functionality
- B. By disguising itself as a legitimate system process visible in the task manager window
- C. By encrypting network traffic to prevent antivirus from scanning its communications data
- D. By running only during times when antivirus software is scheduled for daily scan updates

Answer: A

Q702. What is a watering hole attack?

- A. Placing malware on USB drives and leaving them in public places for people to find
- B. Flooding a server with HTTP requests to make the website unavailable to all users
- C. Sending mass phishing emails to random recipients hoping some will click the links
- D. Compromising websites frequently visited by a target group to infect their devices

Answer: D

Q703. What is the difference between spear phishing and regular phishing?

- A. Spear phishing only targets executives while regular phishing only targets entry-level employees
- B. Spear phishing uses phone calls while regular phishing exclusively uses email messages for targeting
- C. Spear phishing targets specific individuals with personalized messages while phishing is generic mass emails
- D. Spear phishing is legal while regular phishing is considered illegal under all global cyber laws

Answer: C

Q704. What is a logic bomb in malware?

- A. An encryption algorithm that self-destructs after a single use to prevent key reuse
- B. A type of virus that destroys computer hardware by overloading electrical components
- C. A network attack that sends logically malformed packets to crash target router devices
- D. Malicious code that activates when specific conditions or triggers are met at a set time

Answer: D

Q705. How does a drive-by download attack work?

- A. An attacker physically inserts a USB drive with malware into the target computer
- B. An attacker drives near a building and exploits weak WiFi to download sensitive data
- C. Malware is automatically downloaded when a user visits a compromised website page
- D. A user intentionally downloads pirated software that contains hidden malware files

Answer: C

Q706. What is credential stuffing and how does it differ from brute force?

- A. Credential stuffing is slower than brute force because it validates each credential manually
- B. Credential stuffing targets only email accounts while brute force targets all application types
- C. Credential stuffing uses previously stolen username-password pairs while brute force guesses randomly
- D. Credential stuffing requires physical access while brute force is conducted entirely remotely

Answer: C

Q707. What is a supply chain attack in cyber security?

- A. Intercepting purchase orders to redirect payments to attacker-controlled bank accounts
- B. Stealing shipping information to track the location of valuable technology equipment
- C. Attacking the physical supply chain to prevent delivery of hardware to data centers
- D. Compromising a vendor or supplier to gain access to their customers' systems and data

Answer: D

Q708. What distinguishes a zero-day exploit from a known exploit?

- A. A zero-day exploit is less dangerous because it has not been tested in real attacks
- B. A zero-day exploit targets a vulnerability for which no patch or fix currently exists
- C. A zero-day exploit can only be used once before it becomes completely ineffective
- D. A zero-day exploit requires zero technical knowledge to successfully execute remotely

Answer: B

Q709. What is DNS spoofing?

- A. Registering domain names similar to legitimate ones to trick users into visiting them
- B. Encrypting DNS traffic to prevent eavesdropping on domain resolution network queries
- C. Corrupting DNS cache to redirect domain lookups to malicious IP addresses for phishing
- D. Blocking DNS queries to prevent users from resolving any domain names on the network

Answer: C

Q710. What is a fileless malware attack?

- A. An attack that only targets files stored on removable USB drives and external storage
- B. An attack that hides malware inside image files that appear normal to the viewer
- C. An attack that operates entirely in memory without writing any files to the disk drive
- D. An attack that deletes all files on the system before demanding a ransom payment

Answer: C

Q711. How does SAML enable single sign-on across different web applications?

- A. By encrypting browser cookies so they can be shared between different website domains
- B. By requiring users to create identical credentials across every participating application
- C. By storing all user passwords in a shared centralized database accessible to all applications
- D. By exchanging authentication and authorization data between identity and service providers

Answer: D

Q712. What is the difference between authentication and authorization?

- A. Authentication determines access rights while authorization verifies the user's identity
- B. Authentication verifies who the user is while authorization determines what they can access
- C. Authentication and authorization are identical processes with different names in security
- D. Authentication is used for external users while authorization is only for internal users

Answer: B

Q713. What is the primary advantage of hardware security tokens over SMS-based two-factor authentication?

- A. Hardware tokens are resistant to SIM swapping and interception attacks that affect SMS codes
- B. Hardware tokens are cheaper to deploy than SMS-based authentication across all organizations
- C. Hardware tokens do not require any batteries or charging to function and generate codes
- D. Hardware tokens provide faster authentication than SMS because they work without networks

Answer: A

Q714. What is OAuth 2.0 primarily designed for?

- A. Providing a universal username and password standard across all web applications
- B. Delegated authorization allowing third-party applications limited access to user resources
- C. Encrypting all data stored in databases used by web applications and mobile apps
- D. Replacing traditional firewalls with application-level access control mechanisms

Answer: B

Q715. What is the purpose of an identity provider (IdP) in federated identity management?

- A. Developing custom authentication protocols for each individual web application
- B. Creating and managing user identities and authenticating users for service providers
- C. Providing internet connectivity to all users within a federated network system
- D. Storing all application data and user files in a centralized cloud storage system

Answer: B

Q716. What is privileged access management (PAM)?

- A. A tool for creating user accounts in bulk during organizational onboarding events
- B. A method of granting all employees administrator access for operational efficiency
- C. A protocol for resetting forgotten passwords through automated self-service portals
- D. A framework for managing and monitoring accounts with elevated system access privileges

Answer: D

Q717. How does passwordless authentication work?

- A. It uses very short passwords of one or two characters for faster login processes
- B. It stores passwords in the browser so users do not need to type them each time
- C. It automatically generates and fills in passwords without any user involvement
- D. It replaces passwords with alternatives like biometrics, security keys, or magic links

Answer: D

Q718. What is the purpose of access reviews in identity management?

- A. Automatically granting additional permissions to users who request them through helpdesk
- B. Periodically reviewing and validating user access rights to ensure they remain appropriate
- C. Removing all user access permissions at the end of each fiscal quarter for renewal
- D. Creating new user accounts for employees who join the organization each calendar month

Answer: B

Q719. What is the FIDO2 authentication standard?

- A. A protocol for encrypting data stored in cloud computing environments and databases
- B. A framework for managing firewall rules across multiple network security appliances
- C. A specification for designing user interfaces in mobile banking applications only
- D. A standard enabling strong passwordless authentication using public key cryptography

Answer: D

Q720. What is just-in-time (JIT) access provisioning?

- A. Pre-provisioning all possible access rights to users during their first day of employment
- B. Granting elevated access only when needed and automatically revoking it after a set period
- C. Allowing users to permanently keep any elevated access they have ever been granted
- D. Manually reviewing and approving every access request within ninety business days

Answer: B

Q721. How does dynamic application security testing (DAST) differ from SAST?

- A. DAST requires source code access while SAST only needs the compiled application binary files
- B. DAST is only for mobile applications while SAST works exclusively on web-based applications
- C. DAST tests running applications from the outside while SAST analyzes source code without execution
- D. DAST analyzes source code while SAST tests running applications by sending network requests

Answer: C

Q722. What is threat modeling in secure software development?

- A. Creating 3D models of physical security threats for visual presentation to stakeholders
- B. Systematically identifying and evaluating potential threats to an application during design
- C. Documenting all past security incidents in a database for historical trending analysis
- D. Monitoring production environments for active threats and ongoing security incidents

Answer: B

Q723. What is the STRIDE threat model used for?

- A. Rating software developers on their security knowledge, training, and certification levels
- B. Measuring the speed, throughput, reliability, integrity, durability, and efficiency of systems
- C. Categorizing threats into Spoofing, Tampering, Repudiation, Information Disclosure, DoS, and Elevation
- D. Defining six phases of the software development lifecycle from planning through deployment

Answer: C

Q724. What is software composition analysis (SCA)?

- A. Identifying vulnerabilities in third-party libraries and open-source components used in software
- B. Analyzing the performance composition of software to optimize memory and CPU utilization
- C. Composing new software by combining multiple commercial off-the-shelf application products
- D. Creating detailed architectural diagrams of software systems for documentation purposes

Answer: A

Q725. What is the purpose of security requirements in the SDLC?

- A. Specifying the programming languages and frameworks that developers must use exclusively
- B. Documenting the team structure and roles assigned for the software development project effort
- C. Defining specific security features and controls that the software must implement from the start
- D. Listing the hardware requirements needed to run the software on production server environments

Answer: C

Q726. What is the concept of 'shift left' in DevSecOps?

- A. Integrating security testing earlier in the development pipeline rather than only at the end
- B. Moving all development operations to a different time zone for round-the-clock development
- C. Shifting security responsibility entirely from developers to the dedicated security team
- D. Postponing security testing until after deployment to avoid slowing down development speed

Answer: A

Q727. What is a secure coding standard?

- A. A set of rules and guidelines for writing code that minimizes security vulnerabilities
- B. A certification program for developers who complete an advanced security course
- C. A programming language designed specifically for writing security-focused applications
- D. A testing framework that automatically generates test cases for security validation

Answer: A

Q728. What is fuzzing in security testing?

- A. Testing software performance under heavy load conditions to find stability issues
- B. Blurring sensitive data in screenshots to prevent information disclosure publicly
- C. Reviewing code documentation for unclear or ambiguous security requirement specifications
- D. Providing random, unexpected, or malformed input to software to discover vulnerabilities

Answer: D

Q729. What is the OWASP Secure Coding Practices guide?

- A. A programming language designed for building security-focused web applications
- B. A proprietary software tool for scanning applications for known security defects
- C. A certification exam that developers must pass before writing production code
- D. A checklist of secure coding practices to help developers avoid common vulnerabilities

Answer: D

Q730. What is the purpose of a security champion program in development teams?

- A. Replacing developers with security professionals to write all application source code
- B. Hiring external consultants to perform all security testing on behalf of the team
- C. Creating a competition between teams to determine which team has the fewest bugs
- D. Embedding security-focused developers within teams to promote secure coding practices

Answer: D

Q731. How does WPA3 improve security over WPA2?

- A. WPA3 uses shorter encryption keys for faster wireless communication performance
- B. WPA3 uses Simultaneous Authentication of Equals preventing offline dictionary attacks
- C. WPA3 requires physical USB connections between all devices on the wireless network
- D. WPA3 eliminates the need for passwords entirely by using open authentication mode

Answer: B

Q732. What is the KRACK attack and which protocol does it target?

- A. An attack that cracks WEP encryption using statistical analysis of captured packets
- B. An attack exploiting NFC protocols to steal contactless payment card information
- C. An attack targeting Bluetooth pairing to intercept data between connected devices
- D. An attack exploiting WPA2's four-way handshake to reinstall encryption keys

Answer: D

Q733. What is the purpose of network access control (NAC) for wireless networks?

- A. Increasing wireless signal strength throughout the building for better coverage
- B. Ensuring only compliant and authorized devices can connect to the wireless network
- C. Providing guest users with permanent unrestricted access to corporate resources
- D. Reducing the number of available wireless channels to minimize signal interference

Answer: B

Q734. What is mobile application containerization?

- A. Compressing mobile applications to reduce their download size from app stores
- B. Isolating work applications and data from personal content on a mobile device
- C. Grouping multiple applications into a single installation package for convenience
- D. Converting desktop applications into mobile-compatible versions automatically

Answer: B

Q735. What is a deauthentication attack in wireless networks?

- A. Changing the wireless network password to lock out all currently connected users
- B. Preventing new devices from authenticating by overloading the DHCP lease pool
- C. Sending forged deauth frames to disconnect clients from a wireless access point
- D. Disabling the authentication server to prevent all users from logging in remotely

Answer: C

Q736. What security feature does certificate pinning provide in mobile applications?

- A. Encrypting stored application data using the server's public certificate key
- B. Creating self-signed certificates for development testing environments only
- C. Automatically updating expired certificates without requiring user interaction
- D. Preventing man-in-the-middle attacks by validating the server's specific certificate

Answer: D

Q737. What is the security concern with NFC (Near Field Communication) on mobile devices?

- A. NFC signals interfere with WiFi and Bluetooth connections on the same device
- B. Attackers in close proximity can intercept or modify NFC communications and data
- C. NFC technology is incompatible with modern mobile operating system versions
- D. NFC connections consume excessive battery power causing devices to shut down

Answer: B

Q738. What is the purpose of a wireless intrusion prevention system (WIPS)?

- A. Encrypting all wireless traffic between access points and the core network
- B. Managing the distribution of IP addresses for all wireless connected devices
- C. Monitoring wireless networks for unauthorized access points and malicious activity
- D. Increasing the range and signal strength of legitimate wireless access points

Answer: C

Q739. How does mobile threat defense (MTD) protect enterprise mobile devices?

- A. By detecting device, network, and application-level threats on mobile endpoints continuously
- B. By requiring users to connect to the corporate VPN before enabling any application
- C. By replacing the mobile operating system with a custom hardened version for security
- D. By blocking all internet access on mobile devices to prevent any possible threats

Answer: A

Q740. What is SIM swapping and why is it a security concern?

- A. An attacker uses a modified SIM card to intercept all cellular traffic in a geographic area
- B. An attacker convinces a carrier to transfer a victim's number to a new SIM for account takeover
- C. An attacker physically removes the SIM card from a device to prevent it from making calls
- D. An attacker clones a SIM card to use two identical phone numbers simultaneously on the network

Answer: B

Q741. How does the security responsibility differ between IaaS, PaaS, and SaaS?

- A. Customer responsibility decreases from IaaS to SaaS as the provider manages more layers
- B. Customer responsibility increases from IaaS to SaaS as applications become more complex
- C. Security responsibility is identical across all three models with no practical differences
- D. The cloud provider has no security responsibility in any of the three service model types

Answer: A

Q742. What is cloud security posture management (CSPM)?

- A. A training program that teaches employees how to securely use cloud applications at work
- B. A certification framework for evaluating the security capabilities of cloud service providers
- C. A physical security system protecting the data center buildings of cloud service providers
- D. Tools that continuously monitor cloud environments for misconfigurations and compliance violations

Answer: D

Q743. What is the security risk of excessive IAM permissions in cloud environments?

- A. Over-privileged accounts automatically trigger account suspension by the cloud provider
- B. Too many permissions cause billing increases because each permission has a monthly fee
- C. Excessive permissions make cloud services run slower due to additional authorization checks
- D. Over-privileged accounts increase blast radius if credentials are compromised by attackers

Answer: D

Q744. What is a cloud workload protection platform (CWPP)?

- A. A tool that monitors employee productivity when working on cloud-based applications
- B. A platform for managing cloud service billing and cost optimization across accounts
- C. A framework for distributing workloads evenly across multiple cloud provider regions
- D. A solution protecting server workloads including VMs, containers, and serverless functions

Answer: D

Q745. What is the purpose of cloud-native application protection platforms (CNAPP)?

- A. Protecting only legacy applications that have been migrated from on-premises to the cloud
- B. Managing the deployment and scaling of cloud-native applications across multiple regions
- C. Providing network connectivity between different cloud providers for hybrid deployments
- D. Unifying CSPM, CWPP, and other cloud security capabilities into a single integrated platform

Answer: D

Q746. What is the security concern with serverless computing?

- A. Serverless functions cannot be encrypted so all data processed is exposed in plaintext
- B. Reduced visibility into the execution environment makes traditional security monitoring difficult
- C. Serverless computing eliminates all security risks because the provider manages everything
- D. Serverless functions always run with unlimited permissions that cannot be restricted at all

Answer: B

Q747. How does cloud key management service (KMS) enhance security?

- A. By providing unlimited free encryption keys to all users across cloud provider accounts
- B. By eliminating the need for encryption by using secure physical storage in data centers
- C. By automatically encrypting all data without requiring any configuration from customers
- D. By centrally managing encryption keys with access controls, rotation, and audit logging

Answer: D

Q748. What is the purpose of VPC peering in cloud networking security?

- A. Sharing virtual machines between different cloud provider accounts for collaboration
- B. Duplicating VPC configurations across regions for disaster recovery backup purposes
- C. Enabling private network connectivity between VPCs without traversing the public internet
- D. Creating public endpoints for cloud services accessible from anywhere on the internet

Answer: C

Q749. What is a container escape vulnerability?

- A. A configuration error that causes containers to restart repeatedly without stopping
- B. A flaw allowing a process inside a container to break out and access the host system
- C. A storage limitation that prevents containers from saving data to persistent volumes
- D. A network issue that prevents containers from communicating with external services

Answer: B

Q750. What is the role of security groups in cloud environments?

- A. Creating user groups for shared access to cloud management console dashboards
- B. Grouping security team members for project management and task assignment purposes
- C. Acting as virtual firewalls controlling inbound and outbound traffic to cloud resources
- D. Organizing cloud resources into categories for billing and cost allocation tracking

Answer: C

Q751. What is the order of volatility and why is it important in forensics?

- A. A ranking of hard drives by age to determine which should be imaged first during analysis
- B. A guideline for collecting evidence from most volatile to least volatile to preserve critical data
- C. A method for organizing case files by the severity of the crime being investigated
- D. A standard for labeling evidence bags based on the type of device they contain

Answer: B

Q752. What is file carving in digital forensics?

- A. Converting file formats from proprietary to open standards for analysis compatibility
- B. Splitting large files into smaller pieces for easier transfer between forensic systems
- C. Removing metadata from files to protect the privacy of investigation subjects
- D. Recovering files from unallocated disk space using file signatures and structure analysis

Answer: D

Q753. What is timeline analysis in digital forensics?

- A. Scheduling forensic team meetings and deadlines for completing investigation milestones
- B. Measuring how long it takes to complete each step of the forensic investigation process
- C. Creating a chronological sequence of events from multiple evidence sources to reconstruct activities
- D. Tracking the expiration dates of forensic software licenses and tool subscriptions

Answer: C

Q754. How does live forensics differ from traditional dead-box forensics?

- A. Live forensics uses automated tools while dead-box forensics requires entirely manual analysis
- B. Live forensics is performed in real-time during an attack while dead-box is done months afterward
- C. Live forensics analyzes running systems to capture volatile data while dead-box analyzes powered-off devices
- D. Live forensics only applies to servers while dead-box forensics is exclusively for desktop computers

Answer: C

Q755. What is steganographic analysis in digital forensics?

- A. Recovering encrypted files by testing common passwords against encryption algorithms
- B. Scanning network traffic logs for unusual patterns indicating data exfiltration attempts
- C. Analyzing the physical damage to storage media caused by fire, water, or impact events
- D. Detecting and extracting hidden data concealed within images, audio, or other media files

Answer: D

Q756. What is the role of email header analysis in forensic investigations?

- A. Tracing the origin and path of emails by examining routing information in headers
- B. Counting the total number of emails sent by a suspect during a specific period
- C. Formatting email evidence for presentation in court using standardized templates
- D. Reading the content of encrypted emails without needing the decryption password

Answer: A

Q757. What is the significance of slack space in disk forensics?

- A. Slack space is where the operating system stores its temporary cache files during boot
- B. Slack space is reserved by the file system for future use and always contains zero data
- C. Slack space is unused memory that should be defragmented to improve disk performance
- D. Slack space may contain remnants of previously deleted files that can serve as evidence

Answer: D

Q758. What is network forensics?

- A. Managing network device configurations and firmware updates across the organization
- B. Designing and building secure network architectures for enterprise organizations
- C. Optimizing network performance by identifying and resolving bandwidth bottlenecks
- D. Capturing and analyzing network traffic to investigate security incidents and gather evidence

Answer: D

Q759. What is the purpose of a forensic report?

- A. Providing real-time alerts about ongoing security incidents to the response team
- B. Documenting findings, methodology, and conclusions in a clear format for stakeholders
- C. Automatically generating remediation steps for all security vulnerabilities discovered
- D. Creating backups of all evidence files for long-term archival storage purposes

Answer: B

Q760. What is the purpose of a forensic toolkit?

- A. A collection of specialized software tools used for evidence collection and analysis
- B. A physical kit containing screwdrivers and tools for disassembling computer hardware
- C. A set of legal documents required before beginning any forensic investigation
- D. A training manual for new forensic investigators joining the digital forensics team

Answer: A

Q761. How should an incident response team handle evidence during a ransomware attack?

- A. Immediately paying the ransom to recover files before collecting any forensic evidence
- B. Deleting the ransom note to prevent panic among employees and organizational leadership
- C. Formatting all affected systems right away to eliminate the ransomware infection quickly
- D. Preserving encrypted files and ransom notes as evidence before attempting any recovery

Answer: D

Q762. What is the difference between short-term and long-term containment?

- A. Short-term involves rebuilding systems while long-term involves monitoring for threat recurrence
- B. Short-term takes days to implement while long-term is applied within the first few minutes
- C. Short-term stops immediate damage while long-term implements durable fixes allowing investigation
- D. Short-term is performed by management while long-term is handled by the technical team

Answer: C

Q763. What is the purpose of tabletop exercises in incident response?

- A. Testing network bandwidth by sending large amounts of simulated attack traffic
- B. Simulating incident scenarios through discussion to test plans and identify gaps
- C. Training new employees on how to use desktop computers and office applications
- D. Reviewing the physical layout of office furniture to optimize workspace ergonomics

Answer: B

Q764. What is an indicator of compromise (IoC)?

- A. A financial indicator showing the cost of recovering from a security breach event
- B. Observable evidence that a system or network may have been breached by an attacker
- C. A metric measuring the overall compromise between security and user convenience
- D. A performance metric tracking system uptime percentage across all business hours

Answer: B

Q765. When should law enforcement be notified during an incident response?

- A. When the incident involves criminal activity, data breaches with legal requirements, or national security
- B. Only after the incident has been fully resolved and all systems have been restored to normal
- C. Law enforcement should never be notified as it could compromise the internal investigation
- D. Immediately for every security event regardless of severity or type of incident detected

Answer: A

Q766. What is the purpose of incident severity classification?

- A. Categorizing incidents by impact level to determine appropriate response and resource allocation
- B. Ranking incident response team members by their seniority within the organizational hierarchy
- C. Sorting incident tickets alphabetically by the name of the person who reported them
- D. Classifying incidents by the time of day they occurred for shift scheduling optimization

Answer: A

Q767. What is the role of threat intelligence in incident response?

- A. Providing context about threats, attacker tactics, and IoCs to improve detection and response
- B. Generating marketing reports about the organization's security posture for stakeholders
- C. Replacing the need for incident response teams by using artificial intelligence for defense
- D. Automatically blocking all threats without requiring any human intervention or analysis

Answer: A

Q768. What is a playbook in incident response?

- A. A predefined set of procedures for handling specific types of security incidents consistently
- B. A report template used for documenting the final results of completed investigations
- C. A catalog of all security tools and their license keys used by the response team
- D. A training manual for new employees covering general workplace policies and procedures

Answer: A

Q769. What is the purpose of an incident response retainer with a third-party provider?

- A. Hiring permanent full-time security staff through a third-party recruitment agency partner
- B. Ensuring expert incident response resources are available on-demand when a major incident occurs
- C. Purchasing cyber insurance to cover the financial costs of future data breach events
- D. Outsourcing all daily security monitoring operations to a third-party managed service provider

Answer: B

Q770. What is a war room in incident response?

- A. A dedicated space where the response team coordinates and manages a major incident together
- B. A secure server room where all critical infrastructure is housed during normal operations
- C. A virtual meeting room used for daily stand-up meetings about project status and updates
- D. A training facility where employees practice responding to simulated physical emergencies

Answer: A

Q771. How does user and entity behavior analytics (UEBA) enhance security monitoring?

- A. By monitoring only privileged user accounts and ignoring standard employee activities
- B. By replacing SIEM systems entirely with machine learning based threat detection tools
- C. By blocking all user activities that deviate from a predefined list of approved actions
- D. By establishing behavioral baselines and detecting anomalies that indicate potential threats

Answer: D

Q772. What is the difference between signature-based and anomaly-based detection?

- A. Signature-based matches known patterns while anomaly-based detects deviations from normal behavior
- B. Signature-based requires more computing resources than anomaly-based detection methods do
- C. Signature-based is used for network monitoring while anomaly-based is used only for endpoints
- D. Signature-based detects unknown threats while anomaly-based only detects previously seen attacks

Answer: A

Q773. What is the purpose of log correlation in SIEM systems?

- A. Connecting related events from different sources to identify complex attack patterns
- B. Converting all logs into a single standardized format for easier visual reading
- C. Backing up log data to secondary storage for disaster recovery preparedness
- D. Removing duplicate log entries to reduce storage space requirements on servers

Answer: A

Q774. What is the role of threat intelligence feeds in security monitoring?

- A. Automatically patching all vulnerabilities found across the organization's infrastructure
- B. Generating monthly compliance reports for regulatory auditors and board presentations
- C. Providing up-to-date indicators of compromise to enhance detection rules and alert accuracy
- D. Replacing human analysts by automatically investigating and resolving all security alerts

Answer: C

Q775. What is a honeypot used for in security monitoring?

- A. Storing encrypted passwords securely in a protected database for authentication use
- B. Providing backup internet connectivity when the primary connection fails completely
- C. Distributing security patches to all endpoints across the organizational network
- D. Attracting attackers to a decoy system to study their techniques and detect intrusions

Answer: D

Q776. What is the challenge of alert fatigue in security operations?

- A. Alerts being displayed in colors that are difficult for analysts to distinguish on screens
- B. Alerts arriving too slowly for analysts to respond before the threat window has closed
- C. Overwhelming analysts with too many alerts causing them to miss genuine threats in the volume
- D. Alert systems consuming too much network bandwidth reducing application performance speed

Answer: C

Q777. What is the purpose of a security baseline in monitoring?

- A. Establishing normal behavior patterns so deviations can be identified as potential threats
- B. Setting the minimum hardware requirements for running security monitoring software tools
- C. Defining the maximum number of security alerts allowed per day before system shutdown
- D. Creating a template for security reports that must be submitted to regulatory authorities

Answer: A

Q778. What is the role of packet capture (PCAP) in security monitoring?

- A. Recording raw network packets for detailed analysis of network communications and attacks
- B. Compressing network traffic to reduce bandwidth usage across organizational network links
- C. Encrypting network packets to protect data confidentiality during transit between systems
- D. Blocking malicious packets from reaching their intended destination on the network

Answer: A

Q779. What is extended detection and response (XDR)?

- A. A unified platform integrating detection across endpoints, network, cloud, and email sources
- B. A compliance tool that extends audit reporting to cover international regulatory frameworks
- C. An advanced firewall that extends network protection to include wireless access points
- D. A backup system that extends data retention periods beyond standard policy requirements

Answer: A

Q780. What is the purpose of dark web monitoring for organizations?

- A. Purchasing threat intelligence tools that are exclusively available on dark web forums
- B. Blocking all organizational users from accessing dark web sites through network filters
- C. Detecting if organizational data, credentials, or assets appear on dark web marketplaces
- D. Monitoring the dark web to recruit security researchers who operate on hidden platforms

Answer: C

Q781. How does GDPR's 'right to be forgotten' impact organizations?

- A. Organizations can ignore deletion requests if the data is stored in encrypted backup archives
- B. Organizations must delete all data about a person including legally required financial records
- C. Organizations must delete personal data upon request when there is no legitimate reason to retain it
- D. The right to be forgotten only applies to social media companies and not other organizations

Answer: C

Q782. What is the difference between data privacy and data security?

- A. Privacy is a technical control while security is a legal requirement enforced by regulators
- B. Privacy only applies to digital data while security only applies to physical document records
- C. Privacy governs how data should be handled while security implements protections against unauthorized access
- D. Privacy and security are identical concepts with no meaningful distinction between them in practice

Answer: C

Q783. What is a Data Protection Impact Assessment (DPIA)?

- A. An assessment evaluating privacy risks of data processing activities before implementation
- B. A financial audit calculating the total cost of data protection measures for budgeting
- C. A training evaluation assessing employee knowledge of organizational privacy policies
- D. A performance test measuring how fast data protection systems process user requests

Answer: A

Q784. What is the Computer Fraud and Abuse Act (CFAA)?

- A. An international treaty establishing global standards for ethical hacking practices
- B. A state-level law in California that requires breach notification within 24 hours
- C. A European regulation that mandates encryption for all cross-border data transfers
- D. A US federal law that criminalizes unauthorized access to protected computer systems

Answer: D

Q785. What is the role of a Data Protection Officer (DPO)?

- A. Managing the physical security of the organization's data center facilities full time
- B. Selling the organization's data to third parties for approved marketing purposes
- C. Overseeing compliance with data protection regulations and serving as a point of contact
- D. Developing and maintaining all software applications used by the organization directly

Answer: C

Q786. What is the purpose of PCI DSS?

- A. Establishing security standards for organizations that handle credit card payment data
- B. Regulating the manufacturing and distribution of personal computer hardware devices
- C. Defining standards for physical security of government classified information systems
- D. Setting requirements for environmental sustainability in data center construction

Answer: A

Q787. What is the concept of data minimization in privacy law?

- A. Collecting only the minimum personal data necessary for the specified legitimate purpose
- B. Minimizing the number of employees who know about the organization's privacy practices
- C. Compressing personal data to minimize the storage space required on organizational servers
- D. Deleting all personal data immediately after collection regardless of business requirements

Answer: A

Q788. What are bug bounty programs and their legal considerations?

- A. Programs that penalize employees financially for introducing bugs into production code
- B. Programs that pay hackers to attack competitor organizations to find their security weaknesses
- C. Programs rewarding researchers for finding vulnerabilities with clear legal scope and authorization
- D. Programs offering rewards for reporting software bugs related to user interface design only

Answer: C

Q789. What is the purpose of HIPAA in the United States?

- A. Establishing standards for highway infrastructure protection against cyber attacks
- B. Protecting the privacy and security of patient health information in healthcare organizations
- C. Defining rules for international trade of technology products between countries
- D. Regulating the hiring practices of technology companies across the United States market

Answer: B

Q790. What is the legal concept of 'reasonable security measures'?

- A. Security controls that are appropriate and proportionate to the risk and data sensitivity involved
- B. Security measures dictated entirely by the IT department without legal guidance or oversight
- C. The maximum possible security measures regardless of cost or impact on business operations
- D. Only the most basic security measures such as passwords without any additional protections

Answer: A

Q791. How does the convergence of IT and OT create new security challenges?

- A. OT systems were designed for isolated networks and connecting them to IT increases their attack surface
- B. IT and OT convergence eliminates all security risks by combining the strengths of both domains
- C. IT and OT convergence only affects manufacturing companies and has no impact on other sectors
- D. OT systems are inherently more secure than IT systems and improve the overall security posture

Answer: A

Q792. What is the security concern with large language models (LLMs) in cyber security?

- A. LLMs can generate phishing content, discover vulnerabilities, and create malicious code at scale
- B. LLMs eliminate all social engineering threats by educating users automatically at all times
- C. LLMs cannot process security-related queries and are completely useless for cyber attacks
- D. LLMs are only used for defensive security purposes and have no offensive applications whatsoever

Answer: A

Q793. What is Secure Access Service Edge (SASE)?

- A. A physical security perimeter around data centers using advanced biometric access controls
- B. A software development methodology focused on creating secure edge computing applications
- C. A network protocol designed to replace TCP/IP with more secure communication mechanisms
- D. A cloud architecture combining network security and WAN capabilities into a single service model

Answer: D

Q794. How does 5G technology introduce new cyber security challenges?

- A. 5G networks are slower than 4G making them less attractive targets for cyber attackers
- B. 5G only affects mobile phone security and has no impact on enterprise network security
- C. Increased device density, network slicing complexity, and expanded attack surface from IoT proliferation
- D. 5G eliminates all existing cyber threats through its built-in advanced security architecture

Answer: C

Q795. What is the concept of security mesh architecture?

- A. A physical mesh network used for building security camera systems in large corporate offices
- B. A decorative pattern used in security badge design for employee identification purposes
- C. A distributed security approach that extends controls around individuals and assets regardless of location
- D. A single centralized security platform that manages all security controls from one location

Answer: C

Q796. What is the security impact of digital twins in industrial environments?

- A. Digital twins can be used to simulate attacks and test defenses without risking physical systems
- B. Digital twins automatically patch vulnerabilities in physical systems they are connected to
- C. Digital twins are identical copies of malware used by security researchers in sandboxed environments
- D. Digital twins eliminate all physical security requirements for industrial control systems completely

Answer: A

Q797. What is privacy-enhancing technology (PET)?

- A. Software that automatically deletes all browsing history and cookies after each web session
- B. Physical barriers installed around servers to prevent unauthorized personnel from viewing screens
- C. Technologies that enable data processing while minimizing personal data exposure and privacy risks
- D. Traditional encryption methods that have been rebranded with a more modern marketing name

Answer: C

Q798. How does autonomous vehicle security present unique cyber security challenges?

- A. Autonomous vehicles use isolated systems with no network connectivity eliminating all cyber risks
- B. Autonomous vehicle security is identical to traditional computer security with no unique requirements
- C. Compromised vehicle systems could endanger physical safety, requiring real-time security with minimal latency
- D. Autonomous vehicles are immune to cyber attacks because they use proprietary operating systems

Answer: C

Q799. What is the concept of cyber-physical systems security?

- A. Securing the physical locks and barriers that protect server rooms from unauthorized entry
- B. Protecting integrated systems where digital components interact with and control physical processes
- C. Installing physical firewalls in the walls of data centers to prevent electromagnetic leaks
- D. Developing software applications that run on both mobile phones and desktop computers

Answer: B

Q800. What is the emerging concept of security-as-code?

- A. Defining and managing security policies and controls as code within the development pipeline
- B. Replacing all security hardware appliances with software-only solutions running on general servers
- C. Writing malware using programming languages instead of using automated malware creation tools
- D. Teaching security professionals how to write code for developing custom vulnerability scanners

Answer: A

Q801. What is the primary distinction between a vulnerability scan and a penetration test?

- A. Vulnerability scans are illegal while penetration tests are legal
- B. Vulnerability scans identify weaknesses while penetration tests actively exploit them
- C. Vulnerability scans are more expensive than penetration tests
- D. Penetration tests only check for viruses

Answer: B

Q802. In the context of cyber security, what is 'due diligence'?

- A. Ignoring minor security incidents
- B. The practice of thoroughly investigating and understanding security risks before making decisions
- C. Using only open-source security tools
- D. Delegating all security to third parties

Answer: B

Q803. What distinguishes a 'grey hat' hacker from a 'white hat' and 'black hat' hacker?

- A. Grey hats only hack government systems
- B. Grey hats may find vulnerabilities without authorization but do not have malicious intent
- C. Grey hats only use automated tools
- D. Grey hats exclusively perform social engineering

Answer: B

Q804. What is the purpose of a cyber security maturity model?

- A. To measure internet speed across an organization
- B. To assess an organization's security capabilities and identify areas for improvement
- C. To rank employees by technical skill
- D. To determine which antivirus to purchase

Answer: B

Q805. How does 'defense in depth' differ from relying on a single security control?

- A. Defense in depth uses only physical controls
- B. Defense in depth employs multiple layers of security so that if one fails, others still protect
- C. Single security controls are always more effective
- D. Defense in depth only applies to cloud environments

Answer: B

Q806. What role does asset management play in cyber security?

- A. It only tracks the financial value of equipment
- B. It helps organizations identify what needs protecting and ensures proper security controls are applied
- C. It replaces the need for access controls
- D. It is only relevant for physical security

Answer: B

Q807. What is the difference between a proactive and reactive security approach?

- A. Proactive security responds after an incident while reactive security prevents them
- B. Proactive security anticipates and prevents threats before they occur, while reactive responds after an incident
- C. There is no difference between the two approaches
- D. Reactive security is always more cost-effective

Answer: B

Q808. What is the significance of compliance frameworks like ISO 27001 in cyber security?

- A. They eliminate all cyber threats automatically
- B. They provide standards and guidelines for establishing information security management systems
- C. They are only applicable to government organizations
- D. They replace the need for technical security controls

Answer: B

Q809. What is a security control in the context of cyber security?

- A. A device used to lock server rooms only
- B. A safeguard or countermeasure to avoid, detect, or minimize security risks
- C. A type of malware detection algorithm
- D. A firewall brand name

Answer: B

Q810. Why is third-party risk management important in cyber security?

- A. Third parties never pose security risks
- B. Vendors and partners with access to your systems can introduce vulnerabilities if not properly managed
- C. It is only needed for financial institutions
- D. Third-party risks only affect large enterprises

Answer: B

Q811. What is the difference between a corrective and a compensating security control?

- A. They are identical in function
- B. A corrective control restores systems after an incident, while a compensating control provides an alternative when a primary control cannot be implemented
- C. Corrective controls prevent incidents while compensating controls detect them
- D. Compensating controls are always more expensive than corrective controls

Answer: B

Q812. How does the principle of 'defense in depth' apply to data protection?

- A. By using only one very strong encryption algorithm
- B. By applying multiple layers of security controls such as encryption, access controls, monitoring, and backups
- C. By relying solely on physical security measures
- D. By giving all employees access to all data for transparency

Answer: B

Q813. What is the purpose of a business impact analysis (BIA) in information security?

- A. To determine the salary of security staff
- B. To identify critical business functions and the impact of disruptions on the organization
- C. To select antivirus software vendors
- D. To audit employee attendance records

Answer: B

Q814. What distinguishes a technical control from an administrative control in security?

- A. Technical controls are always more effective
- B. Technical controls are implemented through technology while administrative controls are implemented through policies and procedures
- C. Administrative controls are not used in modern security
- D. Technical controls do not require any maintenance

Answer: B

Q815. What is the concept of 'security by design' in information security?

- A. Adding security features only after deployment
- B. Integrating security considerations into every phase of system design and development from the beginning
- C. Designing systems with no security to keep them simple
- D. Using only open-source security tools in design

Answer: B

Q816. How does quantitative risk analysis differ from qualitative risk analysis?

- A. Quantitative analysis uses numerical values and financial metrics while qualitative uses descriptive categories like high, medium, and low
- B. Qualitative analysis is always more accurate
- C. Quantitative analysis does not consider financial impact
- D. They produce identical results using different terminology

Answer: A

Q817. What is the purpose of a data retention policy?

- A. To keep all data indefinitely for compliance
- B. To define how long data should be stored, when it should be deleted, and how disposal should occur
- C. To prevent any data from being deleted
- D. To maximize storage utilization on servers

Answer: B

Q818. What is the 'need to know' principle and how does it complement least privilege?

- A. It means everyone needs to know everything for transparency
- B. It restricts access to information based on job necessity, working alongside least privilege to limit both permissions and information access
- C. It only applies to classified government information
- D. It replaces the need for access control lists

Answer: B

Q819. What is the role of a security operations center (SOC) in an organization?

- A. To develop new software applications
- B. To centrally monitor, detect, analyze, and respond to security incidents in real time
- C. To manage human resources records
- D. To handle customer billing inquiries

Answer: B

Q820. What is the difference between data at rest, data in transit, and data in use?

- A. They all refer to the same state of data
- B. Data at rest is stored, data in transit is being transmitted, and data in use is actively being processed in memory
- C. Data in use is always more secure than data at rest
- D. Data in transit cannot be encrypted

Answer: B

Q821. What is the primary difference between a block cipher and a stream cipher in terms of data processing?

- A. Block ciphers are always more secure
- B. Block ciphers encrypt fixed-size blocks of data while stream ciphers encrypt data one bit or byte at a time
- C. Stream ciphers cannot use keys
- D. Block ciphers do not use initialization vectors

Answer: B

Q822. What is a certificate revocation list (CRL) used for?

- A. To list all valid certificates in use
- B. To publish a list of digital certificates that have been revoked before their expiration date
- C. To generate new cryptographic keys
- D. To store encrypted email messages

Answer: B

Q823. How does the Diffie-Hellman key exchange enable secure communication without prior shared secrets?

- A. It sends the encryption key in plaintext
- B. It allows two parties to jointly establish a shared secret over an insecure channel using mathematical operations
- C. It requires a trusted third party to distribute keys
- D. It only works on encrypted networks

Answer: B

Q824. What is the purpose of a message digest in cryptography?

- A. To compress messages for faster transmission
- B. To produce a fixed-length fingerprint of data that can verify its integrity
- C. To encrypt messages using symmetric keys
- D. To generate random encryption keys

Answer: B

Q825. What is the Online Certificate Status Protocol (OCSP) and how does it differ from CRL?

- A. OCSP replaces encryption entirely
- B. OCSP provides real-time certificate validation by querying the CA directly, unlike CRLs which are periodically published lists
- C. OCSP is a type of symmetric encryption
- D. OCSP and CRL are identical mechanisms

Answer: B

Q826. Why is AES considered more secure than DES for modern encryption needs?

- A. AES uses shorter keys than DES
- B. AES supports key lengths of 128, 192, and 256 bits while DES uses only 56-bit keys, making DES vulnerable to brute-force attacks
- C. DES is newer than AES
- D. AES is an asymmetric algorithm while DES is symmetric

Answer: B

Q827. What is key escrow and what security concern does it raise?

- A. Key escrow means destroying keys after use
- B. Key escrow involves storing copies of encryption keys with a trusted third party, raising concerns about unauthorized access to those keys
- C. Key escrow is a method of key generation
- D. Key escrow only applies to physical locks

Answer: B

Q828. What is a chosen-plaintext attack in cryptanalysis?

- A. An attack where the attacker can only observe ciphertext
- B. An attack where the attacker can choose arbitrary plaintexts and obtain their corresponding ciphertexts to deduce the key
- C. An attack that only targets hash functions
- D. An attack that requires physical access to the encryption device

Answer: B

Q829. What is the role of an initialization vector (IV) in encryption?

- A. It replaces the encryption key
- B. It adds randomness to the encryption process so that identical plaintexts produce different ciphertexts
- C. It determines the key length
- D. It is used only in hash functions

Answer: B

Q830. What is envelope encryption and where is it commonly used?

- A. Encrypting physical mail envelopes
- B. A technique where data is encrypted with a data key, and the data key is then encrypted with a master key, commonly used in cloud services
- C. Using two different symmetric algorithms simultaneously
- D. Encrypting only email attachments

Answer: B

Q831. What is the difference between a network-based IDS and a host-based IDS?

- A. They monitor the same traffic in the same way
- B. A network-based IDS monitors network traffic on a network segment, while a host-based IDS monitors activity on an individual host system
- C. Host-based IDS is always more effective
- D. Network-based IDS can only monitor wireless traffic

Answer: B

Q832. How does network address translation (NAT) provide a layer of security?

- A. NAT encrypts all outgoing traffic
- B. NAT hides internal IP addresses from external networks, making it harder for attackers to directly target internal devices
- C. NAT replaces the need for firewalls entirely
- D. NAT increases network bandwidth

Answer: B

Q833. What is the purpose of a network demilitarized zone (DMZ) architecture?

- A. To eliminate the need for internal firewalls
- B. To place public-facing services in an isolated network segment between the external and internal networks
- C. To directly expose internal servers to the internet
- D. To increase internal network speeds

Answer: B

Q834. What is SSL stripping and how does it compromise network security?

- A. Removing physical cables from network equipment
- B. A man-in-the-middle technique that downgrades HTTPS connections to HTTP, exposing data in plaintext
- C. A method of improving SSL performance
- D. A technique for updating SSL certificates

Answer: B

Q835. What is the purpose of network flow analysis in security monitoring?

- A. To increase network throughput
- B. To collect and analyze metadata about network connections (source, destination, ports, protocols, volume) to detect anomalies and threats
- C. To encrypt all network traffic
- D. To replace packet capture entirely

Answer: B

Q836. How does a proxy firewall (application-level gateway) differ from a packet-filtering firewall?

- A. Proxy firewalls are always slower and less secure
- B. A proxy firewall inspects traffic at the application layer and acts as an intermediary, while a packet filter only examines packet headers
- C. Packet-filtering firewalls provide more granular inspection
- D. Proxy firewalls cannot filter web traffic

Answer: B

Q837. What is MAC address spoofing and why is it a security concern?

- A. It is a method of encrypting MAC addresses
- B. An attacker changes their device's MAC address to impersonate another device, potentially bypassing MAC-based access controls
- C. MAC addresses cannot be changed
- D. MAC spoofing only affects wireless networks

Answer: B

Q838. What is the security benefit of implementing network segmentation with micro-segmentation?

- A. It eliminates the need for any firewalls
- B. It limits lateral movement by creating granular security zones around individual workloads, even within the same network segment
- C. Micro-segmentation only applies to cloud environments
- D. It increases network latency without security benefits

Answer: B

Q839. What is DNS cache poisoning and what risk does it pose?

- A. Clearing the DNS cache for performance
- B. An attack that inserts fraudulent DNS entries into a resolver's cache, redirecting users to malicious websites
- C. A method of improving DNS performance
- D. DNS caches cannot be manipulated

Answer: B

Q840. What is the purpose of implementing 802.1Q VLAN tagging?

- A. To encrypt VLAN traffic
- B. To identify which VLAN a frame belongs to as it traverses trunk links between switches, enabling logical network separation
- C. To replace IP addressing on the network
- D. To increase the maximum number of physical ports on a switch

Answer: B

Q841. What is the primary difference between CBC and CTR modes of operation in block ciphers?

- A. CBC encrypts blocks independently while CTR chains them together
- B. CTR mode turns a block cipher into a stream cipher allowing parallel encryption
- C. CBC is faster than CTR in all cases
- D. CTR requires a larger key size than CBC

Answer: B

Q842. What is the purpose of key derivation functions like HKDF?

- A. To compress encryption keys for storage
- B. To derive one or more cryptographic keys from a source of key material
- C. To generate public-private key pairs
- D. To verify digital signatures

Answer: B

Q843. How does authenticated encryption differ from standard encryption?

- A. Authenticated encryption is slower but uses smaller keys
- B. Authenticated encryption requires two separate algorithms always
- C. Authenticated encryption provides both confidentiality and integrity/authenticity in a single operation
- D. Authenticated encryption only works with asymmetric algorithms

Answer: C

Q844. What is the security implication of using a predictable initialization vector (IV)?

- A. It has no security impact
- B. It can enable chosen-plaintext attacks and leak information about encrypted data
- C. It only affects performance, not security
- D. It makes decryption impossible

Answer: B

Q845. What is Galois/Counter Mode (GCM) and why is it preferred for modern encryption?

- A. A hash function that replaces SHA-256
- B. An authenticated encryption mode providing both confidentiality and integrity with high performance
- C. A key exchange protocol replacing Diffie-Hellman
- D. A digital signature algorithm

Answer: B

Q846. What is the purpose of a cryptographic nonce and how does it differ from an IV?

- A. They are identical concepts with no difference
- B. A nonce is used only once and may not need to be random, while an IV typically needs to be unpredictable
- C. A nonce is always longer than an IV
- D. An IV is used in symmetric encryption while a nonce is for asymmetric only

Answer: B

Q847. What is the Diffie-Hellman key exchange vulnerable to without authentication?

- A. Brute force attacks
- B. Man-in-the-middle attacks
- C. Replay attacks
- D. Dictionary attacks

Answer: B

Q848. What is the role of a Merkle tree in cryptography?

- A. To store encryption keys hierarchically
- B. To efficiently verify data integrity using a tree of hashes
- C. To generate random numbers
- D. To perform key exchange between multiple parties

Answer: B

Q849. Why is SHA-1 considered deprecated for security purposes?

- A. It produces output that is too long
- B. It is too slow for modern systems
- C. Practical collision attacks have been demonstrated against it
- D. It only works with symmetric encryption

Answer: C

Q850. What is the concept of crypto-agility?

- A. The ability to switch between encryption algorithms quickly without major system changes
- B. Using multiple encryption algorithms simultaneously on all data
- C. Encrypting data faster than real-time
- D. The ability to break any encryption algorithm

Answer: A

Q851. What is the role of AppArmor in Linux security?

- A. A firewall management tool
- B. A mandatory access control system that restricts programs' capabilities based on per-program profiles
- C. A disk encryption utility
- D. A network monitoring tool

Answer: B

Q852. What is the difference between discretionary access control (DAC) and role-based access control (RBAC)?

- A. DAC and RBAC are identical
- B. In DAC, the resource owner sets permissions; in RBAC, permissions are assigned based on organizational roles
- C. DAC is more secure than RBAC in all cases
- D. RBAC allows users to share permissions freely

Answer: B

Q853. What is a privilege escalation attack and what are its two types?

- A. An attack that only affects network devices, divided into internal and external
- B. An attack that gains elevated access, divided into horizontal (accessing other users' resources) and vertical (gaining higher privileges)
- C. An attack that targets passwords only, divided into online and offline
- D. An attack on physical security, divided into tailgating and piggybacking

Answer: B

Q854. What is Windows BitLocker and what does it protect against?

- A. A network firewall for Windows
- B. Full disk encryption that protects data on lost or stolen devices by encrypting the entire volume
- C. An antivirus program
- D. A password manager

Answer: B

Q855. How does the Linux sudo command enhance security compared to logging in as root?

- A. sudo is less secure than root login
- B. sudo provides temporary elevated privileges with logging and accountability, reducing the risk of accidental system-wide damage
- C. sudo grants permanent root access
- D. sudo disables all security features

Answer: B

Q856. What is the purpose of System Integrity Protection (SIP) in macOS?

- A. To encrypt user files
- B. To restrict the root account from modifying protected system files and processes
- C. To manage network connections
- D. To install system updates

Answer: B

Q857. What are cgroups in Linux and how do they contribute to security?

- A. A type of firewall rule
- B. Control groups that limit, account for, and isolate the resource usage of process groups
- C. A user authentication mechanism
- D. A file encryption method

Answer: B

Q858. What is the security purpose of UEFI Secure Boot?

- A. To speed up the boot process
- B. To ensure only trusted, digitally signed software can run during the boot process, preventing bootkits
- C. To encrypt the hard drive
- D. To manage user accounts at boot time

Answer: B

Q859. What is the purpose of audit policies in Windows operating systems?

- A. To track software licenses
- B. To define which security events are recorded in the event log for monitoring and forensics
- C. To manage desktop appearances
- D. To control internet bandwidth

Answer: B

Q860. How does memory protection through W^X (Write XOR Execute) policy improve OS security?

- A. It doubles available memory
- B. It ensures memory pages are either writable or executable but never both, preventing code injection attacks
- C. It encrypts all memory contents
- D. It speeds up memory access

Answer: B

Q861. What is the security purpose of the X-Content-Type-Options header?

- A. To encrypt content between browser and server
- B. To prevent browsers from MIME-sniffing a response away from the declared content-type, stopping content-type confusion attacks
- C. To compress HTTP responses
- D. To cache web pages more effectively

Answer: B

Q862. How does a web application firewall (WAF) differ from a traditional network firewall?

- A. They are identical in function
- B. A WAF inspects HTTP/HTTPS traffic at the application layer to detect attacks like XSS and SQL injection, while network firewalls operate at lower layers
- C. A WAF only blocks IP addresses
- D. Network firewalls can inspect application-layer content better

Answer: B

Q863. What is the security risk of exposing source maps in a production web application?

- A. Source maps slow down the server
- B. Source maps reveal original source code structure and logic, aiding attackers in finding vulnerabilities
- C. Source maps corrupt the database
- D. Source maps have no security impact

Answer: B

Q864. What is a Server-Side Template Injection (SSTI) attack?

- A. Injecting CSS into server templates
- B. Injecting malicious code into server-side template engines that gets executed on the server, potentially leading to remote code execution
- C. Injecting JavaScript into the client browser
- D. Modifying HTML templates in a text editor

Answer: B

Q865. What is the purpose of the SameSite cookie attribute and what values can it take?

- A. It determines cookie expiration time with values of Short, Medium, and Long
- B. It controls when cookies are sent with cross-site requests, with values Strict, Lax, and None, helping prevent CSRF attacks
- C. It sets the cookie's encryption strength
- D. It determines which users can see the cookie

Answer: B

Q866. What is an XML External Entity (XXE) attack?

- A. A cross-site scripting attack using XML
- B. An attack exploiting XML parsers that process external entity references, potentially leading to file disclosure, SSRF, or denial of service
- C. An attack that replaces XML with JSON
- D. A SQL injection through XML queries

Answer: B

Q867. What is the security concern with JWT tokens stored in localStorage?

- A. localStorage is too small for JWT tokens
- B. JWTs in localStorage are accessible to JavaScript and vulnerable to XSS attacks, unlike HttpOnly cookies
- C. localStorage automatically encrypts JWTs
- D. There are no concerns with this approach

Answer: B

Q868. What is a subdomain takeover vulnerability?

- A. When a subdomain loads too slowly
- B. When an attacker claims a subdomain that points to an external service the organization no longer uses, enabling phishing or cookie theft
- C. When a subdomain uses a different TLS certificate
- D. When subdomains share the same database

Answer: B

Q869. How does Content Security Policy (CSP) report-only mode help in deploying security headers?

- A. It blocks all violations immediately
- B. It logs policy violations without enforcing them, allowing developers to test and refine CSP rules before enforcement
- C. It only works in development environments
- D. It disables all scripts on the page

Answer: B

Q870. What is a broken access control vulnerability in web applications?

- A. When the login page does not load properly
- B. When users can access resources or perform actions beyond their intended permissions due to improper authorization checks
- C. When CSS styles are not applied correctly
- D. When the server cannot handle multiple users

Answer: B

Q871. What is credential stuffing and how does it differ from brute-force attacks?

- A. They are the same type of attack
- B. Credential stuffing uses stolen username-password pairs from data breaches to try on other services, while brute-force tries random combinations
- C. Credential stuffing only targets email accounts
- D. Brute-force uses stolen credentials while credential stuffing does not

Answer: B

Q872. How does TOTP (Time-Based One-Time Password) authentication work?

- A. It sends a password via email each time
- B. It uses a shared secret and current time to generate a unique code that is valid for a short period, typically 30 seconds
- C. It uses facial recognition technology
- D. It generates a permanent second password

Answer: B

Q873. What is the difference between RBAC and ABAC in terms of flexibility?

- A. RBAC is more flexible than ABAC
- B. ABAC is more flexible because it can use multiple attributes (user, resource, environment, action) for access decisions, while RBAC is based only on predefined roles
- C. They are equally flexible
- D. Neither is flexible

Answer: B

Q874. What is the purpose of a challenge-response authentication protocol?

- A. To test the user's typing speed
- B. To verify identity by requiring the client to correctly respond to a random challenge from the server, proving knowledge without transmitting the secret
- C. To challenge users with security questions only
- D. To respond to denial-of-service attacks

Answer: B

Q875. What is identity federation and what problem does it solve?

- A. Creating multiple identities for one user
- B. Allowing users to authenticate across multiple organizations using a single identity from a trusted identity provider, reducing credential sprawl
- C. Combining all passwords into one
- D. Disabling authentication for external users

Answer: B

Q876. What is the security risk of using SMS-based two-factor authentication?

- A. SMS 2FA has no security risks
- B. SMS messages can be intercepted through SIM swapping, SS7 protocol vulnerabilities, or phone number porting, compromising the second factor
- C. SMS 2FA is the most secure method available
- D. SMS 2FA only works on smartphones

Answer: B

Q877. What is the concept of separation of duties in access control?

- A. Giving all duties to one person
- B. Dividing critical tasks among multiple users so no single person has complete control over a sensitive process, preventing fraud and errors
- C. Separating work and personal accounts
- D. Dividing the network into segments

Answer: B

Q878. What is an access token in the context of API authentication?

- A. A physical token used to enter a building
- B. A digital credential that represents the authorization granted to a client to access specific API resources on behalf of a user
- C. A password stored in plain text
- D. A certificate installed on the server

Answer: B

Q879. How does multi-factor authentication using a hardware security key work?

- A. The key stores the user's password
- B. The key uses public-key cryptography to create a unique credential per site, signing challenges without transmitting secrets
- C. The key generates random passwords
- D. The key connects to the internet to verify identity

Answer: B

Q880. What is the purpose of an identity governance and administration (IGA) platform?

- A. To manage website content
- B. To automate identity lifecycle management, access certifications, and policy enforcement across an organization
- C. To encrypt all user data
- D. To monitor network bandwidth

Answer: B

Q881. What is the DREAD threat rating model used for?

- A. Ranking programming languages by security
- B. Assessing and prioritizing threats based on Damage, Reproducibility, Exploitability, Affected Users, and Discoverability
- C. Rating the performance of security tools
- D. Evaluating developer skill levels

Answer: B

Q882. What is the purpose of a security code review and how does it differ from a regular code review?

- A. They are identical processes
- B. A security code review specifically focuses on identifying security vulnerabilities, insecure patterns, and compliance issues in addition to regular code quality concerns
- C. Security code reviews only look at encryption code
- D. Regular code reviews already cover all security concerns

Answer: B

Q883. What is the principle of fail-secure versus fail-open in software design?

- A. They mean the same thing
- B. Fail-secure denies access when a failure occurs, maintaining security; fail-open allows access during failures, prioritizing availability over security
- C. Fail-secure means the system never fails
- D. Fail-open means the system is always secure

Answer: B

Q884. How does a software composition analysis (SCA) tool help secure development?

- A. By writing code automatically
- B. By identifying known vulnerabilities in open-source dependencies and third-party libraries used in the project
- C. By testing network security
- D. By managing development team schedules

Answer: B

Q885. What is the security benefit of implementing Content Security Policy headers during development?

- A. CSP speeds up page loading
- B. CSP restricts which resources the browser can load, reducing the impact of XSS attacks by preventing execution of unauthorized scripts
- C. CSP encrypts HTTP traffic
- D. CSP manages user sessions

Answer: B

Q886. What is the role of a CI/CD security gate and what checks should it include?

- A. It blocks all deployments permanently
- B. It enforces automated security checks in the build pipeline including SAST, dependency scanning, and secret detection before code can be deployed
- C. It only checks code formatting
- D. It is a physical security checkpoint

Answer: B

Q887. What is the OWASP Application Security Verification Standard (ASVS)?

- A. A programming language standard
- B. A framework providing a basis for testing web application security controls through three levels of security verification requirements
- C. A network protocol standard
- D. A hardware security specification

Answer: B

Q888. What is the security concern with using eval() or similar dynamic code execution functions?

- A. eval() is too slow for production use
- B. eval() executes strings as code, and if user input reaches eval(), it enables arbitrary code execution attacks
- C. eval() only works in specific languages
- D. eval() causes memory leaks

Answer: B

Q889. What is the security purpose of input allowlisting versus blocklisting?

- A. Blocklisting is always more secure than allowlisting
- B. Allowlisting defines what input IS permitted (more secure), while blocklisting defines what IS NOT permitted (prone to bypasses through new attack patterns)
- C. Allowlisting blocks all input
- D. They provide identical security

Answer: B

Q890. What is the purpose of security regression testing in the SDLC?

- A. Testing if the application runs on older systems
- B. Verifying that previously fixed security vulnerabilities have not been reintroduced by new code changes
- C. Testing the application's speed under load
- D. Checking if old features still work

Answer: B

Q891. What is the security difference between WPA2-Personal and WPA2-Enterprise?

- A. There is no security difference
- B. WPA2-Personal uses a shared passphrase for all users while WPA2-Enterprise uses individual credentials authenticated through a RADIUS server, providing better accountability
- C. WPA2-Personal is more secure
- D. WPA2-Enterprise is only for home networks

Answer: B

Q892. What is a captive portal and what security risks does it present?

- A. A portal that captures all network traffic permanently
- B. A web page that requires user interaction before granting network access; risks include credential theft on fake portals and unencrypted initial connections
- C. A firewall configuration tool
- D. A type of network encryption

Answer: B

Q893. How does mobile application sandboxing protect the operating system?

- A. It runs all apps in a virtual sandbox visual environment
- B. Each app runs in its own isolated environment with restricted access to system resources and other apps' data, containing any malicious behavior
- C. It prevents apps from using the internet
- D. It only protects pre-installed apps

Answer: B

Q894. What is a wireless deauthentication attack and how does it work?

- A. An attack that changes WiFi passwords
- B. An attack that sends forged deauthentication frames to disconnect clients from a WiFi network, often used as a precursor to evil twin attacks
- C. An attack that disables Bluetooth only
- D. An attack that overheats wireless routers

Answer: B

Q895. What is the security purpose of mobile device attestation?

- A. Verifying the user's identity
- B. Cryptographically verifying the integrity and security state of a mobile device (that it has not been rooted, jailbroken, or tampered with) before granting access to resources
- C. Attesting to the device's warranty status
- D. Checking the device's battery health

Answer: B

Q896. What is the role of 802.1X in wireless network security?

- A. It defines WiFi frequency bands
- B. It provides port-based network access control requiring devices to authenticate before being granted network access, using EAP methods and a RADIUS server
- C. It sets maximum WiFi speeds
- D. It determines wireless signal range

Answer: B

Q897. What are the security implications of using public USB charging stations for mobile devices?

- A. Public USB chargers are always safe
- B. They pose a risk of 'juice jacking' where modified charging stations can install malware or steal data through the USB data connection
- C. They charge devices faster than private chargers
- D. They only affect older devices

Answer: B

Q898. What is WiFi Protected Setup (WPS) and why is it considered a security risk?

- A. WPS provides the highest WiFi security available
- B. WPS simplifies WiFi setup but its PIN-based method is vulnerable to brute-force attacks due to a design flaw that splits the 8-digit PIN into two halves
- C. WPS only works with WPA3
- D. WPS is required for all modern routers

Answer: B

Q899. How does a mobile VPN protect data on public WiFi networks?

- A. It makes the WiFi network faster
- B. It creates an encrypted tunnel between the mobile device and the VPN server, protecting all data from eavesdropping on the local network
- C. It blocks all incoming calls
- D. It only protects web browsing, not apps

Answer: B

Q900. What is the purpose of MAC address randomization on mobile devices?

- A. To improve WiFi speed
- B. To prevent tracking of the device across different WiFi networks by changing the hardware address used during scanning and connecting
- C. To authenticate to WiFi networks
- D. To encrypt wireless traffic

Answer: B

Q901. What is the principle of shared responsibility in cloud security and how does it vary by service model?

- A. The cloud provider is responsible for all security
- B. In IaaS, customers manage OS and application security; in PaaS, the provider manages the platform; in SaaS, the provider handles nearly everything except user access and data classification
- C. The customer is responsible for all security
- D. Shared responsibility means both parties do the same tasks

Answer: B

Q902. What is the security risk of overly permissive S3 bucket policies in AWS?

- A. S3 buckets cannot be made public
- B. Overly permissive policies can expose sensitive data to the public internet, leading to data breaches as seen in numerous high-profile incidents
- C. Permissive policies improve performance
- D. S3 bucket policies do not affect security

Answer: B

Q903. What is cloud-native logging and monitoring and why is it critical for security?

- A. It is the same as traditional on-premises logging
- B. Cloud-native logging aggregates logs from distributed cloud services (API calls, access events, resource changes) for security monitoring, threat detection, and compliance auditing
- C. It only monitors server uptime
- D. It replaces the need for security controls

Answer: B

Q904. What is a service mesh in cloud-native architecture and how does it enhance security?

- A. A physical network of cloud servers
- B. An infrastructure layer that handles service-to-service communication with built-in security features like mutual TLS, access policies, and traffic encryption without application code changes
- C. A type of cloud storage
- D. A mesh WiFi network for cloud data centers

Answer: B

Q905. What are cloud access security brokers (CASBs) and what four pillars of functionality do they provide?

- A. CASBs only block cloud access
- B. CASBs provide visibility, data security, threat protection, and compliance for cloud services by acting as intermediaries between users and cloud providers
- C. CASBs are a type of cloud storage
- D. CASBs only monitor bandwidth usage

Answer: B

Q906. What is infrastructure as code (IaC) security scanning and why is it important?

- A. It scans physical infrastructure for defects
- B. It analyzes IaC templates (Terraform, CloudFormation) for security misconfigurations before deployment, preventing insecure infrastructure from being provisioned
- C. It scans code for programming errors only
- D. It is only used for cost optimization

Answer: B

Q907. What is the blast radius concept in cloud architecture and how can it be minimized?

- A. The physical distance a cloud data center's explosion can reach
- B. The extent of damage that can result from a security incident; minimized through network segmentation, separate accounts, least privilege, and micro-segmentation
- C. The WiFi signal range of cloud servers
- D. The number of users affected by a cloud outage

Answer: B

Q908. What is the security purpose of a cloud-native web application firewall versus a traditional WAF?

- A. They are identical in functionality
- B. Cloud-native WAFs are designed for dynamic cloud environments with auto-scaling, API protection, bot management, and integration with cloud provider services and DevOps pipelines
- C. Cloud-native WAFs are less capable than traditional WAFs
- D. Traditional WAFs work better in cloud environments

Answer: B

Q909. What is a landing zone in cloud architecture and how does it relate to security?

- A. A physical location where cloud servers are installed
- B. A pre-configured, secure baseline environment with account structure, networking, IAM, logging, and security controls that serves as the foundation for cloud deployments
- C. A backup storage location
- D. A cloud provider's customer support portal

Answer: B

Q910. What is the security risk of using long-lived access keys for cloud service accounts?

- A. Long-lived keys are the most secure option
- B. Long-lived access keys that are not rotated can be leaked or stolen, providing persistent access; short-lived credentials and role-based access with temporary tokens are preferred
- C. Long-lived keys automatically rotate
- D. There is no risk with long-lived keys

Answer: B

Q911. What is the role of timestamps (MAC times) in forensic timeline reconstruction?

- A. Timestamps are unreliable and never used
- B. Modified, Accessed, and Created timestamps on files help establish when events occurred, reconstruct user activities, and build a chronological timeline of events
- C. Timestamps only show the current date
- D. MAC times refer to Apple computer forensics only

Answer: B

Q912. What is the difference between volatile and non-volatile evidence, and why does collection order matter?

- A. All evidence has the same priority for collection
- B. Volatile evidence (RAM, running processes, network connections) is lost when power is removed and must be collected first; non-volatile evidence (hard drives) persists and can be collected later
- C. Non-volatile evidence must always be collected first
- D. The collection order does not affect the investigation

Answer: B

Q913. How does file carving work and when is it necessary in forensic analysis?

- A. File carving is a wood-working term not related to forensics
- B. File carving recovers files from unallocated disk space by searching for file signatures (headers and footers) without relying on file system metadata, useful when files are deleted or file system is damaged
- C. File carving only works on image files
- D. File carving requires the original file names

Answer: B

Q914. What is a forensic triage and when is it performed?

- A. The final step of a forensic investigation
- B. A rapid preliminary assessment of digital evidence to identify the most relevant devices and data, prioritize investigation efforts, and determine the scope of a full forensic examination
- C. A medical procedure unrelated to forensics
- D. A method of permanently deleting evidence

Answer: B

Q915. What is registry analysis in Windows forensics and what type of evidence can it reveal?

- A. Registry analysis only shows software license keys
- B. Windows Registry analysis reveals system configuration, recently accessed files, USB device history, installed software, user accounts, network connections, and user activity patterns
- C. The registry contains no useful forensic data
- D. Registry analysis is only possible on Linux systems

Answer: B

Q916. What is the purpose of using EnCase or FTK in digital forensics?

- A. They are video editing software
- B. They are comprehensive digital forensic suites that provide evidence acquisition, analysis, search, bookmark, and reporting capabilities for forensic investigations
- C. They are antivirus programs
- D. They are network monitoring tools

Answer: B

Q917. What is network packet capture (PCAP) analysis in forensic investigations?

- A. A method of capturing physical evidence
- B. Analysis of recorded network traffic data to identify communications, data transfers, malware command-and-control traffic, and evidence of network-based attacks
- C. A technique for capturing screenshots
- D. A method of printing network diagrams

Answer: B

Q918. What is the significance of examining browser artifacts in a forensic investigation?

- A. Browser artifacts only show bookmarks
- B. Browser artifacts including history, cookies, cache, downloads, saved passwords, autofill data, and session data can reveal user activities, communications, and intent
- C. Browser data is automatically deleted and cannot be recovered
- D. Only the last visited website can be determined

Answer: B

Q919. What is the role of hash sets in forensic analysis?

- A. Hash sets encrypt evidence for storage
- B. Known hash sets (like NSRL) help forensic investigators quickly identify and filter known-good files (operating system files) and known-bad files (malware) from the evidence
- C. Hash sets are random number generators
- D. Hash sets delete duplicate files

Answer: B

Q920. What is the difference between a forensic investigation for civil versus criminal cases?

- A. There is no difference in forensic methodology
- B. Criminal cases require higher standards of evidence (beyond reasonable doubt), stricter chain of custody, and may involve law enforcement; civil cases use preponderance of evidence and may be conducted by private firms
- C. Civil cases require more evidence than criminal cases
- D. Digital forensics is only used in criminal cases

Answer: B

Q921. What is the role of threat intelligence in improving incident response capabilities?

- A. Threat intelligence is not useful for incident response
- B. Threat intelligence provides context about adversary tactics, techniques, and procedures (TTPs) that helps IR teams detect, attribute, and respond to incidents more effectively
- C. Threat intelligence only helps with marketing
- D. Threat intelligence replaces the need for incident response

Answer: B

Q922. How do tabletop exercises improve incident response preparedness?

- A. They test physical security of table-mounted servers
- B. They simulate incident scenarios through discussion-based exercises, testing communication, decision-making, and plan effectiveness without disrupting actual systems
- C. They are physical fitness exercises for the IR team
- D. They only test technical controls

Answer: B

Q923. What is the importance of establishing communication channels before an incident occurs?

- A. Communication planning is not important for incident response
- B. Pre-established out-of-band communication channels ensure the IR team can communicate securely during an incident, even if the attacker has compromised normal communication systems
- C. Only email communication is needed
- D. Communication should only be planned during the incident

Answer: B

Q924. What is the difference between eradication and remediation in incident response?

- A. They are identical concepts
- B. Eradication removes the immediate threat (malware, attacker access) while remediation addresses the underlying vulnerability or weakness that allowed the incident to occur
- C. Eradication is a permanent solution while remediation is temporary
- D. Remediation happens before eradication

Answer: B

Q925. What is a cyber incident severity classification system and why is it important?

- A. It classifies incidents by the day of the week they occur
- B. A framework that categorizes incidents by impact level to determine response priority, resource allocation, escalation requirements, and notification obligations
- C. It only applies to government organizations
- D. Severity classification slows down the response

Answer: B

Q926. How should evidence be handled during incident response to ensure it is admissible in court?

- A. Evidence handling procedures do not matter for legal proceedings
- B. Evidence must be collected following proper procedures including maintaining chain of custody, using write-protected copies, documenting all actions, and preserving original evidence integrity
- C. Only digital copies of evidence are needed in court
- D. Evidence should be modified to make it clearer for the judge

Answer: B

Q927. What is the role of a crisis communications plan in incident response?

- A. It handles technical remediation steps
- B. It defines how the organization communicates about the incident to stakeholders, media, customers, and regulators, managing reputation and meeting notification obligations
- C. It is only needed for public companies
- D. Communication is not part of incident response

Answer: B

Q928. What is the purpose of maintaining an asset inventory for incident response?

- A. To track office furniture
- B. An up-to-date inventory of all IT assets helps IR teams quickly identify affected systems, understand their criticality, determine blast radius, and prioritize response efforts
- C. Asset inventories are only for accounting
- D. It is not relevant to incident response

Answer: B

Q929. What is the concept of 'containment before eradication' in incident response?

- A. Eradication should always come before containment
- B. The incident must be contained (preventing further spread or damage) before eradication begins, because premature eradication may alert the attacker and lose evidence
- C. Containment and eradication happen simultaneously always
- D. Neither containment nor eradication is necessary

Answer: B

Q930. What is log normalization and why is it important for SIEM effectiveness?

- A. Log normalization deletes unnecessary log entries
- B. Log normalization converts logs from different sources into a consistent format, enabling effective correlation, analysis, and searching across heterogeneous data sources in a SIEM
- C. Log normalization encrypts log data
- D. Normalization is only needed for network logs

Answer: B

Q931. What is the MITRE ATT&CK framework's role in creating SIEM detection rules?

- A. MITRE ATT&CK is not related to SIEM detection
- B. ATT&CK provides a comprehensive matrix of adversary tactics and techniques that can be mapped to specific SIEM detection rules, helping identify coverage gaps and prioritize detection development
- C. ATT&CK replaces SIEM systems entirely
- D. ATT&CK only covers network-level attacks

Answer: B

Q932. How does network flow analysis (NetFlow/sFlow) complement full packet capture for security monitoring?

- A. Flow analysis and packet capture provide identical information
- B. Flow analysis provides metadata about communications (who talked to whom, when, how much data) at scale, while full packet capture provides content details; together they provide both breadth and depth of visibility
- C. Flow analysis replaces the need for packet capture
- D. Packet capture is always superior to flow analysis

Answer: B

Q933. What is the concept of security monitoring use cases and how should they be prioritized?

- A. Use cases are user stories for software development
- B. Security monitoring use cases define specific threat scenarios to detect (e.g., brute force, lateral movement), and should be prioritized based on threat likelihood, business impact, and available log sources
- C. All use cases should be implemented simultaneously
- D. Use cases are only needed for compliance

Answer: B

Q934. What is the difference between agent-based and agentless monitoring for endpoints?

- A. They produce identical results with no trade-offs
- B. Agent-based monitoring installs software on endpoints for deep visibility and real-time detection, while agentless monitoring uses network traffic or remote queries with less endpoint impact but potentially less visibility
- C. Agent-based monitoring is always less secure
- D. Agentless monitoring provides better endpoint visibility

Answer: B

Q935. What is a SIEM correlation rule and how does it reduce false positives?

- A. Correlation rules only forward raw logs
- B. Correlation rules combine multiple individual events across different sources to identify attack patterns, reducing false positives by requiring multiple conditions to be met before triggering an alert
- C. Correlation rules increase false positives
- D. Correlation is only used for compliance reporting

Answer: B

Q936. What is threat hunting and how does it differ from traditional alert-based security monitoring?

- A. Threat hunting and alert-based monitoring are identical
- B. Threat hunting is a proactive approach where analysts actively search for hidden threats that may have evaded automated detection, using hypotheses and data analysis rather than waiting for alerts
- C. Threat hunting only looks for known malware signatures
- D. Threat hunting replaces the need for SIEM

Answer: B

Q937. What is the purpose of a SOC analyst playbook for specific alert types?

- A. Playbooks are scripts that automatically fix all security issues
- B. Playbooks provide standardized investigation and response procedures for specific alert types, ensuring consistent analysis, reducing response time, and enabling junior analysts to handle complex alerts
- C. Playbooks are only used for documentation
- D. Playbooks replace the need for skilled analysts

Answer: B

Q938. What is the significance of DNS monitoring in security operations?

- A. DNS monitoring only measures internet speed
- B. DNS monitoring detects malicious domains, DNS tunneling for data exfiltration, domain generation algorithms (DGA) used by malware, and DNS-based command-and-control communications
- C. DNS is too basic to be useful for security
- D. DNS monitoring only applies to web servers

Answer: B

Q939. What is the role of data loss prevention (DLP) in security monitoring?

- A. DLP only prevents accidental file deletion
- B. DLP monitors and controls data movement to detect and prevent unauthorized transmission of sensitive data outside the organization, covering email, web, endpoints, and cloud services
- C. DLP replaces encryption
- D. DLP only works for physical documents

Answer: B

Q940. What is the extraterritorial jurisdiction challenge in cybercrime prosecution?

- A. Cybercrime only occurs within one country
- B. Cybercrimes often cross national boundaries, creating challenges in determining which country's laws apply, obtaining evidence from foreign jurisdictions, and securing cooperation between different legal systems
- C. Extraterritorial jurisdiction is not relevant to cybercrime
- D. All countries have identical cybercrime laws

Answer: B

Q941. What is the difference between data controller and data processor under data protection laws?

- A. They are the same entity with different names
- B. A data controller determines the purposes and means of data processing, while a data processor processes data on behalf of the controller according to their instructions
- C. A data processor has more authority than a controller
- D. Only data controllers are regulated by law

Answer: B

Q942. What is the principle of data minimization in privacy law?

- A. Minimizing the quality of data collected
- B. Organizations should only collect and process the minimum amount of personal data necessary for the specified purpose, avoiding excessive data collection
- C. Collecting as much data as possible for future use
- D. Minimizing the storage cost of data

Answer: B

Q943. What obligations do organizations have under breach notification laws?

- A. Organizations have no obligation to disclose breaches
- B. Organizations must notify affected individuals and/or regulatory authorities within specified timeframes when personal data breaches occur, with details about the breach and mitigation measures
- C. Only breaches involving financial data require notification
- D. Notification is only required if data is published publicly

Answer: B

Q944. What is the legal concept of reasonable security measures in cybersecurity?

- A. There is no legal standard for security measures
- B. The legal standard requiring organizations to implement security measures appropriate to the sensitivity of data they handle and the risks they face, considered in negligence and liability cases
- C. Only government agencies must implement security measures
- D. Reasonable security means using the most expensive solutions

Answer: B

Q945. What is the safe harbor provision in cybersecurity law?

- A. A physical safe for storing digital evidence
- B. Legal protections that shield organizations from liability if they follow specific security frameworks, standards, or practices as defined by applicable law
- C. A harbor where cyber criminals are safe from prosecution
- D. A type of data encryption standard

Answer: B

Q946. What are the legal implications of employee monitoring in the workplace?

- A. Employers can monitor employees without any restrictions
- B. Employee monitoring must balance the employer's legitimate security interests with employee privacy rights, often requiring notice, consent, and proportionality under applicable labor and privacy laws
- C. Employee monitoring is illegal in all jurisdictions
- D. Only IT staff can be monitored

Answer: B

Q947. What is the legal status of ethical hacking and penetration testing?

- A. All hacking is illegal regardless of intent
- B. Ethical hacking and penetration testing are legal when performed with explicit written authorization from the system owner, defining scope, methods, and timeframes
- C. Ethical hackers never need permission
- D. Penetration testing is illegal in all countries

Answer: B

Q948. What are the regulatory requirements of PCI DSS for organizations handling payment card data?

- A. PCI DSS is optional for all organizations
- B. PCI DSS mandates security requirements including network segmentation, encryption of cardholder data, access controls, vulnerability management, and regular security testing for organizations that process, store, or transmit payment card data
- C. PCI DSS only applies to banks
- D. PCI DSS only covers physical card transactions

Answer: B

Q949. What is the role of the Federal Investigation Agency (FIA) in cybercrime enforcement in Pakistan?

- A. FIA has no authority over cybercrimes
- B. FIA's National Response Centre for Cyber Crime (NR3C) is responsible for investigating and prosecuting cybercrimes under PECA 2016, including unauthorized access, data theft, online fraud, and cyber harassment
- C. FIA only handles physical crimes
- D. FIA delegates all cybercrime cases to provincial police

Answer: B

Q950. How does Secure Access Service Edge (SASE) architecture change traditional network security?

- A. SASE replaces the internet with private networks
- B. SASE converges network security functions (firewall, CASB, ZTNA, SWG) with WAN capabilities into a cloud-delivered service, providing security at the edge closer to users and devices regardless of location
- C. SASE is a type of VPN
- D. SASE only works for on-premises networks

Answer: B

Q951. What is the security concern with generative AI tools like large language models in enterprise environments?

- A. Generative AI has no security concerns
- B. Risks include data leakage through prompts, prompt injection attacks, generation of convincing phishing content, AI hallucination producing incorrect security advice, and employees sharing sensitive data with external AI services
- C. Generative AI is perfectly secure for all enterprise use
- D. Only the AI developer faces security risks

Answer: B

Q952. What is the concept of software-defined networking (SDN) and its security implications?

- A. SDN is a type of physical network cable
- B. SDN separates the network control plane from the data plane, enabling programmatic network management; security implications include centralized security policy enforcement but also a single point of failure if the controller is compromised
- C. SDN has no impact on network security
- D. SDN eliminates the need for firewalls

Answer: B

Q953. How do deepfakes pose a threat to cybersecurity beyond misinformation?

- A. Deepfakes are only used for entertainment
- B. Deepfakes enable advanced social engineering attacks including voice cloning for CEO fraud, video impersonation for identity verification bypass, and synthetic media for blackmail and manipulation
- C. Deepfakes can only create images, not video or audio
- D. Deepfakes are easily detected by all security systems

Answer: B

Q954. What is the security challenge of operational technology (OT) and IT convergence?

- A. OT and IT convergence has no security challenges
- B. OT systems were designed for isolated environments with different priorities (availability over confidentiality), and connecting them to IT networks exposes them to cyber threats they were not designed to withstand, while patching and updating is difficult due to uptime requirements
- C. OT systems are naturally more secure than IT systems
- D. IT security tools work identically in OT environments

Answer: B

Q955. What is the 'harvest now, decrypt later' threat in quantum computing?

- A. A quantum computing farming technique
- B. Adversaries are collecting and storing encrypted data today with the expectation that future quantum computers will be able to decrypt it, threatening long-term data confidentiality
- C. A data backup strategy
- D. A machine learning training approach

Answer: B

Q956. What is extended detection and response (XDR) and how does it evolve beyond EDR?

- A. XDR is just a marketing term for EDR
- B. XDR extends EDR by correlating and analyzing security data across multiple security layers (endpoints, network, cloud, email) to provide unified threat detection and response with broader visibility than endpoint-only EDR
- C. XDR only monitors network traffic
- D. XDR replaces all other security tools

Answer: B

Q957. What is the emerging concept of cybersecurity mesh architecture (CSMA)?

- A. A physical mesh network for security devices
- B. CSMA is a composable architecture approach where security controls are deployed as modular, interoperable services that can be distributed across the IT environment, enabling consistent security policy regardless of asset location
- C. CSMA is a type of encryption
- D. CSMA only applies to wireless networks

Answer: B

Q958. What is the role of threat intelligence platforms (TIPs) in modern cybersecurity?

- A. TIPs generate cyber attacks for testing
- B. TIPs aggregate, correlate, and operationalize threat intelligence from multiple sources, enabling organizations to contextualize threats, share intelligence, and integrate IoCs into security controls for proactive defense
- C. TIPs only provide news articles about cyber threats
- D. TIPs replace the need for security analysts

Answer: B

Q959. What is the security implication of edge computing for data protection?

- A. Edge computing eliminates all security concerns
- B. Edge computing processes data closer to the source, creating challenges in securing distributed processing nodes, ensuring consistent security policies across many edge locations, and protecting data at locations with potentially less physical security
- C. Edge computing is more secure than cloud computing
- D. Edge computing does not handle sensitive data

Answer: B

Q960. What is the purpose of conducting a gap analysis of the incident response plan?

- A. To find physical gaps in server room security
- B. To identify weaknesses and areas for improvement in the current IR plan by comparing it against best practices and frameworks
- C. To measure network bandwidth gaps
- D. To find gaps in employee schedules

Answer: B

Cyber Security — Hard

480 questions

Q961. In the STRIDE threat model, what does the 'E' stand for?

- A. Elevation of Privilege
- B. Exploitation in security contexts
- C. Enumeration in security contexts
- D. Data-at-rest symmetric encryption

Answer: A

Q962. Which is a key difference between ISO 27001 and ISO 27002?

- A. ISO 27001 specifies requirements for an ISMS, ISO 27002 provides implementation guidelines
- B. ISO 27001 is for networks, ISO 27002 is for applications
- C. ISO 27001 is outdated within enterprise security environments
- D. They are identical standards with no meaningful distinction between them

Answer: A

Q963. What is the primary purpose of the MITRE ATT&CK framework?

- A. To design websites for enterprise computing environments
- B. To develop antivirus software across enterprise computing systems
- C. To provide a knowledge base of adversary tactics, techniques, and procedures
- D. To generate comprehensive network topology diagrams

Answer: C

Q964. In cyber security, what does APT stand for?

- A. Active Port Tunneling
- B. Application Protocol Transfer
- C. Advanced Persistent Threat
- D. Automated Penetration Testing

Answer: C

Q965. Which best describes the concept of 'zero trust' architecture?

- A. Using zero encryption within enterprise security environments
- B. Trusting all internal network traffic regardless of the specific situation or context
- C. Never trusting any entity by default and always verifying regardless of location
- D. Blocking all external connections regardless of the specific situation or context

Answer: C

Q966. What is the purpose of a threat intelligence platform (TIP)?

- A. To manage employee records within enterprise computing infrastructure
- B. To create backup systems for enterprise computing environments
- C. To collect, aggregate, and analyze threat data from multiple sources
- D. To design and build user-facing application interfaces

Answer: C

Q967. Which regulation specifically addresses data protection for EU citizens?

- A. PCI DSS
- B. SOX
- C. HIPAA
- D. GDPR

Answer: D

Q968. What is the 'kill chain' model in cyber security?

- A. A network topology used in enterprise network infrastructure
- B. A method to physically destroy hardware installed within computing infrastructure
- C. A type of encryption algorithm applied to data protection workflows
- D. A framework describing the stages of a cyber attack from reconnaissance to exfiltration

Answer: D

Q969. What is the role of a Security Operations Center (SOC)?

- A. To continuously monitor, detect, analyze, and respond to security incidents
- B. To design network hardware across the enterprise network infrastructure
- C. To develop new software products across enterprise computing systems
- D. To sell security products within enterprise computing infrastructure

Answer: A

Q970. In Pakistan, which law primarily addresses cybercrime?

- A. Industrial Relations Act in security contexts
- B. Prevention of Electronic Crimes Act (PECA) 2016
- C. Foreign Exchange Act within enterprise security environments
- D. Pakistan Penal Code only in security contexts

Answer: B

Q971. What is the Bell-LaPadula security model primarily designed to protect?

- A. Confidentiality
- B. Non-repudiation
- C. Integrity
- D. Availability

Answer: A

Q972. Which security model uses 'no read down, no write up' to protect integrity?

- A. Brewer-Nash
- B. Clark-Wilson
- C. Biba Model
- D. Bell-LaPadula

Answer: C

Q973. What is the Clark-Wilson integrity model based on?

- A. Discretionary access control in security contexts
- B. Mandatory access control within enterprise security environments
- C. Well-formed transactions and separation of duties
- D. Role-based access control in security contexts

Answer: C

Q974. In security, what is a 'rainbow table'?

- A. A precomputed table for reversing cryptographic hash functions
- B. A type of firewall configuration for security management
- C. A network topology map used in enterprise network infrastructure
- D. A colorful network diagram used in enterprise network infrastructure

Answer: A

Q975. What is the purpose of 'salting' in password hashing?

- A. Making passwords taste better within enterprise security environments
- B. Compressing the password within enterprise security environments
- C. Adding random data to the password before hashing to prevent rainbow table attacks
- D. Encrypting the password twice within enterprise security environments

Answer: C

Q976. What is the Kerckhoffs principle in cryptography?

- A. Encryption is unnecessary within enterprise security environments
- B. A system should be secure even if everything except the key is public knowledge
- C. Keys should be shared publicly within enterprise security environments
- D. All algorithms should be secret regardless of the specific situation or context

Answer: B

Q977. What is the difference between DAC and MAC access control models?

- A. MAC allows users to share freely within enterprise security environments
- B. DAC allows owners to set permissions, MAC enforces system-wide policies based on classification labels
- C. They are fundamentally identical concepts with no meaningful differences
- D. DAC is more restrictive than MAC within enterprise security environments

Answer: B

Q978. Which describes 'security through obscurity'?

- A. Relying on secrecy of implementation rather than sound design for security
- B. Implementing zero trust within enterprise security environments
- C. Using strong encryption within enterprise security environments
- D. Using multiple firewalls within enterprise security environments

Answer: A

Q979. What is a Common Vulnerability Scoring System (CVSS)?

- A. A cryptographic algorithm used for protecting sensitive data
- B. A standard network communication protocol for data transfer
- C. A type of malware scanner designed for compromising system integrity
- D. A standardized framework for rating the severity of security vulnerabilities

Answer: D

Q980. What does TCB (Trusted Computing Base) refer to?

- A. An automated backup solution designed for enterprise-level data recovery and restoration
- B. A physical network cable used for connecting devices within the network infrastructure
- C. The totality of protection mechanisms within a computer system responsible for enforcing security policy
- D. A cloud computing platform used for hosting and managing distributed enterprise workloads

Answer: C

Q981. What mathematical problem does RSA encryption rely on?

- A. The traveling salesman optimization problem
- B. The difficulty of factoring large prime numbers
- C. The discrete logarithm mathematical problem
- D. Elliptic curve mathematical computation

Answer: B

Q982. What is an Initialization Vector (IV) in block cipher modes?

- A. A type of hash function within modern computing environments
- B. A random value used with a key to ensure identical plaintexts produce different ciphertexts
- C. The first encryption key within enterprise security environments
- D. A decryption algorithm within enterprise computing environments

Answer: B

Q983. Which block cipher mode is considered insecure for encrypting multiple blocks?

- A. CBC (Cipher Block Chaining)
- B. ECB (Electronic Codebook)
- C. GCM (Galois/Counter Mode)
- D. CTR (Counter) in security contexts

Answer: B

Q984. What is elliptic curve cryptography (ECC)?

- A. A data compression technique that reduces file sizes within enterprise computing environments
- B. A cryptographic hashing algorithm used within enterprise computing environments for integrity
- C. A symmetric block cipher algorithm widely used for protecting sensitive data at rest
- D. An asymmetric approach using elliptic curves over finite fields for smaller key sizes with equivalent security

Answer: D

Q985. What is a man-in-the-middle attack on key exchange?

- A. A brute force attack targeting enterprise network infrastructure used in enterprise computing environments
- B. A dictionary attack targeting enterprise network infrastructure used in enterprise computing environments
- C. An attacker intercepting and altering communication between two parties who believe they are communicating directly
- D. Physically stealing keys within enterprise security environments

Answer: C

Q986. What is perfect forward secrecy (PFS)?

- A. Using only symmetric encryption without any additional considerations needed
- B. Using the same key forever with no meaningful distinction between them
- C. Never changing encryption keys in any deployment scenario or context
- D. A property ensuring that compromise of long-term keys does not compromise past session keys

Answer: D

Q987. What is the purpose of HMAC?

- A. To reduce data sizes through compression methods and management
- B. To verify both data integrity and authenticity using a hash function and secret key
- C. To encrypt messages using established cryptographic standards
- D. To generate random numbers within enterprise computing infrastructure

Answer: B

Q988. Which is a post-quantum cryptographic approach?

- A. Lattice-based cryptography
- B. AES-128 symmetric block cipher
- C. SHA-256 cryptographic hashing
- D. RSA-2048 key exchange scheme

Answer: A

Q989. What is the purpose of key stretching algorithms like PBKDF2 and bcrypt?

- A. To speed up encryption using established cryptographic standards
- B. To compress encryption keys using established cryptographic standards
- C. To make brute-force attacks on passwords computationally expensive by adding iterations
- D. To make keys shorter within enterprise computing infrastructure

Answer: C

Q990. What is homomorphic encryption?

- A. Encryption that only works on text without any additional considerations needed
- B. A method to compress encrypted data for managing enterprise data resources
- C. Encryption that allows computations on ciphertext without decrypting it first
- D. A type of symmetric encryption applied to data protection workflows

Answer: C

Q991. What is a next-generation firewall (NGFW)?

- A. Same as a traditional firewall with no meaningful distinction between them regardless of the deployment context or scenario
- B. Only works with IPv6 without any additional considerations needed
- C. Combines traditional firewall with application awareness, deep packet inspection, and integrated threat intelligence
- D. Does not filter traffic in any deployment scenario or context regardless of the deployment context or scenario

Answer: C

Q992. What is the purpose of 802.1X network authentication?

- A. To increase and optimize overall network throughput speed and management
- B. To create and deploy enterprise wireless network infrastructure
- C. To provide port-based network access control requiring authentication before granting access
- D. To compress network traffic across the enterprise network infrastructure

Answer: C

Q993. What is deep packet inspection (DPI)?

- A. Checking only packet headers without any additional considerations needed
- B. Examining the full content of network packets including the data payload for threats
- C. A type of compression within modern computing environments
- D. A routing protocol used in enterprise network infrastructure

Answer: B

Q994. What is network forensics?

- A. Designing and deploying new network topologies across enterprise security infrastructure
- B. The capture, recording, and analysis of network events to discover the source of security attacks
- C. Configuring and managing enterprise routers across enterprise security infrastructure
- D. Installing and configuring network cables across enterprise security infrastructure environments

Answer: B

Q995. What is the purpose of a SIEM system in network security?

- A. To speed up network traffic across the enterprise network infrastructure
- B. To provide real-time analysis of security alerts from network hardware and applications
- C. To compress log files within the data management framework and management
- D. To generate comprehensive network topology diagrams and management

Answer: B

Q996. What is microsegmentation in network security?

- A. A type of network cable used in enterprise network infrastructure
- B. A standard network communication protocol for data transmission
- C. Creating fine-grained security zones to isolate workloads and apply granular security policies
- D. Dividing networks into very large segments within enterprise security environments

Answer: C

Q997. What is a network tap used for in security monitoring?

- A. A network routing device for managing enterprise traffic flows
- B. A way to steal water within modern computing environments
- C. A passive device that copies network traffic for monitoring without affecting data flow
- D. A wireless access point within modern computing environments

Answer: C

Q998. What is BGP hijacking?

- A. A legitimate routing practice within modern computing environments
- B. A symmetric key data encryption algorithm applied to data protection workflows
- C. An attack where an AS falsely announces ownership of IP prefixes to redirect internet traffic
- D. A next-generation firewall for filtering network traffic flows

Answer: C

Q999. What is the purpose of NetFlow in security analysis?

- A. To encrypt traffic using established cryptographic standards and management
- B. To collect and record IP traffic flow data for analysis of patterns and anomaly detection
- C. To block websites within enterprise computing infrastructure and management
- D. To increase available bandwidth for improved download speeds

Answer: B

Q1000. What is a zero-day vulnerability?

- A. A vulnerability that has been patched within modern computing environments
- B. A previously unknown vulnerability exploited before the vendor has released a fix
- C. A vulnerability that only affects old networks used in enterprise computing environments
- D. A vulnerability in outdated software used in enterprise computing environments

Answer: B

Q1001. What is a kernel-level rootkit?

- A. A next-generation firewall security appliance deployed across enterprise environments
- B. A legitimate administrative system utility tool used in enterprise computing environments
- C. A standard user-space application for management within the identity management system
- D. Malware that operates at the kernel level, hiding its presence and having unrestricted system access

Answer: D

Q1002. What is Trusted Platform Module (TPM)?

- A. A wired network interface adapter for connectivity used in enterprise computing environments
- B. A next-generation firewall for filtering network traffic flows
- C. A hardware-based security chip that provides cryptographic functions and secure key storage
- D. A type of volatile random access memory module used in enterprise computing environments

Answer: C

Q1003. What is Secure Boot?

- A. A UEFI feature ensuring only signed and trusted software loads during boot, preventing bootkits
- B. A symmetric key data encryption algorithm applied to data protection workflows
- C. An automated feature for creating data backups for managing enterprise data resources
- D. An option to boot the system more quickly within modern computing environments

Answer: A

Q1004. What is the purpose of Control Flow Integrity (CFI)?

- A. To ensure program execution follows the intended control flow graph, preventing code-reuse attacks
- B. To manage file systems within the data management framework and management
- C. To control user access within enterprise computing infrastructure and management
- D. To control traffic flow on networks across the enterprise network infrastructure

Answer: A

Q1005. What is a race condition vulnerability in OS security?

- A. A concurrency competition between processes commonly found within modern computing environments
- B. A type of scheduling algorithm used to manage processes within enterprise computing environments
- C. A flaw where the outcome depends on timing, allowing exploitation of the window between check and use
- D. A fast parallel processing technique used within enterprise computing environments for tasks

Answer: C

Q1006. What is namespace isolation in Linux containers?

- A. Naming convention for files within enterprise security environments in security contexts
- B. A kernel feature providing isolated views of system resources for processes, forming the basis of container security
- C. A DNS configuration within modern computing environments for security management
- D. A programming concept used in enterprise computing environments for security management

Answer: B

Q1007. What is integrity measurement architecture (IMA) in Linux?

- A. Detecting unauthorized file changes by maintaining and verifying cryptographic hashes
- B. Measuring and reporting available disk storage capacity on server devices
- C. Improving overall system performance throughput and resource utilization
- D. Compressing files to reduce disk storage consumption across system drives

Answer: A

Q1008. What is credential guard in Windows 10/11?

- A. A standard software application used for routine enterprise computing operation tasks
- B. A user account type with specifically assigned and managed access permissions
- C. A feature using virtualization-based security to isolate and protect credential secrets from theft
- D. A next-generation firewall solution for filtering and monitoring network traffic

Answer: C

Q1009. What is eBPF and how is it used in OS security?

- A. A standard network communication protocol for data transmission
- B. A symmetric key data encryption algorithm applied to data protection workflows
- C. An in-kernel virtual machine allowing sandboxed programs for security monitoring without modifying kernel source
- D. A file format used for structured data storage for managing enterprise data resources

Answer: C

Q1010. What is the difference between Type 1 and Type 2 hypervisors in security?

- A. Type 1 runs directly on hardware providing better isolation, Type 2 runs on a host OS adding extra attack surface
- B. Type 1 is software-only without any additional considerations needed
- C. Type 2 is more secure within enterprise security environments across computing environments
- D. No security difference in any deployment scenario or context or notable impact in any deployment scenario

Answer: A

Q1011. What is Server-Side Request Forgery (SSRF)?

- A. A type of load balancing technique for traffic distribution used in enterprise computing environments
- B. A server performance optimization technique for scaling within modern computing environments
- C. An attack that abuses server functionality to make requests to internal resources the attacker cannot directly access
- D. A server configuration method for managing network access used in enterprise network infrastructure

Answer: C

Q1012. What is a JSON Web Token (JWT) and its security considerations?

- A. A database management system for storing structured records
- B. A compact token format for secure claims transfer requiring proper validation and strong algorithms
- C. A CSS preprocessor for advanced stylesheet development used in enterprise computing environments
- D. A web browser application for accessing internet-hosted content

Answer: B

Q1013. What is the purpose of Subresource Integrity (SRI)?

- A. To improve and optimize page load speed performance and management
- B. To compress and minify JavaScript and CSS scripts within enterprise computing infrastructure
- C. To allow browsers to verify fetched resources have not been tampered with using cryptographic hashes
- D. To organize and manage web page static resources within enterprise computing infrastructure

Answer: C

Q1014. What is insecure deserialization?

- A. A database schema problem causing data integrity errors for managing enterprise data resources
- B. A vulnerability where untrusted data abuses deserialization logic, potentially leading to remote code execution
- C. Slow and inefficient data loading performance issues within enterprise security environments
- D. A data format compatibility issue between API versions for managing enterprise data resources

Answer: B

Q1015. What is a race condition vulnerability in web applications?

- A. A performance optimization for improving application speed
- B. A flaw where concurrent requests exploit timing gaps in security checks
- C. A caching problem causing stale data to be served
- D. A load balancing issue affecting resource distribution

Answer: B

Q1016. What is HTTP Request Smuggling?

- A. A load balancing technique used for efficiently distributing incoming network traffic
- B. A server-side caching mechanism designed for improving web page loading performance
- C. A content delivery method used for efficiently distributing static web assets globally
- D. An attack exploiting differences in how front-end and back-end servers parse HTTP requests

Answer: D

Q1017. What is a blind SQL injection attack?

- A. A properly secured and parameterized SQL database query for managing enterprise data resources
- B. An SQL injection where the attacker infers information by observing application behavior rather than direct output
- C. A database backup technique for creating data snapshots for managing enterprise data resources
- D. A standard SQL query that returns no data results for managing enterprise data resources

Answer: B

Q1018. What is the purpose of the HSTS header?

- A. To compress and reduce the size of HTTP request headers
- B. To force browsers to only use HTTPS, preventing protocol downgrade attacks
- C. To block and prevent all incoming and outgoing traffic
- D. To improve and speed up web connection establishment

Answer: B

Q1019. What is DOM-based XSS?

- A. It is identical in behavior to reflected cross-site scripting across computing environments
- B. Reflected XSS is exclusively a client-side vulnerability within enterprise security environments
- C. DOM-based XSS is a server-side vulnerability entirely within enterprise security environments
- D. XSS that executes entirely in the browser by manipulating the DOM, unlike reflected XSS which involves the server

Answer: D

Q1020. What is the purpose of OAuth 2.0 in web security?

- A. To provide delegated authorization allowing third-party apps limited access without sharing credentials
- B. To compress and reduce the size of transmitted web data and management
- C. To completely block all unauthorized users from access and management
- D. To encrypt all web traffic using symmetric cipher algorithms and management

Answer: A

Q1021. What is a polymorphic virus?

- A. A virus designed for running on multiple hardware platforms
- B. A virus that infects only once and does not replicate further
- C. Malware that changes its code signature each time it replicates to evade detection
- D. A benign virus that enhances system security posture overall

Answer: C

Q1022. What is a supply chain attack?

- A. A written record or logbook for tracking daily operational activities
- B. An attack compromising a trusted vendor or software update mechanism to distribute malware
- C. A procurement process for purchasing IT equipment used in enterprise computing environments
- D. Attacking physical supply chains and logistics routes

Answer: B

Q1023. What is an APT lifecycle?

- A. A system update cycle for deploying maintenance patches within modern computing environments
- B. A backup schedule for creating regular recovery snapshots for managing enterprise data resources
- C. A multi-stage attack process involving reconnaissance, initial access, persistence, lateral movement, and exfiltration
- D. A standard software development lifecycle methodology used in enterprise computing environments

Answer: C

Q1024. What is living-off-the-land (LOTL) technique?

- A. Environmental computing initiatives for reducing power usage
- B. A green IT practice focused on sustainable computing used in enterprise computing environments
- C. Agricultural computing systems for managing farming operations
- D. Using legitimate system tools already present on the target system to carry out malicious activities

Answer: D

Q1025. What is a logic bomb?

- A. Malicious code that remains dormant until triggered by a specific condition like a date or event
- B. A physical hardware component used in computing infrastructure
- C. A written record or logbook for tracking daily operational activities
- D. A software application used for enterprise computing operations

Answer: A

Q1026. What is DNS tunneling used for in attacks?

- A. A DNS caching technique for reducing query response time used in enterprise computing environments
- B. Encoding data within DNS queries to establish a covert communication channel bypassing firewalls
- C. A DNS optimization technique for faster name resolution used in enterprise computing environments
- D. Improving DNS resolution speed and performance latency

Answer: B

Q1027. What is steganography in cyber attacks?

- A. A type of prehistoric dinosaur classification study used in enterprise computing environments
- B. A printing technique for high-resolution image output used in enterprise computing environments
- C. Hiding malicious data or communications within ordinary-looking files like images or audio
- D. A photography method for capturing wide-angle images used in enterprise computing environments

Answer: C

Q1028. What is a metamorphic virus?

- A. A virus that migrates and spreads between different host operating systems and platforms
- B. A virus that executes only once on a system and then immediately self-terminates
- C. Malware that completely rewrites its own code with each iteration while maintaining functionality
- D. A benign virus variant that improves overall system security posture and performance

Answer: C

Q1029. What is the MITRE ATT&CK technique T1055 (Process Injection)?

- A. A standard software development pattern used for organizing code into modules
- B. A performance optimization technique for accelerating overall application speed
- C. A technique where malicious code is injected into the address space of a legitimate process
- D. A runtime debugging technique used for identifying and resolving code defects

Answer: C

Q1030. What is command and control (C2) infrastructure?

- A. A network management tool used in enterprise network infrastructure
- B. The communication framework attackers use to maintain control over compromised systems
- C. A software development platform for building applications
- D. A military headquarters facility for coordinating operations

Answer: B

Q1031. What is Kerberos authentication protocol?

- A. A software application used for enterprise computing operations
- B. A ticket-based network authentication protocol using symmetric key cryptography and a trusted third party
- C. A standard web browser for accessing internet sites within modern computing environments
- D. A standard network communication protocol for data transmission

Answer: B

Q1032. What is the difference between RBAC and ABAC in scalability?

- A. Neither is scalable within enterprise security environments
- B. RBAC is always more scalable regardless of the specific situation or context
- C. ABAC is more scalable for complex environments as it uses attributes rather than requiring new roles
- D. They scale identically with no meaningful distinction between them

Answer: C

Q1033. What is the FIDO2 authentication standard?

- A. A passwordless authentication standard using public-key cryptography with hardware authenticators
- B. A standard network communication protocol for data transmission
- C. A credential management mechanism for user authentication flows
- D. A file format used for structured data storage for managing enterprise data resources

Answer: A

Q1034. What is just-in-time (JIT) access provisioning?

- A. A type of SSO used across multiple services within enterprise computing environments
- B. Permanent access granted for all users regardless of the specific situation or context
- C. Fast login systems used for authentication within enterprise security environments
- D. Granting privileged access only when needed for a specific duration, then automatically revoking it

Answer: D

Q1035. What is the OAuth 2.0 authorization code flow?

- A. A flow where the client receives an authorization code exchanged for tokens via a back-channel
- B. A file transfer process within enterprise computing environments
- C. A simple password exchange within the identity management system
- D. A database query method for managing enterprise data resources

Answer: A

Q1036. What is the purpose of Privileged Access Management (PAM)?

- A. To create user accounts for enterprise computing environments
- B. To manage email accounts within enterprise computing infrastructure
- C. To secure, manage, and monitor privileged accounts and access to critical systems
- D. To design websites for enterprise computing environments and management

Answer: C

Q1037. What is a pass-the-hash attack?

- A. A hashing algorithm used for data integrity within enterprise computing environments
- B. Sharing passwords openly across teams within enterprise security environments for access
- C. A credential management mechanism designed for enterprise user authentication flows
- D. An attack where captured password hashes are used directly to authenticate without knowing the password

Answer: D

Q1038. What is continuous authentication?

- A. Logging into the system repeatedly at regular intervals across computing environments
- B. A type of multi-factor authentication using OTP tokens used in enterprise computing environments
- C. Continuously verifying identity throughout a session using behavioral biometrics and contextual signals
- D. A session timeout feature that logs users out within the identity management system

Answer: C

Q1039. What is claim-based identity?

- A. A type of insurance policy covering liability damages used in enterprise computing environments
- B. Using assertions about a user made by a trusted identity provider for authentication and authorization
- C. A file management system for organizing stored documents used in enterprise computing environments
- D. Filing legal claims in a court of law or tribunal within enterprise security environments

Answer: B

Q1040. What is certificate-based authentication?

- A. A training certification program for development used in enterprise computing environments
- B. Using digital certificates (X.509) to verify identity, preferred for machine-to-machine communication
- C. A structured file format used for data storage and exchange
- D. Using paper certificates for physical identity verification

Answer: B

Q1041. What is the Microsoft Security Development Lifecycle (SDL)?

- A. A Microsoft commercial product for endpoint security used in enterprise computing environments
- B. A testing framework for automated unit testing of code used in enterprise computing environments
- C. A high-level language used for development within modern computing environments
- D. A process integrating security into every SDLC phase including training, design, and verification

Answer: D

Q1042. What is a Software Bill of Materials (SBOM)?

- A. A project plan for scheduling development milestones used in enterprise computing environments
- B. A software application used for enterprise computing operations
- C. A software license for managing application access rights
- D. A comprehensive inventory of all components, libraries, and dependencies in software

Answer: D

Q1043. What is Interactive Application Security Testing (IAST)?

- A. Manual code review conducted by human security reviewers
- B. A software application used for enterprise computing operations
- C. A hybrid combining SAST and DAST using instrumentation agents within the running application
- D. A type of unit test for verifying individual code functions used in enterprise computing environments

Answer: C

Q1044. What is Runtime Application Self-Protection (RASP)?

- A. A programming pattern used in enterprise computing environments
- B. Security technology built into an application that detects and prevents attacks in real-time
- C. A self-replicating strain of malicious software targeting systems
- D. A software application used for enterprise computing operations

Answer: B

Q1045. What is the purpose of fuzz testing (fuzzing)?

- A. Load testing application performance under heavy traffic
- B. Providing random, unexpected, or malformed data as input to discover vulnerabilities
- C. Testing user experience through usability research sessions
- D. Testing application inputs with soft expected normal values

Answer: B

Q1046. What is taint analysis in secure development?

- A. Tracking untrusted data through an application to identify where it reaches sensitive operations unsanitized
- B. A color analysis tool for evaluating interface themes used in enterprise computing environments
- C. Source code formatting and style standardization practices across computing environments
- D. A code obfuscation method for protecting intellectual property used in enterprise computing environments

Answer: A

Q1047. What is the concept of 'secure defaults'?

- A. Having no security settings or controls enabled at all regardless of the deployment context or scenario
- B. Using default vendor passwords for all system accounts across computing environments
- C. Shipping applications with the most secure configuration, requiring users to explicitly reduce security
- D. Applications require no configuration at all to operate regardless of the deployment context or scenario

Answer: C

Q1048. What is a security champion program in DevSecOps?

- A. Embedding security-trained developers within teams to advocate for security practices
- B. A competitive security skills competition between teams used in enterprise computing environments
- C. A hiring initiative for recruiting security professionals
- D. A certification program for professional security credentials

Answer: A

Q1049. What is the purpose of a bug bounty program?

- A. Paying developers to write application source code faster and more efficiently
- B. Incentivizing external researchers to find and responsibly report vulnerabilities for rewards
- C. An automated software testing service for quality assurance and code validation
- D. A code quality improvement program focused on long-term maintainability goals

Answer: B

Q1050. What is DevSecOps?

- A. Integrating security into DevOps pipelines through automation, making security a shared responsibility
- B. A software application used for enterprise computing operations
- C. A next-generation firewall security appliance deployed across enterprise environments
- D. A high-level language used for development within modern computing environments

Answer: A

Q1051. What is the KRACK vulnerability?

- A. A software application used for enterprise computing operations
- B. A defect in wireless network adapter hardware components
- C. A type of physical screen damage on mobile device displays
- D. A vulnerability in WPA2's four-way handshake allowing nonce reuse and traffic decryption

Answer: D

Q1052. What is 802.11w (Protected Management Frames)?

- A. A wireless speed standard for improving data throughput used in enterprise computing environments
- B. A physical network cable used for connecting infrastructure devices
- C. A physical hardware component used in computing infrastructure
- D. An amendment providing integrity protection for management frames, preventing deauthentication attacks

Answer: D

Q1053. What is mobile application wrapping?

- A. Gift wrapping services for physical phone accessories and mobile device peripherals
- B. Encrypting app names for the purpose of obscuring installed applications on devices
- C. Adding a management layer around apps to enforce security policies without modifying source code
- D. Compressing mobile apps for reduced storage space consumption on user devices

Answer: C

Q1054. What is BLE security vulnerability tracking?

- A. A fitness tracker feature for monitoring health activity
- B. Monitoring BLE implementations for vulnerabilities like BlueBorne, KNOB, and BIAS
- C. A battery monitoring tool for optimizing power consumption
- D. Tracking the location of lost Bluetooth headphones

Answer: B

Q1055. What is certificate pinning in mobile applications?

- A. A user account type with specifically assigned access permissions
- B. Physically pinning paper certificates to a bulletin board
- C. A file attachment method for sending documents via email used in enterprise computing environments
- D. Hardcoding a host's expected certificate to prevent MITM attacks using fraudulent certificates

Answer: D

Q1056. What is the security risk of SS7 protocol vulnerabilities?

- A. There is no risk at all from this vulnerability in any deployment scenario or context
- B. A physical hardware component used in computing infrastructure
- C. A mobile device application for portable computing operations
- D. SS7 flaws allow interception of calls, SMS messages, and location tracking of mobile users globally

Answer: D

Q1057. What is containerization in enterprise mobile security?

- A. A storage method for archiving large data file volumes used in enterprise computing environments
- B. Physical shipping containers for transporting goods
- C. A charging technique for optimizing battery power cycles used in enterprise computing environments
- D. Creating isolated encrypted workspaces separating corporate and personal data on devices

Answer: D

Q1058. What is a man-in-the-disk attack on Android?

- A. An attack exploiting insecure external storage use by apps to maliciously modify data
- B. A network attack exploiting wireless protocol vulnerabilities
- C. A disk cleaning tool for removing temporary cached files
- D. A hardware attack targeting physical disk drive components

Answer: A

Q1059. What is a Wireless Intrusion Prevention System (WIPS)?

- A. A mobile device application used for managing and configuring Wi-Fi network settings
- B. A system monitoring wireless spectrum for unauthorized access and automatically taking countermeasures
- C. A network routing device used for managing enterprise traffic flows across infrastructure
- D. A Bluetooth peripheral device used for short-range data communication between systems

Answer: B

Q1060. What is the security significance of baseband processor vulnerabilities?

- A. A battery concern related to excessive power consumption used in enterprise computing environments
- B. A display issue causing visual glitches and rendering problems used in enterprise computing environments
- C. Baseband vulnerabilities can allow remote code execution bypassing all OS-level security controls
- D. Supply chain attacks have no significance in software security regardless of the deployment context or scenario

Answer: C

Q1061. What is a confused deputy problem in cloud security?

- A. A confused employee who cannot follow security procedures used in enterprise computing environments
- B. A next-generation firewall for filtering network traffic flows
- C. A privilege escalation where a trusted service is tricked into performing unauthorized actions
- D. A DNS issue causing incorrect domain name resolution used in enterprise computing environments

Answer: C

Q1062. What is CNAPP?

- A. A web development framework designed for building scalable application frontend interfaces
- B. An integrated security platform combining CSPM, CWPP, and application security for cloud-native apps
- C. A mobile application store platform for downloading and distributing software packages
- D. A cloud computing platform used for hosting and managing distributed enterprise workloads

Answer: B

Q1063. What is the NIST Cloud Computing Reference Architecture?

- A. A framework defining five major actors and their roles in cloud computing
- B. A cloud provider offering infrastructure as a service
- C. A security tool for scanning and detecting vulnerabilities
- D. A programming standard for writing cloud application code

Answer: A

Q1064. What are the security implications of serverless computing?

- A. No security is needed because there are no significant threats regardless of the deployment context or scenario
- B. This is completely secure by default without configuration across computing environments
- C. Same security profile as traditional server hosting with no meaningful distinction between them
- D. Reduced OS attack surface but increased function-level risks including injection and insecure dependencies

Answer: D

Q1065. What is cloud key management service (KMS)?

- A. A managed service for creating and managing cryptographic keys for cloud data encryption
- B. A music streaming service platform for audio content used in enterprise computing environments
- C. A file management tool for organizing cloud storage resources
- D. A network routing service for directing cloud data traffic

Answer: A

Q1066. What is cloud infrastructure entitlements management (CIEM)?

- A. Managing cloud storage volumes and tracking data capacity limits across regions
- B. Creating and provisioning new cloud user accounts across multiple environments
- C. Managing cloud billing and implementing cost optimization strategies for budgets
- D. Detecting and remediating excessive or unused permissions across multi-cloud environments

Answer: D

Q1067. What is CSA STAR certification?

- A. An astronomy certification for studying celestial objects used in enterprise computing environments
- B. A web browser application for accessing internet-hosted content
- C. A movie rating classification for content age suitability used in enterprise computing environments
- D. A third-party assessment of a cloud provider's security posture based on the Cloud Controls Matrix

Answer: D

Q1068. What is the blast radius concept in cloud security?

- A. A network speed measurement using bandwidth throughput used in enterprise computing environments
- B. The potential scope of damage if a security component is compromised, minimized through segmentation
- C. A storage capacity metric for measuring disk volume sizes used in enterprise computing environments
- D. A physical security measure for protecting building infrastructure

Answer: B

Q1069. What are immutable infrastructure security benefits?

- A. There are no security benefits from this approach regardless of the deployment context or scenario
- B. Faster server performance through hardware optimization
- C. Permanent servers that never change with no security benefit regardless of the deployment context or scenario
- D. Servers are never modified after deployment, reducing configuration drift and persistent threats

Answer: D

Q1070. What is the security challenge of multi-cloud environments?

- A. Managing consistent security policies and visibility across multiple providers with different tools
- B. Multi-cloud eliminates all security risks through redundancy
- C. There are no challenges with this approach at all in any deployment scenario or context
- D. Multi-cloud is inherently more secure than single-cloud

Answer: A

Q1071. What is a YARA rule?

- A. A next-generation firewall for filtering network traffic flows
- B. A programming language rule used in enterprise computing environments
- C. A database management system for storing structured records
- D. A pattern-matching tool for describing malware families based on textual or binary patterns

Answer: D

Q1072. What is the Sleuth Kit and Autopsy platform?

- A. A security scanner deployed across enterprise environments for security management
- B. An open-source digital forensics framework for file system analysis, timeline creation, and keyword searching
- C. A detective game within modern computing environments for security management
- D. A network monitor used in enterprise network infrastructure for security management

Answer: B

Q1073. What is file carving in digital forensics?

- A. Physically cutting a hard drive within enterprise security environments
- B. Wood carving with a computer within enterprise security environments
- C. Editing file names within enterprise security environments
- D. Recovering files from unallocated space by identifying file headers and footers without metadata

Answer: D

Q1074. What is the difference between live and dead forensics?

- A. Live forensics analyzes a running system for volatile data; dead forensics examines a powered-off system
- B. One involves living organisms within enterprise security environments
- C. There is no difference in any deployment scenario or context or notable impact in any deployment scenario
- D. One is more dangerous within enterprise security environments

Answer: A

Q1075. What is Windows Registry forensics?

- A. Fixing Registry errors and resolving system conflicts within enterprise security environments
- B. Registering Windows products and managing software licenses within enterprise security environments
- C. Creating Registry backups and restoration points within enterprise security environments
- D. Analyzing the Registry to extract evidence of user activities, installed software, and USB history

Answer: D

Q1076. What is the significance of slack space?

- A. Unused RAM within enterprise security environments across computing environments
- B. Empty space on a disk within enterprise security environments
- C. Space between end of a file and end of disk cluster, which may contain remnants of previous data
- D. A software application used for enterprise computing operations

Answer: C

Q1077. What is artifact analysis in mobile forensics?

- A. Examining mobile-specific data like SMS databases, call logs, app caches, and location data
- B. Analyzing phone hardware within enterprise security environments
- C. Repairing mobile devices within enterprise security environments
- D. Studying ancient artifacts within enterprise security environments

Answer: A

Q1078. What is the Volatility framework?

- A. A database tool used in enterprise computing environments
- B. A network scanner used in enterprise network infrastructure
- C. Measuring system performance within enterprise security environments
- D. An advanced memory forensics framework for extracting digital artifacts from RAM dumps

Answer: D

Q1079. What is evidence spoliation?

- A. A chain of custody step within modern computing environments
- B. Intentional or negligent destruction or alteration of evidence, leading to legal sanctions
- C. Proper evidence handling within enterprise security environments
- D. A forensic technique within enterprise computing environments

Answer: B

Q1080. What is cloud forensics and its challenges?

- A. A standard software application used for enterprise computing operations and workflows
- B. Weather analysis and atmospheric monitoring within enterprise security environments
- C. An automated backup solution designed for enterprise data recovery and restoration needs
- D. Forensic investigation in cloud environments facing data distribution, multi-jurisdiction, and limited access

Answer: C

Q1081. How does the SANS incident response framework differ from NIST?

- A. SANS has only 3 steps within enterprise security environments without additional considerations
- B. SANS emphasizes Identification vs. NIST's Detection and Analysis, and is more practitioner-oriented
- C. They are fundamentally identical concepts with no meaningful differences between them at all
- D. NIST has 10 distinct steps within enterprise security environments for incident handling

Answer: B

Q1082. What is a CSIRT and its key functions?

- A. A specialized team handling incident monitoring, triage, analysis, coordination, and advisory services
- B. A software application used for enterprise computing operations
- C. A user account type with specifically assigned access permissions
- D. A network administration team managing infrastructure used in enterprise computing environments

Answer: A

Q1083. What is cyber insurance?

- A. An automated backup solution designed for enterprise data recovery and restoration needs
- B. Insurance covering financial losses from cyber incidents including forensics, legal fees, and notification
- C. Car insurance for covering vehicle damage and repair costs resulting from accidents
- D. Health insurance covering IT staff medical expenses, benefits, and wellness programs

Answer: B

Q1084. What is the role of threat intelligence in incident response?

- A. A type of self-replicating malicious software designed for system infiltration attacks
- B. Generating and distributing threats against specifically targeted computing systems
- C. Providing context about threat actors, TTPs, and IOCs to accelerate detection and guide response
- D. A compliance reporting tool designed for generating audit documentation and reports

Answer: C

Q1085. What is SOAR?

- A. A next-generation firewall security appliance deployed across enterprise environments
- B. A software application used for enterprise computing operations
- C. A technology stack integrating security tools, automating response tasks, and orchestrating workflows
- D. A music platform for streaming and listening to audio used in enterprise computing environments

Answer: C

Q1086. What are the legal considerations during incident response?

- A. Evidence preservation, regulatory notification requirements, privacy laws, and law enforcement coordination
- B. Only criminal law applies without any specific legislation across computing environments
- C. Only privacy laws matter during incident response without any additional considerations needed
- D. No legal considerations exist during incident response regardless of the deployment context or scenario

Answer: A

Q1087. What is the purpose of tabletop exercises?

- A. Playing board games for team building recreation purposes across computing environments
- B. Discussion-based exercises to test plans, identify gaps, and improve coordination without operational impact
- C. A physical security measure for protecting building infrastructure
- D. Testing table strength and durability for furniture quality across computing environments

Answer: B

Q1088. What is the 'assume breach' concept?

- A. Accepting all risks without any mitigation efforts across computing environments
- B. Planning with the assumption that breaches will occur, focusing on detection, response, and recovery
- C. Disabling all security controls for maximum performance across computing environments
- D. Ignoring security entirely without any protective measures

Answer: B

Q1089. What is the role of digital forensics in incident response?

- A. An activity reserved only for law enforcement investigators used in enterprise computing environments
- B. A separate unrelated discipline with no connection within modern computing environments
- C. Completely replacing incident response with alternatives across computing environments
- D. Providing scientific examination of evidence to understand attack vector, scope, timeline, and attribution

Answer: D

Q1090. What is the difference between an incident and a data breach?

- A. They are functionally identical frameworks with no differences in any deployment scenario or context
- B. A breach is less serious than a standard incident within modern computing environments
- C. An incident only involves malware and no other threats designed for compromising system integrity
- D. An incident is any event threatening security; a breach specifically involves confirmed unauthorized access to sensitive data

Answer: D

Q1091. What is a SOC maturity model?

- A. A project management tool for tracking development milestones
- B. A software development model for iterative feature delivery
- C. A framework for assessing and improving SOC capabilities across people, processes, and technology
- D. A hiring framework for recruiting security team personnel used in enterprise computing environments

Answer: C

Q1092. What is SOAR in security monitoring?

- A. A software application used for enterprise computing operations
- B. A flight control system used in aviation navigation within modern computing environments
- C. A platform integrating security tools, automating response workflows, and orchestrating incident handling
- D. A cloud computing platform for hosting enterprise workloads

Answer: C

Q1093. What is MITRE ATT&CK's role in security monitoring?

- A. A next-generation firewall security appliance deployed across enterprise environments
- B. A vulnerability scanner for identifying system weaknesses used in enterprise computing environments
- C. A software application used for enterprise computing operations
- D. Providing a matrix of adversary tactics and techniques to guide detection rule creation and coverage assessment

Answer: D

Q1094. What is Extended Detection and Response (XDR)?

- A. A display resolution specification for mobile devices installed within computing infrastructure
- B. A unified platform integrating data from endpoints, networks, cloud, and email for holistic threat detection
- C. A network routing device for managing enterprise traffic flows
- D. A physical hardware component used in computing infrastructure

Answer: B

Q1095. What is deception technology?

- A. Social media manipulation campaigns designed for spreading misinformation and propaganda
- B. A magic show performance with illusion tricks for entertainment and audience engagement
- C. A VPN technology used for creating encrypted tunnel connections across networks securely
- D. Using decoys, breadcrumbs, and traps to detect attackers who bypassed perimeter defenses

Answer: D

Q1096. What are MTTD and MTTR?

- A. Time to reboot systems within enterprise security environments
- B. Average time to build software within enterprise security environments
- C. Time to install updates regardless of the specific situation or context
- D. Key SOC metrics measuring average time to detect (MTTD) and to contain and resolve (MTTR) incidents

Answer: D

Q1097. What is a purple team exercise?

- A. A standard software application commonly used for enterprise computing operations and tasks
- B. A collaborative exercise where offensive and defensive teams work together to improve detection
- C. A training exercise specifically designed for improving team coordination and basic skills
- D. A color-themed social event designed for team-building and recreational group activities

Answer: B

Q1098. What is network detection and response (NDR)?

- A. A network routing device for managing enterprise traffic flows
- B. A DNS service for resolving domain names to IP addresses used in enterprise computing environments
- C. A solution monitoring network traffic using behavioral analytics and ML to detect and respond to threats
- D. A network speed test for measuring bandwidth throughput used in enterprise computing environments

Answer: C

Q1099. What is the role of machine learning in security monitoring?

- A. Replacing all human security analysts with automation across computing environments
- B. Machine learning is used only for spam filtering without any additional considerations needed
- C. Enhancing detection of novel threats, reducing false positives, and automating large-scale data analysis
- D. Making systems slower through additional processing overhead

Answer: C

Q1100. What is a detection engineering program?

- A. Building physical detectors for environmental monitoring
- B. An electrical engineering course covering power distribution used in enterprise computing environments
- C. A systematic approach to developing, testing, and measuring effectiveness of security detection rules
- D. A physical hardware component used in computing infrastructure

Answer: C

Q1101. How does PECA 2016 address cyber terrorism?

- A. Only monetary fines are imposed for cyber terrorism under PECA across computing environments
- B. Only verbal warnings are issued for cyber terrorism under PECA across computing environments
- C. Section 10 criminalizes using information systems for terrorism with penalties up to 14 years and 50 million rupees
- D. PECA does not address terrorism in any of its provisions in any deployment scenario or context

Answer: C

Q1102. What is the Budapest Convention on Cybercrime?

- A. A European trade agreement governing international commerce used in enterprise computing environments
- B. A military alliance treaty for coordinating national defense used in enterprise computing environments
- C. The first international treaty addressing cybercrime with a framework for international cooperation
- D. A climate agreement for addressing global environmental change used in enterprise computing environments

Answer: C

Q1103. What are the challenges of cross-border cybercrime investigation?

- A. There are no challenges with this approach at all or notable impact in any deployment scenario
- B. Only language barriers present challenges for investigation
- C. Only time zone differences present challenges for investigation
- D. Jurisdictional conflicts, varying laws, evidence preservation across borders, and MLAT delays

Answer: D

Q1104. What is the legal framework for digital evidence in Pakistan?

- A. Qanun-e-Shahadat Order 1984, PECA 2016, and Electronic Transactions Ordinance 2002 collectively govern it
- B. Only criminal law applies without any specific legislation
- C. No legal framework exists for digital evidence in Pakistan regardless of the deployment context or scenario
- D. Only PECA applies without any other relevant legislation across computing environments

Answer: A

Q1105. What is the legal principle of proportionality in cyber security law?

- A. Minimum security is always sufficient regardless of threat across computing environments
- B. No security is needed because there are no significant threats regardless of the deployment context or scenario
- C. Maximum security should always be applied regardless of risk
- D. Security measures and penalties should be proportionate to the actual risk, balancing security with rights

Answer: D

Q1106. What is the legal status of encryption in different jurisdictions?

- A. Encryption is mandatory everywhere in the world without exception
- B. Legality varies: some countries mandate it, others restrict or require key escrow
- C. There are no laws about encryption in any country worldwide
- D. Encryption is illegal everywhere in the world without exception

Answer: B

Q1107. What is the NIS2 Directive?

- A. A physical hardware component widely used in enterprise computing infrastructure systems
- B. A standard software application used for routine enterprise computing operations tasks
- C. A standard network communication protocol used for transmitting data across systems
- D. An EU directive expanding cybersecurity requirements with risk management and incident reporting mandates

Answer: D

Q1108. What is 'duty of care' in cybersecurity?

- A. The legal obligation to implement reasonable security measures to protect stakeholders from foreseeable harm
- B. Only governments have a duty of care not private organizations regardless of the deployment context or scenario
- C. Duty of care is only a medical concept with no legal relevance regardless of the deployment context or scenario
- D. Having no responsibility for security is the accepted standard regardless of the deployment context or scenario

Answer: A

Q1109. What are the implications of the Schrems II decision?

- A. There are no implications for this technology whatsoever regardless of the deployment context or scenario
- B. It invalidated EU-US Privacy Shield and requires assessment of destination country data protection adequacy
- C. It only affects social media platforms and no other services without impacting any other system components
- D. It only affects banking institutions and no other organizations without impacting any other system components

Answer: B

Q1110. What is Pakistan's Personal Data Protection Bill?

- A. An education policy for governing school curriculum standards used in enterprise computing environments
- B. It has been fully enacted and is currently in force as law across computing environments
- C. A banking regulation governing financial institution operations used in enterprise computing environments
- D. A proposed legislation establishing data protection rights, consent requirements, and a regulatory authority

Answer: D

Q1111. What is post-quantum cryptography?

- A. Encryption methods used after quantum computers become obsolete across computing environments
- B. A standard cryptographic algorithm used for protecting sensitive data in transit and at rest
- C. Algorithms resistant to quantum attacks; NIST is standardizing lattice-based algorithms like CRYSTALS-Kyber
- D. Current encryption methods that are already widely deployed across enterprise computing environments

Answer: C

Q1112. What is confidential computing?

- A. A privacy policy document for governing data handling used in enterprise computing environments
- B. Hardware-based technology protecting data in use through trusted execution environments (TEEs)
- C. Keeping computer use private from other household members
- D. A secret language used for encoding confidential communications

Answer: B

Q1113. What is the significance of supply chain attacks in software?

- A. Compromising trusted development tools or dependencies to distribute malware through legitimate channels
- B. Supply chain attacks only affect physical goods and shipping without impacting any other system components
- C. Supply chain attacks have no significance in software security regardless of the deployment context or scenario
- D. Supply chain attacks only affect shipping and transportation without impacting any other system components

Answer: A

Q1114. What is cyber-physical systems security?

- A. A gym security system for monitoring facility access points used in enterprise computing environments
- B. Securing physical buildings with locks and access controls across computing environments
- C. Protecting systems where computational elements control physical processes, where attacks cause physical damage
- D. A fitness tracker for monitoring personal health activity used in enterprise computing environments

Answer: C

Q1115. What is adversarial machine learning?

- A. Training ML models faster using distributed computing across computing environments
- B. Techniques exploiting ML model vulnerabilities through adversarial inputs causing misclassification
- C. A competitive tournament between different ML implementations used in enterprise computing environments
- D. A type of supervised learning for classification tasks used in enterprise computing environments

Answer: B

Q1116. What is digital sovereignty in cybersecurity?

- A. Owning a domain name for a personal website registration across computing environments
- B. A nation's ability to control its digital infrastructure, data, and technology without foreign dependency
- C. Online voting rights for participating in digital elections
- D. A type of digital currency for online financial transactions used in enterprise computing environments

Answer: B

Q1117. What is the role of STIX/TAXII in threat intelligence?

- A. Email services for sending and receiving electronic messages
- B. Social media platforms for sharing posts and content within enterprise security environments
- C. Standardized formats and protocols for sharing structured cyber threat intelligence between organizations
- D. File sharing platforms for distributing documents and media

Answer: C

Q1118. What are the security implications of quantum key distribution (QKD)?

- A. QKD weakens security by introducing new attack vectors across computing environments
- B. There are no implications for this technology whatsoever regardless of the deployment context or scenario
- C. QKD uses quantum mechanics for theoretically unbreakable key exchange, as eavesdropping is detectable
- D. QKD replaces all existing encryption methods entirely across computing environments

Answer: C

Q1119. What is security chaos engineering?

- A. Creating chaos and deliberate disruption in security team operations and workflow
- B. Disabling all security controls to test overall system resilience under attack
- C. Intentionally injecting security failures to proactively identify weaknesses and validate controls
- D. Random penetration testing performed without any defined methodology or test plan

Answer: C

Q1120. What is the emerging threat of AI-powered social engineering?

- A. AI improving social skills through virtual communication across computing environments
- B. Using AI for convincing phishing at scale, deepfake impersonation, and personalized attacks harder to detect
- C. AI creating automated social media accounts for marketing across computing environments
- D. AI managing social media accounts for content scheduling across computing environments

Answer: B

Q1121. An organization discovers that an employee has been exfiltrating data for months undetected. Which security principle failure is MOST directly responsible?

- A. Failure in conducting regular penetration testing against external perimeters
- B. Failure in deploying sufficient network bandwidth for threat detection systems
- C. Failure in implementing proper security monitoring and audit logging controls
- D. Failure in applying timely software patches to all organizational endpoints

Answer: C

Q1122. A company implements multiple overlapping security controls. Which cyber security concept does this BEST demonstrate?

- A. Defense in depth using layered protections across multiple levels
- B. Least privilege limiting individual user permissions and access
- C. Security through obscurity relying on hidden system implementations
- D. Single point of failure with centralized management of controls

Answer: A

Q1123. During a risk assessment, the team identifies a threat with low probability but catastrophic impact. How should this risk be classified?

- A. Medium risk since probability and impact partially offset each other
- B. Low risk because the probability of occurrence is minimal
- C. High risk because catastrophic impact outweighs the low probability
- D. Negligible risk as low probability events should be safely ignored

Answer: C

Q1124. An attacker exploits a zero-day vulnerability before the vendor releases a patch. Which response strategy is MOST effective immediately?

- A. Waiting for the vendor to release an official patch before taking action
- B. Completely disconnecting all systems from every network until fully resolved
- C. Implementing compensating controls like network segmentation and monitoring
- D. Reinstalling the operating system on every affected endpoint in the network

Answer: C

Q1125. A security team must justify budget for new security tools to executive leadership. Which approach is MOST effective?

- A. Presenting the total number of blocked attacks from the previous fiscal year
- B. Highlighting technical specifications and vendor certifications of new tools
- C. Quantifying risk reduction in financial terms using established risk frameworks
- D. Comparing the organization's tool count with those of industry competitors

Answer: C

Q1126. In a mature security program, how should threat intelligence be integrated into operations?

- A. Shared exclusively with executive management during annual board review meetings
- B. Only used for generating quarterly reports for compliance and audit purposes
- C. Stored in offline archives for reference during post-incident forensic investigations
- D. Fed into SIEM, vulnerability management, and incident response processes continuously

Answer: D

Q1127. A nation-state actor targets critical infrastructure using advanced persistent threats. What classification does this attacker fall under?

- A. Opportunistic criminal seeking quick financial gain from random targets
- B. Hactivist motivated primarily by political or social activist causes
- C. Script kiddie with limited technical skill and basic attack toolkits
- D. Advanced Persistent Threat actor with significant resources and funding

Answer: D

Q1128. When implementing a security program, why is it critical to align security objectives with business goals?

- A. To reduce the overall complexity of network architecture within the organization systems
- B. To guarantee full compliance with every international security regulation and standard
- C. To ensure security spending directly supports organizational risk tolerance and priorities
- D. To minimize the total number of security staff required across all business departments

Answer: C

Q1129. A company wants to measure the effectiveness of its security program. Which metric is MOST meaningful?

- A. Mean time to detect and respond to confirmed security incidents over time
- B. Number of security certifications held by individual team members on staff
- C. Total count of firewall rules currently configured across all network devices
- D. Total annual budget allocated to the information security department budget

Answer: A

Q1130. An organization adopts a risk transfer strategy for a specific cyber threat. What does this typically involve?

- A. Purchasing cyber insurance to shift the financial burden of a potential breach
- B. Implementing additional technical controls to fully eliminate the identified risk
- C. Accepting the risk without any changes to current operations and security posture
- D. Ignoring the identified risk entirely because the likelihood is considered minimal

Answer: A

Q1131. An organization implements RBAC but discovers that a junior analyst can access executive financial reports. What is the MOST likely cause?

- A. The analyst discovered a zero-day exploit in the access control subsystem
- B. The RBAC system has a software bug that randomly grants elevated permissions
- C. Network segmentation failure allows unauthorized traffic to reach the server
- D. Role definitions are too broad and include permissions beyond job requirements

Answer: D

Q1132. A security architect must choose between preventive and detective controls for a limited budget. What is the BEST strategy?

- A. Invest entirely in preventive controls since prevention eliminates the need for detection
- B. Defer all investment until the budget allows for implementing both control types fully
- C. Balance both control types since prevention reduces incidents and detection catches gaps
- D. Invest entirely in detective controls since detection is always cheaper than prevention

Answer: C

Q1133. When applying the Clark-Wilson integrity model, what mechanism ensures data consistency?

- A. Periodic deletion of old data items to maintain storage efficiency across servers
- B. Well-formed transactions through constrained data items and transformation procedures
- C. Unrestricted access to data items by any authenticated user in the system
- D. Mandatory encryption of all data items at rest and in transit across all systems

Answer: B

Q1134. An organization finds that employees share credentials to bypass access restrictions. Which combination of controls BEST addresses this?

- A. Increasing the maximum number of concurrent sessions allowed for each user account
- B. Multi-factor authentication combined with behavioral analytics and session monitoring
- C. Stronger password complexity rules combined with more frequent password rotation policies
- D. Implementing longer password expiration periods with reduced login attempt thresholds

Answer: B

Q1135. In the Bell-LaPadula model, what does the 'no write down' rule prevent?

- A. Users from reading data at a higher classification level than their clearance
- B. Applications from creating temporary files on shared network storage volumes
- C. Users from writing data to a lower classification level causing data leakage
- D. Administrators from deleting files that are marked as system-critical assets

Answer: C

Q1136. A company merges with another organization that uses different security standards. What is the FIRST step for security integration?

- A. Merging all user accounts into one directory without modifying access permissions
- B. Immediately applying the stricter security standards across both organizations
- C. Conducting a gap analysis to identify differences in security controls and policies
- D. Replacing all security tools in both organizations with a single vendor solution

Answer: C

Q1137. Which scenario BEST illustrates a violation of the separation of duties principle?

- A. A security analyst monitors alerts while an incident responder handles confirmed breaches
- B. A developer writes code while a separate quality team reviews and approves the deployment
- C. A project manager assigns tasks while team leads independently track milestone progress
- D. A single administrator both creates user accounts and approves their own access requests

Answer: D

Q1138. An organization's quantitative risk analysis calculates an ALE of \$500,000 for a specific threat. The proposed countermeasure costs \$600,000 annually. What is the BEST decision?

- A. Defer the decision until the annualized loss expectancy increases above the threshold
- B. Implement the countermeasure and increase the security budget to cover the difference
- C. Implement the countermeasure because any protection is worth the investment overall
- D. Reject the countermeasure because its cost exceeds the annualized loss expectancy

Answer: D

Q1139. What is the primary advantage of using attribute-based access control over role-based access control?

- A. ABAC provides finer-grained access decisions based on multiple contextual attributes
- B. ABAC is simpler to implement and requires significantly less administrative overhead
- C. ABAC requires fewer computing resources and scales better on minimal hardware
- D. ABAC eliminates the need for any user authentication before granting system access

Answer: A

Q1140. How does the Biba integrity model differ from the Bell-LaPadula confidentiality model?

- A. Biba focuses on availability while Bell-LaPadula focuses on authentication mechanisms
- B. Biba allows unrestricted access while Bell-LaPadula enforces mandatory access controls
- C. Biba is only applicable to military systems while Bell-LaPadula is used commercially
- D. Biba uses no read down and no write up to protect integrity of data classifications

Answer: D

Q1141. An organization must protect data against future quantum computing attacks. Which cryptographic approach should they adopt?

- A. Switching from AES-256 to AES-128 for faster encryption processing speed
- B. Implementing post-quantum cryptographic algorithms like lattice-based schemes
- C. Increasing RSA key sizes to 4096 bits for stronger conventional encryption
- D. Using triple DES instead of AES for backward compatibility with systems

Answer: B

Q1142. A developer uses ECB mode to encrypt user profile images. What vulnerability does this introduce?

- A. ECB mode is too slow for image encryption causing application timeouts
- B. ECB mode requires keys that are too long for practical image encryption use
- C. ECB mode produces identical ciphertext for identical plaintext blocks revealing patterns
- D. ECB mode cannot handle binary data formats like images and only works with text

Answer: C

Q1143. During a TLS handshake, what is the purpose of the server's digital certificate?

- A. To establish the symmetric session key used for bulk data encryption directly
- B. To authenticate the server's identity and provide its public key to the client
- C. To encrypt all subsequent data transmitted between the client and the server
- D. To verify the client's identity before allowing any connection to be established

Answer: B

Q1144. An attacker performs a birthday attack against a hash function. What is the attacker exploiting?

- A. Side-channel information leaked through timing of the hash computation
- B. Buffer overflow vulnerabilities in the hash function software implementation
- C. Weak password policies that allow users to set easily guessable passwords
- D. The probability that two different inputs produce the same hash output value

Answer: D

Q1145. Why is key stretching used in password-based key derivation functions like PBKDF2?

- A. To reduce the length of the derived key for efficient storage in databases
- B. To make brute-force attacks slower by increasing computational cost per attempt
- C. To convert asymmetric keys into symmetric keys for hybrid encryption schemes
- D. To enable multiple users to share the same derived key without coordination

Answer: B

Q1146. In a hybrid encryption scheme, how are symmetric and asymmetric encryption combined?

- A. Symmetric encryption is used for authentication while asymmetric handles data integrity
- B. Asymmetric encryption secures the symmetric key which then encrypts the bulk data
- C. Both algorithms encrypt the same data independently for redundancy and verification
- D. Asymmetric encryption protects the bulk data while symmetric encryption protects the keys

Answer: B

Q1147. What is homomorphic encryption and why is it significant for cloud computing?

- A. Encryption that allows computations on ciphertext producing results matching plaintext operations
- B. Encryption that only works within a single cloud provider's proprietary infrastructure
- C. Encryption that uses identical keys across all cloud servers for simplified key management
- D. Encryption that automatically scales key sizes based on the current computational workload

Answer: A

Q1148. A security team discovers that their certificate authority's private key has been compromised. What is the MOST critical immediate action?

- A. Notifying end users to temporarily ignore certificate warnings in their browsers
- B. Revoking all certificates issued by the compromised CA and reissuing from a new CA
- C. Generating new public keys for all end-user certificates across the organization
- D. Increasing the key length of the compromised CA certificate to strengthen security

Answer: B

Q1149. What is the primary security concern with using MD5 for integrity verification today?

- A. MD5 is vulnerable to collision attacks making it unreliable for integrity checks
- B. MD5 requires excessive computational resources compared to newer hash algorithms
- C. MD5 can only hash text data and cannot process binary files or media content
- D. MD5 produces output that is too long for efficient storage in modern databases

Answer: A

Q1150. How does a hardware security module (HSM) enhance cryptographic key management?

- A. By converting all asymmetric keys to symmetric keys for faster processing speeds
- B. By automatically distributing encryption keys to all devices on the local network
- C. By compressing encryption keys to reduce the required storage space significantly
- D. By storing keys in tamper-resistant hardware that prevents unauthorized key extraction

Answer: D

Q1151. An organization detects lateral movement within their network after an initial breach. Which network control would MOST effectively limit this?

- A. Implementing micro-segmentation with strict inter-segment access control policies
- B. Upgrading the external firewall to a newer model with higher throughput capacity
- C. Deploying additional wireless access points throughout the building for coverage
- D. Increasing the bandwidth of the core network switches to handle more traffic loads

Answer: A

Q1152. During a DDoS attack, the security team notices traffic from thousands of spoofed source IPs. Which mitigation is MOST effective?

- A. Disabling all network interfaces until the attack subsides and traffic normalizes
- B. Increasing server hardware resources to absorb the additional malicious traffic load
- C. Using upstream provider scrubbing services with BGP flowspec to filter attack traffic
- D. Blocking each individual spoofed IP address manually in the perimeter firewall rules

Answer: C

Q1153. A penetration tester discovers that VLAN hopping is possible on the network. What misconfiguration is MOST likely responsible?

- A. Spanning tree protocol enabled with default priority values on all switch devices
- B. Trunk ports configured with native VLAN matching user access VLAN assignments
- C. DHCP snooping disabled on trunk ports connecting distribution and core switches
- D. Access ports configured with the correct VLAN but incorrect speed auto-negotiation

Answer: B

Q1154. What is the primary advantage of implementing a zero trust network architecture?

- A. Eliminating the need for any network firewalls or intrusion prevention systems
- B. Reducing network complexity by consolidating all traffic through a single gateway
- C. Automatically encrypting all data at rest without requiring application-level changes
- D. Verifying every access request regardless of whether it originates inside the network

Answer: D

Q1155. An analyst notices DNS tunneling in network logs. What is the attacker MOST likely trying to achieve?

- A. Improving network performance by bypassing the organization's proxy server
- B. Speeding up DNS resolution by caching responses on the local DNS server
- C. Testing the resilience of the DNS infrastructure against denial of service
- D. Exfiltrating data or establishing command-and-control channels via DNS queries

Answer: D

Q1156. A company implements network detection and response (NDR). What capability does NDR primarily add beyond traditional IDS?

- A. NDR replaces the need for endpoint detection and response tools on workstations
- B. NDR provides faster packet forwarding speeds than traditional IDS appliances can
- C. NDR uses behavioral analytics and machine learning to detect unknown threat patterns
- D. NDR eliminates false positive alerts entirely through deterministic rule matching

Answer: C

Q1157. What is BGP hijacking and why is it dangerous?

- A. Exploiting BGP to advertise false routes, redirecting internet traffic through attacker-controlled networks
- B. Attacking the BGP daemon process to crash border routers and cause widespread network service outages
- C. Intercepting BGP keepalive messages to prevent routers from maintaining their established peer sessions
- D. Using BGP messages to flood internal routing tables causing memory exhaustion on core switch devices

Answer: A

Q1158. An organization uses deep packet inspection (DPI) but encrypted traffic is increasing. What is the BEST approach to maintain visibility?

- A. Deploying TLS inspection proxies that decrypt, inspect, and re-encrypt the traffic
- B. Downgrading encryption standards to weaker ciphers that are easier to inspect
- C. Banning all encrypted traffic on the corporate network to maintain full inspection
- D. Switching entirely to metadata analysis and abandoning payload inspection efforts

Answer: A

Q1159. What is the security risk of running network services on default ports?

- A. Default ports cannot be protected by firewalls due to protocol specification limits
- B. Attackers can easily identify and target services running on well-known default ports
- C. Default ports consume more bandwidth than non-standard ports on network devices
- D. Default ports are inherently less secure than non-standard ports at the protocol level

Answer: B

Q1160. How does software-defined networking (SDN) improve network security management?

- A. SDN centralizes control enabling dynamic, programmable security policy enforcement
- B. SDN automatically encrypts all traffic without requiring any additional configuration
- C. SDN replaces all traditional security tools including firewalls and IDS solutions
- D. SDN eliminates the need for physical network infrastructure entirely in all cases

Answer: A

Q1161. A Linux server is compromised via a kernel exploit that bypasses all userspace security controls. Which hardening mechanism would have BEST mitigated this?

- A. Enabling Kernel Address Space Layout Randomization and Secure Boot with signed kernels
- B. Installing additional antivirus software with real-time scanning enabled on the server
- C. Configuring iptables with strict ingress and egress filtering rules on all interfaces
- D. Implementing disk quotas to limit the amount of storage each user account can consume

Answer: A

Q1162. An attacker uses a return-oriented programming (ROP) attack. Which OS security mechanisms specifically counter this technique?

- A. Disk encryption and secure boot processes that verify operating system integrity
- B. File permission checks and user account control prompts before executing programs
- C. Network firewall rules that block incoming connections from suspicious IP addresses
- D. Control Flow Integrity and Shadow Stack implementations in the processor and compiler

Answer: D

Q1163. A containerized application escapes its container and gains host access. What is the MOST effective prevention mechanism?

- A. Using rootless containers with user namespaces and mandatory access control policies
- B. Running containers with root privileges to ensure all security modules are accessible
- C. Disabling all container networking to prevent communication with the host system
- D. Allocating more CPU and memory resources to the container for improved isolation

Answer: A

Q1164. An organization discovers a rootkit embedded in their server OS. What makes rootkits particularly dangerous compared to other malware?

- A. Rootkits only target mobile operating systems and cannot infect desktop computers
- B. Rootkits spread faster than worms across networks using standard network protocols
- C. Rootkits modify the OS kernel or firmware to hide their presence from security tools
- D. Rootkits are larger in file size than other malware making them consume more disk space

Answer: C

Q1165. What is the security advantage of using immutable infrastructure for OS deployments?

- A. Immutable systems are faster because they do not need to load any security modules
- B. Immutable systems never require updates because they are inherently secure by design
- C. Immutable systems are replaced entirely rather than patched, preventing persistent compromise
- D. Immutable systems allow administrators to modify configurations at any time without risk

Answer: C

Q1166. A system administrator needs to detect fileless malware that operates entirely in memory. Which approach is MOST effective?

- A. Scanning network traffic logs for unusual DNS queries originating from external servers
- B. Using endpoint detection and response tools with memory analysis and behavioral monitoring
- C. Checking file integrity using checksums against a known-good baseline of all system files
- D. Running traditional signature-based antivirus scans on all files stored on disk

Answer: B

Q1167. What is the security risk of symlink race conditions in Unix-like operating systems?

- A. They consume excessive disk space by duplicating files across multiple directory paths
- B. They prevent legitimate users from accessing their home directories during peak usage
- C. They allow attackers to redirect privileged operations to unintended files via symbolic links
- D. They cause system slowdowns by creating circular references in the filesystem structure

Answer: C

Q1168. How does Windows Credential Guard protect authentication credentials?

- A. By storing credentials in an encrypted file on the system drive protected by BitLocker
- B. By requiring all users to change their passwords every seven days automatically
- C. By isolating credential processes in a virtualization-based security container from the OS
- D. By disabling all remote authentication protocols to prevent credential interception

Answer: C

Q1169. A security team must secure a system running legacy software that cannot be patched. What is the BEST compensating strategy?

- A. Replacing all other systems on the network with the same legacy software for consistency
- B. Isolating the system with strict network segmentation, monitoring, and application whitelisting
- C. Ignoring the risk since the legacy software has been stable for many years without issues
- D. Granting administrator access to all users so they can manually address security concerns

Answer: B

Q1170. What is the security significance of Trusted Platform Module (TPM) in OS security?

- A. TPM increases CPU processing speed to enable faster encryption and decryption operations
- B. TPM automatically installs operating system updates without requiring administrator input
- C. TPM replaces the need for software-based firewalls by filtering traffic in hardware only
- D. TPM provides a hardware root of trust for secure boot, key storage, and platform integrity

Answer: D

Q1171. A web application uses JWT tokens for authentication. An attacker changes the algorithm header to 'none'. What vulnerability does this exploit?

- A. Insecure session management allowing unlimited concurrent sessions for each user account
- B. Cross-site request forgery enabling unauthorized actions on behalf of authenticated users
- C. JWT algorithm confusion attack bypassing signature verification when none is accepted
- D. Server-side template injection allowing arbitrary code execution on the web application

Answer: C

Q1172. A penetration tester discovers a blind SQL injection vulnerability. How does blind SQL injection differ from standard SQL injection?

- A. Blind SQL injection requires physical access to the database server for exploitation
- B. Blind SQL injection infers data from application behavior rather than direct query output
- C. Blind SQL injection cannot extract any data from the database regardless of technique
- D. Blind SQL injection only works against NoSQL databases and not relational databases

Answer: B

Q1173. An application is vulnerable to server-side template injection (SSTI). What is the primary risk?

- A. Attackers can execute arbitrary code on the server through template engine exploitation
- B. Users can modify the visual layout and CSS styling of the application pages
- C. The application performance degrades because templates are rendered on every request
- D. Template files are exposed to search engines causing information disclosure publicly

Answer: A

Q1174. A security team discovers that their web application's deserialization process is vulnerable. What is the MOST severe potential impact?

- A. Remote code execution through crafted serialized objects processed by the application
- B. Loss of CSS styling information causing visual degradation on mobile device browsers
- C. Slightly increased response times when processing large serialized data payloads
- D. The application displays incorrect date formats in the user interface components

Answer: A

Q1175. How does a race condition vulnerability manifest in web applications?

- A. Two concurrent requests exploit timing gaps to bypass validation or duplicate operations
- B. Page loading speed varies based on the number of concurrent users on the platform
- C. The application crashes when multiple users attempt to log in at the same exact time
- D. Database queries return incorrect results when executed during high traffic periods

Answer: A

Q1176. A web application uses OAuth 2.0 for authentication. What is the primary risk of the 'open redirect' vulnerability in this context?

- A. It allows attackers to steal authorization codes by redirecting to attacker-controlled URLs
- B. It causes the OAuth server to crash when processing redirect URLs with special characters
- C. It exposes the client secret in the URL parameters visible in the browser address bar
- D. It prevents legitimate users from completing the authentication flow on mobile browsers

Answer: A

Q1177. What is a prototype pollution vulnerability in JavaScript web applications?

- A. An attack that modifies the Object prototype to inject properties affecting all objects
- B. An issue where deprecated JavaScript functions cause compatibility errors in browsers
- C. A vulnerability where JavaScript files are served without proper content type headers
- D. A memory leak caused by creating too many JavaScript objects without cleanup routines

Answer: A

Q1178. A web application allows file uploads. Which combination of controls BEST prevents malicious file upload attacks?

- A. Limiting file size to one megabyte and allowing only uploads during business hours daily
- B. Validating file type, scanning content, storing outside webroot, and renaming uploaded files
- C. Only checking the file extension on the client side before allowing the upload to proceed
- D. Compressing all uploaded files into ZIP archives and storing them on a shared network drive

Answer: B

Q1179. What is the security impact of HTTP request smuggling?

- A. It exploits discrepancies between frontend and backend HTTP parsing to inject requests
- B. It downgrades HTTPS connections to HTTP by modifying the request protocol header field
- C. It causes web pages to load slower because requests are processed in an incorrect order
- D. It prevents the server from processing legitimate user requests during peak traffic hours

Answer: A

Q1180. How does Subresource Integrity (SRI) protect web applications that use CDN-hosted scripts?

- A. By caching CDN resources locally so they load faster regardless of CDN server status
- B. By verifying that fetched resources match expected cryptographic hashes before execution
- C. By automatically updating CDN-hosted scripts when new versions become available online
- D. By encrypting the communication channel between the browser and CDN server endpoints

Answer: B

Q1181. An organization is hit by ransomware that also exfiltrates data before encryption. What is this attack technique called?

- A. Crypto-jacking where computing resources are hijacked to mine cryptocurrency secretly
- B. Triple extortion where customers of the victim are also contacted for ransom payments
- C. Double extortion where data is stolen and encrypted to pressure victims into paying
- D. Single extortion where data is encrypted and a ransom is demanded for decryption

Answer: C

Q1182. A threat actor uses living-off-the-land techniques (LOLBins) during an attack. Why is this approach effective?

- A. Because using legitimate system tools avoids triggering signature-based security detections
- B. Because system binaries can be modified without administrator privileges on any system
- C. Because custom malware tools are faster than legitimate system binaries for attacks
- D. Because legitimate tools provide encryption capabilities not available in custom malware

Answer: A

Q1183. An APT group maintains persistence on a compromised system for months. Which technique is MOST commonly used for long-term persistence?

- A. Physically accessing the server room weekly to reinstall their malware on the server
- B. Continuously running brute force attacks against the same system to maintain access
- C. Creating scheduled tasks and modifying registry run keys to survive system reboots
- D. Sending daily phishing emails to the same users to re-establish compromised access

Answer: C

Q1184. During malware analysis, a sample is identified as metamorphic. How does metamorphic malware differ from polymorphic malware?

- A. Metamorphic malware only changes file names while polymorphic malware changes the entire file structure
- B. Metamorphic malware is exclusive to mobile devices while polymorphic malware targets desktop systems
- C. Metamorphic malware spreads faster than polymorphic malware across networks due to its smaller size
- D. Metamorphic malware completely rewrites its own code while polymorphic malware only encrypts its payload

Answer: D

Q1185. A company discovers that attackers used a compromised software update mechanism to distribute malware. What type of attack is this?

- A. Man-in-the-middle attack intercepting downloads between the vendor and customer networks
- B. Supply chain attack leveraging trusted software distribution channels for malware delivery
- C. Watering hole attack targeting employees who visit the software vendor's support website
- D. Brute force attack against the software vendor's authentication system and user accounts

Answer: B

Q1186. An incident responder identifies that malware communicates using DNS TXT records for command and control. Why is this technique effective?

- A. DNS TXT records can carry arbitrary data and DNS traffic is often allowed through firewalls
- B. DNS TXT records are not logged by any DNS server making forensic investigation impossible
- C. DNS TXT records travel faster than other DNS record types due to priority routing protocols
- D. DNS TXT records are encrypted by default making them invisible to all security monitoring tools

Answer: A

Q1187. What is the MITRE ATT&CK framework and how does it help in understanding malware attacks?

- A. A vulnerability scanner that automatically detects and patches malware on infected systems
- B. A firewall configuration standard that blocks all known malware communication protocols
- C. An encryption standard used by government agencies to protect classified communications
- D. A knowledge base of adversary tactics and techniques based on real-world attack observations

Answer: D

Q1188. An organization's EDR detects PowerShell executing encoded commands at 3 AM. What type of attack behavior does this MOST likely indicate?

- A. A scheduled system maintenance script running automated update and cleanup operations
- B. A developer testing deployment scripts in the production environment during quiet hours
- C. Malicious activity using encoded PowerShell for obfuscation during low-monitoring periods
- D. An automated backup process that requires elevated PowerShell execution permissions

Answer: C

Q1189. What is process injection and why do attackers use it?

- A. Injecting malicious code into a legitimate running process to evade detection and gain its privileges
- B. Creating new malicious processes with random names to confuse system administrators during analysis
- C. Modifying the process scheduler to give malware higher priority than legitimate system applications
- D. Injecting monitoring code into malware processes to observe their behavior for research purposes

Answer: A

Q1190. A threat intelligence report describes malware that uses domain generation algorithms (DGA). What is the purpose of DGA?

- A. Creating legitimate websites that appear in search engine results to attract potential victims
- B. Generating secure domain names for the attacker's personal websites and email services
- C. Automatically registering domain names that are about to expire for resale profit purposes
- D. Generating random domain names for C2 communication making it difficult to block all domains

Answer: D

Q1191. An organization implements a zero trust model. How should authentication change compared to traditional perimeter-based security?

- A. Authentication is only required once at the network perimeter and then trusted throughout
- B. Authentication is only required for external users while internal users are auto-trusted
- C. Authentication is eliminated entirely because the zero trust model makes it unnecessary
- D. Continuous authentication and verification is required for every resource access request

Answer: D

Q1192. A company discovers that an attacker used a golden ticket attack in Active Directory. What does this attack achieve?

- A. It disables all group policies allowing unrestricted software installation on endpoints
- B. It modifies DNS records to redirect all domain traffic to attacker-controlled servers
- C. It resets all user passwords in the domain forcing everyone to create new credentials
- D. It creates a forged Kerberos TGT granting unrestricted access to any domain resource

Answer: D

Q1193. What is the security risk of using bearer tokens without additional binding in OAuth 2.0?

- A. Bearer tokens cannot be revoked once issued requiring server restarts to invalidate
- B. Bearer tokens expire too quickly making them impractical for long-running sessions
- C. Bearer tokens are too large in size causing performance issues in network requests
- D. Anyone possessing the token can use it, making stolen tokens immediately exploitable

Answer: D

Q1194. An attacker performs a pass-the-hash attack. What makes this attack possible?

- A. LDAP queries return password hashes to any authenticated user who requests directory data
- B. NTLM authentication allows authentication using the password hash without the actual password
- C. DNS service records expose password hashes when queried with specific malformed requests
- D. Kerberos tickets are stored in plaintext on the domain controller's hard drive permanently

Answer: B

Q1195. How does risk-based adaptive authentication improve security?

- A. It eliminates all authentication requirements to improve user experience and speed
- B. It dynamically adjusts authentication requirements based on contextual risk signals
- C. It uses the same authentication method regardless of the risk level of the access
- D. It requires maximum authentication strength for all access requests at all times

Answer: B

Q1196. An organization deploys FIDO2 hardware keys for authentication. What is the primary attack vector this eliminates?

- A. Physical theft attacks because FIDO2 keys cannot be used by anyone other than the owner
- B. Phishing attacks because FIDO2 keys are bound to specific origins and resist credential theft
- C. Malware attacks because FIDO2 keys include built-in antivirus scanning for connected devices
- D. Denial of service attacks because FIDO2 keys provide faster authentication to all services

Answer: B

Q1197. What is the security challenge of implementing SSO across an organization?

- A. SSO cannot work with cloud applications and is limited to on-premises systems only
- B. A compromised SSO credential grants access to all connected applications simultaneously
- C. SSO does not support multi-factor authentication and relies only on password access
- D. SSO requires separate passwords for each application making it complex to manage

Answer: B

Q1198. How does certificate-based authentication provide stronger security than password-based authentication?

- A. Certificates are shorter than passwords making them faster to process during authentication
- B. Certificates use asymmetric cryptography and cannot be phished or replayed like passwords
- C. Certificates work without any network connection making them suitable for offline access
- D. Certificates do not expire so they never need to be renewed or rotated by administrators

Answer: B

Q1199. An organization needs to implement step-up authentication. In what scenario is this MOST appropriate?

- A. When a system administrator performs routine maintenance on non-critical services
- B. When a user initially logs in to access their basic email and calendar application
- C. When a new employee creates their initial account during the onboarding process
- D. When an already authenticated user attempts to access a highly sensitive resource

Answer: D

Q1200. What is Kerberoasting and how does it exploit Active Directory authentication?

- A. Modifying Kerberos configuration files to disable ticket encryption on the domain controllers
- B. Intercepting Kerberos tickets in transit by performing ARP spoofing on the network segment
- C. Requesting service tickets for service accounts and cracking them offline to obtain passwords
- D. Flooding the Key Distribution Center with ticket requests to cause a denial of service

Answer: C

Q1201. A development team discovers a critical vulnerability in a third-party library used across all their applications. What is the MOST effective response?

- A. Rewriting all applications from scratch without using any third-party library dependencies
- B. Assessing impact, applying patches or mitigations, and implementing SCA in the CI/CD pipeline
- C. Removing the library immediately without considering functionality impact on applications
- D. Ignoring the vulnerability because third-party libraries are the vendor's sole responsibility

Answer: B

Q1202. How does interactive application security testing (IAST) combine benefits of SAST and DAST?

- A. IAST only works on interpreted languages while SAST and DAST work on compiled languages
- B. IAST replaces both SAST and DAST entirely by using artificial intelligence for scanning
- C. IAST runs SAST and DAST tools simultaneously in separate processes for speed improvement
- D. IAST instruments the application to analyze code execution during runtime testing activities

Answer: D

Q1203. A CI/CD pipeline has no security gates. What risks does this introduce and what should be implemented?

- A. Vulnerable code reaches production; implement SAST, DAST, SCA, and secret scanning gates
- B. No risk since CI/CD automation inherently eliminates all security vulnerabilities in code
- C. Only compliance risks exist; implement audit logging at the deployment stage exclusively
- D. Only performance risks exist; implement load testing gates to validate application speed

Answer: A

Q1204. What is the security risk of hardcoded secrets in source code?

- A. Hardcoded secrets only pose a risk if the application is written in interpreted languages
- B. Hardcoded secrets make the code run slower because encryption is computed at every startup
- C. Secrets in code are exposed in version control history even after removal from current files
- D. Secrets in code are automatically protected by the compiler during the build optimization

Answer: C

Q1205. How does runtime application self-protection (RASP) provide security?

- A. RASP scans source code before compilation to identify potential security vulnerabilities early
- B. RASP monitors network traffic outside the application to filter malicious request patterns
- C. RASP instruments the application to detect and block attacks from within the running application
- D. RASP encrypts all application data at rest to prevent unauthorized access to stored files

Answer: C

Q1206. A development team uses microservices architecture. What unique security challenges does this introduce?

- A. No unique challenges since microservices security is identical to monolithic application security
- B. Only performance challenges exist since microservices do not introduce any new security concerns
- C. Increased attack surface from inter-service communications requiring mutual authentication and encryption
- D. Reduced attack surface because microservices are inherently more secure than monolithic applications

Answer: C

Q1207. What is the security significance of Software Bill of Materials (SBOM)?

- A. SBOM is a financial document listing the costs of all software licenses for budget management
- B. SBOM provides a complete inventory of software components enabling vulnerability tracking and response
- C. SBOM is a testing report showing the results of all security scans performed on the software
- D. SBOM is a deployment guide listing the steps required to install software on production servers

Answer: B

Q1208. A penetration tester finds that the application is vulnerable to mass assignment. What is this vulnerability?

- A. An attacker modifies object properties by injecting unexpected parameters into API requests
- B. An attacker assigns administrator roles to all users by exploiting a database trigger flaw
- C. An attacker overwhelms the server by sending thousands of simultaneous assignment operations
- D. An attacker copies large amounts of data by exploiting a flaw in the file download feature

Answer: A

Q1209. What is the purpose of security regression testing?

- A. Verifying that previously fixed security vulnerabilities are not reintroduced by new code changes
- B. Scanning the production environment for new zero-day vulnerabilities not found during development
- C. Testing the application performance under heavy load to identify potential denial of service risks
- D. Reviewing the application source code for compliance with organizational coding style guidelines

Answer: A

Q1210. How does a security-focused infrastructure as code (IaC) scanning tool improve deployment security?

- A. It encrypts all infrastructure configuration files to prevent unauthorized access to server settings
- B. It detects misconfigurations and security policy violations in infrastructure definitions before deployment
- C. It replaces the need for network firewalls by implementing security rules within the code itself
- D. It automatically fixes all application code vulnerabilities found during the build compilation process

Answer: B

Q1211. An organization discovers that employees' mobile devices are connecting to an evil twin access point in their parking lot. What is the MOST comprehensive solution?

- A. Simply changing the corporate WiFi password and notifying employees via company-wide email
- B. Banning all wireless device usage on the company premises to eliminate wireless attack surface
- C. Deploying WIPS with rogue AP detection, enforcing VPN usage, and implementing 802.1X authentication
- D. Installing signal jammers in the parking lot to prevent all wireless connections from working

Answer: C

Q1212. A mobile application stores sensitive data in shared preferences without encryption. What is the primary attack vector this enables?

- A. Data extraction from the device's unencrypted shared preferences via backup or root access
- B. Credential theft through man-in-the-middle attacks on the application's API communications
- C. Denial of service against the application by filling shared preferences with random data
- D. Remote exploitation of the application through the mobile operating system's network stack

Answer: A

Q1213. What is the security implication of 5G network slicing from a cyber security perspective?

- A. 5G slicing reduces bandwidth making traditional wireless attacks completely ineffective
- B. 5G slicing eliminates all wireless security concerns by design through quantum encryption
- C. 5G slicing is only relevant for consumer devices and has no enterprise security implications
- D. Compromise of one slice could potentially affect other slices if isolation controls are insufficient

Answer: D

Q1214. An attacker uses a Stingray device near a corporate office. What type of attack is this?

- A. A Bluetooth sniffing attack that captures data from nearby paired mobile device connections
- B. An IMSI catcher attack that intercepts mobile communications by impersonating a cell tower
- C. A WiFi jamming attack that disrupts all wireless network signals within a building area
- D. A GPS spoofing attack that provides false location data to mobile navigation applications

Answer: B

Q1215. A company implements a zero trust approach for mobile devices. What is the MOST critical component?

- A. Allowing unrestricted access once the device passes an initial security check at enrollment
- B. Trusting all devices on the corporate WiFi network and blocking only external connections
- C. Requiring biometric authentication only for the initial device unlock at the start of day
- D. Continuous device posture assessment and risk-based conditional access for every request

Answer: D

Q1216. What is the security risk of mobile deep links and how can they be exploited?

- A. Deep links expose the application source code to reverse engineering through URL scheme analysis
- B. Deep links increase application size making mobile devices slower and consuming excessive storage
- C. Attackers craft malicious deep links to redirect users to phishing pages or trigger unauthorized actions
- D. Deep links prevent the application from receiving push notifications from the backend server

Answer: C

Q1217. How does a Wi-Fi pineapple device facilitate wireless attacks?

- A. It generates random WiFi networks to confuse nearby devices and prevent them from connecting
- B. It automates evil twin, deauth, and man-in-the-middle attacks on wireless networks for penetration testing
- C. It increases WiFi signal range allowing attackers to connect to distant corporate wireless networks
- D. It blocks all wireless signals in an area creating a denial of service for all wireless devices

Answer: B

Q1218. An enterprise deploys EMM (Enterprise Mobility Management). How does EMM differ from basic MDM?

- A. EMM only manages device hardware while MDM manages both hardware and installed applications
- B. EMM replaces all security controls while MDM supplements existing security infrastructure
- C. EMM includes MDM plus mobile application management, content management, and identity management
- D. EMM is exclusively for iOS devices while MDM supports only Android mobile operating systems

Answer: C

Q1219. What is the security concern with wireless mesh networks in IoT deployments?

- A. Each mesh node is a potential attack point and compromise can propagate through the entire network
- B. Mesh networks provide too much bandwidth making it easy for attackers to exfiltrate large data
- C. Mesh networks cannot be encrypted because the routing protocol requires plaintext communication
- D. Mesh networks are immune to all wireless attacks because they use frequency hopping technology

Answer: A

Q1220. A mobile app development team implements runtime application self-protection. How does this specifically benefit mobile security?

- A. RASP detects and prevents tampering, debugging, and exploitation attempts within the running app
- B. RASP automatically updates the mobile application without requiring app store review approval
- C. RASP replaces the need for encryption by preventing all unauthorized access at the OS level
- D. RASP eliminates the need for any server-side security controls by handling everything on device

Answer: A

Q1221. An organization discovers that an attacker accessed their cloud environment using stolen API keys from a public GitHub repository. What series of actions is MOST appropriate?

- A. Banning all developers from using version control systems to prevent future key exposure
- B. Migrating all workloads to a different cloud provider to prevent similar future incidents
- C. Immediately rotating compromised keys, auditing access logs, implementing secret scanning, and reviewing damage
- D. Only deleting the exposed keys from the GitHub repository and continuing normal operations

Answer: C

Q1222. How does a cloud-native SIEM differ from a traditional on-premises SIEM?

- A. Cloud SIEM is identical to traditional SIEM except it runs on virtual machines in the cloud
- B. Cloud SIEM provides less security visibility because it cannot access on-premises network traffic
- C. Cloud SIEM only works with a single cloud provider and cannot ingest logs from other sources
- D. Cloud SIEM scales elastically, integrates with cloud APIs, and reduces infrastructure management overhead

Answer: D

Q1223. What is the primary security risk of cross-account access in multi-account cloud environments?

- A. Cross-account access always requires physical network connections between cloud data centers
- B. Cross-account access is inherently secure because cloud providers enforce strict isolation
- C. Cross-account access eliminates all identity management because users share a single identity
- D. Overly permissive trust policies between accounts can enable lateral movement across the organization

Answer: D

Q1224. A company uses Kubernetes in production. What is the MOST critical security concern for the cluster?

- A. Disabling all logging to improve cluster performance and reduce storage costs for log files
- B. Allowing all pods to run as root to simplify application deployment and avoid permission issues
- C. Ensuring the cluster uses the maximum number of nodes to handle any distributed denial of service
- D. Securing the API server, implementing RBAC, network policies, and scanning container images for vulnerabilities

Answer: D

Q1225. What is the confused deputy problem in cloud security?

- A. A network routing issue where traffic intended for one VPC is accidentally sent to a different one
- B. A vulnerability where a trusted service is tricked into performing unauthorized actions on behalf of an attacker
- C. A bug in cloud billing systems that charges the wrong customer account for resource consumption
- D. A situation where multiple cloud administrators have conflicting security policies applied simultaneously

Answer: B

Q1226. How does infrastructure as code (IaC) introduce security risks in cloud deployments?

- A. IaC only affects development environments and has no impact on production security posture
- B. IaC templates can codify misconfigurations that are then consistently deployed across all environments
- C. IaC templates are automatically scanned by cloud providers preventing any security misconfigurations
- D. IaC eliminates all security risks because automated deployments are inherently more secure

Answer: B

Q1227. An organization implements a multi-cloud strategy. What is the PRIMARY security challenge this introduces?

- A. Reducing the total cost of cloud services by negotiating volume discounts with providers
- B. Synchronizing time zones between different cloud provider regions for consistent logging
- C. Maintaining consistent security policies and visibility across different cloud provider environments
- D. Ensuring all cloud providers use identical hardware specifications in their data centers

Answer: C

Q1228. What is the security implication of instance metadata services (IMDS) in cloud environments?

- A. IMDS is only accessible from the cloud provider's management console and not from instances
- B. SSRF attacks can access IMDS to retrieve instance credentials and escalate privileges in the cloud
- C. IMDS automatically rotates all credentials making them useless even if they are intercepted
- D. IMDS provides no security-relevant information and is only used for instance performance metrics

Answer: B

Q1229. What is the purpose of cloud security benchmarks like CIS Benchmarks for cloud providers?

- A. Comparing the pricing and cost structures of competing cloud service provider offerings
- B. Providing prescriptive security configuration guidelines specific to each cloud provider's services
- C. Measuring the network performance and latency of different cloud provider data center regions
- D. Rating the customer support quality of different cloud providers based on user surveys

Answer: B

Q1230. A cloud architect needs to protect sensitive workloads from the cloud provider's own administrators. What technology BEST addresses this?

- A. Cloud provider managed keys which give the provider full access to encrypted data
- B. Standard server-side encryption which protects data at rest from external attackers only
- C. Confidential computing using hardware-based trusted execution environments for data in use
- D. Network encryption using TLS for all data in transit between cloud services and users

Answer: C

Q1231. A forensic investigator encounters full disk encryption on a suspect's device. What is the MOST effective approach to access the data?

- A. Ignoring the encrypted drive and relying solely on other unencrypted evidence sources
- B. Capturing memory while the device is running to extract encryption keys before shutdown
- C. Using brute force to crack the encryption key which is guaranteed to succeed quickly
- D. Contacting the encryption software vendor to request a master decryption backdoor key

Answer: B

Q1232. During a forensic investigation, anti-forensic techniques are discovered. Which technique makes timestamp analysis unreliable?

- A. Timestomping which modifies file timestamps to mislead investigators about event timing
- B. Disk wiping which securely overwrites all data making recovery completely impossible
- C. File encryption which prevents investigators from reading the content of seized files
- D. Steganography which hides data within image files to prevent content-based detection

Answer: A

Q1233. A forensic team must analyze a cloud-based system. What unique challenge does cloud forensics present compared to traditional forensics?

- A. Cloud forensics requires no specialized tools because standard disk imaging tools always work
- B. Limited physical access to infrastructure, multi-tenancy, and jurisdiction complexities across regions
- C. Cloud systems always store data in plaintext making the analysis process significantly faster
- D. Cloud providers automatically preserve all forensic evidence eliminating the need for imaging

Answer: B

Q1234. An investigator analyzes Windows registry hives. What type of forensic evidence can be found in the NTUSER.DAT file?

- A. System-wide network configuration settings including IP addresses and DNS server addresses
- B. User-specific data including recent documents, typed URLs, mounted devices, and application history
- C. Operating system installation date, service pack level, and licensed software product keys
- D. Hardware inventory data including processor model, memory capacity, and disk serial numbers

Answer: B

Q1235. How does memory forensics help in analyzing fileless malware attacks?

- A. Memory analysis only works on Linux systems and cannot be applied to Windows-based attacks
- B. Memory forensics is unnecessary for fileless malware since it leaves traces on the hard drive
- C. Memory analysis can reveal injected code, suspicious processes, and network connections invisible on disk
- D. Memory forensics requires the suspect device to be powered off before analysis can begin

Answer: C

Q1236. A forensic investigation involves analyzing a compromised Docker container. What specific challenges does container forensics present?

- A. Container ephemeral nature means evidence is lost when containers are destroyed or recreated
- B. Containers cannot be compromised because they run in completely isolated security sandboxes
- C. Containers always persist their complete state to disk making forensic imaging straightforward
- D. Container forensics is identical to virtual machine forensics with no additional complexity

Answer: A

Q1237. What is the forensic significance of the Windows Event Log entry ID 4624?

- A. It records successful logon events including logon type, source address, and authentication details
- B. It records when new user accounts are created by administrators in the domain directory
- C. It records when a user's account is locked out after exceeding maximum login attempt limits
- D. It records failed logon attempts with the incorrect password that was entered by the user

Answer: A

Q1238. An investigator uses the Volatility framework for memory analysis. What primary capability does this tool provide?

- A. Creating forensic images of hard drives using bit-level copying with hash verification
- B. Generating automated forensic reports from evidence files without any manual analysis
- C. Analyzing memory dumps to extract processes, network connections, registry hives, and malware artifacts
- D. Scanning network traffic captures for signatures of known malware communication patterns

Answer: C

Q1239. How does browser forensics help in investigating cyber crimes?

- A. Browser data is automatically deleted when the browser is closed making forensics impossible
- B. Browser forensics requires the user's cooperation to provide their browsing history voluntarily
- C. Browser artifacts reveal browsing history, cached pages, cookies, downloaded files, and stored credentials
- D. Browser forensics can only recover the URLs of websites visited within the last twenty-four hours

Answer: C

Q1240. What is the Locard's Exchange Principle and how does it apply to digital forensics?

- A. All digital evidence must be exchanged between prosecution and defense before trial begins
- B. Evidence can only be collected from the suspect's primary device and not from any other source
- C. Every digital interaction leaves traces on both systems, providing evidence of contact and activity
- D. Digital evidence loses its value over time and must be analyzed within seventy-two hours

Answer: C

Q1241. During a major breach, the IR team discovers the attacker has compromised Active Directory. What is the MOST critical containment action?

- A. Resetting the KRBTGT account password twice and all privileged account passwords immediately
- B. Disconnecting the domain controller from the network and waiting for the vendor's support
- C. Simply changing the domain administrator password and continuing the ongoing investigation
- D. Rebuilding all workstations from scratch before addressing the domain controller compromise

Answer: A

Q1242. An organization experiences a supply chain compromise affecting a widely used library. How should the IR team prioritize their response?

- A. Waiting for the library vendor to release a patch before taking any response actions at all
- B. Only investigating systems that have already shown signs of compromise and ignoring others
- C. Immediately removing the library from all systems regardless of the impact on business operations
- D. Identifying all systems using the compromised library, assessing exposure, and implementing mitigations

Answer: D

Q1243. During incident response, the team suspects the attacker is monitoring their communications. What is the BEST approach?

- A. Switching to out-of-band communication channels not accessible to the attacker for coordination
- B. Broadcasting false information through compromised channels to mislead and confuse the attacker
- C. Stopping all communication until the attacker has been fully removed from all systems
- D. Continuing to use normal email and chat systems while being careful about what is discussed

Answer: A

Q1244. How should an IR team handle a situation where the attacker has deployed multiple backdoors across the network?

- A. Removing backdoors one at a time over several weeks to minimize disruption to operations
- B. Mapping all identified backdoors, planning simultaneous removal, and monitoring for re-establishment
- C. Negotiating with the attacker to voluntarily remove the backdoors from compromised systems
- D. Ignoring the backdoors and focusing solely on strengthening perimeter firewall rule sets

Answer: B

Q1245. An organization faces a data breach requiring notification under GDPR. What is the maximum timeframe for notifying the supervisory authority?

- A. Within 30 calendar days of discovering the personal data breach event
- B. Within 72 hours of becoming aware of the personal data breach event
- C. Within 7 business days of completing the full investigation of the breach
- D. Within 24 hours of the breach occurring regardless of discovery timing

Answer: B

Q1246. What is the primary benefit of implementing SOAR (Security Orchestration, Automation, and Response) in IR?

- A. Replacing all human analysts with automated systems that handle every type of incident alone
- B. Providing unlimited storage for security event logs without any cost or capacity limitations
- C. Eliminating the need for incident response plans because automation handles everything directly
- D. Automating repetitive response tasks and orchestrating tools to reduce response time and errors

Answer: D

Q1247. An IR team must decide whether to monitor an attacker or immediately contain the breach. What factors should guide this decision?

- A. Ongoing risk to data, attacker's current activity level, legal obligations, and intelligence value
- B. Only the cost of the monitoring equipment needed to track the attacker's network activities
- C. Whether the attack is occurring during business hours or after regular working time periods
- D. The personal preferences of the incident response team lead regarding investigation approach

Answer: A

Q1248. How does a threat hunting program complement traditional incident response?

- A. Replacing incident response entirely with continuous proactive threat elimination processes
- B. Proactively searching for threats that evade automated detection before they cause incidents
- C. Only hunting for threats after an incident has been detected and reported by monitoring
- D. Focusing exclusively on external threats while ignoring all potential insider threat actors

Answer: B

Q1249. During a breach investigation, the IR team finds the initial access occurred six months ago. What does this indicate about the organization's detection capabilities?

- A. Six months of undetected access is normal and expected for all organizations of similar size
- B. Significant detection gaps exist since the attacker maintained undetected access for an extended period
- C. The detection systems were working correctly but were intentionally configured to delay alerts
- D. The organization's detection capabilities are excellent because they eventually found the breach

Answer: B

Q1250. An organization's IR plan includes both technical and executive response procedures. Why is executive involvement critical during major incidents?

- A. Executives perform the technical analysis because they have the deepest technical expertise
- B. Executives authorize resource allocation, make business decisions, and manage external communications
- C. Executives are only needed to sign off on the final incident report after full resolution
- D. Executives should not be involved because their involvement slows down the response process

Answer: B

Q1251. A SOC team needs to reduce mean time to detect (MTTD) threats. Which approach is MOST effective?

- A. Increasing the alert threshold to generate fewer alerts so analysts can review each one carefully
- B. Hiring more analysts to manually review every single log entry generated across all systems
- C. Implementing automated detection with tuned correlation rules, ML models, and threat intelligence integration
- D. Reducing the number of monitored data sources to focus only on firewall logs for simplicity

Answer: C

Q1252. How does deception technology enhance security monitoring beyond traditional honeypots?

- A. Encrypting all network traffic to prevent attackers from understanding the network architecture
- B. Blocking all network traffic from unknown sources before it reaches any internal network system
- C. Replacing all production systems with decoy systems to eliminate the real attack surface completely
- D. Deploying decoy assets throughout the network that generate high-fidelity alerts on any interaction

Answer: D

Q1253. A SIEM generates thousands of alerts daily but only a small fraction are true positives. What is the BEST approach to improve alert quality?

- A. Tuning detection rules, implementing alert scoring, and using contextual enrichment to prioritize alerts
- B. Increasing the alert volume by adding more detection rules without evaluating existing ones
- C. Disabling all alert rules and relying entirely on manual log review by security analysts
- D. Forwarding all alerts directly to management without any analyst review or prioritization

Answer: A

Q1254. What is the primary challenge of monitoring encrypted traffic in a SOC?

- A. Encrypted traffic cannot be captured by any network monitoring tools available on the market today
- B. Encrypted traffic hides payload content requiring alternative approaches like metadata analysis or decryption
- C. Encrypted traffic is identical to unencrypted traffic from a monitoring perspective in all cases
- D. Encrypted traffic always indicates malicious activity and should be blocked by default at all times

Answer: B

Q1255. How should a SOC measure the effectiveness of its detection capabilities?

- A. Counting the total number of alerts generated per day regardless of their quality or accuracy
- B. Measuring only the number of incidents that resulted in data loss across the entire organization
- C. Tracking the number of security tools deployed without evaluating their actual effectiveness
- D. Using metrics like MTTD, MTTR, detection coverage mapped to ATT&CK, and false positive rates

Answer: D

Q1256. An organization implements a data lake for security analytics. What advantage does this provide over traditional SIEM?

- A. Eliminating the need for any structured queries because data lakes organize data automatically
- B. Providing real-time alerting capabilities that are not possible with any traditional SIEM tools
- C. Storing massive volumes of raw data at lower cost enabling long-term analysis and threat hunting
- D. Replacing all security analysts with automated systems that process data without human input

Answer: C

Q1257. What is the concept of 'purple teaming' and how does it improve security monitoring?

- A. Collaboration between red and blue teams to improve detection rules and response procedures together
- B. A compliance framework that requires organizations to conduct biannual security assessments only
- C. A third-party team that replaces both offensive and defensive teams with automated testing tools
- D. A training program exclusively for security operations center analysts to improve their soft skills

Answer: A

Q1258. A SOC analyst notices beaconing behavior in network traffic. What does this MOST likely indicate?

- A. A legitimate application checking for software updates from the vendor's update server
- B. Malware communicating with a command-and-control server at regular intervals for instructions
- C. Network equipment sending heartbeat messages to confirm device connectivity and status
- D. A user streaming video content which creates periodic network traffic burst patterns

Answer: B

Q1259. How does security monitoring need to adapt for serverless and containerized environments?

- A. Monitoring must shift to application-level telemetry, API calls, and runtime behavior analysis
- B. Serverless environments require no monitoring because cloud providers handle all security aspects
- C. Traditional host-based monitoring tools work identically in serverless and container environments
- D. Container monitoring only needs to check container image signatures at deployment time only

Answer: A

Q1260. What is the role of AI and machine learning in modern security monitoring?

- A. Completely replacing all human security analysts with fully autonomous threat response systems
- B. Detecting complex patterns and anomalies at scale that human analysts and rules cannot identify
- C. Only generating marketing reports about security posture for executive board presentations
- D. Slowing down threat detection because ML models require extensive computation time always

Answer: B

Q1261. An organization operates across multiple countries with different data protection laws. What is the MOST effective approach to compliance?

- A. Storing all data in a single country to avoid dealing with multiple legal jurisdictions entirely
- B. Ignoring all data protection laws because it is impossible to comply with all of them simultaneously
- C. Implementing the strictest applicable standard as a baseline and adding jurisdiction-specific requirements
- D. Complying only with the laws of the country where the organization's headquarters is located

Answer: C

Q1262. A security researcher discovers a critical vulnerability in widely-used software. Under what circumstances could disclosure become legally problematic?

- A. If the researcher is employed by a competing software company in the same industry sector
- B. If the vulnerability affects more than one hundred users across multiple different countries
- C. If the researcher uses automated scanning tools rather than performing manual code review only
- D. If disclosure occurs without vendor notification, outside a bug bounty scope, or violates CFAA provisions

Answer: D

Q1263. How does the concept of 'privacy by design' differ from traditional approaches to data privacy?

- A. Privacy by design embeds privacy into systems from the start rather than adding it after development
- B. Privacy by design only applies to mobile applications while traditional approaches cover all systems
- C. Privacy by design eliminates the need for privacy policies by making systems inherently compliant
- D. Privacy by design requires hiring a dedicated privacy team while traditional approaches do not

Answer: A

Q1264. An organization faces a conflict between a government's lawful data access request and GDPR data protection requirements. How should this be handled?

- A. Deleting all requested data to prevent both the government and regulators from accessing it
- B. Automatically complying with the government request and ignoring all data protection requirements
- C. Refusing all government requests regardless of their legal basis or jurisdictional authority
- D. Engaging legal counsel to evaluate the request, applying appropriate legal frameworks, and documenting decisions

Answer: D

Q1265. What is the legal significance of the Schrems II ruling for international data transfers?

- A. It created a new encryption standard that must be used for all international data transfers
- B. It permanently banned all data transfers between the European Union and the United States
- C. It established that cloud providers are solely responsible for all international data compliance
- D. It invalidated Privacy Shield and requires additional safeguards for EU-US personal data transfers

Answer: D

Q1266. How do whistleblower protection laws intersect with cyber security?

- A. They mandate that all security incidents be reported to the media for public transparency
- B. They require whistleblowers to have formal cyber security certifications before making reports
- C. They protect individuals who report security violations or illegal activities from employer retaliation
- D. They only protect whistleblowers in government organizations and not in private sector companies

Answer: C

Q1267. What legal challenges does the use of AI in cyber security decision-making present?

- A. AI in cyber security is universally prohibited by international law across all jurisdictions globally
- B. AI decisions may lack transparency, create liability issues, and potentially introduce discriminatory bias
- C. AI decisions are always legally defensible because they are based on objective data processing only
- D. AI decision-making eliminates all legal liability because responsibility shifts to the algorithm vendor

Answer: B

Q1268. An employee uses company resources to conduct unauthorized security testing on a third party. What legal liability does the organization face?

- A. The organization is only liable if the testing successfully finds vulnerabilities in the third party
- B. The organization may be vicariously liable for the employee's unauthorized actions using company resources
- C. The organization is automatically protected by its acceptable use policy regardless of circumstances
- D. The organization has no liability because unauthorized actions are solely the employee's responsibility

Answer: B

Q1269. How does the concept of 'digital sovereignty' affect cloud computing and data storage decisions?

- A. Digital sovereignty eliminates the need for data encryption because national borders provide protection
- B. Digital sovereignty only applies to government data and has no impact on commercial organizations
- C. Digital sovereignty is a theoretical concept with no practical legal implications for organizations
- D. Countries require data to be stored within their borders, affecting cloud architecture and provider selection

Answer: D

Q1270. What is the legal framework for active cyber defense measures and 'hack back' by private organizations?

- A. Hack-back is legal as long as the organization documents the attacker's IP address beforehand
- B. Private sector hack-back is generally illegal under most jurisdictions, with very limited exceptions
- C. International law provides blanket immunity for all defensive cyber operations by organizations
- D. Organizations have unrestricted legal authority to attack systems used to breach their networks

Answer: B

Q1271. How does quantum key distribution (QKD) fundamentally change secure communication?

- A. QKD eliminates the need for key management because quantum keys never expire or need rotation
- B. QKD replaces all existing encryption algorithms with faster quantum-based alternatives for performance
- C. QKD only works over distances shorter than one meter making it impractical for any real deployment
- D. QKD uses quantum mechanics to detect any eavesdropping attempt making interception theoretically detectable

Answer: D

Q1272. An organization considers implementing confidential computing. What specific security problem does this technology solve?

- A. Improving network transmission security by implementing quantum-resistant encryption protocols everywhere
- B. Protecting data during processing in hardware-isolated enclaves even from privileged system administrators
- C. Encrypting data at rest using stronger algorithms than currently available standard encryption methods
- D. Preventing physical theft of servers by implementing advanced biometric access controls on hardware

Answer: B

Q1273. How does the concept of 'shift everywhere' extend beyond traditional 'shift left' in DevSecOps?

- A. Security responsibility shifts entirely to developers with no involvement from security teams
- B. Security is applied at every stage including design, development, deployment, and runtime operations
- C. All security testing is moved to the earliest possible development stage and never done elsewhere
- D. Security testing shifts to production only because that is where real threats actually manifest

Answer: B

Q1274. What is the security implication of neuromorphic computing for future cyber threats?

- A. Neuromorphic computing has no relevance to cyber security and only affects medical applications
- B. Neuromorphic computers cannot be programmed for malicious purposes due to hardware constraints
- C. Brain-inspired computing could enable more sophisticated AI attacks that mimic human decision patterns
- D. Neuromorphic computing eliminates all cyber threats by processing data like a biological brain

Answer: C

Q1275. How does the concept of software-defined perimeter (SDP) improve upon traditional VPN security?

- A. SDP hides infrastructure from unauthorized users entirely and provides least-privilege application access
- B. SDP replaces all firewalls and IDS systems with a single unified network security appliance
- C. SDP provides faster network speeds than VPN by eliminating all encryption overhead completely
- D. SDP is only applicable to on-premises networks and cannot be used with cloud environments

Answer: A

Q1276. What are the security challenges of Web3 and decentralized applications?

- A. Decentralized applications are inherently secure because they have no single point of failure
- B. Web3 security is identical to Web2 security with no new threat models or attack surfaces
- C. Smart contract vulnerabilities, private key management, and irreversible transactions create unique risks
- D. Web3 eliminates all security risks through decentralization making traditional attacks impossible

Answer: C

Q1277. How does adversarial machine learning pose a threat to AI-powered security systems?

- A. Adversarial ML only affects image recognition systems and has no impact on security applications
- B. Adversarial ML improves security by training models to be more robust against future attacks
- C. Attackers craft inputs that deceive ML models into misclassifying malicious activity as benign traffic
- D. ML models are immune to adversarial inputs because they learn from millions of training examples

Answer: C

Q1278. What is the 'harvest now, decrypt later' threat in the context of quantum computing?

- A. Organizations harvest threat intelligence now to decrypt future attack patterns using AI analysis
- B. Security teams collect malware samples now for decryption and analysis at a later scheduled date
- C. Adversaries collect encrypted data now planning to decrypt it when quantum computers become available
- D. Cloud providers harvest encryption keys now to simplify key management processes in the future

Answer: C

Q1279. How does the concept of cyber threat intelligence platforms evolve with emerging technologies?

- A. Platforms only focus on historical threat data and cannot incorporate real-time intelligence feeds
- B. Platforms integrate AI for automated analysis, predict attack trends, and enable proactive defense strategies
- C. Platforms become obsolete because automated security tools eliminate the need for threat intelligence
- D. Platforms are replaced entirely by blockchain-based solutions that distribute intelligence automatically

Answer: B

Q1280. What is the emerging concept of cyber insurance and how does it impact organizational security posture?

- A. Cyber insurance only covers physical damage to hardware and does not cover data breach expenses
- B. Cyber insurance premiums are identical regardless of the organization's security maturity level
- C. Cyber insurance transfers financial risk but requires demonstrating minimum security controls for coverage
- D. Cyber insurance eliminates the need for security controls because all losses are financially covered

Answer: C

Q1281. An organization implements a risk-based approach to cyber security. How should residual risk be handled after applying security controls?

- A. Residual risk should be completely eliminated before proceeding
- B. Residual risk should be documented, accepted by management, and continuously monitored
- C. Residual risk does not exist if proper controls are applied
- D. Residual risk is only relevant during the initial risk assessment

Answer: B

Q1282. How does the Diamond Model of intrusion analysis differ from the Cyber Kill Chain?

- A. The Diamond Model focuses only on malware classification
- B. The Diamond Model analyzes intrusions using four core features (adversary, infrastructure, capability, victim) while the Kill Chain describes attack phases sequentially
- C. The Cyber Kill Chain is only applicable to insider threats
- D. The Diamond Model is a compliance framework

Answer: B

Q1283. What is the primary challenge of implementing a cyber security governance framework in a multinational organization?

- A. Multinational organizations are too small for governance frameworks
- B. Balancing global security standards with varying local regulations, cultural differences, and regional threat landscapes
- C. Governance frameworks are only designed for single-country operations
- D. The only challenge is language translation of policies

Answer: B

Q1284. In cyber security risk quantification, what does the FAIR (Factor Analysis of Information Risk) model provide that qualitative approaches do not?

- A. FAIR eliminates the need for risk assessments entirely
- B. FAIR provides a quantitative framework for measuring risk in financial terms using probability and loss magnitude
- C. FAIR only measures physical security risks
- D. FAIR replaces all other security frameworks

Answer: B

Q1285. An organization's CISO must present the return on security investment (ROSI) to the board. Which approach BEST demonstrates security value?

- A. Listing all security tools purchased
- B. Correlating security spending with reduction in incident frequency, severity, and potential financial losses avoided
- C. Comparing the number of employees in the security team to other departments
- D. Only showing compliance checkbox completions

Answer: B

Q1286. What is the key difference between cyber security and cyber resilience?

- A. They are identical concepts with different names
- B. Cyber security focuses on preventing and protecting against attacks, while cyber resilience encompasses the ability to continue operating during and recover quickly after an attack
- C. Cyber resilience only applies to cloud systems
- D. Cyber security is outdated and replaced by cyber resilience

Answer: B

Q1287. When conducting a cyber security gap analysis, what should be the primary benchmark?

- A. The security budget of competing organizations
- B. Industry-specific frameworks and regulatory requirements mapped against current security posture
- C. The number of security tools deployed
- D. Employee satisfaction surveys about security policies

Answer: B

Q1288. How does the concept of 'security debt' parallel 'technical debt' in software development?

- A. Security debt only affects open-source projects
- B. Deferred security fixes and shortcuts accumulate over time, increasing the cost and complexity of future remediation, similar to technical debt
- C. Security debt refers to financial debts from purchasing security tools
- D. Security debt is only relevant during audits

Answer: B

Q1289. What challenge does the convergence of Information Technology (IT) and Operational Technology (OT) create for cyber security programs?

- A. IT and OT systems have identical security requirements
- B. OT systems often have different lifecycles, patching constraints, and availability requirements than IT systems, requiring tailored security approaches
- C. Convergence eliminates all security risks
- D. OT systems are inherently more secure than IT systems

Answer: B

Q1290. A board of directors asks for a cyber risk appetite statement. What should this statement primarily define?

- A. The exact number of acceptable security incidents per year
- B. The level and types of cyber risk the organization is willing to accept in pursuit of its strategic objectives
- C. A list of all security tools the organization uses
- D. The cyber security team's annual training schedule

Answer: B

Q1291. An organization performs a quantitative risk analysis and calculates an SLE of \$200,000 and an ARO of 0.5. What is the Annual Loss Expectancy (ALE) and how should it guide security investment?

- A. ALE is \$400,000 and any control costing less should be implemented
- B. ALE is \$100,000 and security controls costing less than this amount may be justified
- C. ALE is \$200,000 because ARO is irrelevant
- D. ALE cannot be calculated from SLE and ARO alone

Answer: B

Q1292. How does the Brewer-Nash (Chinese Wall) model address conflicts of interest in access control?

- A. It prevents users from reading any data
- B. It dynamically restricts access based on a user's previous access history to prevent conflicts of interest
- C. It only applies to financial institutions
- D. It uses mandatory access control labels exclusively

Answer: B

Q1293. What is the security principle of 'open design' and why is it considered stronger than security through obscurity?

- A. Open design means making all passwords public
- B. Open design means a system's security should not depend on the secrecy of its implementation, as peer review strengthens security
- C. Open design prohibits the use of encryption
- D. Open design means using only open-source operating systems

Answer: B

Q1294. When implementing a defense-in-depth strategy, why is diversity of controls important across layers?

- A. Diverse controls are cheaper to implement
- B. Using different types of controls at each layer ensures that a single vulnerability or exploit technique cannot bypass multiple layers simultaneously
- C. Diverse controls are easier to manage
- D. Using identical controls at each layer provides better consistency

Answer: B

Q1295. An organization must choose between risk avoidance, risk mitigation, risk transfer, and risk acceptance. Under what circumstances is risk acceptance the most appropriate strategy?

- A. When the cost of mitigation exceeds the potential impact and the risk falls within the organization's risk appetite
- B. Risk acceptance is never appropriate in cyber security
- C. When regulatory requirements mandate specific controls
- D. When the risk involves critical infrastructure systems

Answer: A

Q1296. How does the Graham-Denning model address the secure creation and deletion of objects and subjects in access control?

- A. It only defines read and write permissions
- B. It defines eight protection rules governing how subjects and objects can be securely created, deleted, and how access rights can be transferred
- C. It is limited to network access control only
- D. It replaces the need for authentication mechanisms

Answer: B

Q1297. A security architect must evaluate the trade-off between usability and security. Which approach BEST balances these competing requirements?

- A. Always maximize security regardless of usability impact
- B. Implement risk-based controls that apply stronger security to higher-risk activities while minimizing friction for low-risk operations
- C. Always maximize usability regardless of security impact
- D. Apply the same level of security to all operations regardless of risk

Answer: B

Q1298. What is the Harrison-Ruzzo-Ullman (HRU) model's contribution to access control theory?

- A. It only addresses network segmentation
- B. It provides a formal framework for analyzing whether a given access control system can reach an unsafe state, proving that the general safety problem is undecidable
- C. It replaces all other access control models
- D. It only applies to physical access control systems

Answer: B

Q1299. Why is the concept of 'trust but verify' being replaced by 'never trust, always verify' in modern security architectures?

- A. Because trust but verify is easier to implement
- B. Because the traditional perimeter-based model assumed internal network trust, which modern threats like insider attacks and lateral movement have proven inadequate
- C. Because never trust always verify is less expensive
- D. Because trust but verify only applies to physical security

Answer: B

Q1300. An organization discovers that its security controls address individual risks but fail to account for compound or cascading risks. What concept addresses this gap?

- A. Single point of failure analysis
- B. Systemic risk analysis, which evaluates how individual risks can combine and cascade to create larger, emergent risks across interconnected systems
- C. Simple vulnerability scanning
- D. Individual risk acceptance for each identified risk

Answer: B

Q1301. A security architect must choose between RSA-2048 and ECC P-256 for a resource-constrained IoT deployment. Which is more suitable and why?

- A. RSA-2048 because it uses larger keys
- B. ECC P-256 because it provides equivalent security to RSA-2048 with much smaller keys and lower computational overhead
- C. Neither can be used in IoT devices
- D. RSA-2048 because it is newer than ECC

Answer: B

Q1302. What is the security vulnerability of using AES in GCM mode with a reused nonce?

- A. It has no effect on security
- B. Nonce reuse in AES-GCM can reveal the authentication key and allow forgery of ciphertexts, catastrophically breaking both confidentiality and authenticity
- C. It only affects performance, not security
- D. It causes the encryption to become stronger

Answer: B

Q1303. How does lattice-based cryptography provide resistance against quantum computing attacks?

- A. It uses longer RSA keys
- B. Lattice-based problems like Learning With Errors (LWE) are believed to be hard for both classical and quantum computers to solve
- C. It relies on the same mathematical problems as RSA
- D. It only works on quantum computers

Answer: B

Q1304. What is the primary security risk when using CBC mode encryption without a MAC (encrypt-only)?

- A. CBC mode cannot encrypt data properly
- B. It is vulnerable to padding oracle attacks, where an attacker can decrypt ciphertext by manipulating padding and observing error responses
- C. CBC mode without MAC uses too much bandwidth
- D. There is no security risk with encrypt-only CBC

Answer: B

Q1305. Why is the Dual_EC_DRBG random number generator considered compromised, and what broader lesson does it teach?

- A. It generates numbers that are too random
- B. It contained suspected NSA backdoor constants that could allow prediction of output, teaching that cryptographic standards must be transparent and verifiable
- C. It was too slow for practical use
- D. It only worked on specific hardware

Answer: B

Q1306. In a TLS 1.3 handshake, what fundamental change was made compared to TLS 1.2 regarding cipher suites?

- A. TLS 1.3 no longer uses encryption
- B. TLS 1.3 removed support for static RSA key exchange and non-AEAD ciphers, mandating forward secrecy and authenticated encryption
- C. TLS 1.3 only supports symmetric encryption
- D. TLS 1.3 increased the handshake from two to four round trips

Answer: B

Q1307. What is the security implication of Shor's algorithm for current public key cryptographic systems?

- A. It makes symmetric encryption obsolete
- B. It can efficiently factor large numbers and compute discrete logarithms on a quantum computer, breaking RSA and ECC
- C. It only affects hash functions
- D. It strengthens existing public key systems

Answer: B

Q1308. What is the difference between deterministic and probabilistic encryption, and why does it matter for security?

- A. They produce identical outputs
- B. Deterministic encryption always produces the same ciphertext for the same plaintext and key, while probabilistic encryption includes randomness producing different ciphertexts, which prevents pattern leakage
- C. Deterministic encryption is always more secure
- D. Probabilistic encryption cannot be decrypted

Answer: B

Q1309. What is the security purpose of key commitment in authenticated encryption, and what attacks does its absence enable?

- A. Key commitment is only about key storage efficiency
- B. Without key commitment, an attacker can craft a ciphertext that decrypts to valid but different plaintexts under different keys, enabling invisible salamander attacks
- C. Key commitment prevents brute force attacks
- D. Key commitment is required only for symmetric ciphers

Answer: B

Q1310. A cryptographic implementation passes all functional tests but is found vulnerable to a timing side-channel attack. What is the root cause and mitigation?

- A. The algorithm itself is broken and must be replaced
- B. The implementation takes different amounts of time for different inputs, leaking information. Mitigation requires constant-time implementation of cryptographic operations
- C. Timing attacks only affect hardware, not software
- D. Using a faster processor eliminates timing attacks

Answer: B

Q1311. An attacker performs a VLAN hopping attack using double tagging. What network misconfiguration enables this attack?

- A. Using encrypted VLANs
- B. Having the native VLAN of trunk ports be the same as a user VLAN, allowing double-tagged frames to traverse VLAN boundaries
- C. Using too many VLANs on the network
- D. Having DHCP enabled on the network

Answer: B

Q1312. How does a network detection and response (NDR) solution differ from traditional IDS/IPS systems?

- A. NDR only monitors email traffic
- B. NDR uses behavioral analytics and machine learning to detect unknown threats based on network behavior patterns, while traditional IDS/IPS primarily relies on signatures
- C. NDR replaces firewalls entirely
- D. Traditional IDS/IPS is always more effective than NDR

Answer: B

Q1313. What is the security risk of running recursive DNS resolvers open to the internet?

- A. Recursive DNS resolvers cannot be accessed from the internet
- B. Open recursive resolvers can be abused for DNS amplification DDoS attacks and DNS cache poisoning, and may leak information about internal queries
- C. Open resolvers improve DNS performance globally
- D. Recursive DNS resolvers are inherently secure

Answer: B

Q1314. A company implements TLS inspection (SSL decryption) on their network firewall. What privacy and security trade-offs must be considered?

- A. TLS inspection has no trade-offs
- B. TLS inspection enables detection of encrypted threats but breaks end-to-end encryption, requires careful certificate management, and raises privacy concerns for employee communications
- C. TLS inspection improves end-to-end encryption
- D. TLS inspection eliminates the need for endpoint security

Answer: B

Q1315. What makes a Slowloris attack effective against web servers, and how does it differ from volumetric DDoS attacks?

- A. Slowloris uses high bandwidth to overwhelm servers
- B. Slowloris keeps many connections open by sending partial HTTP requests slowly, exhausting the server's connection pool with minimal bandwidth, unlike volumetric attacks that flood with traffic
- C. Slowloris only targets DNS servers
- D. Slowloris attacks are always detectable by firewalls

Answer: B

Q1316. How does the QUIC protocol change the traditional network security monitoring approach?

- A. QUIC uses unencrypted UDP packets that are easier to monitor
- B. QUIC encrypts transport metadata that was previously visible in TCP+TLS, reducing the effectiveness of traditional network monitoring that relied on inspecting TCP headers and TLS handshake details
- C. QUIC only works on IPv4 networks
- D. QUIC eliminates the need for any encryption

Answer: B

Q1317. What is the security implication of IPv6 dual-stack network configurations if IPv6 security controls are not properly implemented?

- A. IPv6 is inherently secure and needs no additional controls
- B. Attackers can use IPv6 to bypass IPv4 security controls if the network has IPv6 enabled but not properly monitored or firewalled, creating a shadow network path
- C. Dual-stack configurations always improve security
- D. IPv6 cannot coexist with IPv4 on the same network

Answer: B

Q1318. What is a network covert channel and how can DNS be used to create one?

- A. A covert channel is an encrypted VPN tunnel
- B. A covert channel uses legitimate network protocols in unintended ways to secretly communicate. DNS tunneling encodes data within DNS queries and responses to bypass network controls
- C. Covert channels require dedicated hardware
- D. DNS cannot be used for covert communication

Answer: B

Q1319. An organization implements a software-defined perimeter (SDP). How does this differ from traditional VPN-based remote access?

- A. SDP and VPN are identical technologies
- B. SDP authenticates users before granting network visibility and creates individual, encrypted connections to specific resources, while VPNs grant broad network access after authentication
- C. SDP requires no authentication
- D. VPNs provide more granular access control than SDP

Answer: B

Q1320. What is the security risk of protocol downgrade attacks in network communications?

- A. Downgrade attacks improve security by using older, more tested protocols
- B. Attackers force communications to use weaker, older protocol versions with known vulnerabilities, bypassing the security of newer protocols even when both endpoints support them
- C. Downgrade attacks only affect physical layer protocols
- D. Modern protocols are immune to downgrade attacks

Answer: B

Q1321. In a padding oracle attack against CBC mode encryption, what information does the attacker exploit?

- A. The encryption key leaked through side channels
- B. The server's response indicating whether decrypted padding is valid or invalid
- C. The predictability of the initialization vector
- D. The weakness of the block cipher itself

Answer: B

Q1322. What is the security advantage of X25519 over traditional ECDH for key exchange?

- A. X25519 uses a larger key size
- B. X25519 is resistant to timing attacks by design and avoids common implementation pitfalls
- C. X25519 uses symmetric encryption internally
- D. X25519 does not require any computation

Answer: B

Q1323. What is the primary challenge of implementing threshold cryptography in a distributed system?

- A. It requires quantum computers
- B. The need to securely distribute key shares while ensuring no single party can reconstruct the full key alone
- C. It only works with RSA encryption
- D. It doubles the key length required

Answer: B

Q1324. How does a commitment scheme work in cryptographic protocols, and what two properties must it satisfy?

- A. Speed and efficiency
- B. Hiding (the committed value is secret) and binding (the committer cannot change the value after committing)
- C. Encryption and decryption
- D. Authentication and authorization

Answer: B

Q1325. What is the security risk of using AES-CBC with a MAC computed over the ciphertext versus over the plaintext?

- A. There is no difference between the two approaches
- B. MAC-then-encrypt can be vulnerable to padding oracle attacks, while encrypt-then-MAC is generally secure
- C. MAC over ciphertext is always insecure
- D. Both approaches are equally vulnerable

Answer: B

Q1326. What is the significance of the discrete logarithm problem in modern cryptography?

- A. It is only used in hash functions
- B. It is easily solvable with modern computers
- C. It provides the mathematical hardness underlying Diffie-Hellman, DSA, and elliptic curve cryptography
- D. It has been proven to be unsolvable

Answer: C

Q1327. A system uses ECDSA for digital signatures. What catastrophic failure occurs if the same random nonce k is used for two different signatures with the same private key?

- A. The signatures become too large to verify
- B. Nothing happens as the nonce does not affect security
- C. The private key can be mathematically recovered from the two signatures
- D. The public key becomes invalid

Answer: C

Q1328. What is the fundamental difference between zero-knowledge proofs and traditional authentication mechanisms?

- A. Zero-knowledge proofs are faster than password authentication
- B. Zero-knowledge proofs allow a prover to demonstrate knowledge of a secret without revealing any information about the secret itself
- C. Zero-knowledge proofs require biometric data
- D. Zero-knowledge proofs eliminate the need for cryptographic keys

Answer: B

Q1329. In post-quantum cryptography, why are lattice-based schemes considered promising candidates for standardization?

- A. They use simple XOR operations only
- B. They offer strong security proofs from worst-case hardness assumptions and support diverse functionalities including encryption, signatures, and FHE
- C. They require no computational resources
- D. They are identical to RSA but with larger keys

Answer: B

Q1330. What is format-preserving encryption (FPE) and in what scenario is it particularly useful?

- A. Encryption that always produces the same output format as the input, useful for encrypting data that must retain its format such as credit card numbers in legacy databases
- B. Encryption that compresses data before encrypting
- C. Encryption that formats output as JSON
- D. Encryption that preserves file extensions

Answer: A

Q1331. How does Kernel Address Space Layout Randomization (KASLR) differ from standard ASLR, and what are its limitations?

- A. KASLR and ASLR are identical with no differences
- B. KASLR randomizes kernel memory addresses but can be bypassed through information disclosure vulnerabilities that leak kernel pointers
- C. KASLR only applies to user-space programs
- D. KASLR provides perfect protection against all kernel attacks

Answer: B

Q1332. What is a TOCTOU (Time of Check to Time of Use) race condition, and how can it be exploited in OS security?

- A. A type of brute-force attack on passwords
- B. A vulnerability where the state of a resource changes between checking its properties and using it, allowing privilege escalation through symlink attacks
- C. A network timing attack on encryption
- D. A denial-of-service attack using timestamps

Answer: B

Q1333. A security team must harden a Linux server. What is the security benefit of implementing seccomp-bpf?

- A. It encrypts network traffic
- B. It filters system calls available to processes, reducing the kernel attack surface by restricting which syscalls a program can make
- C. It manages user passwords
- D. It monitors disk usage

Answer: B

Q1334. What is the Dirty COW vulnerability and what class of OS security issue does it represent?

- A. A malware variant targeting IoT devices
- B. A race condition in the Linux kernel's copy-on-write mechanism that allowed unprivileged users to write to read-only memory mappings for privilege escalation
- C. A Windows registry corruption bug
- D. A DNS spoofing technique

Answer: B

Q1335. How does hardware-based memory tagging (like ARM MTE) improve OS security compared to software-only solutions?

- A. It replaces the need for an operating system
- B. It assigns tags to memory regions and pointers, detecting use-after-free and buffer overflow at hardware speed with minimal performance overhead
- C. It only protects GPU memory
- D. It eliminates the need for encryption

Answer: B

Q1336. What is the security architecture of a microkernel-based OS and why is it considered more secure than a monolithic kernel?

- A. Microkernels are faster and thus more secure
- B. Microkernels run minimal code in kernel mode with drivers and services in user space, reducing the trusted computing base and limiting the impact of vulnerabilities
- C. Microkernels use more encryption
- D. Microkernels eliminate all vulnerabilities

Answer: B

Q1337. An organization detects a sophisticated rootkit using Direct Kernel Object Manipulation (DKOM). What makes this technique particularly dangerous?

- A. It corrupts the bootloader permanently
- B. DKOM modifies kernel data structures directly in memory to hide processes and connections without modifying files on disk, evading traditional detection methods
- C. It only affects network interfaces
- D. It can be detected easily by antivirus software

Answer: B

Q1338. What is the security purpose of Intel SGX (Software Guard Extensions) and what attacks has it proven vulnerable to?

- A. SGX accelerates encryption with no known weaknesses
- B. SGX creates hardware-encrypted enclaves for sensitive computation but has been shown vulnerable to side-channel attacks like Spectre, Foreshadow, and power analysis
- C. SGX is a network security protocol
- D. SGX only protects data at rest

Answer: B

Q1339. How does mandatory integrity control (MIC) in Windows protect against privilege escalation?

- A. It encrypts all system files
- B. It assigns integrity levels to processes and objects, preventing lower-integrity processes from writing to higher-integrity objects even with appropriate DACLs
- C. It disables all non-Microsoft software
- D. It monitors network traffic

Answer: B

Q1340. What is a Spectre-class vulnerability and why is it fundamentally difficult to mitigate at the OS level?

- A. It is a simple buffer overflow fixed by input validation
- B. Spectre exploits speculative execution in modern CPUs to leak data across security boundaries, and full mitigation requires microcode updates and significant performance trade-offs because the vulnerability is in hardware design
- C. It only affects 32-bit operating systems
- D. It is a type of malware that can be removed by antivirus

Answer: B

Q1341. What is a deserialization gadget chain and how is it used in insecure deserialization attacks?

- A. A blockchain verification mechanism
- B. A sequence of existing classes in an application whose methods can be chained together during deserialization to achieve arbitrary code execution
- C. A series of database queries optimized for performance
- D. A method of encrypting serialized objects

Answer: B

Q1342. How can a web cache poisoning attack be executed, and what makes it particularly dangerous?

- A. By deleting the server's cache files
- B. By manipulating unkeyed HTTP headers or parameters to store malicious responses in the cache, which are then served to all subsequent users
- C. By overloading the cache with too many requests
- D. By encrypting cached content

Answer: B

Q1343. What is a second-order SQL injection and why is it harder to detect than standard SQL injection?

- A. SQL injection performed by two attackers simultaneously
- B. Malicious input stored in the database and later executed when retrieved and used in another SQL query without sanitization
- C. SQL injection that only works on the second attempt
- D. A SQL injection that targets backup databases

Answer: B

Q1344. A web application implements CSP with 'unsafe-inline' for scripts. What attack vectors does this leave open and how can nonce-based CSP improve security?

- A. It has no impact on security
- B. 'unsafe-inline' allows any inline scripts to execute, defeating XSS protection. Nonce-based CSP allows only scripts with a matching server-generated nonce, blocking injected scripts
- C. Nonce-based CSP is less secure than unsafe-inline
- D. unsafe-inline only affects CSS, not scripts

Answer: B

Q1345. What is a Server-Side Request Forgery (SSRF) attack through DNS rebinding, and how does it bypass common SSRF mitigations?

- A. It exploits DNS caching to speed up the server
- B. The attacker controls a domain that first resolves to an allowed IP, passes validation, then resolves to an internal IP on the actual request, bypassing allowlist/blocklist checks
- C. It only works against DNS servers themselves
- D. It replaces the server's DNS configuration

Answer: B

Q1346. How does a prototype pollution vulnerability in JavaScript enable remote code execution in server-side applications?

- A. By crashing the JavaScript engine
- B. By modifying Object.prototype to inject properties that alter application logic, which in template engines or child process spawning can lead to command injection
- C. By changing the JavaScript language syntax
- D. By corrupting the package.json file

Answer: B

Q1347. What is HTTP request smuggling via CL.TE desync and what infrastructure conditions make it possible?

- A. It is a type of DDoS attack
- B. It occurs when a front-end server uses Content-Length and a back-end uses Transfer-Encoding, causing them to disagree on request boundaries, allowing an attacker to smuggle requests
- C. It requires the attacker to have server access
- D. It only affects HTTP/2 connections

Answer: B

Q1348. What is the security implication of GraphQL introspection being enabled in production and how can it be exploited?

- A. GraphQL introspection improves security by documenting the API
- B. Introspection reveals the entire API schema including types, fields, mutations, and relationships, giving attackers a complete map for targeted attacks on sensitive operations
- C. Introspection only affects performance
- D. GraphQL does not support introspection

Answer: B

Q1349. A security researcher finds a web application vulnerable to blind XSS. How does blind XSS differ from reflected XSS and why is it particularly dangerous?

- A. Blind XSS is less dangerous than reflected XSS
- B. Blind XSS payloads are stored and execute in a different application context (like an admin panel) that the attacker cannot directly access, potentially compromising privileged accounts
- C. Blind XSS only works in Internet Explorer
- D. Blind XSS requires physical access to the server

Answer: B

Q1350. What is a WebSocket hijacking attack and how does the lack of same-origin policy enforcement for WebSockets create a vulnerability?

- A. WebSockets cannot be hijacked due to built-in encryption
- B. Cross-site WebSocket hijacking occurs because WebSocket connections do not enforce same-origin policy; a malicious page can initiate a WebSocket connection to a vulnerable server using the victim's cookies
- C. WebSocket hijacking requires modifying network hardware
- D. WebSockets automatically validate all origins

Answer: B

Q1351. How does the WebAuthn protocol provide phishing-resistant authentication?

- A. By using longer passwords
- B. By binding credentials to the origin (domain) so that credentials created for one site cannot be used on a phishing site with a different domain
- C. By requiring users to type the URL manually
- D. By encrypting the login page

Answer: B

Q1352. What is a token binding and how does it prevent token theft attacks?

- A. Binding tokens to a physical location
- B. Cryptographically binding security tokens to the TLS connection so that stolen tokens cannot be used from a different TLS session
- C. Using stronger passwords for token generation
- D. Limiting the number of tokens a user can have

Answer: B

Q1353. An organization implements a zero trust architecture. How should the authentication model change from traditional perimeter-based security?

- A. Authentication is only needed at the perimeter
- B. Every access request must be authenticated, authorized, and continuously validated regardless of network location, assuming no implicit trust
- C. Internal users should not need authentication
- D. Authentication frequency should be reduced for convenience

Answer: B

Q1354. What is a silver ticket attack in Active Directory and how does it differ from a golden ticket attack?

- A. Silver tickets are less harmful versions of golden tickets
- B. A silver ticket forges a service ticket for a specific service using the service account's hash, while a golden ticket forges a TGT using the KRBTGT hash granting access to all services
- C. Silver tickets target Linux systems while golden tickets target Windows
- D. Silver tickets expire faster but golden tickets never expire

Answer: B

Q1355. How does continuous adaptive risk and trust assessment (CARTA) improve access control beyond static policies?

- A. By removing all access controls
- B. By continuously evaluating risk signals during sessions and dynamically adjusting access levels based on changing context and behavior, not just at the point of initial authentication
- C. By requiring re-authentication every minute
- D. By using only biometric authentication

Answer: B

Q1356. What are the security implications of the OAuth 2.0 implicit flow and why was it deprecated for SPAs?

- A. Implicit flow is the most secure OAuth flow
- B. The implicit flow returns tokens in the URL fragment, exposing them to browser history, referrer headers, and potential token leakage; PKCE-enhanced authorization code flow is now recommended for SPAs
- C. Implicit flow was deprecated because it is too slow
- D. SPAs cannot use OAuth at all

Answer: B

Q1357. What is the security risk of JWT algorithm confusion attacks and how do they work?

- A. They slow down JWT verification
- B. An attacker changes the JWT header algorithm from RS256 to HS256, causing the server to verify the signature using the RSA public key as the HMAC secret, allowing token forgery
- C. They only affect expired tokens
- D. They require access to the private key

Answer: B

Q1358. How does decentralized identity (DID) change the traditional identity management paradigm?

- A. By centralizing all identity data in one database
- B. By giving individuals control over their own identity credentials through self-sovereign identity, eliminating dependence on centralized identity providers
- C. By removing the need for authentication entirely
- D. By using only username and password

Answer: B

Q1359. What is a relay attack on NFC-based authentication and how does it bypass physical proximity requirements?

- A. It blocks all NFC signals within an area
- B. An attacker uses two devices to relay the NFC communication between a legitimate token and reader over a distance, making the reader believe the token is nearby when it is not
- C. It speeds up NFC communication
- D. It only works against Bluetooth devices

Answer: B

Q1360. What is the principle behind passwordless authentication using FIDO2 and why is it considered more secure than traditional passwords?

- A. It uses simpler passwords that are easier to remember
- B. FIDO2 replaces passwords with public-key cryptography where the private key never leaves the device, eliminating password databases, phishing, replay attacks, and credential stuffing
- C. It stores passwords in the cloud instead of locally
- D. It only works on specific operating systems

Answer: B

Q1361. What is a supply chain attack through a compromised build pipeline, and how does code signing mitigate it?

- A. An attack on physical supply logistics
- B. An attacker compromises the CI/CD pipeline to inject malicious code into builds; code signing creates a cryptographic proof of build integrity, allowing consumers to verify the software was not tampered with
- C. Code signing prevents all supply chain attacks completely
- D. Build pipelines cannot be compromised

Answer: B

Q1362. How does taint tracking analysis work in identifying security vulnerabilities?

- A. It tracks paint colors in the UI
- B. It marks user-controlled data as 'tainted' and traces its flow through the application to detect if it reaches security-sensitive operations (sinks) without proper sanitization
- C. It monitors CPU temperature
- D. It tracks changes in version control

Answer: B

Q1363. What is the SLSA (Supply-chain Levels for Software Artifacts) framework and what problem does it address?

- A. A programming language specification
- B. A framework defining increasing levels of supply chain security for software artifacts, addressing integrity and provenance from source to build to distribution
- C. A type of encryption standard
- D. A software testing methodology

Answer: B

Q1364. A development team discovers a deserialization vulnerability in their Java application. What makes Java deserialization attacks particularly dangerous?

- A. Java deserialization is slower than other languages
- B. Java's classpath often contains libraries with exploitable gadget chains that allow arbitrary code execution when crafted objects are deserialized
- C. Java does not support serialization
- D. Deserialization only affects database operations

Answer: B

Q1365. What is the security significance of reproducible builds and how do they protect the software supply chain?

- A. Reproducible builds make software faster
- B. Reproducible builds ensure that compiling the same source code always produces identical binary output, allowing independent verification that distributed binaries match the source and have not been tampered with
- C. Reproducible builds only work for open-source software
- D. They are identical to version control

Answer: B

Q1366. How does mutation testing evaluate the effectiveness of a security test suite?

- A. By testing the application with random inputs
- B. By introducing small changes (mutations) to the source code and checking whether the test suite detects them; surviving mutants indicate gaps in test coverage
- C. By mutating the database schema
- D. By changing the programming language

Answer: B

Q1367. What is a confused deputy vulnerability in application development and how can capability-based security prevent it?

- A. When a developer is confused about requirements
- B. A privilege escalation where a trusted program (deputy) is tricked into misusing its authority; capability-based security prevents it by requiring explicit unforgeable tokens for each resource access
- C. A type of race condition
- D. A denial-of-service attack on the development server

Answer: B

Q1368. What is the security impact of improper certificate validation in mobile and desktop applications?

- A. It only affects loading speed
- B. Improper certificate validation can allow man-in-the-middle attacks by accepting invalid, expired, or self-signed certificates, enabling interception of encrypted communications
- C. It improves security by accepting more certificates
- D. It only affects web browsers

Answer: B

Q1369. What are the security implications of using WebAssembly (WASM) in web applications and what unique threats does it introduce?

- A. WASM has no security implications
- B. WASM can be used to obfuscate malicious code, perform cryptomining, bypass traditional JavaScript security analysis, and introduce memory safety vulnerabilities from languages like C/C++
- C. WASM is inherently more secure than JavaScript
- D. WASM cannot interact with web APIs

Answer: B

Q1370. How does a security-focused chaos engineering approach differ from traditional chaos engineering?

- A. They are identical practices
- B. Security chaos engineering deliberately introduces security-relevant failures (certificate expiration, credential leaks, firewall rule removal) to verify that security controls respond correctly under adverse conditions
- C. Security chaos engineering only tests network outages
- D. It replaces the need for penetration testing

Answer: B

Q1371. How does the KRACK (Key Reinstallation Attack) exploit the WPA2 four-way handshake, and what specifically is compromised?

- A. It cracks the WiFi password directly
- B. It forces reinstallation of an already-in-use key by manipulating handshake messages, causing nonce reuse that allows decryption of packets and potential injection in TKIP and GCMP modes
- C. It disables the access point's encryption hardware
- D. It only works against WEP encryption

Answer: B

Q1372. What is the FragAttacks vulnerability class in WiFi security and why does it affect virtually all WiFi devices?

- A. It only affects devices manufactured before 2010
- B. FragAttacks comprises design flaws in WiFi's frame aggregation and fragmentation features present in the IEEE 802.11 standard since 1997, along with widespread implementation flaws, affecting nearly all WiFi devices
- C. It is specific to one router manufacturer
- D. It only works when WiFi encryption is disabled

Answer: B

Q1373. How do Stingray/IMSI catcher devices operate, and what defense mechanisms exist against them?

- A. They simply block cell signals
- B. They impersonate legitimate cell towers to force phones to connect, intercepting calls, texts, and tracking location; 5G standalone networks with SUPI concealment help mitigate identification attacks
- C. They only work within 1 meter range
- D. They are only used for WiFi attacks

Answer: B

Q1374. What is the security risk of Bluetooth Low Energy (BLE) tracking and how have recent attacks like BLESAs exploited BLE reconnection?

- A. BLE is too short-range to pose security risks
- B. BLESAs (BLE Spoofing Attacks) exploits the reconnection process where authentication is optional, allowing an attacker to impersonate a previously paired device and send spoofed data
- C. BLE attacks require physical contact with the device
- D. BLE tracking only works indoors

Answer: B

Q1375. What are the unique security challenges of eSIM compared to traditional physical SIM cards?

- A. eSIMs are identical to physical SIMs in security
- B. eSIMs introduce risks of remote provisioning attacks, profile manipulation, and remove the physical possession requirement, but also eliminate SIM swapping at carrier stores
- C. eSIMs cannot be used for authentication
- D. eSIMs are less secure in every aspect

Answer: B

Q1376. How does the Dragonblood attack affect WPA3-SAE (Simultaneous Authentication of Equals) security?

- A. WPA3 is completely immune to all attacks
- B. Dragonblood reveals side-channel vulnerabilities in WPA3-SAE's dragonfly handshake including timing-based and cache-based attacks that can recover the WiFi password, plus downgrade attacks to WPA2
- C. Dragonblood only affects WPA2 networks
- D. Dragonblood is a physical attack on routers

Answer: B

Q1377. What is a man-in-the-disk attack on Android and why is it effective despite Android's sandboxing?

- A. An attack that fills up the device storage
- B. An attack exploiting Android apps that use external storage (shared by all apps) for receiving data, allowing a malicious app to modify files before the target app processes them
- C. It bypasses device encryption
- D. It only works on rooted devices

Answer: B

Q1378. What security mechanisms does 5G introduce to address the vulnerabilities found in 4G LTE networks?

- A. 5G has no security improvements over 4G
- B. 5G introduces subscriber identity encryption (SUCI), enhanced key hierarchy, network slicing isolation, and edge computing security, addressing IMSI catching, false base stations, and subscriber tracking
- C. 5G only improves speed without security changes
- D. 5G uses the same security as WiFi

Answer: B

Q1379. What is a baseband processor attack and why is it considered one of the most severe mobile security threats?

- A. It is a type of app-level malware that can be removed easily
- B. Baseband attacks target the separate processor handling cellular communications, which runs its own RTOS with high privileges, potentially allowing call interception, SMS theft, and persistent access that survives factory resets
- C. It only affects WiFi connectivity
- D. It requires the attacker to have physical access to the device

Answer: B

Q1380. How does a WiFi direct connection present security risks different from traditional WiFi, and what attack scenarios are unique to WiFi Direct?

- A. WiFi Direct is more secure than traditional WiFi in all cases
- B. WiFi Direct creates peer-to-peer connections that bypass network security controls; attacks include rogue group owner impersonation, unauthorized device discovery, and exploiting WPS vulnerabilities in the setup process
- C. WiFi Direct cannot transfer data
- D. WiFi Direct uses wired connections

Answer: B

Q1381. What is a cloud instance metadata service (IMDS) attack and how did it contribute to the Capital One breach?

- A. IMDS attacks target cloud provider's headquarters
- B. An attacker exploited an SSRF vulnerability to query the instance metadata service at 169.254.169.254, obtaining temporary IAM credentials that gave access to S3 buckets containing customer data
- C. IMDS only contains non-sensitive information
- D. The Capital One breach was caused by a phishing attack

Answer: B

Q1382. How does a cross-account confused deputy attack work in cloud environments?

- A. It confuses cloud administrators about their account credentials
- B. A malicious customer tricks a cloud service (deputy) into using the service's cross-account permissions to access resources in a victim's account by providing the victim's account ID as the target
- C. It only affects personal cloud accounts
- D. It is a type of DDoS attack

Answer: B

Q1383. What is the security challenge of Kubernetes RBAC in multi-tenant cloud environments?

- A. Kubernetes does not support RBAC
- B. Kubernetes RBAC can be complex to configure correctly for multi-tenancy, with risks of overly permissive ClusterRoles, insecure default service accounts, and namespace escape through cluster-wide resources
- C. Kubernetes RBAC is identical to OS-level access control
- D. Multi-tenancy is not possible in Kubernetes

Answer: B

Q1384. What is the security purpose of confidential computing in cloud environments and how does it differ from traditional encryption?

- A. It is the same as encrypting data at rest
- B. Confidential computing protects data during processing (in use) using hardware-based trusted execution environments, ensuring not even the cloud provider can access data during computation
- C. It only encrypts network traffic
- D. It provides encryption for backup data only

Answer: B

Q1385. How does cloud infrastructure entitlements management (CIEM) address the permissions gap problem?

- A. CIEM manages cloud infrastructure hardware
- B. CIEM continuously analyzes actual permission usage versus granted permissions, identifying overly permissive access and recommending least-privilege policies based on real behavior patterns
- C. CIEM replaces IAM entirely
- D. CIEM only manages storage permissions

Answer: B

Q1386. What are the security risks of serverless function event injection and how does it differ from traditional injection attacks?

- A. Serverless functions cannot be attacked
- B. Serverless event injection exploits the various event sources (API Gateway, S3, SNS, DynamoDB) that trigger functions; each source can carry malicious payloads in different formats, expanding the attack surface beyond traditional HTTP input
- C. Serverless injection is identical to web injection
- D. Only API Gateway events can be exploited

Answer: B

Q1387. What is the security significance of cloud provider control plane versus data plane, and how should monitoring differ for each?

- A. Control plane and data plane are the same in cloud environments
- B. The control plane manages cloud resources (API calls, configuration) while the data plane handles actual data traffic; monitoring must cover both: control plane for unauthorized changes and data plane for exfiltration and anomalies
- C. Only the data plane needs monitoring
- D. The control plane is automatically secure

Answer: B

Q1388. How can an attacker exploit insecure cloud function dependencies to achieve persistent access in a serverless environment?

- A. Serverless environments prevent all dependency attacks
- B. An attacker can compromise a dependency in the function's package or inject a malicious layer, gaining code execution each time the function is invoked, achieving persistence without maintaining infrastructure
- C. Dependencies do not affect function security
- D. Only first-party code can be compromised

Answer: B

Q1389. What is a cloud-native application protection platform (CNAPP) and how does it unify previously separate security tools?

- A. CNAPP is a new cloud provider service
- B. CNAPP combines CSPM, CWPP, CIEM, and IaC scanning into a unified platform, providing end-to-end visibility across the application lifecycle from code to cloud runtime, reducing tool sprawl and alert correlation gaps
- C. CNAPP only provides antivirus for cloud
- D. CNAPP replaces the need for IAM

Answer: B

Q1390. What security challenges arise from cloud provider managed keys versus customer managed keys (CMK) for encryption?

- A. There is no difference between managed and customer keys
- B. Provider-managed keys simplify operations but give the provider access to key material; CMKs give customers control but require managing key rotation, access policies, and disaster recovery, with the risk of permanent data loss if keys are lost
- C. Customer keys are always less secure
- D. Provider keys cannot be used for compliance

Answer: B

Q1391. How does SSD wear leveling and TRIM complicate traditional forensic data recovery?

- A. SSDs are easier to forensically analyze than HDDs
- B. TRIM instructs the SSD to erase deleted data blocks for performance, while wear leveling moves data across cells, making traditional recovery of deleted files unreliable as data may be erased or relocated without notice
- C. SSDs store data identically to traditional hard drives
- D. TRIM and wear leveling only affect SSD performance, not forensics

Answer: B

Q1392. What are the forensic challenges of investigating incidents in containerized environments?

- A. Containers are identical to virtual machines for forensics
- B. Containers are ephemeral with short lifespans, share the host kernel, use layered filesystems, and may be automatically destroyed and replaced, making traditional disk imaging impractical and requiring real-time log collection and orchestrator-level evidence
- C. Container forensics uses the same tools as disk forensics
- D. Containers always preserve all evidence

Answer: B

Q1393. How does memory forensics using tools like Volatility detect rootkits that hide from traditional analysis?

- A. Memory forensics cannot detect rootkits
- B. Volatility analyzes RAM dumps to find discrepancies between OS-reported process lists and actual memory structures, detect hidden processes, injected code, and hooks that rootkits use to hide from live system tools
- C. Memory forensics only finds cached passwords
- D. Volatility is a network analysis tool

Answer: B

Q1394. What is the forensic significance of Windows Prefetch files and how do they aid in malware analysis?

- A. Prefetch files only speed up application loading
- B. Prefetch files record which executables were run, when they were last executed, run count, and which DLLs and files were loaded, providing evidence of malware execution even after the malware binary is deleted
- C. Prefetch files are encrypted and cannot be analyzed
- D. Prefetch files exist only on Linux systems

Answer: B

Q1395. What are the legal and technical challenges of performing forensics on encrypted devices with full disk encryption?

- A. Encryption has no impact on forensic investigation
- B. Full disk encryption prevents traditional dead-box forensic imaging of readable data, requiring either live acquisition while the device is running and unlocked, key recovery from memory or key escrow, or legal compulsion to provide the decryption key
- C. Encrypted devices can always be easily decrypted by forensic tools
- D. Encryption only protects deleted files

Answer: B

Q1396. How does NTFS journal forensics (\$UsnJrnl and \$LogFile) provide evidence that regular file analysis cannot?

- A. NTFS journals contain no useful forensic data
- B. The USN Journal records all file system changes (creation, modification, deletion, rename) with timestamps, and \$LogFile records transactional metadata, providing evidence of file operations even after files are deleted and overwritten
- C. NTFS journals only exist on Linux systems
- D. Journal analysis requires access to the original user account

Answer: B

Q1397. What is the role of digital forensics in investigating Advanced Persistent Threat (APT) campaigns?

- A. Digital forensics cannot help with APT investigations
- B. Forensics must trace lateral movement across multiple systems, identify initial access vectors, recover C2 communications, analyze custom malware through reverse engineering, and build a comprehensive timeline spanning months of activity
- C. APT investigations only require network monitoring
- D. APTs leave no forensic evidence

Answer: B

Q1398. What is the forensic approach to analyzing fileless malware that operates entirely in memory?

- A. Fileless malware cannot be detected or analyzed
- B. Analysis requires memory acquisition from the live system, examining PowerShell logs, WMI event subscriptions, registry-resident payloads, and system event logs since traditional disk-based forensics finds no malware files
- C. Fileless malware always leaves files on disk
- D. Simply scanning the hard drive is sufficient

Answer: B

Q1399. How does anti-forensic timestamp manipulation (timestomping) work, and what techniques can detect it?

- A. Timestamps cannot be manipulated
- B. Timestomping modifies file timestamps to blend in with legitimate files; detection involves comparing \$STANDARD_INFORMATION timestamps with \$FILE_NAME timestamps in NTFS MFT, analyzing USN Journal entries, and checking for timestamp inconsistencies
- C. Only law enforcement can detect timestamp manipulation
- D. Timestomping permanently destroys all timestamp evidence

Answer: B

Q1400. What are the unique challenges of cloud forensics compared to traditional on-premises forensics?

- A. Cloud forensics is simpler than traditional forensics
- B. Cloud forensics faces challenges including limited physical access, multi-tenant data isolation, jurisdiction across regions, volatile/ephemeral resources, dependency on provider cooperation, and ensuring forensic soundness without hardware control
- C. Cloud environments provide better forensic capabilities
- D. There are no differences between cloud and traditional forensics

Answer: B

Q1401. During a major incident, the IR team suspects the attacker is monitoring their Slack and email. What should the team do?

- A. Continue using Slack and email to avoid tipping off the attacker
- B. Switch to pre-established out-of-band communication channels such as a separate system, encrypted phone calls, or in-person meetings that the attacker has no access to
- C. Post misleading information on Slack to confuse the attacker
- D. Shut down all communication and work independently

Answer: B

Q1402. How should an IR team handle the discovery that a supply chain partner's software update was the initial attack vector?

- A. Immediately remove all partner software from the network
- B. Coordinate with the partner for joint investigation, isolate affected systems without breaking critical dependencies, analyze the compromised update for IoCs, and notify relevant stakeholders and potentially CISA or sector ISACs
- C. Ignore the supply chain aspect and focus only on internal systems
- D. Publicly blame the supply chain partner

Answer: B

Q1403. An organization detects ransomware but the attacker also exfiltrated data before encryption. How does this dual-threat change the incident response approach?

- A. Focus only on decrypting the ransomware
- B. The response must address both the ransomware (restoration from backups, containment) and the data breach (regulatory notification, determining what data was stolen, monitoring for data publication), treating it as a data breach with ransomware
- C. Data exfiltration is not a concern if data is encrypted on the attacker's end
- D. Only law enforcement can handle this type of incident

Answer: B

Q1404. What is the 'dwell time' metric in incident response and what does a high dwell time indicate about an organization's detection capabilities?

- A. Dwell time measures how long it takes to deploy patches
- B. Dwell time is the duration between initial compromise and detection; high dwell time indicates poor detection capabilities, allowing attackers to deeply entrench, move laterally, and accomplish objectives before being discovered
- C. Dwell time measures the time between incidents
- D. Dwell time is only relevant for insider threats

Answer: B

Q1405. How should an IR team determine whether to rebuild compromised systems from scratch versus cleaning and restoring them?

- A. Always clean and restore to save time
- B. Rebuilding is preferred for high-severity compromises where the attacker had deep access (admin/root), persistence mechanisms may be hidden, or the full extent of compromise is uncertain; cleaning is acceptable for low-severity, well-understood incidents
- C. Always rebuild every system regardless of the incident
- D. The decision should be made by the legal department only

Answer: B

Q1406. What is the role of deception technology during active incident response?

- A. Deception has no place in incident response
- B. Deploying decoys and honeypots during active response can detect attacker lateral movement, misdirect the attacker, gather intelligence on TTPs, and confirm when the attacker has been fully removed from the environment
- C. Deception technology only works before an incident occurs
- D. Deception technology replaces the need for containment

Answer: B

Q1407. What are the legal considerations when an IR team discovers evidence of criminal activity by an insider during a routine security investigation?

- A. Legal considerations do not apply during incident response
- B. The team must immediately involve legal counsel, ensure evidence is collected in a legally defensible manner, consider employee privacy rights, determine whether to involve law enforcement, and separate the HR process from the technical investigation
- C. The IR team should confront the employee directly
- D. The team should delete the evidence to avoid complications

Answer: B

Q1408. How does the MITRE ATT&CK framework enhance incident response activities?

- A. MITRE ATT&CK is only useful for threat intelligence, not incident response
- B. ATT&CK provides a common language for describing attacker techniques, helps map observed activity to known threat groups, identifies gaps in detection coverage, guides investigation by predicting likely next steps, and structures post-incident reporting
- C. ATT&CK replaces the need for an incident response plan
- D. ATT&CK only works for nation-state attacks

Answer: B

Q1409. An IR team has contained a breach but the attacker had access for 6 months. What specific challenges does this extended dwell time create for the response?

- A. Extended dwell time does not create additional challenges
- B. Extended access means the attacker likely established multiple persistence mechanisms, moved laterally across many systems, exfiltrated significant data, potentially tampered with logs and backups, and created credentials that must all be identified and remediated
- C. After 6 months the threat is no longer active
- D. Only the initially compromised system needs attention

Answer: B

Q1410. What is the strategic difference between tactical incident response and strategic incident response during a major breach?

- A. There is no difference between tactical and strategic response
- B. Tactical response focuses on immediate technical actions (containment, eradication, recovery) while strategic response addresses business continuity, stakeholder communication, regulatory compliance, reputation management, and long-term security improvements
- C. Tactical response is done by management and strategic by engineers
- D. Strategic response only happens after the incident is resolved

Answer: B

Q1411. How can an attacker evade security monitoring through living-off-the-land techniques, and what detection strategies are effective?

- A. Living-off-the-land techniques are easily detected by antivirus
- B. Attackers use legitimate system tools (PowerShell, WMI, certutil) for malicious purposes, evading signature-based detection; effective detection requires behavioral analysis, command-line logging, and baselining normal administrative tool usage
- C. These techniques only work on Linux systems
- D. Standard SIEM rules detect all living-off-the-land activity

Answer: B

Q1412. What is the detection challenge of encrypted C2 (command and control) traffic, and what network-level indicators can still reveal it?

- A. Encrypted C2 traffic cannot be detected at all
- B. While content is hidden, encrypted C2 traffic can be detected through JA3/JA3S fingerprinting, certificate analysis, beacon interval patterns, traffic volume anomalies, SNI field inspection, and unusual destination characteristics
- C. Simply decrypting all traffic is the only solution
- D. C2 traffic always uses unencrypted protocols

Answer: B

Q1413. How should security monitoring architecture be designed for a hybrid cloud environment?

- A. Use only on-premises monitoring tools for all environments
- B. Implement unified monitoring that aggregates logs from on-premises, cloud provider native logs, and SaaS audit logs into a centralized SIEM with normalized data formats, cloud-specific detection rules, and cross-environment correlation
- C. Cloud environments do not need monitoring
- D. Use separate, unconnected monitoring for each environment

Answer: B

Q1414. What is the OODA loop (Observe, Orient, Decide, Act) and how does it apply to SOC operations?

- A. OODA is a software development methodology
- B. The OODA loop provides a decision-making framework for SOC analysts: Observe (collect data), Orient (analyze context and threat intelligence), Decide (determine response), Act (execute response); faster OODA loops lead to better security outcomes
- C. OODA only applies to military operations
- D. OODA replaces SIEM technology

Answer: B

Q1415. What are the key metrics for measuring SOC effectiveness, and what do they indicate?

- A. Only the number of alerts matters
- B. Key metrics include MTTD (Mean Time To Detect), MTTR (Mean Time To Respond), false positive rate, alert-to-incident ratio, coverage against ATT&CK techniques, and analyst workload; together they measure detection capability, response efficiency, and operational health
- C. SOC effectiveness cannot be measured
- D. Only compliance audit results matter

Answer: B

Q1416. How does security orchestration, automation, and response (SOAR) reduce mean time to respond?

- A. SOAR replaces all human analysts
- B. SOAR automates repetitive investigation steps (enrichment, correlation), orchestrates actions across security tools, and executes predefined playbooks for common alert types, freeing analysts for complex decisions while reducing response time for routine incidents
- C. SOAR only generates reports
- D. SOAR increases response time due to complexity

Answer: B

Q1417. What is the challenge of monitoring east-west traffic in modern data center and cloud environments?

- A. East-west traffic does not exist in modern networks
- B. East-west (lateral) traffic between servers often bypasses traditional perimeter monitoring; detecting threats requires micro-segmentation with inline monitoring, agent-based endpoint telemetry, or network detection and response (NDR) solutions within the data center
- C. East-west traffic is automatically monitored by firewalls
- D. Only north-south traffic matters for security

Answer: B

Q1418. What is the role of machine learning anomaly detection in security monitoring, and what are its limitations?

- A. ML anomaly detection perfectly detects all threats with no false positives
- B. ML anomaly detection identifies deviations from learned baseline behavior, useful for detecting novel threats without signatures; limitations include high false positive rates during baseline changes, susceptibility to adversarial evasion, and need for quality training data
- C. ML makes all other detection methods obsolete
- D. ML anomaly detection only works for network traffic

Answer: B

Q1419. How should a SOC handle monitoring during a zero-day vulnerability disclosure before patches are available?

- A. Zero-day monitoring is impossible
- B. The SOC should implement compensating detection rules based on known exploitation indicators, increase monitoring of likely target systems, deploy virtual patches through WAF/IPS, monitor for proof-of-concept exploit code releases, and coordinate with threat intelligence for IoCs
- C. Simply wait for the vendor patch
- D. Disconnect all vulnerable systems from the network

Answer: B

Q1420. What is the concept of detection-as-code and how does it improve security monitoring operations?

- A. Detection-as-code means writing detection rules in any programming language
- B. Detection-as-code applies software engineering practices (version control, testing, peer review, CI/CD) to detection rule development, enabling systematic rule management, automated testing for false positives, and collaborative detection engineering
- C. Detection-as-code replaces SIEM platforms
- D. It only applies to cloud-native environments

Answer: B

Q1421. How does the CLOUD Act affect international data access for law enforcement, and what tensions does it create with GDPR?

- A. The CLOUD Act has no impact on international data access
- B. The CLOUD Act allows US authorities to compel US-based technology companies to provide data regardless of where it is stored; this creates conflicts with GDPR's data transfer restrictions and other countries' sovereignty concerns over their citizens' data
- C. The CLOUD Act only applies to cloud storage companies
- D. GDPR and the CLOUD Act are fully compatible

Answer: B

Q1422. What are the legal and ethical considerations of 'hack back' or active cyber defense measures by private organizations?

- A. Hack back is legal and encouraged for all organizations
- B. Active cyber defense by private entities raises complex legal issues: most jurisdictions prohibit unauthorized access regardless of motive, there are attribution challenges, risks of escalation, potential collateral damage, and limited legal frameworks distinguishing defensive from offensive actions
- C. Hack back is universally illegal with no exceptions
- D. Private organizations have unlimited rights to defend their networks through offensive measures

Answer: B

Q1423. How does the Budapest Convention on Cybercrime facilitate international cooperation, and what are its limitations?

- A. The Budapest Convention is no longer relevant
- B. The Convention provides a framework for mutual legal assistance, harmonizing cybercrime definitions and procedures; limitations include non-participation by major cyber powers (Russia, China), slow MLAT processes, and challenges adapting to emerging technologies
- C. The Convention automatically resolves all jurisdictional disputes
- D. All countries are members of the Budapest Convention

Answer: B

Q1424. What legal framework governs the use of artificial intelligence in cybersecurity decision-making, and what liability issues arise?

- A. AI in cybersecurity has no legal implications
- B. The legal framework is evolving: AI-powered security decisions raise questions of liability when automated systems cause false positives blocking legitimate users, algorithmic bias in threat detection, accountability for AI errors, and regulatory requirements for explainable AI
- C. AI decisions are always legally defensible
- D. Only developers of AI are liable, not the organizations using it

Answer: B

Q1425. What are the legal challenges of prosecuting state-sponsored cyber attacks under international law?

- A. International law clearly addresses all state-sponsored cyber attacks
- B. State-sponsored attacks face challenges including attribution difficulty, sovereign immunity, lack of binding international cyber norms, the Tallinn Manual's non-binding status, and debates about whether cyber attacks constitute 'armed attacks' triggering self-defense rights
- C. All state-sponsored attacks are considered acts of war
- D. International courts have jurisdiction over all cyber attacks

Answer: B

Q1426. How does the right to encryption conflict with government demands for lawful access, and what are the implications of encryption backdoors?

- A. There is no conflict between encryption and lawful access
- B. The tension lies between individuals' right to privacy through encryption and governments' need for lawful access to investigate crimes; mandated backdoors weaken security for all users and cannot be limited to authorized access only
- C. Encryption backdoors are perfectly secure
- D. Governments do not seek access to encrypted communications

Answer: B

Q1427. What is the legal significance of Pakistan's Prevention of Electronic Crimes Act (PECA) 2016 Section 10 regarding cyber terrorism?

- A. Section 10 only addresses spam emails
- B. Section 10 criminalizes cyber terrorism including using information systems to create fear, panic, or insecurity, interfering with critical infrastructure, or threatening national security, with severe penalties including imprisonment up to fourteen years
- C. PECA 2016 does not address terrorism
- D. Section 10 is about copyright protection

Answer: B

Q1428. What are the legal implications of cross-border data transfers following the Schrems II ruling, and how do organizations ensure compliance?

- A. Schrems II has no practical impact on data transfers
- B. Schrems II invalidated the EU-US Privacy Shield and requires organizations to conduct transfer impact assessments when using Standard Contractual Clauses, implement supplementary measures, and ensure the destination country's surveillance laws do not undermine EU-level data protection
- C. Organizations can freely transfer data between any countries
- D. Only EU organizations are affected by Schrems II

Answer: B

Q1429. How do cybersecurity regulations for critical infrastructure differ from general data protection laws, and what additional obligations do they impose?

- A. Critical infrastructure has fewer regulations than regular businesses
- B. Critical infrastructure regulations impose additional obligations including mandatory incident reporting with shorter deadlines, minimum security standards, regular security assessments, supply chain security requirements, and potential government oversight or certification
- C. Critical infrastructure only needs to follow GDPR
- D. There are no specific regulations for critical infrastructure

Answer: B

Q1430. What is the legal and ethical framework around vulnerability disclosure, and how do coordinated vulnerability disclosure policies reduce legal risk for researchers?

- A. Vulnerability disclosure has no legal framework
- B. Coordinated vulnerability disclosure (CVD) policies provide a legal safe harbor by defining how researchers should report vulnerabilities, establishing timelines for patches, and protecting researchers from prosecution under computer crime laws when they act in good faith within the policy's scope
- C. All vulnerability disclosure is illegal
- D. Only government researchers can disclose vulnerabilities

Answer: B

Q1431. How does the convergence of AI and cybersecurity create an arms race between attackers and defenders?

- A. AI only benefits defenders, not attackers
- B. Both attackers and defenders use AI: defenders for threat detection, anomaly analysis, and automated response; attackers for evading detection, generating polymorphic malware, automating social engineering, and discovering vulnerabilities, creating a continuous escalation cycle
- C. AI has no impact on the attacker-defender dynamic
- D. The arms race is only theoretical and has not materialized

Answer: B

Q1432. What is the concept of security chaos engineering and how does it proactively improve security posture?

- A. Chaos engineering makes security systems more chaotic
- B. Security chaos engineering deliberately introduces controlled security failures (expired certificates, disabled controls, simulated breaches) to verify that detection, alerting, and response mechanisms function correctly, identifying weaknesses before attackers exploit them
- C. It is identical to penetration testing
- D. Chaos engineering only tests system availability

Answer: B

Q1433. What are the unique security challenges of Web3 and decentralized applications (dApps)?

- A. Web3 applications are inherently more secure than Web2
- B. Web3 introduces unique challenges including smart contract vulnerabilities, irreversible transactions on blockchain, private key management, bridge protocol exploits, flash loan attacks, governance token manipulation, and the tension between immutability and the need to patch vulnerabilities
- C. Web3 has no security challenges
- D. Web3 uses the same security model as traditional web applications

Answer: B

Q1434. How does quantum key distribution (QKD) fundamentally differ from post-quantum cryptography (PQC), and what are the practical deployment challenges of each?

- A. QKD and PQC are the same technology
- B. QKD uses quantum physics to distribute keys with information-theoretic security but requires dedicated fiber/satellite links; PQC uses classical algorithms resistant to quantum attacks that work on existing infrastructure; QKD faces distance and cost limitations while PQC faces algorithm maturity and standardization challenges
- C. QKD is software-based while PQC requires hardware
- D. Neither technology addresses quantum threats

Answer: B

Q1435. What is the security significance of software-defined perimeter (SDP) and how does it implement zero trust principles?

- A. SDP is identical to a traditional VPN
- B. SDP creates dynamic, identity-based perimeters by making application infrastructure invisible to unauthorized users through single-packet authorization, mutual TLS, and need-to-know access, implementing zero trust by eliminating the concept of a trusted network
- C. SDP only works for physical security
- D. SDP replaces firewalls entirely

Answer: B

Q1436. What is the emerging threat of adversarial machine learning and how can it compromise AI-based security systems?

- A. Adversarial ML only affects image recognition
- B. Adversarial ML attacks manipulate AI models through techniques like evasion attacks (crafting inputs to bypass detection), poisoning attacks (corrupting training data), model stealing (extracting model parameters), and inversion attacks (extracting training data from model outputs)
- C. AI security systems cannot be manipulated
- D. Adversarial attacks require access to the model's source code

Answer: B

Q1437. How does the concept of cyber-physical systems security extend traditional cybersecurity to address safety-critical environments?

- A. Cyber-physical security is identical to IT security
- B. Cyber-physical security must address the intersection of digital and physical systems where cyber attacks can cause physical harm (e.g., industrial control systems, medical devices, autonomous vehicles), requiring safety engineering integration, real-time constraints, and consequence-based risk assessment
- C. Cyber-physical systems cannot be hacked
- D. Only military systems are cyber-physical

Answer: B

Q1438. What is the concept of digital sovereignty in cybersecurity and how does it influence technology and policy decisions?

- A. Digital sovereignty only means hosting websites domestically
- B. Digital sovereignty encompasses a nation's ability to control its digital infrastructure, data, and technology supply chains, influencing decisions on data localization, domestic technology development, trusted vendor requirements, and resistance to foreign surveillance
- C. Digital sovereignty has no impact on cybersecurity
- D. Only large countries can achieve digital sovereignty

Answer: B

Q1439. What is the security implication of neuromorphic computing and how might it change the cybersecurity landscape?

- A. Neuromorphic computing only affects battery life
- B. Neuromorphic computing mimics brain neural networks in hardware, potentially enabling real-time AI-powered threat detection with extremely low power consumption, but also poses risks of creating more sophisticated AI-driven attacks and may require new security paradigms for protecting brain-inspired architectures
- C. Neuromorphic computing is identical to quantum computing
- D. Neuromorphic computing has no security implications

Answer: B

Q1440. How does the concept of cyber resilience fundamentally differ from traditional cybersecurity, and why is it becoming more important?

- A. Cyber resilience and cybersecurity are identical concepts
- B. Cyber resilience accepts that breaches are inevitable and focuses on an organization's ability to anticipate, withstand, recover from, and adapt to adverse cyber events while maintaining critical operations, whereas traditional cybersecurity primarily focuses on prevention; resilience is critical as attack sophistication outpaces pure prevention
- C. Cyber resilience means giving up on security
- D. Traditional cybersecurity is always sufficient

Answer: B

SUBJECT

Databases

1,530 Multiple Choice Questions

Databases — Easy

510 questions

Q1. What is a database?

- A. A specialized type of programming language
- B. An organized collection of structured data
- C. A specific type of network protocol used
- D. A physical hardware component in a system

Answer: B

Q2. Which of the following is a popular Database Management System (DBMS)?

- A. CSS
- B. HTML
- C. MySQL
- D. FTP

Answer: C

Q3. What does DBMS stand for?

- A. Data Backup Management System
- B. Database Management System
- C. Database Modification Software
- D. Digital Base Management System

Answer: B

Q4. Which of these is an advantage of using a database system?

- A. Data cannot be shared between users
- B. Data redundancy is significantly reduced
- C. Overall storage cost increases greatly
- D. Programs become much harder to write

Answer: B

Q5. What is data redundancy?

- A. Algorithmic compression of data
- B. Unnecessary duplication of data
- C. Permanent loss of stored data
- D. Cryptographic encryption of data

Answer: B

Q6. Who manages and controls access to a database?

- A. Business Systems Analyst
- B. Database Administrator (DBA)
- C. Network Systems Engineer
- D. Web Interface Designer

Answer: B

Q7. Which of the following is NOT a type of database user?

- A. DBA administrator
- B. Database end user
- C. SQL programmer
- D. Compiler program

Answer: D

Q8. What is the primary goal of a DBMS?

- A. To manage and configure physical hardware components in a server rack
- B. To provide a convenient and efficient way to store and retrieve data
- C. To design and develop interactive websites and web applications online
- D. To replace the functionality of the underlying operating system entirely

Answer: B

Q9. In a file-based system, data is stored in:

- A. Random access memory storage only
- B. Separate files for each application
- C. Cloud-based storage systems only
- D. A single centralized data repository

Answer: B

Q10. Data independence means:

- A. Changes to data storage do not affect application programs
- B. Applications and data definitions are treated as the same
- C. Stored data values cannot be modified by any user at all
- D. Data items are not stored permanently on the storage disk

Answer: A

Q11. The two-tier architecture in databases consists of:

- A. Three layer tiers
- B. Client and Server
- C. Only Client tier
- D. Only Server tier

Answer: B

Q12. In a client-server architecture, the database resides on the:

- A. Client side
- B. Neither side
- C. Both equally
- D. Server side

Answer: D

Q13. What is a centralized database system?

- A. A system designed to operate without any DBMS software layer
- B. A system that runs exclusively on cloud computing platforms
- C. A system with data distributed across multiple remote locations
- D. A system where data is stored and managed at a single location

Answer: D

Q14. Which tier in three-tier architecture handles the user interface?

- A. Presentation tier
- B. Database data tier
- C. None of these tiers
- D. Business logic tier

Answer: A

Q15. What does a query processor do?

- A. Backs up data to external storage
- B. Interprets and executes database queries
- C. Stores data on physical disk drives
- D. Manages user accounts and profiles

Answer: B

Q16. The storage manager in a DBMS is responsible for:

- A. Creating graphical user interfaces
- B. Managing how data is stored on disk
- C. Writing and executing SQL queries
- D. Configuring network communication

Answer: B

Q17. Which component of DBMS translates high-level queries into low-level operations?

- A. Transaction manager
- B. Query compiler module
- C. Buffer manager module
- D. File manager module

Answer: B

Q18. In a three-tier architecture, the middle tier is called:

- A. Data warehouse store
- B. Application server
- C. Database server
- D. Web browser client

Answer: B

Q19. What is the role of the buffer manager?

- A. Creates and maintains database schemas and table structures
- B. Manages the transfer of data between disk and main memory
- C. Manages user authentication and role-based access control
- D. Processes and optimizes SQL queries for efficient execution

Answer: B

Q20. A database engine is also known as:

- A. Query language
- B. Storage engine
- C. Front-end application
- D. Report generator

Answer: B

Q21. Which data model organizes data in tables of rows and columns?

- A. Relational model
- B. Object model
- C. Hierarchical model
- D. Network model

Answer: A

Q22. In the hierarchical data model, data is organized as a:

- A. Graph structure
- B. Star structure
- C. Tree structure
- D. Table structure

Answer: C

Q23. An Entity-Relationship (ER) diagram is used to:

- A. Design network topologies for servers
- B. Create application interfaces for users
- C. Model the conceptual design of a database
- D. Write SQL queries against stored tables

Answer: C

Q24. In an ER diagram, a rectangle represents:

- A. An entity type
- B. A relationship
- C. An attribute
- D. A primary key

Answer: A

Q25. An attribute in a data model represents:

- A. A constraint on records
- B. A property of an entity
- C. A table in the database
- D. A relationship between entities

Answer: B

Q26. Which symbol represents a relationship in an ER diagram?

- A. Ellipse
- B. Rectangle
- C. Triangle
- D. Diamond

Answer: D

Q27. A primary key uniquely identifies:

- A. A relationship instance
- B. A single column value
- C. Each record in a table
- D. An entire database schema

Answer: C

Q28. In the network data model, the relationship between records is:

- A. Many-to-many using sets
- B. No relationships allowed
- C. Hierarchical only structure
- D. One-to-one relationships

Answer: A

Q29. The relational data model was proposed by:

- A. Edgar Dijkstra
- B. Charles Bachman
- C. E.F. Codd
- D. Peter Chen

Answer: C

Q30. Which of these is NOT a data model?

- A. Relational model
- B. Hierarchical model
- C. Compiler model
- D. Network model

Answer: C

Q31. In a relational database, a table is also called a:

- A. Tree
- B. Graph
- C. Document
- D. Relation

Answer: D

Q32. A row in a relational table is called a:

- A. Schema
- B. Attribute
- C. Domain
- D. Tuple

Answer: D

Q33. A column in a relational table is called:

- A. An attribute
- B. A relation
- C. A key value
- D. A tuple

Answer: A

Q34. A primary key must be:

- A. Unique and not null
- B. Duplicate in values
- C. Foreign key based
- D. Nullable and optional

Answer: A

Q35. A foreign key is:

- A. A key used only for sorting records in a particular order
- B. The first column of any table regardless of its data content
- C. An attribute that is always unique across all related tables
- D. An attribute that references the primary key of another table

Answer: D

Q36. The domain of an attribute specifies:

- A. The number of rows containing that attribute
- B. The name assigned to identify that attribute
- C. The set of allowed values for that attribute
- D. The table name where the attribute exists

Answer: C

Q37. The degree of a relation refers to:

- A. The number of tables in database
- B. The number of tuples (rows) stored
- C. The number of keys in the schema
- D. The number of attributes (columns)

Answer: D

Q38. The cardinality of a relation refers to:

- A. The number of keys in the table
- B. The number of foreign key links
- C. The number of tuples (rows)
- D. The number of attributes (columns)

Answer: C

Q39. Which of the following is a property of a relational table?

- A. Rows must be ordered
- B. Each row is unique
- C. Columns can be duplicated
- D. Attributes are ordered

Answer: B

Q40. NULL in a database represents:

- A. An empty string character
- B. A zero numeric value only
- C. A space character in data
- D. An unknown or missing value

Answer: D

Q41. The selection operation in relational algebra:

- A. Selects tuples (rows) that satisfy a condition
- B. Selects specific columns from a given relation
- C. Renames a relation or its attribute identifiers
- D. Combines two relations into a single relation

Answer: A

Q42. The projection operation in relational algebra:

- A. Deletes tuples from a stored relation
- B. Selects specific columns from a relation
- C. Joins two relations on common attributes
- D. Selects specific rows from a given relation

Answer: B

Q43. The union of two relations R and S contains:

- A. Only tuples that are common to both R and S
- B. All tuples from both R and S without duplicates
- C. The Cartesian product of relations R and S
- D. Tuples in R but not present in relation S

Answer: B

Q44. Which relational algebra operation renames a relation or its attributes?

- A. Rename (ρ)
- B. Projection (π)
- C. Division ($\div$)
- D. Selection (σ)

Answer: A

Q45. The Cartesian product of relations R and S is denoted as:

- A. $R \bowtie S$
- B. $R \times S$
- C. $R \cup S$
- D. $R \cap S$

Answer: B

Q46. Set difference ($R - S$) returns:

- A. Common tuples found in R and S
- B. Tuples in S that are not in R
- C. Tuples in R that are not in S
- D. All tuples from both R and S

Answer: C

Q47. For union compatibility, two relations must have:

- A. The same number of attributes with compatible domains
- B. The same name for both relations being combined
- C. The same number of tuples stored in both relations
- D. Identical primary keys defined on both data relations

Answer: A

Q48. The intersection of R and S ($R \cap S$) returns:

- A. Tuples that are only in relation R
- B. Tuples that are only in relation S
- C. All tuples from relations R and S
- D. Tuples that are in both R and S

Answer: D

Q49. Which operation is fundamental in relational algebra (cannot be derived from others)?

- A. Intersection
- B. Division
- C. Selection
- D. Natural join

Answer: C

Q50. The natural join operation:

- A. Renames attributes within a relation for query convenience
- B. Returns all tuples from both relations regardless of matching
- C. Combines tuples from two relations based on common attributes
- D. Removes duplicates from a single relation without conditions

Answer: C

Q51. SQL stands for:

- A. Standard Question Language
- B. Structured Query Language
- C. System Query Logic
- D. Simple Query Language

Answer: B

Q52. Which SQL statement is used to retrieve data from a database?

- A. DELETE
- B. INSERT
- C. UPDATE
- D. SELECT

Answer: D

Q53. Which SQL clause is used to filter rows?

- A. HAVING
- B. ORDER BY
- C. WHERE
- D. GROUP BY

Answer: C

Q54. Which SQL statement adds new rows to a table?

- A. UPDATE SET
- B. INSERT INTO
- C. ALTER TABLE
- D. CREATE INDEX

Answer: B

Q55. Which SQL statement modifies existing data?

- A. DROP
- B. INSERT
- C. UPDATE
- D. SELECT

Answer: C

Q56. Which SQL statement removes rows from a table?

- A. ALTER
- B. DELETE
- C. TRUNCATE
- D. DROP

Answer: B

Q57. Which SQL command creates a new table?

- A. CREATE TABLE
- B. NEW TABLE
- C. INSERT TABLE
- D. ADD TABLE

Answer: A

Q58. Which SQL command permanently removes a table?

- A. DESTROY TABLE
- B. REMOVE TABLE
- C. DELETE TABLE
- D. DROP TABLE

Answer: D

Q59. The ORDER BY clause is used to:

- A. Filter rows by value
- B. Join tables on keys
- C. Sort the result set
- D. Group rows together

Answer: C

Q60. Which keyword is used to remove duplicate rows from a result?

- A. FILTER
- B. UNIQUE
- C. REMOVE
- D. DISTINCT

Answer: D

Q61. A view in SQL is:

- A. A stored procedure in the database server
- B. A virtual table based on a SELECT query
- C. A physical copy of a table stored on disk
- D. A type of index on one or more columns

Answer: B

Q62. Which SQL command creates a view?

- A. INSERT VIEW
- B. MAKE VIEW
- C. BUILD VIEW
- D. CREATE VIEW

Answer: D

Q63. A stored procedure is:

- A. A precompiled set of SQL statements stored in the database
- B. An index structure built on one or more columns of a table
- C. A type of view that shows data from multiple source tables
- D. A constraint that enforces data integrity rules on a table

Answer: A

Q64. A trigger in SQL is:

- A. A constraint that restricts the values allowed in a specific column
- B. A type of view that displays filtered data from the underlying table
- C. A manual query written by the user and executed on demand each time
- D. A procedure that automatically executes when a specific event occurs

Answer: D

Q65. What is a cursor in SQL?

- A. A pointer to a table that references its location in data storage
- B. A database object used to retrieve and process rows one at a time
- C. A type of index structure created on one or more table columns
- D. A constraint on a column that restricts the allowed data values

Answer: B

Q66. Which SQL statement is used to grant privileges?

- A. PERMIT
- B. ALLOW
- C. ENABLE
- D. GRANT

Answer: D

Q67. Which SQL statement removes privileges?

- A. DENY
- B. DISABLE
- C. REVOKE
- D. REMOVE

Answer: C

Q68. An assertion in SQL is:

- A. A view that provides a virtual representation of the underlying data
- B. A stored procedure that runs on a scheduled basis within the server
- C. A constraint that specifies a condition the database must always satisfy
- D. A type of trigger that fires automatically when data changes occur

Answer: C

Q69. What is a transaction in SQL?

- A. A type of view that provides a virtual representation of table data
- B. A single SELECT query that retrieves data from the database tables
- C. An index structure that speeds up data retrieval from large tables
- D. A sequence of operations treated as a single logical unit of work

Answer: D

Q70. The COMMIT statement in SQL:

- A. Creates a savepoint within the current active transaction
- B. Starts a new transaction and initializes the session state
- C. Undoes all changes made in the current transaction entirely
- D. Saves all changes made in the current transaction permanently

Answer: D

Q71. Normalization in databases is the process of:

- A. Adding redundant data to speed up query processing
- B. Deleting all data from every table in the database
- C. Creating as many tables as possible in the schema
- D. Organizing data to reduce redundancy and dependency

Answer: D

Q72. First Normal Form (1NF) requires:

- A. All attributes are restricted to numeric
- B. No foreign keys exist within the tables
- C. All attributes contain atomic values only
- D. Tables have only one column of any data

Answer: C

Q73. A functional dependency $X \rightarrow Y$ means:

- A. The value of Y uniquely determines the value of X
- B. X and Y always contain the same identical values
- C. X and Y are completely unrelated to each other
- D. The value of X uniquely determines the value of Y

Answer: D

Q74. Redundancy in a database leads to:

- A. Insertion, deletion, and update anomalies
- B. Better data integrity and consistency
- C. Faster queries and improved performance
- D. Simplified database design structure

Answer: A

Q75. An update anomaly occurs when:

- A. A table is created for the first time in the schema
- B. A new row is inserted into an existing data table
- C. The same data must be changed in multiple places
- D. An index is dropped from a column in the database

Answer: C

Q76. Second Normal Form (2NF) requires:

- A. Only atomic values without any other constraints
- B. 1NF and no partial dependency on the primary key
- C. No functional dependencies of any kind at all
- D. All attributes must serve as keys in the table

Answer: B

Q77. Third Normal Form (3NF) requires:

- A. No attributes defined at all in it
- B. Only 1NF without other constraints
- C. 2NF and no transitive dependencies
- D. All attributes are primary key parts

Answer: C

Q78. A transitive dependency exists when:

- A. Two keys depend on each other in a circular dependency
- B. A key attribute depends on a non-key attribute directly
- C. A non-key attribute depends on another non-key attribute
- D. No dependencies exist between any attributes at all

Answer: C

Q79. A partial dependency exists when:

- A. A non-key attribute depends on the entire composite primary key
- B. A non-key attribute depends on part of a composite primary key
- C. No dependencies exist between any attributes in the data table
- D. All attributes depend on all keys defined in the relation schema

Answer: B

Q80. Denormalization is:

- A. Removing all tables and their data from the entire database
- B. Intentionally introducing redundancy to improve read performance
- C. Deleting all indexes to reduce overall storage requirements
- D. Normalizing the schema to a higher normal form than before

Answer: B

Q81. ACID properties of a transaction stand for:

- A. Atomicity, Concurrency, Isolation, Durability
- B. Addition, Consistency, Integrity, Durability
- C. Atomicity, Consistency, Isolation, Durability
- D. Atomicity, Consistency, Isolation, Dependency

Answer: C

Q82. Atomicity means that a transaction:

- A. Runs in isolation from all other transactions
- B. Is either fully completed or fully rolled back
- C. Is stored permanently on the disk after start
- D. Can be partially completed and then paused

Answer: B

Q83. Consistency ensures that a transaction:

- A. Never fails regardless of the operations within
- B. Is stored permanently on disk after the commit
- C. Runs concurrently with other active transactions
- D. Brings the database from one valid state to another

Answer: D

Q84. Isolation ensures that:

- A. Transactions are never allowed to execute concurrently
- B. All transactions see each other's changes immediately
- C. Only one transaction can exist in the system at a time
- D. Concurrent transactions do not interfere with each other

Answer: D

Q85. Durability ensures that:

- A. Data is stored only in volatile memory and not on stable disk
- B. Changes can be undone at any time even after transaction commit
- C. Committed transaction changes persist even after system failure
- D. Transactions are temporary and their changes are never saved

Answer: C

Q86. The COMMIT operation in a transaction:

- A. Pauses the active transaction
- B. Makes all changes permanent
- C. Undoes all current changes
- D. Starts a new transaction

Answer: B

Q87. The ROLLBACK operation in a transaction:

- A. Undoes all changes made by the transaction
- B. Starts a new transaction in the session
- C. Deletes the database and all stored data
- D. Saves all changes and makes them permanent

Answer: A

Q88. A transaction that has not yet been committed or rolled back is called:

- A. Committed
- B. Active
- C. Aborted
- D. Terminated

Answer: B

Q89. Which of the following is a transaction state?

- A. All of these: Active, Partially committed, Committed, Failed, Aborted
- B. Only Aborted is considered a valid state for a transaction process
- C. Only Active is considered a valid state for a transaction in a DBMS
- D. Only Committed is considered a valid state for a transaction to be

Answer: A

Q90. A schedule in transaction management is:

- A. A single SQL query executed against the database tables
- B. A backup plan for recovering from a system level failure
- C. A type of index structure for speeding up data lookups
- D. A sequence of operations from one or more transactions

Answer: D

Q91. Concurrency control ensures:

- A. Larger storage capacity on the disk drives
- B. Better network speed between client nodes
- C. Faster hardware performance for the system
- D. Correct execution of concurrent transactions

Answer: D

Q92. A lock in database concurrency control is:

- A. A mechanism to control access to a data item
- B. A backup method for protecting against loss
- C. A type of query for retrieving stored data
- D. A type of index for speeding up data search

Answer: A

Q93. A shared lock allows:

- A. Writing by multiple transactions to the same data item
- B. No access at all to the locked data item by any party
- C. Only one transaction to read the data item at any time
- D. Multiple transactions to read the data item simultaneously

Answer: D

Q94. An exclusive lock allows:

- A. Multiple transactions to write to the same data item
- B. Multiple transactions to read from the same data item
- C. Only one transaction to read and write the data item
- D. No operations at all on the locked data item by anyone

Answer: C

Q95. A deadlock in database systems occurs when:

- A. All transactions commit at the same time without any conflicts found
- B. A single transaction is very slow in completing its data operations
- C. Two or more transactions are waiting for each other to release locks
- D. No transactions are running in the database system at the moment

Answer: C

Q96. The Two-Phase Locking (2PL) protocol has two phases:

- A. Lock phase and Unlock phase only
- B. Growing phase and Shrinking phase
- C. Read phase and Write phase only
- D. Start phase and End phase only

Answer: B

Q97. Which of the following can break a deadlock?

- A. Adding more locks to the waiting transactions
- B. Creating more indexes on the locked data items
- C. Aborting one of the deadlocked transactions
- D. Increasing the buffer size for the data manager

Answer: C

Q98. Starvation in concurrency control occurs when:

- A. All transactions complete quickly without any delays or blocking
- B. A transaction waits indefinitely because others keep getting priority
- C. The database runs out of storage space on the physical disk
- D. No locks are used by any of the transactions in the database

Answer: B

Q99. A lock manager is responsible for:

- A. Designing schemas for new data models
- B. Creating tables in the database schema
- C. Writing SQL queries for data retrieval
- D. Granting and releasing locks on data items

Answer: D

Q100. The lock compatibility matrix shows:

- A. How to create new tables in the database schema now
- B. Which lock types can coexist on the same data item
- C. The execution plans for all queries in the database
- D. Index structures for optimizing data retrieval speed

Answer: B

Q101. Database recovery is the process of:

- A. Deleting all data from the tables in the entire database
- B. Restoring the database to a correct state after a failure
- C. Creating new tables and defining their column structures
- D. Optimizing queries by choosing better execution plans

Answer: B

Q102. A transaction log (write-ahead log) records:

- A. Table structures and their column definitions in the database schema
- B. Only SELECT queries that have been executed against the database tables
- C. User login information and authentication credentials stored securely
- D. All changes made by transactions before they are applied to the database

Answer: D

Q103. The redo operation:

- A. Creates a new database from scratch without using any data from the previous state
- B. Re-applies changes of committed transactions that may not have been written to disk
- C. Undoes all changes of committed transactions and reverts to the original data state
- D. Deletes the entire log file and removes all records of previous transaction actions

Answer: B

Q104. The undo operation:

- A. Reverses changes made by uncommitted transactions
- B. Creates a backup of the current database on disk
- C. Optimizes queries for faster execution performance
- D. Applies committed changes to the database on disk

Answer: A

Q105. A checkpoint in database recovery is:

- A. A type of lock used to prevent concurrent access to shared data items
- B. An SQL statement used to query data from one or more database tables
- C. A point at which all modified data is written to disk and recorded in the log
- D. A type of index structure built on columns for faster data retrieval

Answer: C

Q106. The Write-Ahead Log (WAL) protocol requires:

- A. Logs are only written to disk after the transaction has committed fully
- B. No log is needed for any database operations under this recovery protocol
- C. Data must be written to disk before the corresponding log records are made
- D. Log records must be written to disk before the corresponding data changes

Answer: D

Q107. A system crash typically causes:

- A. Loss of data in main memory but not on disk
- B. No data loss of any kind after the crash
- C. Loss of all data everywhere in the system
- D. Loss of data on disk but not in main memory

Answer: A

Q108. A media failure refers to:

- A. Application crash caused by a software defect
- B. Network slowness affecting query response times
- C. Disk or storage device failure causing data loss
- D. Memory overflow due to excessive data loading

Answer: C

Q109. Database backups are important for:

- A. Reducing table size by removing unnecessary data rows
- B. Recovering from media failures and catastrophic events
- C. Creating new databases from the backed up data copies
- D. Improving query speed and reducing execution latency

Answer: B

Q110. An incremental backup stores:

- A. The entire database every single time
- B. Only the database schema definitions
- C. Only the changes since the last backup
- D. Only the transaction log file entries

Answer: C

Q111. An index in a database is used to:

- A. Replace primary key usage
- B. Slow down query results
- C. Speed up data retrieval
- D. Increase storage redundancy

Answer: C

Q112. A primary index is built on:

- A. An unordered file without sorting
- B. A foreign key only in the relation
- C. Any non-key attribute in the table
- D. The primary key of an ordered file

Answer: D

Q113. A dense index has:

- A. An entry for every block in the data file
- B. Only one entry for the entire data file
- C. An entry for every record in the data file
- D. No entries at all in the index structure

Answer: C

Q114. A sparse index has:

- A. An entry for every individual record stored
- B. An entry for each block or group of records
- C. Entries only for deleted records in the file
- D. No entries at all in the entire index file

Answer: B

Q115. A B-tree is a:

- A. Linear list with sequential access
- B. Balanced search tree used for indexing
- C. Binary search tree with two children
- D. Unbalanced tree without constraints

Answer: B

Q116. In a B+ tree, all actual data pointers are stored in:

- A. All nodes equally
- B. Leaf nodes only
- C. Root node only
- D. Internal nodes only

Answer: B

Q117. Hashing is used for:

- A. Direct access to records using a hash function
- B. Sorting data into a specific order on columns
- C. Sequential access to all records in the file
- D. Creating views on top of the base table data

Answer: A

Q118. A secondary index is built on:

- A. The ordering field of the file
- B. The primary key field only
- C. A non-ordering field of a file
- D. A foreign key field only

Answer: C

Q119. A clustered index means:

- A. Data rows are stored randomly without any ordering
- B. Data rows are physically stored in the index order
- C. No physical ordering is implied by the index at all
- D. Multiple indexes exist on the same column of data

Answer: B

Q120. Which data structure is commonly used for implementing indexes?

- A. Linked list nodes
- B. Stack structure
- C. Queue structure
- D. B+ tree structure

Answer: D

Q121. Query processing involves:

- A. Only creating tables without any query functionality
- B. Only writing queries without any execution or output
- C. Only storing data without any retrieval functionality
- D. Parsing, optimizing, and executing database queries

Answer: D

Q122. A query execution plan describes:

- A. The user who wrote the query and their access privilege details
- B. The table creation date and the schema modification history
- C. The sequence of operations the DBMS will perform to execute a query
- D. The query syntax and grammatical structure in the SQL language

Answer: C

Q123. Query optimization aims to:

- A. Find the most efficient way to execute a query
- B. Make queries longer and more complex to write
- C. Increase the number of disk accesses required
- D. Slow down query processing for all operations

Answer: A

Q124. The EXPLAIN command in SQL shows:

- A. The database schema layout
- B. The query execution plan
- C. User access permissions
- D. The table data content

Answer: B

Q125. A full table scan reads:

- A. Only indexed rows stored
- B. Only the last row stored
- C. Every row in the table
- D. Only the first row found

Answer: C

Q126. Using an index for a query is generally:

- A. Always slower than performing a full table scan
- B. Faster than a full table scan for selective queries
- C. The same speed as a full table scan on all data
- D. Not recommended for any database query at all

Answer: B

Q127. Which join algorithm compares every tuple in one relation with every tuple in another?

- A. Nested Loop Join
- B. Sort-Merge Join
- C. Index Nested Join
- D. Hash Join method

Answer: A

Q128. The cost of a query is typically measured in:

- A. Number of SQL statements used
- B. Lines of code in the procedure
- C. Number of disk I/O operations
- D. Number of tables in the schema

Answer: C

Q129. Selectivity of a condition refers to:

- A. The size of the database measured in total bytes
- B. The number of columns selected in the projection
- C. The fraction of tuples that satisfy the condition
- D. The number of tables referenced in the SQL query

Answer: C

Q130. A query optimizer can be:

- A. Only cost-based approach
- B. Rule-based or cost-based
- C. Neither rule nor cost based
- D. Only rule-based approach

Answer: B

Q131. Database security aims to protect the database from:

- A. Faster queries and improved execution performance
- B. Unauthorized access, modification, and destruction
- C. Larger storage capacity on the available disk space
- D. Better indexing and optimized data retrieval speed

Answer: B

Q132. Authentication in database security verifies:

- A. The identity of a user
- B. The number of indexes
- C. The speed of queries
- D. The size of tables

Answer: A

Q133. Authorization in database security determines:

- A. The operating system and hardware details
- B. What actions a user is permitted to perform
- C. The database version and release number
- D. The user password and login credentials

Answer: B

Q134. The SQL GRANT statement is used to:

- A. Create tables in database
- B. Remove privileges from users
- C. Delete data from tables
- D. Give privileges to users

Answer: D

Q135. The SQL REVOKE statement is used to:

- A. Remove previously granted privileges
- B. Insert data into existing table records
- C. Grant new privileges to database users
- D. Create views on top of the base tables

Answer: A

Q136. A database role is:

- A. An index structure built on columns for faster data retrieval
- B. A type of table used to store data records within the database
- C. A named collection of privileges that can be assigned to users
- D. A query optimization technique to improve execution planning

Answer: C

Q137. Encryption in database security is used to:

- A. Compress data to reduce the amount of storage space it requires
- B. Create indexes on columns for faster retrieval of requested records
- C. Convert data into an unreadable format to prevent unauthorized access
- D. Speed up queries by creating cached copies of frequently used results

Answer: C

Q138. An audit trail in a database records:

- A. All operations performed on the database for security monitoring
- B. Only DDL operations like CREATE, ALTER, and DROP table commands
- C. Only user logins and their session authentication event details
- D. Only SELECT queries that have been executed against stored tables

Answer: A

Q139. SQL injection is a type of:

- A. Indexing strategy used for speeding up data retrieval from the tables
- B. Security attack that exploits vulnerabilities in SQL query construction
- C. Database optimization technique that improves overall query performance
- D. Backup method used for creating copies of the database for recovery

Answer: B

Q140. The principle of least privilege means:

- A. All users should have full access to every table in the database
- B. Users should have only the minimum privileges needed for their tasks
- C. No privileges should be granted to any user in the database
- D. Only administrators should be allowed to use the database system

Answer: B

Q141. A distributed database is:

- A. A database stored only in the cloud without any local presence
- B. A database stored on a single machine in one physical location
- C. A database spread across multiple sites connected by a network
- D. A backup copy of a database kept for disaster recovery purposes

Answer: C

Q142. Data fragmentation in distributed databases means:

- A. Duplicating the entire database at every site simultaneously
- B. Compressing data to reduce storage requirements at each site
- C. Losing data during transfer between sites on the network link
- D. Dividing a relation into smaller parts stored at different sites

Answer: D

Q143. Horizontal fragmentation divides a relation by:

- A. Rows and columns both
- B. Rows (tuples) only
- C. Neither rows nor cols
- D. Columns (attributes)

Answer: B

Q144. Vertical fragmentation divides a relation by:

- A. Rows (tuples) only
- B. Columns (attributes)
- C. Neither of the above
- D. Columns and rows both

Answer: B

Q145. Data replication in distributed databases means:

- A. Compressing data to reduce transfer size
- B. Deleting data from all distributed sites
- C. Encrypting data before sending it across
- D. Storing copies of data at multiple sites

Answer: D

Q146. The transparency goal in distributed databases means:

- A. Data cannot be distributed across multiple sites
- B. Users must know all data locations before querying
- C. Only administrators can query distributed data
- D. Users should not be aware that data is distributed

Answer: D

Q147. Location transparency means:

- A. Data is always stored locally on the user's own machine
- B. Users must specify the exact data location in every query
- C. Users do not need to know where data is physically stored
- D. No remote access is possible from any site in the system

Answer: C

Q148. A distributed DBMS (DDBMS) manages:

- A. Only local databases without any distribution
- B. A database distributed across multiple sites
- C. Only a centralized database on one machine
- D. Only cloud databases hosted by third parties

Answer: B

Q149. Which of the following is an advantage of distributed databases?

- A. Slower queries being preferred by users
- B. Higher costs being beneficial overall
- C. Improved reliability and availability
- D. Increased complexity as an advantage

Answer: C

Q150. Fragmentation transparency means:

- A. Users are unaware that data is fragmented
- B. Data cannot be fragmented at any site
- C. Only fragments can be queried directly
- D. Users must reassemble fragments manually

Answer: A

Q151. NoSQL stands for:

- A. Not Only SQL format
- B. Newer SQL approach
- C. Non-Standard Queries
- D. No Structured Queries

Answer: A

Q152. Which of the following is a NoSQL database?

- A. MySQL
- B. PostgreSQL
- C. Oracle DB
- D. MongoDB

Answer: D

Q153. NoSQL databases are designed primarily for:

- A. Replacing all SQL databases entirely in every use case now
- B. Only relational data organized in normalized table schemas
- C. Small structured datasets only stored in relational tables
- D. Scalability and handling large volumes of unstructured data

Answer: D

Q154. A document-oriented database stores data as:

- A. Graphs with nodes and edge links
- B. Fixed-size blocks on the disk drive
- C. Tables with rows and columns only
- D. Documents (typically JSON or BSON)

Answer: D

Q155. A key-value store maps:

- A. Tables to their schemas
- B. Rows to table columns
- C. Queries to result sets
- D. A unique key to a value

Answer: D

Q156. Redis is an example of a:

- A. Key-value store
- B. Column-family store
- C. Graph database
- D. Relational database

Answer: A

Q157. A column-family store organizes data into:

- A. Documents in JSON format
- B. Traditional rows and tables
- C. Columns and column families
- D. Graphs with node and edges

Answer: C

Q158. Apache Cassandra is an example of a:

- A. Relational database
- B. Document database
- C. Graph database
- D. Column-family store

Answer: D

Q159. A graph database stores data as:

- A. Documents in JSON formats
- B. Key-value pairs only stored
- C. Nodes, edges, and properties
- D. Tables and rows of records

Answer: C

Q160. Neo4j is an example of a:

- A. Document database
- B. Key-value store
- C. Relational database
- D. Graph database

Answer: D

Q161. A data warehouse is:

- A. A transactional database designed for processing day-to-day operations in real-time mode
- B. A subject-oriented, integrated, time-variant, non-volatile collection of data for decisions
- C. A backup storage mechanism that periodically copies data for disaster recovery plans
- D. A type of NoSQL database that supports flexible schemas and horizontal scaling systems

Answer: B

Q162. OLAP stands for:

- A. Offline Analytical Processing
- B. Online Application Programming
- C. Online Automated Processing
- D. Online Analytical Processing

Answer: D

Q163. OLTP stands for:

- A. Online Text Processing
- B. Online Transaction Processing
- C. Online Table Processing
- D. Offline Transaction Processing

Answer: B

Q164. The ETL process stands for:

- A. Export, Translate, Link
- B. Enter, Transfer, List
- C. Extract, Transfer, Log
- D. Extract, Transform, Load

Answer: D

Q165. A fact table in a data warehouse contains:

- A. Quantitative measures and foreign keys to dimension tables
- B. Only dimension attributes used for filtering and grouping
- C. Only text descriptions without any numeric measure values
- D. Only primary keys without any quantitative measure data

Answer: A

Q166. A dimension table in a data warehouse contains:

- A. Descriptive attributes used for filtering and grouping
- B. Log entries recording all data warehouse load operations
- C. Transaction records from the operational source systems
- D. Quantitative measures like sales amounts and quantities

Answer: A

Q167. The star schema consists of:

- A. Multiple fact tables connected to each other in it
- B. A central fact table connected to dimension tables
- C. Only dimension tables without any fact table data
- D. Only fact tables without any dimension table data

Answer: B

Q168. Data mining is:

- A. Writing SQL queries for retrieving records from relational tables
- B. Creating database backups for disaster recovery and data safety
- C. The process of discovering patterns and insights from large datasets
- D. Storing data in a warehouse for long-term historical analysis use

Answer: C

Q169. A data mart is:

- A. A type of OLTP system designed for transactional data processing
- B. A subset of a data warehouse focused on a specific business area
- C. A complete data warehouse containing all enterprise data sources
- D. A NoSQL database designed for handling unstructured data at scale

Answer: B

Q170. Which of the following is an OLAP operation?

- A. INSERT
- B. Drill-down
- C. UPDATE
- D. DELETE

Answer: B

Q171. What is a schema in a database context?

- A. A query used to retrieve specific data
- B. A backup copy of the entire database
- C. The overall design or structure of data
- D. A security policy for user permissions

Answer: C

Q172. Which component stores metadata about a database?

- A. Query processor
- B. Storage engine
- C. Data dictionary
- D. Buffer manager

Answer: C

Q173. What does data integrity refer to in databases?

- A. Amount of storage space data occupies
- B. Number of users accessing data at once
- C. Accuracy and consistency of stored data
- D. Speed of data retrieval operations

Answer: C

Q174. Which language is used to define database structure?

- A. DCL
- B. DML
- C. TCL
- D. DDL

Answer: D

Q175. What is a database instance?

- A. The design blueprint of the database schema
- B. The software used to manage the database
- C. The snapshot of data at a given moment
- D. The physical disk where data is stored

Answer: C

Q176. Which of these is a function of a DBMS?

- A. Rendering graphical user interface views
- B. Compiling high-level source code files
- C. Routing packets across network devices
- D. Managing concurrent data access safely

Answer: D

Q177. What is a query in database terminology?

- A. A method to compress stored records
- B. A protocol for network communication
- C. A type of database backup procedure
- D. A request to retrieve or modify data

Answer: D

Q178. Which of these describes a flat-file database?

- A. Data distributed across multiple server nodes
- B. Data stored in a single two-dimensional table
- C. Data organized in a hierarchical tree format
- D. Data linked through complex network pointers

Answer: B

Q179. What does ACID stand for in databases?

- A. Allocation, Computation, Indexing, Deletion
- B. Access, Control, Integrity, Design
- C. Atomicity, Consistency, Isolation, Durability
- D. Authentication, Concurrency, Integration, Data

Answer: C

Q180. What is the role of a query processor?

- A. It encrypts data before storage use
- B. It parses and executes user queries
- C. It backs up data to external drives
- D. It manages physical disk allocation

Answer: B

Q181. What is a client-server database architecture?

- A. A system where all processing happens on one machine
- B. A system that does not use any network connections
- C. A system where data is stored only on client machines
- D. A system where clients request services from a server

Answer: D

Q182. What is the function of the storage manager in a DBMS?

- A. It handles network routing between nodes
- B. It compiles application source code files
- C. It creates user interface forms and reports
- D. It manages data storage on physical media

Answer: D

Q183. Which tier handles the user interface in a three-tier architecture?

- A. Data tier
- B. Logic tier
- C. Storage tier
- D. Presentation tier

Answer: D

Q184. What is a database buffer in the system architecture?

- A. A backup copy of the entire database files
- B. A network connection between two servers
- C. A temporary memory area for data transfer
- D. A permanent storage location on the disk

Answer: C

Q185. What does the query compiler do in a DBMS?

- A. It handles backup and recovery processes
- B. It stores data permanently on disk drives
- C. It translates queries into executable plans
- D. It manages user authentication credentials

Answer: C

Q186. In a two-tier architecture, which components communicate directly?

- A. Two separate database servers communicate
- B. The middleware and storage tier communicate
- C. The client and database server communicate
- D. Two client applications communicate together

Answer: C

Q187. What is the role of the transaction manager?

- A. Ensuring transactions execute correctly and safely
- B. Designing the user interface for applications
- C. Compressing data files for efficient disk storage
- D. Routing network requests between server clusters

Answer: A

Q188. Which memory component holds recently accessed disk pages?

- A. Permanent storage
- B. Cache buffer pool
- C. Network registers
- D. Virtual memory

Answer: B

Q189. What is a centralized database system?

- A. A system spread across many geographic locations
- B. A system with no persistent data storage at all
- C. A system using only cloud-based storage services
- D. A system running on a single computer at one site

Answer: D

Q190. What is the purpose of the data communication component?

- A. To enable communication between DBMS and users
- B. To perform mathematical calculations on data
- C. To define the structure of database tables
- D. To compress data before writing to disk drives

Answer: A

Q191. What is a data model in database design?

- A. A physical storage device for database files
- B. A conceptual framework for organizing data
- C. A network protocol for data transmission
- D. A programming language for writing queries

Answer: B

Q192. Which data model organizes data into tables of rows and columns?

- A. Network model
- B. Hierarchical model
- C. Relational model
- D. Object model

Answer: C

Q193. In the hierarchical data model, data is organized as:

- A. A graph with many-to-many connections
- B. An unstructured collection of documents
- C. A tree with parent-child relationships
- D. A flat two-dimensional table structure

Answer: C

Q194. What is an entity in the ER model?

- A. A constraint applied to a column in a table
- B. A real-world object represented in a database
- C. A relationship between two database tables
- D. A query used to retrieve records from data

Answer: B

Q195. What does an attribute represent in a data model?

- A. A connection between two separate entities
- B. A property or characteristic of an entity
- C. A method for sorting records in a table
- D. A rule that constrains data entry values

Answer: B

Q196. What is a relationship in the ER model?

- A. A method for indexing table records fast
- B. An association between two or more entities
- C. A column in a database table structure
- D. A constraint on a single attribute value

Answer: B

Q197. Which symbol represents an entity in an ER diagram?

- A. Triangle shape
- B. Rectangle shape
- C. Oval shape
- D. Diamond shape

Answer: B

Q198. What does cardinality specify in a relationship?

- A. The number of entity instances in a link
- B. The data type of attributes in an entity
- C. The access permission level for database users
- D. The storage format for records on a disk

Answer: A

Q199. What is a composite attribute in the ER model?

- A. An attribute that can be divided into subparts
- B. An attribute that has multiple values at once
- C. An attribute derived from other stored values
- D. An attribute that cannot be subdivided further

Answer: A

Q200. Which data model uses nodes and edges to represent data?

- A. Graph data model
- B. Columnar model
- C. Relational model
- D. Flat file model

Answer: A

Q201. What is a tuple in a relational database?

- A. A column in a relational database table
- B. A key in a relational database design
- C. A table in a relational database schema
- D. A row in a relational database table

Answer: D

Q202. What is a primary key in a relational table?

- A. A column that always contains null values
- B. A column that uniquely identifies each row
- C. A column that links to an external file
- D. Any column that contains text data values

Answer: B

Q203. What is a foreign key used for?

- A. Calculating aggregate functions on data sets
- B. Encrypting data in a database column field
- C. Linking one table to another related table
- D. Sorting records in ascending order by value

Answer: C

Q204. What does the term relation refer to in relational databases?

- A. A connection between two different databases
- B. A table with rows and columns of data
- C. A constraint applied to a single column
- D. A query that joins multiple data sources

Answer: B

Q205. What is a domain in a relational database?

- A. A type of relationship between two tables
- B. A network address for the database server
- C. A set of allowable values for an attribute
- D. A backup strategy for database recovery use

Answer: C

Q206. What does a candidate key represent?

- A. A key that is always the foreign key value
- B. A key used only for sorting table records
- C. A key that allows duplicate values in rows
- D. A key that might be selected as primary key

Answer: D

Q207. What is the degree of a relation?

- A. The number of constraints on it
- B. The number of rows in the table
- C. The number of keys in the table
- D. The number of columns in the table

Answer: D

Q208. What are null values used to represent?

- A. The empty string in a text data column
- B. Missing or unknown information in a field
- C. The number zero in a numeric column field
- D. A default value assigned to every record

Answer: B

Q209. What is a superkey in relational databases?

- A. A set of attributes that uniquely identifies rows
- B. A key used only in object-oriented databases
- C. A key with exactly one attribute column only
- D. A key that cannot contain any null value at all

Answer: A

Q210. What is an attribute in a relational table?

- A. A column within the relational table structure
- B. A row within the relational table structure
- C. A constraint applied to the entire table
- D. A relationship between two different tables

Answer: A

Q211. What is relational algebra?

- A. A data storage format for flat files
- B. A programming language for web development
- C. A procedural query language for relations
- D. A network protocol for data transfer

Answer: C

Q212. What does the selection operation do?

- A. It combines two relations into one table
- B. It picks specific rows matching a condition
- C. It picks specific columns from a relation
- D. It removes duplicate rows from a relation

Answer: B

Q213. What does the projection operation do?

- A. It selects specific rows based on conditions
- B. It calculates aggregate values for a group
- C. It joins two relations on a common column
- D. It selects specific columns from a relation

Answer: D

Q214. What is the Cartesian product of two relations?

- A. The difference of rows between both tables
- B. All combinations of rows from both tables
- C. Only matching rows from both of the tables
- D. The intersection of rows from both tables

Answer: B

Q215. What does the union operation require?

- A. Relations must be union-compatible to combine
- B. Relations must share a common primary key
- C. Relations must be stored in the same file
- D. Relations must have different attribute counts

Answer: A

Q216. What does the set difference operation produce?

- A. All rows appearing in either of two relations
- B. Only rows common to both of the two relations
- C. A sorted version of the first input relation
- D. Rows in the first relation but not the second

Answer: D

Q217. What is a natural join operation?

- A. A join that only works on single-column tables
- B. A join that matches rows on common attributes
- C. A join that removes all duplicate column values
- D. A join that produces the Cartesian product only

Answer: B

Q218. What does the rename operation do in relational algebra?

- A. It sorts the tuples in a relation by key order
- B. It deletes a relation from the database schema
- C. It creates a backup copy of the existing table
- D. It changes the name of a relation or attribute

Answer: D

Q219. What is the intersection of two relations?

- A. Tuples present in both relations simultaneously
- B. Tuples in the first but not in the second one
- C. All tuples from both relations combined together
- D. Tuples in the second but not in the first one

Answer: A

Q220. Which operation combines selection and join in one step?

- A. Rename operation
- B. Projection operation
- C. Division operation
- D. Theta join operation

Answer: D

Q221. What does SQL stand for?

- A. Structured Query Language
- B. Sequential Query Logic
- C. Standard Query Library
- D. System Query Lookup

Answer: A

Q222. Which SQL command retrieves data from a table?

- A. UPDATE
- B. DELETE
- C. INSERT
- D. SELECT

Answer: D

Q223. Which clause filters rows in a SELECT statement?

- A. ORDER BY
- B. WHERE
- C. GROUP BY
- D. HAVING

Answer: B

Q224. What does the INSERT command do in SQL?

- A. Modifies existing rows in a table
- B. Adds new rows into a table
- C. Deletes rows from a table
- D. Creates a new database table

Answer: B

Q225. Which SQL keyword removes duplicate rows from results?

- A. FILTER
- B. PRIMARY
- C. DISTINCT
- D. UNIQUE

Answer: C

Q226. What does the UPDATE command do in SQL?

- A. Removes rows from the table
- B. Retrieves rows from a table
- C. Modifies existing rows in table
- D. Creates a new table in database

Answer: C

Q227. Which SQL command deletes rows from a table?

- A. CLEAR
- B. DELETE
- C. DROP
- D. REMOVE

Answer: B

Q228. What does the ORDER BY clause do?

- A. Sorts the result set by columns
- B. Groups rows by a specified column
- C. Limits the number of output rows
- D. Filters rows based on a condition

Answer: A

Q229. Which SQL command creates a new table?

- A. CREATE TABLE
- B. BUILD TABLE
- C. MAKE TABLE
- D. NEW TABLE

Answer: A

Q230. What is the wildcard character in SQL SELECT?

- A. Hash symbol
- B. Asterisk symbol
- C. Ampersand symbol
- D. Dollar symbol

Answer: B

Q231. What is a stored procedure in SQL?

- A. A single SQL query saved as a bookmark
- B. A temporary table created during a query
- C. A precompiled set of SQL statements stored
- D. A backup copy of the database schema file

Answer: C

Q232. What is a view in SQL?

- A. A physical table stored on the disk drive
- B. A log file recording all query activity
- C. A virtual table based on a query result
- D. A backup of the original database table

Answer: C

Q233. What is a trigger in SQL databases?

- A. A type of index for faster data retrieval
- B. Code that executes automatically on events
- C. A manual command executed by the user only
- D. A constraint that limits column data types

Answer: B

Q234. What is an index used for in SQL?

- A. Creating backup copies of table records
- B. Defining relationships between two tables
- C. Encrypting data stored in database tables
- D. Speeding up data retrieval from tables

Answer: D

Q235. What does the GRANT command do in SQL?

- A. Gives users specific access permissions
- B. Creates new columns in existing tables
- C. Backs up data to an external location
- D. Removes existing tables from a database

Answer: A

Q236. What is a cursor in SQL?

- A. A method for creating temporary table data
- B. A type of database constraint on columns
- C. A graphical pointer on the user interface
- D. A mechanism to traverse query result rows

Answer: D

Q237. What does the REVOKE command do?

- A. It deletes data from a specified table
- B. It creates new user accounts in database
- C. It removes previously granted permissions
- D. It modifies the structure of a table

Answer: C

Q238. What is a transaction in SQL?

- A. A unit of work with multiple operations
- B. A single SELECT query against one table
- C. A method for defining table constraints
- D. A type of database index on a column

Answer: A

Q239. What does COMMIT do in SQL?

- A. Deletes all data from the current table
- B. Creates a new savepoint within transaction
- C. Saves all changes made in a transaction
- D. Undoes all changes made in a transaction

Answer: C

Q240. What does ROLLBACK do in SQL?

- A. Saves all changes made in the transaction
- B. Grants permissions to a specified database user
- C. Undoes all changes made in the transaction
- D. Creates a new table in the active database

Answer: C

Q241. What is database normalization?

- A. Organizing data to reduce redundancy and dependency
- B. Adding redundant data to improve query speed
- C. Encrypting data for security and privacy protection
- D. Compressing data to save physical disk storage space

Answer: A

Q242. What is the first normal form (1NF)?

- A. All candidate keys are properly identified and used
- B. Each column contains only atomic indivisible values
- C. All non-key attributes depend on the primary key
- D. No transitive dependencies exist in the relation

Answer: B

Q243. What is a functional dependency?

- A. A relationship between two separate database tables
- B. A method for sorting records in ascending order
- C. A constraint where one attribute determines another
- D. A technique for encrypting stored data values

Answer: C

Q244. What is data redundancy in databases?

- A. Data that is encrypted for security purposes
- B. Data that is properly normalized in tables
- C. Unnecessary repetition of data across tables
- D. Data stored in compressed file formats only

Answer: C

Q245. What problem does normalization primarily solve?

- A. Data anomalies from redundancy
- B. Slow network connection speeds
- C. Lack of user interface design
- D. Insufficient hardware capacity

Answer: A

Q246. What is an insertion anomaly?

- A. Inability to add data without other unrelated data
- B. Difficulty deleting unwanted records from tables
- C. Errors when updating records across multiple tables
- D. Problems with reading data from large database files

Answer: A

Q247. What is a deletion anomaly?

- A. Creating circular references between related tables now
- B. Failing to update all copies of redundant data values
- C. Losing unrelated data when deleting a specific record
- D. Adding duplicate records accidentally to a table field

Answer: C

Q248. What is an update anomaly?

- A. Inconsistency from failing to update all data copies
- B. Inability to insert new records into a database table
- C. Failure to create proper indexes on table columns here
- D. Accidentally deleting records from a database table now

Answer: A

Q249. What is denormalization?

- A. Intentionally adding redundancy for performance gain
- B. Converting a relational database to a flat file type
- C. The process of achieving first normal form status
- D. Removing all indexes from database table columns

Answer: A

Q250. What does 2NF require beyond 1NF?

- A. No transitive dependencies on the primary key
- B. Every column must have a unique non-null value
- C. No partial dependencies on the primary key set
- D. All attributes must be of the same data type

Answer: C

Q251. What is a database transaction?

- A. A report generated from database queries
- B. A backup copy of the database schema
- C. A connection between two database servers
- D. A logical unit of work on a database

Answer: D

Q252. What does the Atomicity property ensure?

- A. All or none of the operations are completed
- B. Changes persist even after system failures
- C. Database remains consistent after transaction
- D. Transactions execute in complete isolation always

Answer: A

Q253. What does the Consistency property ensure?

- A. All operations in a transaction execute together
- B. Database moves from one valid state to another
- C. Transactions run without any interference at all
- D. Transaction results are permanently stored safely

Answer: B

Q254. What does the Isolation property ensure?

- A. All operations complete or none complete at all
- B. Changes are saved permanently after a commit
- C. Concurrent transactions do not interfere with others
- D. The database schema remains unchanged after commits

Answer: C

Q255. What does the Durability property ensure?

- A. Transactions can be rolled back at any later time
- B. Multiple transactions can run at the same time now
- C. Data types remain consistent across all table rows
- D. Committed changes persist even after system failure

Answer: D

Q256. What does COMMIT do in transaction processing?

- A. It creates a new transaction from the current state
- B. It saves all changes made during the transaction
- C. It undoes all changes made during the transaction
- D. It temporarily pauses the current active transaction

Answer: B

Q257. What does ROLLBACK do in transaction processing?

- A. Pauses the transaction until the user resumes it later
- B. Saves all pending changes to the database permanently
- C. Undoes all changes since the last commit or savepoint
- D. Creates a backup of the database before rolling back

Answer: C

Q258. What is a schedule in transaction processing?

- A. An ordered sequence of operations from transactions
- B. A configuration file for database server parameters
- C. A backup plan for the database recovery process
- D. A list of all users currently connected to database

Answer: A

Q259. What is a serial schedule?

- A. Transactions execute one after another without overlap
- B. Transactions execute simultaneously on different processors
- C. Transactions execute their operations in interleaved order
- D. Transactions execute in reverse order of their submission

Answer: A

Q260. What is a transaction log used for?

- A. Storing user passwords and authentication credentials
- B. Managing network connections between client and server
- C. Recording all changes for recovery and audit purposes
- D. Displaying query results to the end user interface

Answer: C

Q261. What is concurrency control in databases?

- A. Backing up data to prevent data loss scenarios
- B. Managing simultaneous access by multiple users
- C. Encrypting data for secure transmission purposes
- D. Compressing data to save disk storage space used

Answer: B

Q262. What is a lock in database concurrency control?

- A. A method for compressing database table records
- B. A technique for encrypting sensitive data fields
- C. A mechanism to restrict access to data items
- D. A backup mechanism for database recovery tasks

Answer: C

Q263. What is a shared lock used for?

- A. Preventing all access to the data item entirely
- B. Allowing only one transaction to access the data
- C. Allowing multiple transactions to read the data
- D. Allowing only write access to the data item

Answer: C

Q264. What is an exclusive lock used for?

- A. Giving one transaction sole access for writing
- B. Sharing write access among multiple transactions
- C. Allowing multiple users to read data together
- D. Preventing all transactions from accessing data

Answer: A

Q265. What is a deadlock in database systems?

- A. A backup process that runs during peak usage hours
- B. A transaction that completes faster than expected
- C. Two transactions waiting for each other indefinitely
- D. A query that returns more results than anticipated

Answer: C

Q266. What is starvation in concurrency control?

- A. A transaction completes too quickly for proper logging
- B. A transaction produces incorrect results due to bad data
- C. A transaction waits indefinitely for a resource access
- D. A transaction uses too much memory during its execution

Answer: C

Q267. What does the lock manager do in a DBMS?

- A. It compiles SQL queries into executable code plans
- B. It creates user accounts and manages permissions
- C. It manages physical storage allocation on disk drives
- D. It grants and releases locks on data items safely

Answer: D

Q268. What is a lock table used for?

- A. Maintaining backup copies of all locked data items
- B. Storing the actual data values of locked records
- C. Tracking which locks are held by which transactions
- D. Recording all SQL queries executed by each user

Answer: C

Q269. What is lock compatibility?

- A. Whether two lock types can coexist on same item
- B. Whether a lock persists after system restart event
- C. Whether a lock can be transferred between users
- D. Whether a lock applies to one or multiple tables

Answer: A

Q270. What is the purpose of a timeout in deadlock handling?

- A. To extend the duration of held locks indefinitely
- B. To increase the priority of waiting transactions
- C. To abort a transaction after waiting too long
- D. To speed up slow-running queries automatically

Answer: C

Q271. What is database recovery?

- A. Adding new data records to database tables
- B. Optimizing queries for faster execution speed
- C. Restoring a database after a failure event
- D. Creating user accounts for database access

Answer: C

Q272. What is a system crash in database context?

- A. A routine update of the database software used
- B. A planned maintenance shutdown of the server
- C. An unexpected failure that disrupts operations
- D. A scheduled backup of the database to disk

Answer: C

Q273. What is a transaction log?

- A. A record of all changes made by transactions
- B. A backup copy of the database schema only
- C. A list of all database users and permissions
- D. A configuration file for server parameters

Answer: A

Q274. What is the purpose of a database backup?

- A. To create a copy of data for recovery purposes
- B. To speed up query processing performance times
- C. To encrypt sensitive data for security compliance
- D. To normalize database tables for better design

Answer: A

Q275. What does the REDO operation do in recovery?

- A. Deletes all log records from the transaction log file
- B. Creates a new backup of the database on the disk
- C. Reapplies changes from committed transactions in log
- D. Undoes changes from committed transactions in log

Answer: C

Q276. What does the UNDO operation do in recovery?

- A. Applies changes from committed transactions to data
- B. Compresses the transaction log for storage savings
- C. Reverses changes from uncommitted transactions safely
- D. Creates new checkpoints in the transaction log file

Answer: C

Q277. What is a checkpoint in database recovery?

- A. A lock that prevents concurrent access to data
- B. A point where the database state is synchronized
- C. A user account with limited database permissions
- D. A query that checks for data integrity violations

Answer: B

Q278. What is the difference between volatile and non-volatile storage?

- A. Non-volatile storage is faster but loses data on power loss
- B. Neither type of storage can lose data under any conditions
- C. Volatile storage is faster but loses data on power loss
- D. Both types of storage retain data after power is lost

Answer: C

Q279. What is a disk failure in database context?

- A. A routine defragmentation of the disk drive unit
- B. A malfunction that causes loss of stored disk data
- C. A scheduled cleaning of the disk read-write heads
- D. A planned upgrade of the disk storage hardware

Answer: B

Q280. What is the purpose of a log buffer?

- A. Temporarily holding log records before writing to disk
- B. Encrypting data before it is stored on the disk drive
- C. Backing up the entire database to an external device
- D. Storing user queries before they are executed later

Answer: A

Q281. What is an index in a database?

- A. A backup method for database recovery tasks
- B. A security mechanism for user authentication
- C. A data structure for fast data retrieval access
- D. A complete copy of the database table data

Answer: C

Q282. What is a primary index in a database?

- A. An index that only supports range-based queries
- B. An index built on the primary key of the table
- C. An index built on a non-key column of a table
- D. An index that stores all table data within itself

Answer: B

Q283. What is a secondary index?

- A. An index that is only used for backup purposes
- B. An index that replaces the primary key entirely
- C. An index on the ordering key of the data file
- D. An index on a non-ordering field of a data file

Answer: D

Q284. What is a dense index?

- A. An index that compresses data for storage efficiency
- B. An index that stores data without any key structure
- C. An index with an entry for every data record stored
- D. An index with an entry for some data records only

Answer: C

Q285. What is a sparse index?

- A. An index with entries for only some data records
- B. An index with an entry for every single data record
- C. An index stored entirely in main memory at all times
- D. An index that does not use any key values at all

Answer: A

Q286. What is a B-tree in database indexing?

- A. A linear list of pointers to database records
- B. A balanced tree structure for disk-based indexing
- C. A binary search tree for in-memory data only
- D. A hash-based structure for exact match queries

Answer: B

Q287. What is a hash index used for?

- A. Efficient pattern matching on text data values
- B. Efficient exact match queries on key values
- C. Efficient range queries on sorted data values
- D. Efficient sequential scanning of all records

Answer: B

Q288. What is a clustered index?

- A. An index where data order differs from index order
- B. An index where data order matches the index order
- C. An index created automatically on every new column
- D. An index that works only on temporary table data

Answer: B

Q289. What is sequential file organization?

- A. Records stored in a tree structure on disk blocks
- B. Records stored in a hash table on the disk blocks
- C. Records stored in random order on the disk blocks
- D. Records stored in sorted order on the disk blocks

Answer: D

Q290. What is a heap file organization?

- A. Records are stored in the order they are inserted
- B. Records are stored in sorted order by primary key
- C. Records are stored using a B-tree index structure
- D. Records are stored in a hash table data structure

Answer: A

Q291. What is query processing in a database?

- A. Managing user accounts and access permissions sets
- B. Designing the physical layout of storage on disk
- C. Translating and executing queries against the data
- D. Creating backup copies of database tables regularly

Answer: C

Q292. What is a query execution plan?

- A. A list of all users who can run the query
- B. A log of all previously executed database queries
- C. A backup schedule for the database operations
- D. The strategy used to execute a database query

Answer: D

Q293. What does a query optimizer do?

- A. It creates indexes on all columns automatically now
- B. It translates SQL into a different programming language
- C. It selects the most efficient query execution plan
- D. It encrypts queries before sending to the server

Answer: C

Q294. What is a full table scan?

- A. Reading only indexed rows from the data table
- B. Reading only the first row of the database table
- C. Reading rows from multiple tables at the same time
- D. Reading every row in the table sequentially stored

Answer: D

Q295. What is the purpose of query parsing?

- A. Sorting the query results in alphabetical order
- B. Checking the syntax and validity of SQL queries
- C. Executing the query against the stored database
- D. Backing up the query text for future reference

Answer: B

Q296. What is a cost-based query optimizer?

- A. An optimizer that selects the plan with lowest cost
- B. An optimizer that ignores indexes during planning
- C. An optimizer that always chooses the first valid plan
- D. An optimizer that only works with simple queries

Answer: A

Q297. What are statistics used for in query optimization?

- A. Estimating the cost of different execution plans
- B. Encrypting data before storage on disk drives
- C. Displaying query results to end users on screen
- D. Creating user accounts and password management

Answer: A

Q298. What is an index scan in query processing?

- A. Scanning all rows in a table without using any index
- B. Using an index to locate specific rows efficiently
- C. Scanning multiple tables simultaneously for matches
- D. Scanning only the first page of the data table

Answer: B

Q299. What is query decomposition?

- A. Encrypting query text to prevent unauthorized access
- B. Combining multiple queries into a single large query
- C. Breaking a complex query into simpler sub-operations
- D. Removing queries from the execution queue permanently

Answer: C

Q300. What is the role of the query evaluation engine?

- A. It manages physical storage allocation on the disk
- B. It executes the plan generated by the optimizer
- C. It creates user interfaces for database applications
- D. It designs the database schema and table layouts

Answer: B

Q301. What is database security?

- A. Protecting data from unauthorized access or damage
- B. Compressing data to save storage space on disk
- C. Optimizing queries for faster execution speeds
- D. Sorting data in alphabetical order in all tables

Answer: A

Q302. What is authentication in database security?

- A. Compressing database files for efficient storage
- B. Verifying the identity of a user accessing data
- C. Encrypting data stored in the database tables
- D. Backing up data to prevent loss from failures

Answer: B

Q303. What is authorization in database security?

- A. Determining what actions a user is allowed to do
- B. Verifying user identity before granting access
- C. Encrypting data before storing it in the database
- D. Creating backup copies of the database regularly

Answer: A

Q304. What does the GRANT command do?

- A. Creates a new database on the server system
- B. Removes user access to database objects entirely
- C. Gives specific permissions to users or roles
- D. Deletes data from a specified database table

Answer: C

Q305. What does the REVOKE command do?

- A. Adds new permissions to a user or role account
- B. Encrypts all data in the specified database table
- C. Creates new user accounts in the database system
- D. Removes previously granted permissions from users

Answer: D

Q306. What is data encryption in database security?

- A. Duplicating data across multiple storage devices
- B. Converting data into an unreadable coded format
- C. Sorting data in a specific predetermined order
- D. Compressing data to reduce its storage size need

Answer: B

Q307. What is a database role?

- A. A named collection of related database privileges
- B. A physical server that hosts the database system
- C. A backup strategy for disaster recovery planning
- D. A type of query used for data retrieval purposes

Answer: A

Q308. What is an audit trail in database security?

- A. A method for optimizing database query performance
- B. A record of all access and changes to the database
- C. A technique for compressing data stored on disk drives
- D. A strategy for distributing data across multiple servers

Answer: B

Q309. What is SQL injection?

- A. A method for optimizing SQL query performance time
- B. An attack that inserts malicious SQL code in input
- C. A process for encrypting SQL queries before running
- D. A technique for backing up SQL database regularly

Answer: B

Q310. What is access control in database security?

- A. Compressing database tables to save storage space
- B. Optimizing the speed of data retrieval operations
- C. Regulating who can access what data and operations
- D. Creating indexes for faster query execution speeds

Answer: C

Q311. What is a distributed database?

- A. A database accessible only through a local area network
- B. A database spread across multiple connected computers
- C. A database that only stores metadata and no actual data
- D. A database stored entirely on a single server machine

Answer: B

Q312. What is data fragmentation in distributed databases?

- A. Compressing data to reduce network transfer overhead
- B. Encrypting data before distributing to remote servers
- C. Breaking a relation into smaller parts across sites
- D. Combining data from multiple sources into one table

Answer: C

Q313. What is horizontal fragmentation?

- A. Replicating the entire table across all network sites
- B. Combining multiple tables into a single unified table
- C. Dividing a table by selecting specific rows from it
- D. Dividing a table by selecting specific columns from it

Answer: C

Q314. What is vertical fragmentation?

- A. Dividing a table by selecting specific rows based on value
- B. Dividing a table by selecting specific columns from table
- C. Combining rows from different tables into a unified view
- D. Storing complete copies of a table at every network site

Answer: B

Q315. What is data replication in distributed databases?

- A. Encrypting data before transmission between all sites
- B. Dividing data into fragments across different sites
- C. Compressing data to reduce the overall storage needed
- D. Storing copies of data at multiple sites for access

Answer: D

Q316. What is a distributed DBMS?

- A. Software that manages data across multiple network sites
- B. Software that only handles backup and recovery operations
- C. Software that manages files on a single computer only
- D. Software that designs user interfaces for database access

Answer: A

Q317. What is location transparency?

- A. Users must establish connections to each site manually
- B. Users can only access data stored on their local site
- C. Users access data without knowing its physical location
- D. Users must specify which site stores the requested data

Answer: C

Q318. What is a site in a distributed database?

- A. A computer or node that hosts part of the database
- B. A column in a distributed database table structure
- C. A type of query used in distributed data processing
- D. A backup copy of the distributed database for safety

Answer: A

Q319. What is fragmentation transparency?

- A. Users access data without knowing how it is fragmented
- B. Users must specify fragment names in all their queries
- C. Users can only query one fragment at a time per query
- D. Users must manually join fragments when querying data

Answer: A

Q320. What is a distributed query?

- A. A query that is scheduled to run at a future time date
- B. A query that runs on a single database server only
- C. A query that only retrieves metadata from the catalog
- D. A query that accesses data from multiple network sites

Answer: D

Q321. What does NoSQL stand for?

- A. Not Only SQL databases
- B. Networked Open SQL system
- C. No Structured Query Language
- D. Non-Sequential Query Logic

Answer: A

Q322. Which is an example of a document-based NoSQL database?

- A. MongoDB
- B. MySQL
- C. Redis
- D. Neo4j

Answer: A

Q323. What is a key-value store?

- A. A database that maps unique keys to their values
- B. A database that uses SQL for all data operations
- C. A database that stores data in relational tables
- D. A database that only stores numeric data values

Answer: A

Q324. Which is an example of a graph database?

- A. MongoDB
- B. Neo4j
- C. Cassandra
- D. Redis

Answer: B

Q325. What type of NoSQL database is Redis?

- A. Document store
- B. Column-family store
- C. Graph database store
- D. Key-value store

Answer: D

Q326. What type of NoSQL database is Cassandra?

- A. Key-value store type
- B. Column-family store
- C. Graph database type
- D. Document store type

Answer: B

Q327. What is a document database?

- A. A database storing data in graph node structures
- B. A database storing data in fixed relational tables
- C. A database storing only text document file contents
- D. A database storing data in flexible document formats

Answer: D

Q328. What is horizontal scaling in NoSQL databases?

- A. Adding more resources to a single existing server
- B. Adding more servers to distribute the workload
- C. Upgrading the software on the existing servers
- D. Reducing the number of servers in the cluster

Answer: B

Q329. What is schema flexibility in NoSQL databases?

- A. Schema changes require rebuilding the entire database
- B. Schema must be defined before any data can be inserted
- C. Records in the same collection can have different fields
- D. All records must follow the exact same fixed schema

Answer: C

Q330. What is the main advantage of NoSQL over relational databases?

- A. More mature and established query language support
- B. Better support for complex multi-table join queries
- C. Better scalability for large unstructured data sets
- D. Stronger ACID transaction support across tables

Answer: C

Q331. What is a data warehouse?

- A. A central repository for integrated historical data
- B. A physical building where servers are maintained
- C. A system for real-time transaction processing only
- D. A temporary storage for data during transmission

Answer: A

Q332. What does ETL stand for in data warehousing?

- A. Evaluate, Test, Launch processing
- B. Encrypt, Transfer, Load processing
- C. Extract, Transform, Load processing
- D. Edit, Track, Log processing steps

Answer: C

Q333. What is a fact table in a data warehouse?

- A. A table containing descriptive text about entities
- B. A table containing measurable business event data
- C. A table containing database configuration settings
- D. A table storing user credentials and permissions

Answer: B

Q334. What is a dimension table?

- A. A table storing system configuration parameters
- B. A table providing descriptive context for facts
- C. A table storing numeric measurement values only
- D. A table containing temporary processing results

Answer: B

Q335. What is OLAP in data warehousing?

- A. Online Application Logging and Processing
- B. Online Analytical Processing of stored data
- C. Optimized Linear Analysis of Procedures
- D. Offline Logical Access and Programming

Answer: B

Q336. What is a star schema?

- A. A schema without any fact or dimension tables used
- B. A schema with many fact tables and one dimension
- C. A schema where all tables are fully normalized form
- D. A schema with one fact table and many dimensions

Answer: D

Q337. What is data mining?

- A. Deleting old data from the database storage files
- B. Backing up data to external storage media drives
- C. Manually entering data into database tables one row
- D. Discovering patterns and knowledge from large data

Answer: D

Q338. What is the purpose of data integration?

- A. Splitting data into smaller separate fragments only
- B. Encrypting data before storing in the data warehouse
- C. Combining data from multiple sources into unified view
- D. Deleting duplicate records from all source databases

Answer: C

Q339. What is a data mart?

- A. A complete enterprise-wide data warehouse system
- B. A tool for mining patterns from large data volumes
- C. A method for encrypting warehouse data at rest now
- D. A subset of a warehouse for a specific department

Answer: D

Q340. What is the difference between OLTP and OLAP?

- A. OLTP and OLAP serve the exact same purpose in databases
- B. OLTP is newer technology while OLAP is the older version
- C. OLTP is for analytics while OLAP is for transactions only
- D. OLTP handles transactions while OLAP handles analytics

Answer: D

Q341. Which type of database stores data in key-value pairs?

- A. Relational database
- B. Key-value store
- C. Hierarchical database
- D. Network database

Answer: B

Q342. What is the main purpose of a database management system?

- A. To compile programming code
- B. To manage and organize data efficiently
- C. To design user interfaces
- D. To configure network settings

Answer: B

Q343. Which of the following is an example of a relational DBMS?

- A. MongoDB
- B. Redis
- C. MySQL
- D. Cassandra

Answer: C

Q344. What does DDL stand for in database systems?

- A. Data Description Logic
- B. Data Definition Language
- C. Database Design Layer
- D. Dynamic Data Linking

Answer: B

Q345. Which language is most commonly used to interact with relational databases?

- A. Python
- B. Java
- C. SQL
- D. C++

Answer: C

Q346. What is a record in a database table?

- A. A column header
- B. A single row of data
- C. The table name
- D. A database file

Answer: B

Q347. Which of the following best describes data consistency?

- A. Data is stored in multiple formats
- B. Data follows defined rules and remains accurate
- C. Data is always encrypted
- D. Data is compressed for storage

Answer: B

Q348. What is a field in a database?

- A. An entire table
- B. A single piece of data within a record
- C. A database backup
- D. A type of query

Answer: B

Q349. Which of the following is NOT a benefit of using a DBMS?

- A. Reduced data redundancy
- B. Improved data security
- C. Increased hardware speed
- D. Data backup and recovery

Answer: C

Q350. What does DML stand for in database terminology?

- A. Data Manipulation Language
- B. Data Modeling Logic
- C. Database Management Layer
- D. Dynamic Memory Loading

Answer: A

Q351. What is the simplest form of database architecture?

- A. Three-tier architecture
- B. Distributed architecture
- C. Single-tier architecture
- D. Cloud architecture

Answer: C

Q352. In database architecture, what does the term 'tier' refer to?

- A. A type of database table
- B. A logical or physical layer in the system
- C. A security protocol
- D. A backup mechanism

Answer: B

Q353. Which component interprets and executes database queries?

- A. Storage engine
- B. Query processor
- C. Network adapter
- D. Display driver

Answer: B

Q354. What is the purpose of the disk manager in a DBMS?

- A. To display query results
- B. To manage the allocation of space on disk
- C. To handle user authentication
- D. To optimize network traffic

Answer: B

Q355. In a two-tier architecture, which component sends requests to the server?

- A. Database engine
- B. Client application
- C. Storage manager
- D. Buffer pool

Answer: B

Q356. What does the term 'back-end' refer to in database architecture?

- A. The user interface
- B. The server side that manages the database
- C. The network cable
- D. The display monitor

Answer: B

Q357. Which architecture separates presentation, logic, and data layers?

- A. Single-tier architecture
- B. Two-tier architecture
- C. Three-tier architecture
- D. Peer-to-peer architecture

Answer: C

Q358. What is the role of the authorization manager in a DBMS?

- A. To optimize queries
- B. To verify user permissions before allowing operations
- C. To manage disk space
- D. To handle data backup

Answer: B

Q359. What does a database server primarily do?

- A. Renders web pages
- B. Processes and responds to database requests
- C. Manages email services
- D. Compiles application code

Answer: B

Q360. Which DBMS component temporarily holds data pages in memory?

- A. Query parser
- B. Buffer pool
- C. Log manager
- D. Schema manager

Answer: B

Q361. What is a simple attribute in the ER model?

- A. An attribute that can be divided into sub-parts
- B. An attribute that cannot be divided further
- C. An attribute with multiple values
- D. An attribute derived from other attributes

Answer: B

Q362. What does a diamond shape represent in an ER diagram?

- A. An entity
- B. An attribute
- C. A relationship
- D. A primary key

Answer: C

Q363. What is a key attribute in an ER diagram?

- A. Any attribute of an entity
- B. An attribute that uniquely identifies each entity instance
- C. An attribute that stores encrypted data
- D. An attribute that links to another table

Answer: B

Q364. In the ER model, what is a one-to-many relationship?

- A. One entity instance relates to at most one instance of another entity
- B. One entity instance can relate to many instances of another entity
- C. Many entity instances relate to many instances of another entity
- D. No entity instances are related

Answer: B

Q365. What is the purpose of an ER diagram?

- A. To write SQL queries
- B. To visually represent the database structure and relationships
- C. To configure network settings
- D. To define programming functions

Answer: B

Q366. Which data model represents data as a collection of objects?

- A. Relational model
- B. Hierarchical model
- C. Object-oriented data model
- D. Flat file model

Answer: C

Q367. In the ER model, what is a many-to-many relationship?

- A. Each entity instance relates to exactly one other
- B. Multiple instances of both entities can relate to each other
- C. Only one entity participates in the relationship
- D. The relationship has no cardinality

Answer: B

Q368. What does an oval shape represent in an ER diagram?

- A. An entity
- B. A relationship
- C. An attribute
- D. A constraint

Answer: C

Q369. What is a strong entity in the ER model?

- A. An entity that has a primary key and can exist independently
- B. An entity that depends on another entity for identification
- C. An entity with the most attributes
- D. An entity connected to all other entities

Answer: A

Q370. In the network data model, what type of relationship structure is used?

- A. Tree structure
- B. Graph structure allowing many-to-many relationships
- C. Simple flat tables
- D. Star schema

Answer: B

Q371. What is an alternate key in a relational database?

- A. A key that is always numeric
- B. A candidate key that is not selected as the primary key
- C. A key used only for backup purposes
- D. A key that allows duplicate values

Answer: B

Q372. What is referential integrity in a relational database?

- A. All tables must have the same number of rows
- B. Foreign key values must match existing primary key values or be null
- C. All attributes must have default values
- D. Tables must be stored on the same disk

Answer: B

Q373. What does the term 'relation' mean in the relational model?

- A. A connection between two databases
- B. A two-dimensional table of data
- C. A physical storage device
- D. A type of programming function

Answer: B

Q374. What is entity integrity in a relational database?

- A. Every table must have at least one row
- B. The primary key of a relation cannot contain null values
- C. All attributes must be of the same data type
- D. Each table must have exactly one foreign key

Answer: B

Q375. How many candidate keys can a relation have?

- A. Exactly one
- B. At most two
- C. One or more
- D. None

Answer: C

Q376. What is the purpose of a composite key?

- A. To encrypt data in a column
- B. To use two or more attributes together to uniquely identify a tuple
- C. To link two databases
- D. To store multiple values in one field

Answer: B

Q377. What is a relation schema?

- A. The data stored in a table at a given time
- B. The definition of a relation including its name and attributes
- C. A backup copy of a relation
- D. A query that retrieves data from a relation

Answer: B

Q378. In a relational table, can two rows be identical?

- A. Yes, any number of duplicate rows are allowed
- B. No, each row must be unique
- C. Only if the table has no primary key defined
- D. Only in temporary tables

Answer: B

Q379. What is the role of a foreign key?

- A. To uniquely identify rows in its own table
- B. To establish a link between data in two tables
- C. To encrypt sensitive data
- D. To automatically generate unique IDs

Answer: B

Q380. What does the term 'arity' of a relation refer to?

- A. The number of rows in the relation
- B. The number of columns (attributes) in the relation
- C. The size of the relation in bytes
- D. The number of constraints on the relation

Answer: B

Q381. Which relational algebra operation removes duplicate rows from the result?

- A. Projection
- B. Selection
- C. Cartesian product
- D. Rename

Answer: A

Q382. What does the selection operation filter?

- A. Columns from a relation
- B. Rows from a relation based on a condition
- C. Duplicate tables
- D. Database users

Answer: B

Q383. What does the projection operation do in relational algebra?

- A. Selects specific rows from a table
- B. Selects specific columns from a relation
- C. Joins two tables together
- D. Renames a relation

Answer: B

Q384. What is the result of a Cartesian product of two relations?

- A. Only matching rows from both relations
- B. All possible combinations of rows from both relations
- C. The intersection of both relations
- D. An empty relation

Answer: B

Q385. For the union operation, what must be true about the two relations?

- A. They must have the same name
- B. They must be union compatible with the same number and types of attributes
- C. They must be stored in the same database
- D. They must have the same number of rows

Answer: B

Q386. What does the set difference operation $R - S$ return?

- A. All tuples in both R and S
- B. Tuples in R that are not in S
- C. Tuples in S that are not in R
- D. The union of R and S

Answer: B

Q387. Which operation combines two union-compatible relations and removes duplicates?

- A. Cartesian product
- B. Natural join
- C. Union
- D. Set difference

Answer: C

Q388. What is the rename operation used for in relational algebra?

- A. To delete a relation
- B. To change the name of a relation or its attributes
- C. To create a copy of a relation
- D. To merge two relations

Answer: B

Q389. How many fundamental operations exist in relational algebra?

- A. Three
- B. Five
- C. Six
- D. Eight

Answer: C

Q390. What is the intersection of two relations R and S?

- A. All tuples in R or S
- B. Tuples that appear in both R and S
- C. Tuples that appear in R but not S
- D. An empty set

Answer: B

Q391. Which SQL clause is used to filter rows based on a condition?

- A. ORDER BY
- B. GROUP BY
- C. WHERE
- D. HAVING

Answer: C

Q392. Which SQL statement is used to add new rows to a table?

- A. UPDATE
- B. INSERT INTO
- C. ALTER TABLE
- D. CREATE TABLE

Answer: B

Q393. What does the SELECT DISTINCT keyword do?

- A. Selects all rows including duplicates
- B. Selects only unique rows, removing duplicates
- C. Selects the first row only
- D. Selects rows in a random order

Answer: B

Q394. Which SQL keyword is used to sort the result set?

- A. SORT BY
- B. ARRANGE BY
- C. ORDER BY
- D. FILTER BY

Answer: C

Q395. What does the SQL COUNT() function do?

- A. Counts the number of tables in a database
- B. Counts the number of rows that match a criteria
- C. Counts the number of databases on the server
- D. Counts the number of characters in a column

Answer: B

Q396. Which SQL command is used to remove a table from the database?

- A. DELETE TABLE
- B. REMOVE TABLE
- C. DROP TABLE
- D. DESTROY TABLE

Answer: C

Q397. What does the SQL UPDATE statement do?

- A. Adds new rows to a table
- B. Modifies existing data in a table
- C. Deletes rows from a table
- D. Creates a new table

Answer: B

Q398. Which SQL operator is used to search for a pattern in a column?

- A. MATCH
- B. LIKE
- C. FIND
- D. SEARCH

Answer: B

Q399. What does the SQL DELETE statement do?

- A. Removes a table structure
- B. Removes specific rows from a table
- C. Removes a database
- D. Removes a column from a table

Answer: B

Q400. Which keyword is used to combine results from two SELECT statements?

- A. MERGE
- B. COMBINE
- C. UNION
- D. APPEND

Answer: C

Q401. What is a trigger in SQL?

- A. A button in the database interface
- B. A special stored procedure that automatically executes when a specified event occurs
- C. A type of database index
- D. A command to start the database server

Answer: B

Q402. What is an SQL view?

- A. A graphical display of the database
- B. A virtual table based on the result of a SELECT query
- C. A physical copy of a table
- D. A type of report

Answer: B

Q403. What is the purpose of the GRANT command in SQL?

- A. To create a new table
- B. To give users specific access privileges on database objects
- C. To delete user accounts
- D. To optimize queries

Answer: B

Q404. What does the REVOKE command do in SQL?

- A. Reverses a transaction
- B. Removes access privileges previously granted to users
- C. Restores deleted data
- D. Reverts schema changes

Answer: B

Q405. What is an assertion in SQL?

- A. A statement that prints debug information
- B. A constraint that specifies a condition the database must always satisfy
- C. A type of stored procedure
- D. A SQL comment

Answer: B

Q406. What is the purpose of the COMMIT command?

- A. To start a new transaction
- B. To save all changes made during the current transaction permanently
- C. To undo changes
- D. To lock the database

Answer: B

Q407. What does the ROLLBACK command do?

- A. Moves the database to a new server
- B. Undoes all changes made during the current transaction
- C. Deletes the entire database
- D. Creates a backup of the database

Answer: B

Q408. What is a CHECK constraint in SQL?

- A. A constraint that checks disk space
- B. A constraint that ensures values in a column satisfy a specified condition
- C. A constraint that checks user passwords
- D. A constraint that validates SQL syntax

Answer: B

Q409. What is the main goal of database normalization?

- A. To increase data redundancy
- B. To organize data to reduce redundancy and improve data integrity
- C. To make queries run faster
- D. To add more tables to the database

Answer: B

Q410. What is First Normal Form (1NF)?

- A. A table with no foreign keys
- B. A table where each column contains only atomic (indivisible) values
- C. A table with at least three columns
- D. A table stored in the first file

Answer: B

Q411. What is a repeating group in database design?

- A. A group of users who access the database repeatedly
- B. Multiple values stored in a single column violating 1NF
- C. A frequently executed query
- D. A backup that runs on a schedule

Answer: B

Q412. What does Second Normal Form (2NF) require?

- A. The table must be in 1NF and have no partial dependencies on a composite key
- B. The table must have at least two columns
- C. All columns must be numeric
- D. The table must have exactly two candidate keys

Answer: A

Q413. What is a partial dependency?

- A. A dependency on the entire primary key
- B. A non-key attribute that depends on only part of a composite primary key
- C. A dependency between two foreign keys
- D. A dependency that only sometimes holds

Answer: B

Q414. What does Third Normal Form (3NF) eliminate?

- A. All foreign keys
- B. Transitive dependencies of non-key attributes on the primary key
- C. All null values
- D. All composite keys

Answer: B

Q415. What type of anomaly occurs when deleting a row causes loss of unrelated data?

- A. Insertion anomaly
- B. Update anomaly
- C. Deletion anomaly
- D. Selection anomaly

Answer: C

Q416. Which normal form addresses the issue of repeating groups?

- A. Second Normal Form
- B. Third Normal Form
- C. First Normal Form
- D. Boyce-Codd Normal Form

Answer: C

Q417. What does the Atomicity property of a transaction ensure?

- A. The transaction is very small
- B. Either all operations in the transaction are completed or none are
- C. The transaction uses atomic data types
- D. The transaction runs on a single processor

Answer: B

Q418. What does the Consistency property of a transaction ensure?

- A. All transactions run at the same speed
- B. A transaction brings the database from one valid state to another valid state
- C. All columns have consistent data types
- D. The database is consistent in size

Answer: B

Q419. What does the Isolation property of a transaction ensure?

- A. Transactions are stored in isolated tables
- B. Concurrent transactions do not interfere with each other
- C. The database is isolated from the network
- D. Each transaction runs on a separate server

Answer: B

Q420. What does the Durability property of a transaction ensure?

- A. The transaction runs for a long duration
- B. Once a transaction is committed, its changes persist even after system failures
- C. The database hardware is durable
- D. The transaction cannot be undone

Answer: B

Q421. What are the possible states of a transaction?

- A. Start, Run, Stop
- B. Active, Partially Committed, Committed, Failed, Aborted
- C. Open, Closed, Locked
- D. Read, Write, Delete

Answer: B

Q422. What happens when a transaction is aborted?

- A. Its changes are permanently saved
- B. All changes made by the transaction are undone
- C. The database is shut down
- D. Other transactions are paused

Answer: B

Q423. What is a schedule in the context of transactions?

- A. A time table for database backups
- B. The chronological order in which operations from concurrent transactions are executed
- C. A list of database users
- D. A plan for database design

Answer: B

Q424. What is a read operation in a transaction?

- A. Reading the database schema
- B. Retrieving a data item value from the database
- C. Reading the transaction log
- D. Reading user input

Answer: B

Q425. What is concurrency control in database systems?

- A. A method to control the number of database tables
- B. A mechanism to manage simultaneous access to the database by multiple transactions
- C. A technique to control query speed
- D. A tool to manage database backups

Answer: B

Q426. What is a lock in the context of concurrency control?

- A. A physical security device for the server
- B. A mechanism that restricts access to a data item by other transactions
- C. A type of database table
- D. A method for encrypting data

Answer: B

Q427. What is a shared lock (S-lock)?

- A. A lock shared across multiple databases
- B. A lock that allows multiple transactions to read a data item simultaneously
- C. A lock that allows both reading and writing
- D. A lock that prevents all access

Answer: B

Q428. What is an exclusive lock (X-lock)?

- A. A lock available only to the DBA
- B. A lock that allows only one transaction to both read and write a data item
- C. A lock that blocks all database operations
- D. A lock used during database installation

Answer: B

Q429. What is the purpose of the two-phase locking (2PL) protocol?

- A. To lock data in two different databases
- B. To ensure serializability by dividing a transaction into a growing phase and a shrinking phase
- C. To perform two types of locks on each data item
- D. To execute transactions in two steps

Answer: B

Q430. In the two-phase locking protocol, what is the growing phase?

- A. The phase where the database grows in size
- B. The phase where a transaction acquires locks but cannot release any
- C. The phase where new tables are created
- D. The phase where transactions are added to the schedule

Answer: B

Q431. In the two-phase locking protocol, what is the shrinking phase?

- A. The phase where the database shrinks in size
- B. The phase where a transaction releases locks but cannot acquire new ones
- C. The phase where tables are deleted
- D. The phase where the transaction log is trimmed

Answer: B

Q432. What is starvation in the context of concurrency control?

- A. When a database runs out of storage
- B. When a transaction waits indefinitely because other transactions keep getting priority
- C. When the server runs out of memory
- D. When the network connection is lost

Answer: B

Q433. What is a lock manager in a DBMS?

- A. A person who manages physical locks on the server room
- B. A DBMS component that handles lock requests, grants, and releases
- C. A software for managing encryption keys
- D. A tool for managing database user accounts

Answer: B

Q434. What is a transaction log (journal)?

- A. A list of all database users
- B. A file that records all changes made to the database for recovery purposes
- C. A log of network traffic
- D. A list of all SQL queries ever written

Answer: B

Q435. What is a database backup?

- A. A secondary database server
- B. A copy of the database stored separately for recovery purposes
- C. A redundant index on a table
- D. An extra user account

Answer: B

Q436. What does the redo operation do during recovery?

- A. Re-executes all queries from the beginning
- B. Reapplies committed transactions whose changes may not have been written to disk
- C. Restarts the database server
- D. Recreates all database tables

Answer: B

Q437. What does the undo operation do during recovery?

- A. Removes the database entirely
- B. Reverses the changes of uncommitted transactions to restore consistency
- C. Undoes the last backup
- D. Removes the recovery log

Answer: B

Q438. What is the purpose of the write-ahead log rule?

- A. To ensure writes happen before reads
- B. To ensure log records are written before the corresponding data changes are flushed to disk
- C. To write data ahead of the schedule
- D. To log all user activities before database access

Answer: B

Q439. What is a system failure in database terminology?

- A. A failure in the database design process
- B. A failure that loses the contents of main memory but not disk storage
- C. A failure that destroys all data permanently
- D. A failure in the user interface

Answer: B

Q440. What is a media failure in database terminology?

- A. A failure in multimedia content
- B. A failure that damages the physical storage medium, potentially losing disk data
- C. A failure in the display monitor
- D. A failure in the network media

Answer: B

Q441. What is an incremental backup?

- A. A backup of the entire database
- B. A backup that captures only the changes made since the last backup
- C. A backup performed incrementally faster each time
- D. A backup that increments the database version

Answer: B

Q442. What is a primary index?

- A. The first index created on a table
- B. An index built on the primary key of an ordered file
- C. An index that can only store integers
- D. The most important index in the database

Answer: B

Q443. What is a B-tree?

- A. A binary tree used for sorting
- B. A balanced search tree where all leaf nodes are at the same level, used for indexing
- C. A tree that only stores binary data
- D. A tree with exactly two branches

Answer: B

Q444. What is the advantage of using indexes?

- A. They reduce storage space
- B. They speed up data retrieval operations
- C. They make all queries faster without any trade-off
- D. They eliminate the need for primary keys

Answer: B

Q445. What is the disadvantage of having too many indexes?

- A. Queries become slower to read
- B. Insert, update, and delete operations become slower as indexes must be maintained
- C. The database becomes read-only
- D. Indexes use no additional storage

Answer: B

Q446. How many clustered indexes can a table have?

- A. As many as needed
- B. Exactly one
- C. Exactly two
- D. None

Answer: B

Q447. What is query optimization?

- A. Writing queries in the shortest way possible
- B. The process of selecting the most efficient execution plan for a query
- C. Removing unnecessary queries from the database
- D. Optimizing the database for more queries

Answer: B

Q448. What is the role of the query parser?

- A. To execute the query
- B. To check the syntax and validate the SQL query before optimization
- C. To store query results
- D. To create database tables

Answer: B

Q449. What is a table scan (full table scan)?

- A. Scanning the table structure for errors
- B. Reading every row in a table to find matching records
- C. Scanning only the index
- D. A security scan of the table

Answer: B

Q450. Why is using an index generally faster than a full table scan?

- A. Indexes are stored on faster disks
- B. Indexes narrow down the search to specific records without scanning every row
- C. Indexes compress the data for faster reading
- D. Indexes always return fewer results

Answer: B

Q451. What does the EXPLAIN command show in SQL?

- A. An explanation of SQL syntax
- B. The execution plan the DBMS will use for a query
- C. The database schema
- D. A list of all indexes

Answer: B

Q452. What is a join operation in the context of query optimization?

- A. Joining two databases together
- B. Combining rows from two or more tables based on related columns
- C. Joining user sessions
- D. Adding columns to a table

Answer: B

Q453. What is cost-based optimization?

- A. Optimizing the monetary cost of database licenses
- B. Selecting the execution plan with the lowest estimated resource cost
- C. Minimizing the cost of database hardware
- D. Optimizing query writing costs

Answer: B

Q454. What is a query tree (query plan tree)?

- A. A tree-shaped database
- B. A tree structure representing the relational algebra operations of a query
- C. A decision tree for choosing queries
- D. A directory structure for storing queries

Answer: B

Q455. What are database statistics used for in query optimization?

- A. Generating reports about database usage
- B. Providing the optimizer with data distribution information to estimate operation costs
- C. Tracking user login statistics
- D. Measuring network performance

Answer: B

Q456. What is data encryption?

- A. Compressing data to save space
- B. Converting data into an unreadable format that can only be decoded with a key
- C. Backing up data to a remote location
- D. Indexing data for faster access

Answer: B

Q457. What is an access control list (ACL) in database security?

- A. A checklist for database design
- B. A list specifying which users have access to which database objects and what operations they can perform
- C. A list of all database connections
- D. A list of recently accessed tables

Answer: B

Q458. What is a database audit trail?

- A. A hiking trail near the data center
- B. A record of all activities performed on the database for monitoring and accountability
- C. A trail of deleted data
- D. A path for data migration

Answer: B

Q459. What is a database firewall?

- A. A physical wall around the server room
- B. A security layer that monitors and filters database traffic to prevent unauthorized access
- C. A backup power supply
- D. A type of database index

Answer: B

Q460. What is data masking?

- A. Hiding the database from the network
- B. Replacing sensitive data with fictitious but realistic data to protect privacy
- C. Compressing data to reduce size
- D. Encrypting data at rest

Answer: B

Q461. What is the advantage of data replication?

- A. It saves storage space
- B. It improves data availability and local read performance
- C. It eliminates the need for backups entirely
- D. It speeds up write operations

Answer: B

Q462. What is a homogeneous distributed database?

- A. A database with the same data at all sites
- B. A distributed database where all sites use the same DBMS software
- C. A database that stores only one type of data
- D. A database with uniform table sizes

Answer: B

Q463. What is a heterogeneous distributed database?

- A. A database with mixed data types
- B. A distributed database where different sites may use different DBMS software
- C. A database with unequal table sizes
- D. A database that stores data in different formats

Answer: B

Q464. What is location transparency in distributed databases?

- A. The database hides where data is physically stored from users
- B. The database reveals the location of all data to users
- C. Transparency about the data center location
- D. Showing users the physical address of each server

Answer: A

Q465. What is replication transparency?

- A. Making replication visible to all users
- B. Users are unaware that data is replicated; the system manages copies automatically
- C. Transparent replication logs
- D. Showing users all copies of data

Answer: B

Q466. Which of the following is a document-based NoSQL database?

- A. MySQL
- B. PostgreSQL
- C. MongoDB
- D. Oracle

Answer: C

Q467. Which of the following is an example of a key-value store?

- A. MongoDB
- B. Redis
- C. Neo4j
- D. Cassandra

Answer: B

Q468. What is a column-family database?

- A. A database organized by families of users
- B. A NoSQL database that stores data in columns grouped into column families
- C. A database with family-related data only
- D. A relational database feature

Answer: B

Q469. Which of the following is a graph database?

- A. Redis
- B. Cassandra
- C. Neo4j
- D. HBase

Answer: C

Q470. What is the main advantage of NoSQL databases over relational databases?

- A. They always provide ACID guarantees
- B. They offer better scalability and flexibility for unstructured or semi-structured data
- C. They always have better query capabilities
- D. They require less storage

Answer: B

Q471. What is a document in a document-based NoSQL database?

- A. A text file on disk
- B. A self-contained data record often in JSON or BSON format with flexible schema
- C. A PDF stored in the database
- D. A database manual

Answer: B

Q472. What does horizontal scaling mean in the context of NoSQL?

- A. Adding more columns to tables
- B. Adding more servers to handle increased load
- C. Making the server hardware more powerful
- D. Increasing disk space on one server

Answer: B

Q473. What is vertical scaling?

- A. Adding more servers to the cluster
- B. Increasing the capacity of a single server by adding more CPU, RAM, or storage
- C. Adding more rows to a table
- D. Stacking databases vertically

Answer: B

Q474. What is ETL in data warehousing?

- A. Electronic Transfer Link
- B. Extract, Transform, Load - the process of moving data from sources into the data warehouse
- C. Extended Table Logic
- D. Enterprise Technology Layer

Answer: B

Q475. What is a dimension table in a data warehouse?

- A. A table that stores database dimensions (size)
- B. A table containing descriptive attributes used to filter, group, and label fact data
- C. A table with three-dimensional data
- D. A temporary table

Answer: B

Q476. What is a star schema in data warehousing?

- A. A schema shaped like a star on the database diagram
- B. A schema with a central fact table connected to multiple dimension tables, forming a star pattern
- C. A schema used in astronomy databases
- D. A schema with star-rated data

Answer: B

Q477. What is OLAP?

- A. Online Automatic Processing
- B. Online Analytical Processing - technology for complex analytical queries on multidimensional data
- C. Open Link Access Protocol
- D. Offline Loading and Processing

Answer: B

Q478. What is a measure in a data warehouse fact table?

- A. The physical size of the table
- B. A numeric value that can be aggregated, such as sales amount or quantity
- C. A unit of measurement for storage
- D. A count of dimension tables

Answer: B

Q479. What is a data cube in OLAP?

- A. A cubic-shaped storage device
- B. A multidimensional representation of data allowing analysis across multiple dimensions
- C. A backup cube for data protection
- D. A rubik's cube made of data drives

Answer: B

Q480. What is a DEFAULT constraint in SQL?

- A. The default database setting
- B. A constraint that assigns a default value to a column when no value is specified during insertion
- C. A constraint that defaults to NULL
- D. The default table in a database

Answer: B

Q481. What is the purpose of the UNIQUE constraint in SQL?

- A. To make all rows unique across all columns
- B. To ensure that all values in a column or combination of columns are distinct
- C. To create a unique database
- D. To uniquely identify the database

Answer: B

Q482. What is an unnormalized table?

- A. A table without a name
- B. A table that contains repeating groups or non-atomic values violating 1NF
- C. A table with too many columns
- D. A table that is not indexed

Answer: B

Q483. What is the relationship between normalization and data redundancy?

- A. Normalization increases data redundancy
- B. Normalization reduces data redundancy by organizing data into separate related tables
- C. Normalization has no effect on redundancy
- D. Normalization eliminates all data from the database

Answer: B

Q484. What is a write operation in a transaction?

- A. Writing documentation about the database
- B. Modifying the value of a data item in the database
- C. Writing a query to the console
- D. Creating a new table

Answer: B

Q485. What does it mean when a transaction is in the 'committed' state?

- A. The transaction is waiting to start
- B. The transaction has completed successfully and its changes are permanently recorded
- C. The transaction has been cancelled
- D. The transaction is still running

Answer: B

Q486. Why is concurrency control necessary in database systems?

- A. To make the database run faster
- B. To ensure data consistency and isolation when multiple transactions access data simultaneously
- C. To control the number of database users
- D. To manage disk space

Answer: B

Q487. What is a full database backup?

- A. A backup that only stores schema information
- B. A complete copy of the entire database at a specific point in time
- C. A backup of the backup system
- D. A partial backup of frequently accessed data

Answer: B

Q488. What is the purpose of a log file in database recovery?

- A. To log user login attempts
- B. To record all database modifications so transactions can be redone or undone during recovery
- C. To store query execution plans
- D. To track disk space usage

Answer: B

Q489. What is the search key of an index?

- A. A physical key to unlock the database
- B. The attribute or set of attributes used to look up records in the index
- C. A password for accessing indexed data
- D. The primary key of the indexed table

Answer: B

Q490. Can a table have multiple non-clustered indexes?

- A. No, only one index per table
- B. Yes, a table can have multiple non-clustered indexes on different columns
- C. Only if the table has no primary key
- D. Only two non-clustered indexes are allowed

Answer: B

Q491. What is an index entry composed of?

- A. Only the table name
- B. A search key value and a pointer to the corresponding data record or block
- C. Only the row number
- D. A copy of the entire row

Answer: B

Q492. What happens to indexes when a row is inserted into a table?

- A. Indexes are unaffected by insertions
- B. All relevant indexes must be updated to include entries for the new row
- C. Indexes are automatically deleted
- D. Only the primary index is updated

Answer: B

Q493. Why might a database optimizer choose not to use an available index?

- A. Indexes can never be ignored
- B. When the optimizer estimates a full table scan would be faster, such as when most rows match the query
- C. The optimizer always uses indexes
- D. Indexes are only used for SELECT queries

Answer: B

Q494. What is a query rewrite in the optimization process?

- A. Rewriting the query in a different programming language
- B. Transforming a query into an equivalent but more efficient form before generating execution plans
- C. Manually rewriting queries for better readability
- D. Rewriting queries to fix syntax errors

Answer: B

Q495. What is the purpose of a password policy in database security?

- A. To make passwords decorative
- B. To enforce rules about password strength, expiration, and reuse to prevent unauthorized access
- C. To store passwords in plain text
- D. To share passwords among users

Answer: B

Q496. What is a privilege in database security?

- A. A type of database table
- B. A right or permission granted to a user to perform specific operations on database objects
- C. A special type of data
- D. A database backup method

Answer: B

Q497. What is the difference between system privileges and object privileges?

- A. There is no difference
- B. System privileges allow operations on the system level while object privileges allow operations on specific database objects
- C. System privileges are for hardware while object privileges are for software
- D. Object privileges are more powerful

Answer: B

Q498. Why is it important to regularly update database software?

- A. To change the database color theme
- B. To patch security vulnerabilities and protect against newly discovered threats
- C. To increase storage space
- D. To delete old data

Answer: B

Q499. What is a distributed database management system (DDBMS)?

- A. A DBMS installed on multiple monitors
- B. Software that manages a distributed database, making distribution transparent to users
- C. A mobile database application
- D. A database stored on a USB drive

Answer: B

Q500. What is a site (node) in a distributed database?

- A. A website that hosts the database
- B. A physical or logical location that has its own local database and processing capability
- C. A web page with database information
- D. A user workstation

Answer: B

Q501. What is the benefit of distributing a database across multiple sites?

- A. It makes the database harder to use
- B. It improves reliability, availability, and performance by placing data closer to where it is used
- C. It reduces the total amount of data stored
- D. It eliminates the need for security

Answer: B

Q502. What is a communication network in a distributed database system?

- A. A social network for database users
- B. The network infrastructure that connects different database sites and enables data transfer between them
- C. A chat system within the database
- D. An internal database messaging system

Answer: B

Q503. What is a collection in MongoDB?

- A. A group of database servers
- B. A grouping of documents, similar to a table in relational databases
- C. A collection of SQL queries
- D. A backup set

Answer: B

Q504. What is a node in a graph database?

- A. A network server
- B. An entity or data point represented as a vertex in the graph, such as a person or product
- C. A connection between two databases
- D. A cluster of documents

Answer: B

Q505. What are the four key characteristics of a data warehouse according to Bill Inmon?

- A. Fast, Reliable, Secure, Scalable
- B. Subject-oriented, Integrated, Time-variant, and Non-volatile
- C. Structured, Indexed, Normalized, and Backed-up
- D. Large, Distributed, Real-time, and Encrypted

Answer: B

Q506. What is a data source in the context of data warehousing?

- A. The power source for the data center
- B. An operational system, external file, or other origin from which data is extracted for the warehouse
- C. The database administrator
- D. A query result

Answer: B

Q507. What is a dashboard in business intelligence?

- A. The front panel of a database server
- B. A visual display of key performance indicators and metrics from the data warehouse for decision making
- C. A database configuration page
- D. A network monitoring tool

Answer: B

Q508. What is the staging area in a data warehouse architecture?

A. A theater stage for database presentations

B. An intermediate storage area where extracted data is cleaned and transformed before loading into the warehouse

- C. The main warehouse storage
- D. A backup location

Answer: B

Q509. What does it mean to revoke a privilege from a database user?

- A. To grant additional privileges
- B. To remove a previously granted permission from the user
- C. To delete the user account
- D. To reset the user password

Answer: B

Q510. What is network latency in the context of distributed databases?

- A. The speed of the database processor
- B. The time delay for data to travel between sites across the network
- C. The size of the database
- D. The number of concurrent users

Answer: B

Databases — Medium

510 questions

Q511. Which level of data abstraction describes how data is actually stored?

- A. View level
- B. Logical level
- C. External level
- D. Physical level

Answer: D

Q512. The three-schema architecture was proposed by:

- A. IEEE standards body
- B. ISO organization
- C. W3C consortium
- D. ANSI/SPARC committee

Answer: D

Q513. Which of the following is a disadvantage of DBMS?

- A. Increased data redundancy across tables
- B. Program-data dependence across modules
- C. High initial cost of hardware and software
- D. Data inconsistency between applications

Answer: C

Q514. What is metadata in a database context?

- A. Temporary data
- B. Data about data
- C. Deleted data
- D. Encrypted data

Answer: B

Q515. Which component of DBMS is responsible for ensuring data consistency?

- A. Integrity manager
- B. Authorization manager
- C. Buffer pool manager
- D. File system manager

Answer: A

Q516. A data dictionary stores:

- A. Application programs and their code
- B. Actual user data stored in tables
- C. Metadata about the database structure
- D. Backup files created periodically

Answer: C

Q517. Which of the following describes logical data independence?

- A. Ability to change applications without affecting the underlying data
- B. Ability to change the logical schema without affecting external schemas
- C. None of these options correctly describes logical data independence
- D. Ability to change physical storage without affecting the logical schema

Answer: B

Q518. What is a database instance?

- A. The database management software application
- B. The overall architectural design of the database
- C. A backup copy of the database stored on disk
- D. The collection of data stored at a particular moment

Answer: D

Q519. A database schema is:

- A. A data structure used for indexing records
- B. A language used for querying the database
- C. The overall design or structure of the database
- D. The actual data values stored in the tables

Answer: C

Q520. Which of the following is NOT a function of a DBMS?

- A. Generating source code for applications
- B. Manipulating and querying stored data
- C. Managing concurrent transaction access
- D. Defining and managing data structures

Answer: A

Q521. In the three-schema architecture, which schema defines the user's view?

- A. External schema
- B. Internal schema
- C. Conceptual schema
- D. Physical schema

Answer: A

Q522. Which DBMS component ensures that the database remains in a consistent state after a failure?

- A. Authorization manager
- B. Buffer manager
- C. Recovery manager
- D. Query optimizer

Answer: C

Q523. The data definition language (DDL) compiler processes:

- A. Application programs written in the host language
- B. Schema definitions and stores metadata in the catalog
- C. User-submitted queries for data retrieval operations
- D. Data insertion commands for populating new tables

Answer: B

Q524. In a two-tier architecture, the application logic resides:

- A. Exclusively on the client tier
- B. Exclusively on the server tier
- C. On a separate dedicated middle tier
- D. On the client side or server side

Answer: D

Q525. What is a database catalog?

- A. A system database containing metadata about the database structure
- B. A comprehensive list of all users who have database access rights
- C. An index structure built on primary keys of all related tables
- D. A sequential log file that records all committed transactions

Answer: A

Q526. Which architecture is most suitable for web-based database applications?

- A. Peer-to-peer architecture
- B. Two-tier architecture
- C. Single-tier architecture
- D. Three-tier architecture

Answer: D

Q527. The runtime database processor executes:

- A. Backup and recovery operations only
- B. Data definition language statements only
- C. Schema compilation processes only
- D. Privileged commands and query plans at runtime

Answer: D

Q528. What is the purpose of the DML precompiler?

- A. To compile DDL commands and update the system catalog
- B. To extract DML commands embedded in application programs
- C. To manage concurrency control among active transactions
- D. To optimize queries by selecting the best execution plan

Answer: B

Q529. In a parallel database system, the main goal is:

- A. Reduced overall storage cost on disk
- B. Simplified query writing for users
- C. Eliminating the need for any backups
- D. Improved performance through parallelism

Answer: D

Q530. Shared-nothing architecture means:

- A. All processors share the memory and disk
- B. Only disk is shared among all processors
- C. Each processor has its own memory and disk
- D. Only memory is shared among processors

Answer: C

Q531. In an Enhanced ER (EER) model, specialization is:

- A. Adding new relationships between existing entity types
- B. Combining multiple entity types into a single entity
- C. The process of defining subclasses of an entity type
- D. Removing attributes from entities during refinement

Answer: C

Q532. A weak entity in an ER diagram:

- A. Cannot be uniquely identified by its own attributes alone
- B. Is always fully independent of any other entity type
- C. Cannot have any attributes or properties defined on it
- D. Has a complete primary key defined on its own attributes

Answer: A

Q533. Total participation in an ER diagram means:

- A. The relationship is optional for all entities
- B. No entity participates in the defined relationship
- C. Only some entities participate in the relationship
- D. Every entity must participate in the relationship

Answer: D

Q534. A composite attribute is:

- A. An attribute that holds multiple values for one entity
- B. An attribute that serves as the key for identification
- C. An attribute whose value is derived from other attributes
- D. An attribute that can be divided into smaller sub-attributes

Answer: D

Q535. A multivalued attribute is represented in an ER diagram by:

- A. A single ellipse
- B. A rectangle shape
- C. A dashed ellipse
- D. A double ellipse

Answer: D

Q536. In an ER diagram, a double rectangle represents:

- A. A weak entity
- B. An attribute
- C. A relationship
- D. A strong entity

Answer: A

Q537. Generalization in EER modeling is:

- A. Adding new attributes to an existing entity class
- B. Removing a class entirely from the ER model
- C. Splitting a single class into multiple subclasses
- D. Combining subclasses into a generalized superclass

Answer: D

Q538. A derived attribute is one whose value:

- A. Must be entered manually by users
- B. Cannot change once it is set initially
- C. Is always NULL in every record stored
- D. Can be computed from other attributes

Answer: D

Q539. The cardinality ratio in an ER model specifies:

- A. The overall size of the database measured in stored records
- B. The number of attributes defined on a single entity type
- C. The total number of entities present in the entire ER model
- D. The number of relationship instances an entity can participate in

Answer: D

Q540. An object-oriented data model supports:

- A. Only tree structures and parent-child links
- B. Only tables, rows, and column structures
- C. Only flat files and sequential access paths
- D. Encapsulation, inheritance, and polymorphism

Answer: D

Q541. A candidate key is:

- A. Any arbitrary attribute present in the table
- B. The same concept as a foreign key reference
- C. A minimal superkey that uniquely identifies tuples
- D. Always the first column defined in the schema

Answer: C

Q542. A superkey is:

- A. Any set of attributes that uniquely identifies tuples
- B. A key that is permitted to contain NULL values
- C. Always limited to a single attribute in the table
- D. The same concept as a foreign key in a relation

Answer: A

Q543. Referential integrity constraint ensures that:

- A. A foreign key value must match a primary key value in the referenced table or be NULL
- B. All columns across every table in the database must have the same compatible data type
- C. All tables in the database schema must have exactly the same total number of data rows
- D. Primary key values can be duplicated across different rows within the same data table

Answer: A

Q544. Entity integrity constraint states that:

- A. Foreign keys must always be unique in tables
- B. Tables must have at least one row of data
- C. All attributes in every row must have values
- D. No primary key attribute can have a NULL value

Answer: D

Q545. An alternate key is:

- A. A composite key made up of two or more attributes
- B. A candidate key that is not chosen as the primary key
- C. A key that is permitted to contain NULL value entries
- D. The same concept as a foreign key in another table

Answer: B

Q546. A composite key is:

- A. A key with only one single attribute value
- B. A foreign key referencing another table
- C. A key that allows NULL values in rows
- D. A key consisting of two or more attributes

Answer: D

Q547. Which of the following is NOT a relational integrity constraint?

- A. Entity integrity
- B. Sequential constraint
- C. Referential integrity
- D. Domain constraint

Answer: B

Q548. A relation schema defines:

- A. Only the primary key definition of the relation in question
- B. The structure of a relation including attribute names and domains
- C. The actual data values currently stored within the relation
- D. Only the foreign keys that reference other related tables

Answer: B

Q549. In relational algebra, the Cartesian product of two relations R and S produces:

- A. All possible combinations of tuples from R and S
- B. The intersection of tuples in both R and S
- C. Only tuples that match on common attributes
- D. The union of all tuples from both R and S

Answer: A

Q550. A relation in 1NF requires that:

- A. The table has only one column of values
- B. There are no foreign keys in the relation
- C. All values are strictly numeric in type
- D. All attribute values are atomic (indivisible)

Answer: D

Q551. The theta join of relations R and S is:

- A. A rename operation applied to change the relation name
- B. A Cartesian product followed by selection with a condition
- C. A set union combining all tuples from two given relations
- D. A simple projection on selected attributes from the relation

Answer: B

Q552. An equi-join is a special case of theta join where:

- A. No condition is applied to the join result
- B. The result of the join is always an empty set
- C. The condition uses only inequality comparisons
- D. The condition uses only equality comparisons

Answer: D

Q553. The division operation $R \div S$ returns:

- A. The Cartesian product of both relations R and S fully
- B. Tuples in R that are associated with every tuple in S
- C. All tuples from relation S regardless of the condition
- D. All tuples from relation R regardless of the condition

Answer: B

Q554. Left outer join includes:

- A. All tuples from the right relation and matching from the left side
- B. Only tuples that have matching values in both of the joined tables
- C. No tuples at all if there is no matching condition between tables
- D. All tuples from the left relation and matching tuples from the right

Answer: D

Q555. In relational algebra, which operation can be derived from the fundamental operations?

- A. Projection
- B. Intersection
- C. Division
- D. Selection

Answer: B

Q556. The semi-join of R and S returns:

- A. All tuples from both relations R and S combined
- B. The complete Cartesian product of R and S result
- C. Only the tuples from relation S not in relation R
- D. Tuples from R that have matching tuples in S

Answer: D

Q557. Anti-join of R and S returns:

- A. All tuples from R regardless of matching in S
- B. The complete Cartesian product of R and S both
- C. Tuples from R that have no matching tuples in S
- D. All tuples that match between relations R and S

Answer: C

Q558. Relational calculus is:

- A. A procedural query language that specifies how to retrieve data
- B. A data definition language for creating tables and constraints
- C. A physical storage language for defining data file structures
- D. A non-procedural query language that describes what to retrieve

Answer: D

Q559. Tuple relational calculus uses:

- A. Tuple variables that range over relations
- B. Domain variables that range over values
- C. Only table names without any variables
- D. No variables of any kind are required

Answer: A

Q560. Domain relational calculus uses:

- A. SQL queries written in standard syntax format
- B. No variables of any kind within the expression
- C. Domain variables that range over attribute domains
- D. Tuple variables that range over entire relations

Answer: C

Q561. Which SQL clause is used to filter groups in a GROUP BY query?

- A. WHERE
- B. HAVING
- C. ORDER BY
- D. LIMIT

Answer: B

Q562. An INNER JOIN returns:

- A. All rows from the left table with NULLs for right
- B. Only rows that have matching values in both tables
- C. All rows from the right table with NULLs for left
- D. All rows from both the left and the right tables

Answer: B

Q563. Which aggregate function returns the number of rows?

- A. COUNT()
- B. SUM()
- C. AVG()
- D. MAX()

Answer: A

Q564. A LEFT JOIN returns:

- A. The Cartesian product of both tables without any condition
- B. All rows from the right table and matching from the left
- C. Only rows that have matching values in both joined tables
- D. All rows from the left table and matching rows from the right

Answer: D

Q565. Which SQL constraint ensures a column cannot have NULL values?

- A. CHECK
- B. UNIQUE
- C. NOT NULL
- D. DEFAULT

Answer: C

Q566. The LIKE operator is used with which wildcard characters?

- A. # and @
- B. % and _
- C. * and ?
- D. & and ^

Answer: B

Q567. What does the IN operator do in SQL?

- A. Checks if a value matches any value in a list
- B. Performs an inner join between two given tables
- C. Inserts data into a specified target table row
- D. Indexes a column for faster search performance

Answer: A

Q568. Which SQL clause limits the number of rows returned?

- A. LIMIT (or TOP in SQL Server)
- B. GROUP BY (aggregation clause)
- C. HAVING (group filter clause)
- D. WHERE (filter condition clause)

Answer: A

Q569. A subquery is:

- A. A query that modifies existing data
- B. A query that creates new table objects
- C. A stored procedure in the database
- D. A query nested inside another query

Answer: D

Q570. The BETWEEN operator selects values:

- A. Within a given inclusive range
- B. Outside a specified range only
- C. That are NULL valued entries
- D. Equal to a single specific value

Answer: A

Q571. A materialized view differs from a regular view in that it:

- A. Cannot be refreshed or updated at all
- B. Physically stores the result set on disk
- C. Is always virtual and never materialized
- D. Does not use SELECT queries for data

Answer: B

Q572. Window functions in SQL are used to:

- A. Delete rows from tables based on specified filtering conditions
- B. Create new tables by defining their schema and column structures
- C. Perform calculations across a set of rows related to the current row
- D. Modify table structure by adding or removing column definitions

Answer: C

Q573. The ROW_NUMBER() function in SQL:

- A. Deletes duplicate rows from the result set of the current query
- B. Returns the primary key value from the underlying data table
- C. Assigns a unique sequential number to each row within a partition
- D. Counts the total number of rows stored in the entire data table

Answer: C

Q574. A Common Table Expression (CTE) is:

- A. A trigger that fires automatically in response to a data change
- B. A stored procedure that encapsulates reusable database SQL logic
- C. A permanent table that persists after the query has been executed
- D. A temporary named result set used within a single SQL statement

Answer: D

Q575. The RANK() function differs from ROW_NUMBER() in that:

- A. RANK() never produces ties in the ranking of any result set rows
- B. RANK() assigns the same rank to ties and skips subsequent ranks
- C. ROW_NUMBER() handles ties by assigning them the same rank value
- D. They are identical and produce exactly the same results always

Answer: B

Q576. A recursive CTE is used for:

- A. Querying hierarchical or recursive data structures
- B. Creating indexes on one or more table columns
- C. Optimizing joins between tables in the query
- D. Dropping tables from the database permanently

Answer: A

Q577. The SAVEPOINT statement in SQL:

- A. Commits the transaction and makes all changes fully permanent
- B. Starts a new transaction and initializes all session variables
- C. Creates a point within a transaction to which you can roll back
- D. Drops a table and removes it permanently from the data schema

Answer: C

Q578. An INSTEAD OF trigger:

- A. Executes after the triggering action ends
- B. Never executes regardless of the trigger
- C. Executes before the triggering action runs
- D. Executes in place of the triggering action

Answer: D

Q579. The MERGE statement in SQL combines:

- A. Only DDL operations such as CREATE, ALTER, and DROP tables
- B. Only INSERT and DELETE operations without any UPDATE logic
- C. Only SELECT operations for querying data from source tables
- D. INSERT, UPDATE, and DELETE operations based on a condition

Answer: D

Q580. What is a sequence in SQL?

- A. A type of trigger that fires on data modification
- B. An object that generates sequential numeric values
- C. An ordered view of data filtered by a set condition
- D. A sorted table arranged by a specific column value

Answer: B

Q581. Boyce-Codd Normal Form (BCNF) is stricter than 3NF because:

- A. Every determinant must be a candidate key
- B. It allows partial dependencies on some keys
- C. It requires no functional dependencies at all
- D. It requires only atomic values in each field

Answer: A

Q582. A relation in BCNF is always in:

- A. 4NF
- B. 1NF
- C. 5NF
- D. 3NF

Answer: D

Q583. A multivalued dependency X '!' Y exists when:

- A. Y is a primary key attribute that uniquely identifies every tuple
- B. A set of values of Y is determined by X, independent of other attributes
- C. X functionally determines Y through a standard functional dependency
- D. No dependency exists between the attributes X and Y in the relation

Answer: B

Q584. Fourth Normal Form (4NF) requires:

- A. No functional dependencies of any kind in the schema
- B. BCNF and no non-trivial multivalued dependencies
- C. All attributes serve as keys in the relation structure
- D. Only 1NF without any additional constraints applied

Answer: B

Q585. A lossless join decomposition ensures that:

- A. Some data is always permanently lost during the decomposition process
- B. Joining the decomposed relations reproduces the original relation exactly
- C. Extra spurious tuples are generated when joining decomposed relations
- D. Only key attributes are preserved in the decomposed relation fragments

Answer: B

Q586. Which normal form deals with join dependencies?

- A. Third Normal Form (3NF)
- B. First Normal Form (1NF)
- C. Second Normal Form (2NF)
- D. Fifth Normal Form (5NF)

Answer: D

Q587. Armstrong's axiom of reflexivity states:

- A. If X determines Y, then XZ determines YZ for any Z
- B. If Y is a subset of X, then X determines Y
- C. If X determines Y and Y determines Z, then X determines Z
- D. If X determines Y, then Y also determines X always

Answer: B

Q588. Armstrong's axiom of augmentation states:

- A. If X determines Y and Y determines Z, then X determines Z
- B. If X determines Y, then Y also determines X in reverse
- C. If X determines Y, then XZ determines YZ for any attribute Z
- D. If Y is a subset of X, then X functionally determines Y

Answer: C

Q589. The canonical (minimal) cover of a set of FDs:

- A. Is always empty regardless of the original dependency set
- B. Contains all possible FDs derivable from the original set
- C. Has maximum redundancy among all its functional dependencies
- D. Has no redundant dependencies and no extraneous attributes

Answer: D

Q590. A dependency-preserving decomposition ensures that:

- A. All original functional dependencies can be enforced without joining decomposed tables
- B. All tables in the decomposition must have exactly the same schema and attributes
- C. Joins are never needed between any of the resulting decomposed relation fragments
- D. Some functional dependencies are always lost during the process of decomposition

Answer: A

Q591. A serial schedule is one where:

- A. Only one transaction is allowed in the database system total
- B. Operations are randomly ordered across multiple transactions
- C. Transactions execute one after another without interleaving
- D. Transactions execute concurrently with operations interleaved

Answer: C

Q592. A serializable schedule is:

- A. A schedule whose result is equivalent to some serial schedule
- B. A schedule where all transactions run simultaneously at once
- C. A schedule that is always a serial schedule by definition
- D. A schedule with no transactions present within its timeline

Answer: A

Q593. Conflict serializability is tested using:

- A. Hash table (bucket-based structure)
- B. B-tree index (balanced tree structure)
- C. ER diagram (entity relationship model)
- D. Precedence graph (serialization graph)

Answer: D

Q594. Two operations conflict if:

- A. They belong to different transactions, access the same data item, and at least one is a write
- B. They are both read operations accessing the same data item from within different transactions
- C. They belong to the same transaction and access different data items within the transaction
- D. They access different data items from different transactions without any write operations

Answer: A

Q595. A dirty read occurs when:

- A. No read operation occurs between any concurrent transactions
- B. A transaction reads data that it has written by itself before
- C. A transaction reads data written by an uncommitted transaction
- D. A transaction reads data that has already been fully committed

Answer: C

Q596. A non-repeatable read occurs when:

- A. No read operation is performed during the entire transaction
- B. All reads from a transaction return the same consistent value
- C. A transaction reads a value once and never accesses it again
- D. A transaction reads the same item twice and gets different values

Answer: D

Q597. A phantom read occurs when:

- A. A transaction re-executes a query and finds new rows added by another transaction
- B. A transaction reads no rows at all from the table during its entire execution
- C. A transaction creates a table and populates it with data from another data source
- D. A transaction deletes rows from the table and those changes are then committed

Answer: A

Q598. The READ UNCOMMITTED isolation level allows:

- A. No anomalies of any kind during transaction processing
- B. Dirty reads, non-repeatable reads, and phantom reads
- C. Only phantom reads but prevents all other read anomalies
- D. Only non-repeatable reads but prevents all dirty reads

Answer: B

Q599. The SERIALIZABLE isolation level provides:

- A. Complete isolation from other transactions
- B. Only dirty read prevention mechanism
- C. Faster performance than other levels
- D. No isolation between any transactions

Answer: A

Q600. A cascading rollback occurs when:

- A. No rollback is needed because every transaction completes without any failures
- B. Only one transaction rolls back while all other transactions continue normally
- C. One transaction's rollback forces other dependent transactions to also roll back
- D. All transactions commit successfully without any errors during their execution

Answer: C

Q601. Strict Two-Phase Locking requires that:

- A. No locks are needed for any transaction under this locking protocol
- B. Only shared locks are used and exclusive locks are never acquired
- C. All exclusive locks are held until the transaction commits or aborts
- D. Locks can be released at any time during the transaction processing

Answer: C

Q602. Rigorous Two-Phase Locking requires:

- A. Locks are released immediately after each individual use
- B. All locks (shared and exclusive) are held until commit or abort
- C. Only exclusive locks are held until the transaction commits
- D. No growing phase exists in this locking protocol at all

Answer: B

Q603. A wait-die deadlock prevention scheme:

- A. Allows younger transactions to wait; older ones are rolled back
- B. Allows older transactions to wait; younger ones are rolled back
- C. Prevents all waiting by immediately aborting every transaction
- D. Kills all transactions to resolve and prevent deadlock situations

Answer: B

Q604. A wound-wait deadlock prevention scheme:

- A. All transactions wait without any preemption at all now
- B. No preemption occurs under any circumstances in the system
- C. Older transactions preempt younger ones; younger ones wait
- D. Younger transactions preempt older ones; older ones wait

Answer: C

Q605. A deadlock detection algorithm uses:

- A. A hash table index
- B. A wait-for graph
- C. An ER diagram model
- D. A B-tree structure

Answer: B

Q606. Timestamp-based concurrency control assigns:

- A. Random access to data items without any order
- B. Locks to all data items accessed by transactions
- C. A unique timestamp to each transaction for ordering
- D. Priority based on the size of each transaction

Answer: C

Q607. In timestamp ordering, if a transaction tries to write a data item with a later read timestamp:

- A. The data item gets deleted
- B. The transaction is rolled back
- C. The read timestamp is updated
- D. The write proceeds normally

Answer: B

Q608. Optimistic concurrency control assumes that:

- A. Conflicts are frequent and locks must always be acquired
- B. All transactions conflict and none can proceed at once
- C. No validation is needed at any point during processing
- D. Conflicts are rare, so validation is done at commit time

Answer: D

Q609. The three phases of optimistic concurrency control are:

- A. Start, Process, End
- B. Read, Validation, Write
- C. Read, Write, Commit
- D. Lock, Execute, Unlock

Answer: B

Q610. Lock granularity refers to:

- A. The number of transactions in the system
- B. The type of query being currently executed
- C. The size of the data item that is locked
- D. The speed of acquiring locks on the item

Answer: C

Q611. Deferred update (no-undo/redo) recovery technique:

- A. Requires neither undo nor redo operations during the recovery process
- B. Writes changes immediately and may need undo for uncommitted changes
- C. Does not write changes to disk until commit, so undo is never needed
- D. Does not use a log file for recording any of the transaction changes

Answer: C

Q612. Immediate update recovery technique:

- A. Does not use a log file for any of the transaction record entries
- B. Writes changes only after commit, requiring only redo operations
- C. Requires only redo operations during the crash recovery process
- D. Writes changes to disk before commit, requiring both undo and redo

Answer: D

Q613. Shadow paging recovery technique:

- A. Maintains two page tables: current and shadow
- B. Requires checkpoints at all times to function
- C. Uses only a log file for the recovery process
- D. Cannot handle any system crashes that may occur

Answer: A

Q614. After a crash, which transactions need to be redone?

- A. All transactions regardless of their commit state now
- B. Committed transactions whose changes may not be on disk
- C. Uncommitted transactions that were active during crash
- D. Only aborted transactions that failed before the crash

Answer: B

Q615. After a crash, which transactions need to be undone?

- A. All transactions regardless of their state at the time of the crash
- B. No transactions need to be undone during the recovery after a crash
- C. Transactions that were active (not committed) at the time of crash
- D. Committed transactions that completed successfully before the crash

Answer: C

Q616. The ARIES recovery algorithm stands for:

- A. Automated Recovery and Integrity Enforcement System
- B. Advanced Recovery and Index Enhancement System
- C. Algorithm for Redo and Integrity Evaluation System
- D. Algorithm for Recovery and Isolation Exploiting Semantics

Answer: D

Q617. ARIES uses which three phases during recovery?

- A. Scan, Parse, Execute
- B. Analysis, Redo, Undo
- C. Lock, Execute, Unlock
- D. Read, Write, Commit

Answer: B

Q618. A log sequence number (LSN) is:

- A. The total number of active transactions
- B. A type of lock used in concurrency control
- C. The size of the log file on stable storage
- D. A unique identifier for each log record

Answer: D

Q619. Fuzzy checkpointing differs from standard checkpointing in that:

- A. No log entry is written to record the occurrence of the checkpoint
- B. The database does not need to stop processing during the checkpoint
- C. All processing must stop completely during the checkpoint operation
- D. It occurs only once during the entire lifetime of the database

Answer: B

Q620. Physiological logging means:

- A. Log records describe changes at the page level with logical operations within the page
- B. Only logical operation names are logged without any reference to the physical pages
- C. Only physical byte-level changes are logged for each modification to the data pages
- D. No logging occurs at all for any modification made to the database during processing

Answer: A

Q621. In a B+ tree of order n , an internal node can have at most:

- A. $n+1$ keys in the node
- B. n keys and n children
- C. n children and $n-1$ keys
- D. $2n$ children in total

Answer: C

Q622. The height of a B+ tree with n records and order p is approximately:

- A. $p \cdot \log(n)$
- B. $n \cdot p$
- C. $\log_p(n)$
- D. n / p

Answer: C

Q623. Static hashing has which disadvantage?

- A. Fixed number of buckets may lead to overflow as data grows
- B. It requires B+ trees to function as the base structure
- C. It is too slow for any practical database use case at all
- D. It cannot handle any data stored in the database tables

Answer: A

Q624. Extendible hashing handles growth by:

- A. Deleting old records to make room for new data entries
- B. Using a B+ tree instead of hashing for the index data
- C. Increasing the number of hash functions used for keys
- D. Doubling the directory and splitting buckets as needed

Answer: D

Q625. Linear hashing handles growth by:

- A. Using a fixed number of buckets without any changes
- B. Splitting all buckets simultaneously at the same time
- C. Splitting buckets in a linear order, one at a time
- D. Never splitting any buckets regardless of the load

Answer: C

Q626. A composite index is an index on:

- A. A stored procedure reference
- B. Two or more columns combined
- C. A view of the base table data
- D. A single column of the table

Answer: B

Q627. A covering index is one that:

- A. Indexes all tables in the database regardless of whether they need indexing
- B. Has no columns included in it and serves no purpose for query optimization
- C. Covers only the primary key column and does not include any other attributes
- D. Contains all columns needed to satisfy a query without accessing the base table

Answer: D

Q628. A bitmap index is most efficient for:

- A. Primary key columns that uniquely identify each row
- B. Columns with unique values that never repeat at all
- C. Columns with high cardinality (many distinct values)
- D. Columns with low cardinality (few distinct values)

Answer: D

Q629. An index-organized table (IOT) stores:

- A. Data separately from all index structures created
- B. Data only in memory without disk-based persistence
- C. Data directly in the index structure (B+ tree)
- D. Data in flat files without any indexing structures

Answer: C

Q630. The term fan-out in a B+ tree refers to:

- A. The number of records stored in the entire data file
- B. The maximum number of pointers (children) in a node
- C. The number of leaf nodes at the bottom of the tree
- D. The total height of the tree from root to leaf level

Answer: B

Q631. Sort-Merge Join works by:

- A. Using a full table scan on both relations without any index use
- B. Sorting both relations on the join attribute and then merging them
- C. Creating a hash table on one relation and probing with the other
- D. Using nested loops to compare every tuple in both the relations

Answer: B

Q632. Hash Join works by:

- A. Nested iteration comparing every pair of tuples from relations
- B. Sorting both relations on the join attribute and merging results
- C. Using indexes on both tables to find matching join attributes
- D. Building a hash table on one relation and probing with the other

Answer: D

Q633. An index nested loop join uses:

- A. No index at all on either of the two relations being joined here
- B. An index on the inner relation to find matching tuples efficiently
- C. An index on the outer relation only for scanning its tuple values
- D. A full scan of both relations without using any index structures

Answer: B

Q634. Pushing selections down in a query tree:

- A. Increases the size of intermediate results
- B. Removes selections entirely from the plan
- C. Has no effect on the intermediate results
- D. Reduces the size of intermediate results

Answer: D

Q635. Materialization in query processing means:

- A. Not storing any intermediate results and processing everything in one single pass
- B. Using only indexes to process the query without accessing the base table data
- C. Computing and storing the result of each operation before passing it to the next
- D. Only computing the final result without evaluating any intermediate operations

Answer: C

Q636. Pipelining in query processing means:

- A. Processing operations one at a time in a strict sequential fashion
- B. Storing all intermediate results on disk before the next operation
- C. Passing tuples from one operation to the next as they are produced
- D. Using only a single CPU core for all database query computations

Answer: C

Q637. The catalog statistics used by a cost-based optimizer include:

- A. Only the table name without any other statistical information
- B. Only the primary key definition of each table in the catalog
- C. Number of tuples, number of distinct values, number of blocks
- D. Only column names without cardinality or distribution details

Answer: C

Q638. Heuristic optimization rules include:

- A. Perform joins before selections and projections always
- B. Always use Cartesian products first before other joins
- C. Perform selections and projections as early as possible
- D. Avoid using indexes for any data retrieval operations

Answer: C

Q639. The join ordering problem in query optimization refers to:

- A. Ordering rows within a single table by a specific column
- B. Ordering columns in a table based on their data type size
- C. Determining the most efficient order to join multiple tables
- D. Sorting query results in ascending or descending sequence

Answer: C

Q640. Dynamic programming is used in query optimization for:

- A. Creating tables and defining column data types
- B. Parsing SQL queries into abstract syntax trees
- C. Managing locks in the concurrency control layer
- D. Finding the optimal join order for multiple tables

Answer: D

Q641. Mandatory Access Control (MAC) is based on:

- A. Security labels and clearance levels assigned to subjects and objects
- B. User discretion in deciding who can access which database objects
- C. No access control at all with every user having full open access
- D. Role assignments only without any additional security label checks

Answer: A

Q642. Discretionary Access Control (DAC) allows:

- A. Data owners to grant and revoke access to other users
- B. No access control with all data open to every user
- C. Access based on security clearance only in the system
- D. Only the DBA to control access to all database objects

Answer: A

Q643. Role-Based Access Control (RBAC) assigns privileges to:

- A. Individual users only without grouping
- B. Queries only without role management
- C. Tables only without any user control
- D. Roles, which are then assigned to users

Answer: D

Q644. The Bell-LaPadula model enforces:

- A. Only read access controls without write restrictions
- B. No read-up and no write-down rules for confidentiality
- C. No rules at all for controlling information access
- D. Only write access controls without read restrictions

Answer: B

Q645. Data masking in database security:

- A. Encrypts entire databases for secure data storage
- B. Deletes all data from the tables in the database
- C. Replaces sensitive data with realistic but fake data
- D. Creates copies of data for backup and recovery use

Answer: C

Q646. A view can provide security by:

- A. Deleting sensitive records from the database tables
- B. Encrypting all data stored in the underlying tables
- C. Creating new tables with restricted column schemas
- D. Restricting which columns and rows a user can see

Answer: D

Q647. Virtual Private Database (VPD) provides:

- A. Only audit logging without any enforcement of access policies
- B. Column-level encryption only without any row-level restrictions
- C. No security features of any kind for the database application
- D. Row-level security by automatically adding conditions to queries

Answer: D

Q648. Transparent Data Encryption (TDE) encrypts:

- A. Query results only sent to clients
- B. Log files only on the disk storage
- C. Data in transit only on the network
- D. Data at rest on the storage media

Answer: D

Q649. Database activity monitoring (DAM) is used to:

- A. Optimize storage allocation across available disk partitions
- B. Speed up queries by caching frequently accessed data in memory
- C. Create backups of the database on a scheduled periodic basis
- D. Monitor and record database activity in real-time for security

Answer: D

Q650. The GRANT OPTION in SQL allows:

- A. Creating tables and views in the database schema at will
- B. Revoking all privileges from every user in the database
- C. Deleting databases and all their data from the server
- D. A grantee to further grant the same privilege to other users

Answer: D

Q651. The two-phase commit (2PC) protocol ensures:

- A. Only read operations complete
- B. No atomicity guarantees at all
- C. Only local transaction atomicity
- D. Atomicity of distributed transactions

Answer: D

Q652. In the 2PC protocol, the coordinator first sends:

- A. A prepare (vote) message to all participants
- B. A commit message directly to all the sites
- C. An abort message to cancel the transaction
- D. A query to all participants for data first

Answer: A

Q653. If any participant votes No in 2PC, the coordinator:

- A. Ignores the vote and proceeds with commit
- B. Sends an abort message to all participants
- C. Retries the transaction automatically again
- D. Sends a commit message to all participants

Answer: B

Q654. A distributed query involves:

- A. No data access from any site at all
- B. Only local data access on one site
- C. Accessing data from a single table
- D. Accessing data from multiple sites

Answer: D

Q655. The CAP theorem states that a distributed system can guarantee at most:

- A. Two of three: Consistency, Availability, Partition Tolerance
- B. None of these properties can be guaranteed in any system
- C. Only one property out of the three possible in the theorem
- D. All three properties at the same time without any trade-offs

Answer: A

Q656. Semi-join optimization in distributed queries:

- A. Increases data transfer by sending all records across the link
- B. Uses full Cartesian products between all fragments in the join
- C. Reduces data transfer by sending only relevant join attributes
- D. Eliminates all joins from the distributed query execution plan

Answer: C

Q657. Mixed fragmentation combines:

- A. No fragmentation is used at all
- B. Only vertical fragmentation types
- C. Horizontal and vertical fragmentation
- D. Only horizontal fragmentation types

Answer: C

Q658. Full replication means:

- A. Every site has a complete copy of the entire database
- B. No replication exists at any site in the system at all
- C. Only one site has the data and others have no copies
- D. Data is partially replicated at only a few of sites

Answer: A

Q659. The blocking problem in 2PC occurs when:

- A. No failures occur during the execution of the two-phase commit protocol
- B. The coordinator fails after sending prepare but before sending the decision
- C. All participants commit successfully without any failures in the protocol
- D. The network is fast and reliable with no message delays or lost packets

Answer: B

Q660. Distributed deadlock detection is more complex because:

- A. All transactions are local and never span multiple sites
- B. Deadlocks can involve transactions at multiple sites
- C. Deadlocks never occur in distributed database systems
- D. Only one site has transactions running at any time

Answer: B

Q661. The BASE properties in NoSQL stand for:

- A. Batch Analysis, Stable Environment for data processing
- B. Binary Access, Sequential Evaluation of stored records
- C. Basic Access, Secure Entry with authentication support
- D. Basically Available, Soft state, Eventually consistent

Answer: D

Q662. NoSQL databases typically favor:

- A. Strong consistency over availability in all distributed environments
- B. Availability and partition tolerance over strong consistency (AP in CAP)
- C. Neither availability nor consistency for distributed data systems
- D. Only consistency without any consideration for availability at all

Answer: B

Q663. Sharding in NoSQL databases refers to:

- A. Compressing data to reduce storage on each server node
- B. Horizontally partitioning data across multiple servers
- C. Vertically splitting a table into separate column groups
- D. Encrypting data before storing it on distributed servers

Answer: B

Q664. MongoDB stores data in which format?

- A. CSV (Comma Separated)
- B. Plain text (Flat file)
- C. XML (Markup Language)
- D. BSON (Binary JSON)

Answer: D

Q665. In MongoDB, a collection is analogous to:

- A. A column in a relational database
- B. A database in a relational system
- C. A table in a relational database
- D. A row in a relational database

Answer: C

Q666. MapReduce in NoSQL is used for:

- A. Managing transactions and ensuring ACID properties on the data
- B. Creating indexes on columns for improving data retrieval speed
- C. Single-record queries for retrieving individual data items fast
- D. Processing large datasets in parallel across distributed systems

Answer: D

Q667. Which NoSQL database type is best for social network relationship queries?

- A. Key-value store
- B. Document database
- C. Column-family store
- D. Graph database

Answer: D

Q668. Denormalization in NoSQL is common because:

- A. It reduces storage requirements on each server in cluster
- B. It enforces referential integrity between related records
- C. It avoids expensive joins and improves read performance
- D. It simplifies updates by reducing the number of writes

Answer: C

Q669. Apache HBase is built on top of:

- A. MongoDB
- B. Redis
- C. Hadoop HDFS
- D. MySQL

Answer: C

Q670. Eventual consistency in NoSQL means:

- A. No consistency is guaranteed between replicas at any point in time whatsoever
- B. Data is never consistent across replicas regardless of the time that passes
- C. All reads always return the latest written value from any replica in the cluster
- D. All replicas will converge to the same state given enough time without new updates

Answer: D

Q671. The snowflake schema extends the star schema by:

- A. Removing dimension tables from the schema completely
- B. Adding more fact tables to handle additional measures
- C. Normalizing dimension tables into sub-dimension tables
- D. Denormalizing all tables into one single flat structure

Answer: C

Q672. Roll-up in OLAP means:

- A. Drilling down to see more detailed level of data
- B. Selecting a specific dimension value for filtering
- C. Aggregating data to a higher level of summarization
- D. Rotating the data cube axes for a different view

Answer: C

Q673. Drill-down in OLAP means:

- A. Rotating the data cube to view a different perspective
- B. Navigating from summarized data to more detailed data
- C. Aggregating data to higher levels of summarization
- D. Removing dimensions from the cube for simplification

Answer: B

Q674. Slice operation in OLAP:

- A. Selects a single value for one dimension, creating a sub-cube
- B. Aggregates all dimensions into one single summary result
- C. Selects multiple values for multiple dimensions simultaneously
- D. Rotates the cube to provide an alternative data perspective

Answer: A

Q675. Dice operation in OLAP:

- A. Selects a single dimension value for data filtering
- B. Rotates the cube axes for an alternative data view
- C. Selects specific values for two or more dimensions
- D. Aggregates data across all available dimension types

Answer: C

Q676. Pivot (rotate) in OLAP:

- A. Deletes a dimension from the cube structure entirely
- B. Rotates the data axes to provide an alternative view
- C. Merges cubes from different source data warehouses
- D. Adds new data records into the existing data cube

Answer: B

Q677. MOLAP (Multidimensional OLAP) stores data in:

- A. A multidimensional array (cube) structure
- B. Key-value stores for fast data retrieval
- C. Flat files stored on the file system disk
- D. Relational tables with rows and columns

Answer: A

Q678. ROLAP (Relational OLAP) stores data in:

- A. Relational database tables
- B. Flat files on disk storage
- C. NoSQL database documents
- D. Multidimensional arrays

Answer: A

Q679. Association rule mining discovers:

- A. Database schemas and their normalization structure in the system
- B. Index structures and their configuration for faster data access
- C. Relationships between items (e.g., items frequently bought together)
- D. SQL queries and their execution plans for optimized performance

Answer: C

Q680. The Apriori algorithm is used in data mining for:

- A. Finding frequent itemsets and association rules
- B. Classification of records into known classes
- C. Regression analysis on continuous value data
- D. Clustering data into groups of similar items

Answer: A

Q681. What distinguishes logical data independence from physical data independence?

- A. Logical is used in NoSQL while physical is used in relational databases only
- B. Logical shields from schema changes while physical shields from storage changes
- C. Logical handles storage changes while physical handles schema modifications
- D. Logical applies to queries while physical applies to transactions exclusively

Answer: B

Q682. Which mapping exists between the external and conceptual schema levels?

- A. External-internal mapping
- B. External-conceptual mapping
- C. Schema-instance mapping
- D. Conceptual-internal mapping

Answer: B

Q683. In the three-schema architecture, the conceptual level describes:

- A. Individual user views of specific data subsets
- B. How data is physically stored on disk devices
- C. The overall logical structure of the database
- D. The network protocols used for data transfer

Answer: C

Q684. Which type of database language handles data access permissions?

- A. Data Control Language handles permissions
- B. Data Definition Language handles permissions
- C. Transaction Control Language handles them
- D. Data Manipulation Language handles them

Answer: A

Q685. What is the difference between a database schema and a database state?

- A. Schema is the data values while state is the structure definition
- B. Schema applies to queries while state applies to transactions only
- C. Schema is the structure definition while state is the current data
- D. Schema is temporary storage while state is permanent data storage

Answer: C

Q686. Which of the following best describes program-data independence?

- A. Program logic is separated from the data structure definitions
- B. Programs and data are stored on the same physical disk drives
- C. Data files are independent of the operating system platform used
- D. Programs automatically generate data without user intervention

Answer: A

Q687. What is the self-describing nature of a database system?

- A. The system generates documentation without any user involvement
- B. The system repairs itself when errors are automatically detected
- C. The system contains both data and its own metadata description
- D. The system creates backups at regular predetermined time intervals

Answer: C

Q688. Which database model was the first to be widely commercially adopted?

- A. Hierarchical data model
- B. Object-oriented model
- C. Network data model
- D. Relational data model

Answer: A

Q689. What is the purpose of a database catalog?

- A. To compress data for efficient disk storage
- B. To maintain metadata about database objects
- C. To store actual user data records in tables
- D. To provide backup copies of all user data

Answer: B

Q690. Which characteristic differentiates a DBMS from a traditional file system?

- A. DBMS cannot handle concurrent access but file systems can do it
- B. DBMS supports data abstraction while file systems do not offer it
- C. DBMS uses flat files while file systems use structured tables
- D. DBMS requires less memory than file systems for similar data sets

Answer: B

Q691. In a three-tier architecture, what role does the application server play?

- A. It serves as intermediary between client and database
- B. It only handles user interface rendering on screen
- C. It manages the operating system kernel processes
- D. It directly stores data in physical database files

Answer: A

Q692. What advantage does a parallel database architecture provide?

- A. Simplified system administration with fewer components
- B. Improved query performance through concurrent processing
- C. Reduced hardware costs compared to single server systems
- D. Elimination of the need for any backup or recovery plan

Answer: B

Q693. Which component of DBMS architecture handles authorization checks?

- A. File manager component
- B. Buffer manager component
- C. Disk space manager component
- D. Authorization manager component

Answer: D

Q694. What is the shared-nothing architecture in parallel databases?

- A. All processors share a common memory and disk pool
- B. Each processor has its own private memory and disk
- C. All processors share memory but have private disk units
- D. Each processor has private memory but shares all disks

Answer: B

Q695. How does connection pooling improve database performance?

- A. By encrypting all data connections between nodes
- B. By compressing data packets sent over the network link
- C. By duplicating queries across multiple server instances
- D. By reusing existing connections instead of creating new

Answer: D

Q696. What distinguishes embedded database systems from server-based ones?

- A. Embedded systems only work with object-oriented languages
- B. Embedded systems require dedicated server hardware always
- C. Embedded systems support more concurrent users overall
- D. Embedded systems run within the application process itself

Answer: D

Q697. In DBMS architecture, what does the disk manager handle?

- A. Allocating and managing space on storage devices
- B. Rendering graphical interfaces for database users
- C. Parsing and optimizing SQL queries for execution
- D. Managing user sessions and authentication tokens

Answer: A

Q698. What is the benefit of a shared-disk architecture in clusters?

- A. Data is automatically replicated to every single node
- B. Any node can access any data on the shared storage
- C. Each node operates with completely independent data
- D. No network communication is needed between the nodes

Answer: B

Q699. Which architectural approach best supports database scalability?

- A. Single monolithic server with maximum hardware resources
- B. Standalone desktop database with local file storage only
- C. Distributed architecture with horizontal partitioning used
- D. Centralized file-based system with sequential data access

Answer: C

Q700. What is the purpose of the catalog manager in DBMS?

- A. To maintain and provide access to system metadata
- B. To manage physical disk read and write operations
- C. To encrypt data before transmission over networks
- D. To handle user interface rendering for the client

Answer: A

Q701. What is a weak entity in the ER model?

- A. An entity that cannot participate in any relationships
- B. An entity that depends on another entity for identity
- C. An entity that contains only a single attribute value
- D. An entity that has its own unique primary key defined

Answer: B

Q702. How does the network data model differ from the hierarchical model?

- A. Network model has no relationships while hierarchical has them
- B. Network model only supports one-to-one cardinality mappings
- C. Network model allows many-to-many while hierarchical does not
- D. Network model uses tables while hierarchical uses trees only

Answer: C

Q703. What is generalization in the Enhanced ER model?

- A. Removing attributes from entities to simplify the schema
- B. Combining entities into a higher-level general entity type
- C. Breaking a general entity into specific subtype entities
- D. Creating relationships between two unrelated entity types

Answer: B

Q704. What is the difference between total and partial participation?

- A. Total is for strong entities while partial is for weak ones
- B. Total applies to attributes while partial applies to entities
- C. Total means some entities participate while partial means all
- D. Total means all entities participate while partial means some

Answer: D

Q705. In the ER model, what is a multivalued attribute?

- A. An attribute that is shared between multiple entities
- B. An attribute that cannot store any null values at all
- C. An attribute with exactly one value for each entity
- D. An attribute with multiple values for a single entity

Answer: D

Q706. What is specialization in the Enhanced ER model?

- A. Removing relationships between entities to reduce complexity
- B. Merging multiple entities into one general supertype entity
- C. Converting attributes into separate independent entity types
- D. Defining subtypes of an entity based on distinguishing traits

Answer: D

Q707. What does a ternary relationship involve in the ER model?

- A. A relationship connecting exactly two entity types together
- B. A relationship with three attributes and no entity types
- C. A relationship connecting an entity type to itself only
- D. A relationship connecting exactly three entity types at once

Answer: D

Q708. What is the object-relational data model?

- A. A model designed exclusively for unstructured text data
- B. A model using only hierarchical tree data structures
- C. A model combining relational and object-oriented features
- D. A model that eliminates the concept of tables entirely

Answer: C

Q709. What is a derived attribute in the ER model?

- A. An attribute that uniquely identifies each entity instance
- B. An attribute whose value is entered manually by users
- C. An attribute that connects two entities in a relationship
- D. An attribute whose value is computed from other attributes

Answer: D

Q710. What does the concept of aggregation address in ER modeling?

- A. Combining multiple attributes into a single composite one
- B. Treating a relationship set as a higher-level entity set
- C. Splitting an entity into multiple independent sub-entities
- D. Removing redundant relationships from the ER diagram view

Answer: B

Q711. What is the difference between a superkey and a candidate key?

- A. A candidate key is a minimal superkey without extra attributes
- B. A superkey cannot contain null values but candidate keys can
- C. A candidate key allows duplicate values but superkeys do not
- D. A superkey is always a single attribute unlike candidate key

Answer: A

Q712. What does entity integrity constraint specifically require?

- A. Primary key attributes cannot contain any null values
- B. Every table must have at least one foreign key defined
- C. All attributes in a table must have default values set
- D. Foreign keys must reference existing primary key values

Answer: A

Q713. What is referential integrity in relational databases?

- A. Each table must have at least two candidate key options
- B. Foreign key values must match existing primary key values
- C. All attributes must have the same domain type definition
- D. Each row must have a unique value in every single column

Answer: B

Q714. What is the cardinality of a relation?

- A. The number of candidate keys defined in the table
- B. The number of tuples or rows currently in the table
- C. The number of attributes or columns in the table
- D. The number of foreign key constraints on the table

Answer: B

Q715. What is a composite key in relational databases?

- A. A key consisting of two or more attributes
- B. A key consisting of a single attribute only
- C. A key that allows null values in all parts
- D. A key that is automatically generated by DBMS

Answer: A

Q716. What is the difference between natural and surrogate keys?

- A. Natural keys are system-generated while surrogate keys are not
- B. Natural keys come from data while surrogate keys are generated
- C. Natural keys are always composite while surrogate ones are not
- D. Natural keys must be numeric while surrogate keys are textual

Answer: B

Q717. What does the closure of a set of functional dependencies represent?

- A. The complete set of all functional dependencies that hold
- B. The total count of candidate keys in a given relation set
- C. The minimal number of attributes needed for a primary key
- D. The maximum number of tuples a relation can store safely

Answer: A

Q718. What is an alternate key in relational databases?

- A. Any attribute that is part of the primary key itself
- B. A foreign key that references a composite primary key
- C. A candidate key that is not chosen as primary key
- D. A key that only exists in temporary database tables

Answer: C

Q719. What does the relational model information principle state?

- A. All information is represented as values in relations
- B. Some data must be stored outside of relation tables
- C. Data can be stored in hierarchical tree structures
- D. Relations can contain duplicate tuples when necessary

Answer: A

Q720. What is a relation schema compared to a relation instance?

- A. Schema defines queries while instance defines storage format
- B. Schema is the data while instance is the structure definition
- C. Schema and instance are always identical in all situations
- D. Schema is the structure while instance is the actual data set

Answer: D

Q721. How does a left outer join differ from a natural join?

- A. Left outer join requires explicit join conditions unlike natural
- B. Left outer join only works on tables with identical schemas
- C. Left outer join excludes all unmatched tuples entirely
- D. Left outer join preserves all tuples from the left relation

Answer: D

Q722. What is the division operation used for in relational algebra?

- A. Splitting a relation into two halves by row count
- B. Calculating arithmetic division on numeric column values
- C. Removing duplicate tuples from a single relation quickly
- D. Finding tuples associated with all values in another set

Answer: D

Q723. What is the equivalence rule for pushing selection through a join?

- A. Selections on attributes of one relation can precede the join
- B. Selections can never be applied before join operations
- C. Selections must always follow projection operations in order
- D. Selections and projections are interchangeable in all cases

Answer: A

Q724. How does semijoin differ from natural join?

- A. Semijoin returns only attributes from one input relation
- B. Semijoin returns attributes from both input relations used
- C. Semijoin does not require any common attributes to operate
- D. Semijoin always produces more tuples than natural join does

Answer: A

Q725. What does the extended projection operation allow?

- A. Arithmetic expressions and renaming in the output columns
- B. Filtering rows based on aggregate function result values
- C. Joining multiple relations without specifying join conditions
- D. Only selecting existing columns from the base relation

Answer: A

Q726. In relational algebra, which operation is NOT a fundamental operation?

- A. Union operation
- B. Selection operation
- C. Intersection operation
- D. Cartesian product

Answer: C

Q727. What is the antijoin operation in relational algebra?

- A. Tuples from left relation that have no match in right
- B. Tuples from right relation that have no match in left
- C. Tuples from both relations that match on join condition
- D. All possible combinations of tuples from both relations

Answer: A

Q728. How is the generalized projection different from standard projection?

- A. It can only select a single column from the relation
- B. It removes the requirement for union compatibility checks
- C. It allows arithmetic operations on attributes in output
- D. It automatically eliminates null values from the results

Answer: C

Q729. What does the grouping operation (gamma) do in extended relational algebra?

- A. It renames the relation and all of its attribute names used
- B. It sorts tuples by a specified attribute in ascending order
- C. It groups tuples and applies aggregate functions to groups
- D. It removes all duplicate tuples from the relation entirely

Answer: C

Q730. What is a full outer join in relational algebra?

- A. A join that only returns tuples from the right side table
- B. A join preserving unmatched tuples from both left and right
- C. A join that produces the Cartesian product without any filter
- D. A join preserving only matching tuples from both relations

Answer: B

Q731. What is the difference between WHERE and HAVING clauses?

- A. WHERE works only with joins while HAVING works with all queries
- B. WHERE is used in DDL while HAVING is used in DML statements
- C. WHERE filters individual rows while HAVING filters groups
- D. WHERE filters after grouping while HAVING filters before it

Answer: C

Q732. Which type of join returns rows only when there is a match in both tables?

- A. Inner join
- B. Right outer join
- C. Left outer join
- D. Full outer join

Answer: A

Q733. What does the GROUP BY clause do in SQL?

- A. Sorts results in descending order by default
- B. Removes duplicate values from a single column only
- C. Limits the number of rows in the result output
- D. Groups rows sharing common values for aggregation

Answer: D

Q734. What is a subquery in SQL?

- A. A query nested inside another SQL query statement
- B. A query that creates a temporary table permanently
- C. A query that runs automatically on a fixed schedule
- D. A query that only works with a single table at once

Answer: A

Q735. What is the purpose of the LIKE operator in SQL?

- A. To perform exact matching on numeric columns only
- B. To perform pattern matching on string column values
- C. To compare two tables for identical structure setup
- D. To join tables based on a common column attribute

Answer: B

Q736. What does the BETWEEN operator do in SQL?

- A. Joins two tables on a specified column attribute
- B. Checks if a value exists in a list of values
- C. Filters values within a specified inclusive range
- D. Sorts results between ascending and descending order

Answer: C

Q737. What is the difference between DELETE and TRUNCATE?

- A. DELETE can use WHERE clause while TRUNCATE removes all rows
- B. DELETE is DDL while TRUNCATE is DML in SQL classification
- C. DELETE is faster than TRUNCATE for large table operations
- D. DELETE removes the table while TRUNCATE only removes rows

Answer: A

Q738. Which aggregate function counts non-null values in a column?

- A. MAX function
- B. AVG function
- C. COUNT function
- D. SUM function

Answer: C

Q739. What does the IN operator do in SQL?

- A. Checks if a value falls within a numeric range
- B. Verifies if a column has null values present
- C. Performs pattern matching on string data values
- D. Checks if a value matches any value in a list

Answer: D

Q740. What is a self-join in SQL?

- A. A join between two completely different database tables
- B. A join that requires no ON clause or condition specified
- C. A join where a table is joined with itself using aliases
- D. A join that automatically matches all columns in tables

Answer: C

Q741. What is the difference between a stored procedure and a function in SQL?

- A. Functions can modify data while procedures cannot modify it
- B. Procedures cannot take parameters while functions always can
- C. Procedures can return values while functions cannot do so
- D. Functions must return a value while procedures need not do

Answer: D

Q742. What is a materialized view and how does it differ from a regular view?

- A. Materialized views store data physically while regular do not
- B. Materialized views cannot be refreshed while regular ones can
- C. Materialized views use less storage than regular views always
- D. Materialized views are always up to date while regular are not

Answer: A

Q743. What are the different types of triggers based on timing?

- A. FAST, MEDIUM, and SLOW trigger performance speed types
- B. BEFORE, AFTER, and INSTEAD OF trigger execution types
- C. START, MIDDLE, and END trigger position location types
- D. FIRST, SECOND, and THIRD trigger priority level types

Answer: B

Q744. What is dynamic SQL and when is it used?

- A. SQL that only works with temporary tables created at runtime
- B. SQL statements constructed and executed at runtime dynamically
- C. SQL that automatically optimizes queries without any user input
- D. SQL embedded directly in application source code at compile

Answer: B

Q745. What is the purpose of a savepoint in SQL transactions?

- A. To lock all tables preventing concurrent access by others
- B. To create a backup of the entire database at that moment
- C. To permanently commit all changes up to that specific point
- D. To mark a point within a transaction for partial rollback

Answer: D

Q746. What is the difference between clustered and non-clustered indexes?

- A. Clustered indexes use more disk space than non-clustered type
- B. Clustered indexes sort table data while non-clustered do not
- C. Clustered indexes are slower than non-clustered for all uses
- D. Clustered indexes cannot be created on primary key columns

Answer: B

Q747. What is a Common Table Expression (CTE) used for?

- A. Creating database users with specific permission levels set
- B. Encrypting sensitive data before storing it in the database
- C. Permanently storing query results in a physical table file
- D. Defining temporary named result sets within a single query

Answer: D

Q748. What does the EXPLAIN command do in SQL?

- A. It describes the structure of a database table in detail
- B. It shows the execution plan for a given SQL query statement
- C. It documents the purpose of stored procedures and functions
- D. It provides help text for all available SQL commands listed

Answer: B

Q749. What is a window function in SQL?

- A. A function that manages database connection windows open
- B. A function that creates graphical windows in applications
- C. A function that performs calculations across row sets
- D. A function that displays query results in separate panels

Answer: C

Q750. What is the purpose of the CASCADE option in DROP TABLE?

- A. It drops the table and all dependent objects together
- B. It only drops the table if it contains no data rows
- C. It converts the table to a view before deletion occurs
- D. It creates a backup before dropping the specified table

Answer: A

Q751. What is a partial dependency in database normalization?

- A. A key attribute depends on a non-key attribute in table
- B. A non-key attribute depends on another non-key attribute
- C. A non-key attribute depends on the entire composite key
- D. A non-key attribute depends on part of a composite key

Answer: D

Q752. What additional requirement does 3NF impose over 2NF?

- A. No transitive dependencies on the primary key column
- B. All tables must have at least two candidate key options
- C. Every attribute must be part of some candidate key set
- D. All attributes must be of the same domain type value

Answer: A

Q753. What is a transitive dependency?

- A. Attribute A depends on a foreign key in another table
- B. Attribute A depends on B which depends on primary key
- C. Attribute A and B both depend on the same primary key
- D. Attribute A depends directly on the primary key only

Answer: B

Q754. How does BCNF differ from 3NF?

- A. BCNF allows transitive dependencies that 3NF does not
- B. BCNF is less restrictive than 3NF for all relations
- C. BCNF requires every determinant to be a candidate key
- D. BCNF only applies to tables with a single candidate key

Answer: C

Q755. What is a multivalued dependency?

- A. An attribute depends on the concatenation of all other ones
- B. One attribute value determines a set of independent values
- C. Two attributes are always equal in every tuple of table
- D. One attribute value determines exactly one other value

Answer: B

Q756. What does the fourth normal form (4NF) address?

- A. Transitive dependencies among non-key attributes used
- B. Join dependencies that cannot be reduced any further
- C. Partial dependencies on the composite primary key
- D. Multivalued dependencies that cause data redundancy

Answer: D

Q757. What is lossless join decomposition?

- A. A decomposition that always results in fewer tables total
- B. A decomposition where some data may be lost permanently
- C. A decomposition where original data can be fully recovered
- D. A decomposition that eliminates all foreign keys from tables

Answer: C

Q758. What is dependency preservation in normalization?

- A. All functional dependencies can be checked within tables
- B. Dependencies are only preserved in the original table form
- C. Dependencies are stored in a separate metadata table file
- D. Some functional dependencies may be lost during the split

Answer: A

Q759. When is denormalization justified in database design?

- A. When read performance is prioritized over write efficiency
- B. When all tables are already in first normal form status
- C. When the database has too few tables for organization
- D. When the database has no functional dependencies at all

Answer: A

Q760. What is the difference between lossless and lossy decomposition?

- A. Lossless recovers original data while lossy creates spurious tuples
- B. Lossless is always faster while lossy requires more processing
- C. Lossless uses more storage while lossy uses compression methods
- D. Lossless adds extra data while lossy preserves original data

Answer: A

Q761. What is a serializable schedule?

- A. A schedule that does not require any concurrency control
- B. A schedule equivalent in result to some serial schedule
- C. A schedule where transactions run in strict serial order
- D. A schedule that allows all types of data conflicts always

Answer: B

Q762. What is the difference between conflict serializability and view serializability?

- A. View serializability allows more schedules than conflict type
- B. Conflict serializability is less restrictive than view type
- C. View serializability requires all transactions to be serial
- D. Conflict serializability only applies to read operations here

Answer: A

Q763. What is a dirty read in transaction processing?

- A. Reading data from a backup file instead of the live database
- B. Reading data written by an uncommitted transaction directly
- C. Reading data from a table that has no index defined on it
- D. Reading data that has been committed by another transaction

Answer: B

Q764. What is a non-repeatable read anomaly?

- A. A transaction reads the same row twice and gets same results
- B. A transaction cannot read any rows because they are all locked
- C. A transaction reads the same row twice and gets different results
- D. A transaction reads rows from a table that does not exist yet

Answer: C

Q765. What is a cascading rollback in transaction management?

- A. Automatically committing all pending transactions at the same time
- B. Rolling back a single transaction without affecting any others
- C. Restarting the database server to clear all active transactions
- D. One transaction rollback causing other transactions to rollback

Answer: D

Q766. What is the write-ahead logging (WAL) protocol?

- A. Neither log records nor data pages need to be written to disk
- B. Log records are written to disk before data pages are flushed
- C. Both log records and data pages are written simultaneously
- D. Data pages are written to disk before log records are made

Answer: B

Q767. What is a recoverable schedule?

- A. A schedule that automatically recovers from all hardware failures
- B. A schedule where no transaction ever needs to be rolled back
- C. A schedule where transactions commit only after dependencies commit
- D. A schedule where all transactions execute in strict serial order

Answer: C

Q768. What are the different transaction isolation levels in SQL?

- A. Read Uncommitted, Read Committed, Repeatable Read, Serializable
- B. Private, Shared, Protected, and Public access levels for data
- C. Fast, Normal, Slow, and Batch processing modes for transactions
- D. Low, Medium, High, and Maximum security levels for transactions

Answer: A

Q769. What is a checkpoint in database recovery?

- A. A point where all current data is deleted from the database
- B. A point where the database schema is backed up to a file
- C. A point where database state is synchronized with the log
- D. A point where new transactions are prevented from starting

Answer: C

Q770. What is a cascadeless schedule?

- A. A schedule allowing transactions to read uncommitted data freely
- B. A schedule where transactions only read data from committed ones
- C. A schedule where all transactions execute without any commits
- D. A schedule that permits cascading rollbacks to occur frequently

Answer: B

Q771. What is the two-phase locking protocol?

- A. Transactions use only shared locks in the first phase entirely
- B. Transactions acquire all locks first then release all at end
- C. Transactions can acquire and release locks in any order freely
- D. Transactions have a growing phase and a shrinking phase

Answer: D

Q772. How does the wait-die scheme handle deadlock prevention?

- A. Younger transactions wait while older ones are rolled back
- B. Older transactions wait while younger ones are rolled back
- C. Transactions are randomly selected for rollback on each conflict
- D. All transactions are rolled back when a conflict is detected

Answer: B

Q773. How does the wound-wait scheme handle deadlock prevention?

- A. Older transactions are rolled back while younger ones wait
- B. Older transactions preempt younger ones while younger wait
- C. Both transactions are always rolled back on any conflict
- D. Transactions are processed strictly in first-come first-served

Answer: B

Q774. What is timestamp-based concurrency control?

- A. Scheduling transactions alphabetically by their name values
- B. Using random numbers to determine transaction execution order
- C. Assigning locks based on the size of each transaction query
- D. Ordering transactions by timestamps to ensure serializability

Answer: D

Q775. What is optimistic concurrency control?

- A. Transactions acquire locks before accessing any data items
- B. Transactions are executed one at a time in serial order only
- C. Transactions always use exclusive locks on all data items used
- D. Transactions proceed without locks and validate before commit

Answer: D

Q776. What is a lock escalation?

- A. Converting an exclusive lock to a shared lock type
- B. Releasing all locks held by a transaction at once
- C. Converting fine-grained locks to coarser-grained locks
- D. Transferring locks from one transaction to another one

Answer: C

Q777. What is the difference between pessimistic and optimistic concurrency control?

- A. Pessimistic uses timestamps while optimistic uses lock-based
- B. Pessimistic locks data preemptively while optimistic validates at commit
- C. Pessimistic allows dirty reads while optimistic prevents all of them
- D. Pessimistic is always faster than optimistic concurrency approaches

Answer: B

Q778. What is intention locking in a multigranularity system?

- A. Locking individual rows before accessing a full table record
- B. Preventing any transaction from acquiring exclusive lock types
- C. Locking the entire database before any transaction can start
- D. Signaling intent to lock at a finer granularity level below

Answer: D

Q779. What is the phantom problem in concurrency control?

- A. A transaction reads data that was never written to database
- B. New rows appear in repeated queries due to concurrent inserts
- C. Data becomes corrupted due to simultaneous write operations
- D. A transaction loses its locks after a network disconnection

Answer: B

Q780. What is multiversion concurrency control (MVCC)?

- A. Maintaining a single version of data with exclusive locking
- B. Maintaining multiple versions of data for concurrent access
- C. Deleting old versions of data immediately after each update
- D. Compressing multiple data versions into a single stored copy

Answer: B

Q781. What is the difference between immediate and deferred database modification?

- A. Immediate writes to log first while deferred writes to disk first
- B. Immediate uses checkpoints while deferred uses log records for data
- C. Immediate only works for read operations while deferred handles writes
- D. Immediate writes changes to disk before commit while deferred waits

Answer: D

Q782. What is shadow paging in database recovery?

- A. Creating shadows of all user queries for later replay at anytime
- B. Maintaining two page tables: current and shadow for safe recovery
- C. Encrypting page data to prevent unauthorized access to records
- D. Duplicating the entire database to a secondary server for backup

Answer: B

Q783. What is the steal/no-steal policy in buffer management?

- A. Steal and no-steal both prevent uncommitted pages from being flushed
- B. No-steal allows uncommitted pages while steal prevents their flushing
- C. Steal allows uncommitted pages to be flushed while no-steal prevents
- D. Steal refers to memory allocation while no-steal refers to disk usage

Answer: C

Q784. What is the force/no-force policy in buffer management?

- A. Force uses redo logging while no-force uses undo logging exclusively
- B. Force writes all pages at checkpoint while no-force writes at commit
- C. Force prevents dirty pages from being flushed while no-force allows
- D. Force requires all pages written at commit while no-force does not

Answer: D

Q785. What is media recovery versus crash recovery?

- A. Media recovery restores from disk failure while crash handles system crash
- B. Media recovery handles software bugs while crash recovery handles disk
- C. Media recovery is faster than crash recovery in all failure situations
- D. Media recovery uses logs while crash recovery uses backups exclusively

Answer: A

Q786. What information does a log record typically contain?

- A. Only the transaction ID and the timestamp of each operation
- B. Only the names of tables that were accessed during the transaction
- C. Only the SQL query text that was executed by the user application
- D. Transaction ID, data item, old value, new value, and operation type

Answer: D

Q787. What is fuzzy checkpointing?

- A. A checkpoint that runs automatically without any administrator control
- B. A checkpoint that allows transactions to continue during the process
- C. A checkpoint that only saves metadata without any actual data pages
- D. A checkpoint that stops all transactions until it completes fully

Answer: B

Q788. What is the purpose of the log sequence number (LSN)?

- A. To assign priority levels to different transactions for scheduling
- B. To uniquely identify and order log records in the transaction log
- C. To measure the physical size of each log record stored on disk
- D. To count the total number of transactions in the database system

Answer: B

Q789. What is a warm backup versus a cold backup?

- A. Cold backup is taken while database is online while warm is offline
- B. Warm backup is taken while database is online while cold is offline
- C. Warm backup only copies logs while cold backup copies entire database
- D. Cold backup is faster than warm backup in all operational scenarios

Answer: B

Q790. What is point-in-time recovery?

- A. Recovering only the tables that were modified since last backup
- B. Recovering the database schema without any of the actual data
- C. Recovering the database to a specific moment before the failure
- D. Recovering the database to its most recent committed state only

Answer: C

Q791. What is the difference between a B-tree and a B+ tree?

- A. B+ tree requires more levels than B-tree for the same data volume
- B. B-tree supports range queries while B+ tree only supports exact match
- C. B+ tree stores data only in leaves while B-tree stores in all nodes
- D. B-tree stores data in leaves only while B+ tree stores in all nodes

Answer: C

Q792. What is a composite index?

- A. An index on a single column of a database table
- B. An index built on two or more columns together
- C. An index that is shared between multiple tables
- D. An index that uses both hash and tree structures

Answer: B

Q793. What is index selectivity and why does it matter?

- A. The time required to create the index on a large data table
- B. The size of the index file compared to the original data file
- C. The ratio of distinct values to total records affecting efficiency
- D. The number of tables that reference a particular index entry

Answer: C

Q794. What is extendible hashing?

- A. A static hash function that never changes bucket count
- B. A hashing technique that eliminates all collision handling
- C. A hash function that only works with string data values
- D. A dynamic hashing scheme that grows and shrinks as needed

Answer: D

Q795. What is a bitmap index best suited for?

- A. Columns with very high cardinality like unique identifiers
- B. Tables that are frequently updated with new data records
- C. Columns with low cardinality like gender or status fields
- D. Queries that require exact match lookups on primary keys

Answer: C

Q796. What is the advantage of a covering index?

- A. It works only with aggregate queries on numeric data types
- B. It automatically updates when the table schema is changed
- C. It satisfies a query entirely without accessing table data
- D. It requires more storage space than a regular index type

Answer: C

Q797. What is linear hashing?

- A. A dynamic hashing scheme that splits buckets incrementally
- B. A hashing scheme that uses a fixed number of buckets always
- C. A hash function that only produces sequential hash values
- D. A hashing technique that stores data in a linear linked list

Answer: A

Q798. What is the difference between static and dynamic hashing?

- A. Static is always faster than dynamic for all query types used
- B. Static uses trees while dynamic uses linear search methods
- C. Static supports range queries while dynamic only supports exact
- D. Static has fixed buckets while dynamic adjusts bucket count

Answer: D

Q799. What is a multi-level index?

- A. An index structure with multiple levels like a tree hierarchy
- B. An index on multiple columns of the same table structure
- C. An index shared across multiple databases on the network
- D. An index that indexes other indexes for circular referencing

Answer: A

Q800. What is the difference between ordered and unordered file organization?

- A. Ordered files use hashing while unordered files use tree indexes
- B. Ordered files cannot be indexed while unordered files always can
- C. Ordered files are stored in RAM while unordered are stored on disk
- D. Ordered files maintain sort order while unordered files do not

Answer: D

Q801. What is the difference between a nested loop join and a hash join?

- A. Hash join only works with equality conditions while nested loop not
- B. Hash join builds a hash table while nested loop iterates through rows
- C. Nested loop is always faster than hash join for all data sizes
- D. Nested loop requires sorted input while hash join does not need it

Answer: B

Q802. What is a merge join and when is it most efficient?

- A. Merge join sorts and merges inputs and is efficient when pre-sorted
- B. Merge join uses nested loops and is efficient for indexed columns
- C. Merge join hashes both inputs and is efficient for small tables
- D. Merge join uses random access and is efficient for all query types

Answer: A

Q803. What is query rewriting in optimization?

- A. Changing the query results to match expected output format
- B. Adding additional columns to the query output automatically
- C. Transforming a query into an equivalent but more efficient form
- D. Replacing all queries with stored procedure calls instead

Answer: C

Q804. What is selectivity estimation in query optimization?

- A. The process of selecting the best join algorithm for queries
- B. The process of selecting which indexes to create on tables
- C. Estimating the total storage required for query result output
- D. Estimating the fraction of rows that satisfy a given condition

Answer: D

Q805. What is the purpose of a histogram in query optimization?

- A. Tracking the number of queries executed over time periods
- B. Displaying query performance metrics to database operators
- C. Providing data distribution information for cost estimation
- D. Graphically showing the structure of the execution plan tree

Answer: C

Q806. What is pipelining in query execution?

- A. Passing results from one operation to the next immediately
- B. Executing all operations in parallel on separate processors
- C. Storing intermediate results to disk before continuing
- D. Repeating operations until the correct result is obtained

Answer: A

Q807. What is materialization in query execution?

- A. Executing all operations in a single pass without stopping
- B. Storing intermediate results temporarily before next step
- C. Deleting temporary tables after query execution completes
- D. Converting query results into a permanent database table

Answer: B

Q808. What is predicate pushdown in query optimization?

- A. Removing all filter conditions to speed up processing time
- B. Duplicating filter conditions across multiple plan operators
- C. Moving filter conditions closer to the data source in plan
- D. Moving filter conditions to the end of the execution plan

Answer: C

Q809. What does the catalog store that helps query optimization?

- A. The source code of all stored procedures and functions used
- B. The network configuration for distributed query processing
- C. Table sizes, column statistics, and index information data
- D. The actual data rows from all database tables and views

Answer: C

Q810. What is a left-deep join tree in query optimization?

- A. A tree where both inputs of each join are intermediate results
- B. A tree where the left input of each join is a base table
- C. A tree where the right input of each join is a base table
- D. A tree where all joins are performed simultaneously in parallel

Answer: C

Q811. What is the difference between discretionary and mandatory access control?

- A. Discretionary allows owner control while mandatory uses labels
- B. Discretionary is managed by system while mandatory by data owner
- C. Discretionary applies to tables while mandatory applies to rows
- D. Discretionary uses encryption while mandatory uses authentication

Answer: A

Q812. What is role-based access control (RBAC)?

- A. Encrypting data based on the role of the user accessing data
- B. Assigning permissions to roles that are then granted to users
- C. Auditing user actions based on their assigned role in system
- D. Assigning permissions directly to each individual user account

Answer: B

Q813. How does parameterized querying prevent SQL injection attacks?

- A. It blocks all user input from being used in SQL queries
- B. It validates user input using regular expressions only here
- C. It encrypts the entire SQL query before sending to server
- D. It separates SQL code from user input preventing injection

Answer: D

Q814. What is transparent data encryption (TDE)?

- A. Encrypting data using user-provided keys for each session
- B. Encrypting data at rest without requiring application changes
- C. Encrypting data only when it is transmitted over a network
- D. Encrypting only the database schema and not the actual data

Answer: B

Q815. What is the principle of least privilege in database security?

- A. Users should share a single account to simplify access control
- B. Users should have only minimum access needed for their tasks
- C. Users should have maximum access to complete tasks efficiently
- D. Users should have temporary access that expires after each day

Answer: B

Q816. What is database masking?

- A. Replacing sensitive data with realistic but fake values shown
- B. Removing all data from a database for security purposes only
- C. Encrypting all database files using symmetric key algorithms
- D. Compressing database tables to obscure the stored data values

Answer: A

Q817. What is a security label in mandatory access control?

- A. A classification assigned to data and users for access rules
- B. A type of encryption key used to protect data at rest only
- C. A log entry recording when data was last accessed by users
- D. A physical tag attached to the database server hardware

Answer: A

Q818. What is the purpose of database activity monitoring?

- A. Managing storage allocation across multiple database disk drives
- B. Optimizing query performance by tracking slow-running queries
- C. Tracking and analyzing database activities for security threats
- D. Scheduling automated backups during off-peak operational hours

Answer: C

Q819. What is column-level security in databases?

- A. Restricting access to specific columns within a table view
- B. Encrypting individual columns using different key algorithms
- C. Restricting access to entire database tables for all users
- D. Compressing specific columns to reduce storage size needed

Answer: A

Q820. What is a virtual private database (VPD)?

- A. A database that is only accessible through a VPN connection
- B. A feature that adds predicates to queries based on user policy
- C. A temporary database created for each user session separately
- D. A database running on a private physical server exclusively

Answer: B

Q821. What is the CAP theorem in distributed databases?

- A. Compression, Authentication, Privacy are the three key goals
- B. Concurrency, Atomicity, Persistence must all be maintained
- C. Consistency, Availability, Performance can all be guaranteed
- D. Consistency, Availability, Partition tolerance: pick two of three

Answer: D

Q822. What is the two-phase commit protocol used for?

- A. Replicating data between primary and secondary server sites
- B. Fragmenting tables into horizontal and vertical data parts
- C. Ensuring atomicity of transactions across distributed sites
- D. Optimizing queries across multiple distributed database sites

Answer: C

Q823. What is the difference between homogeneous and heterogeneous distributed databases?

- A. Homogeneous uses same DBMS at all sites while heterogeneous not
- B. Heterogeneous uses same DBMS at all sites while homogeneous not
- C. Heterogeneous has more data than homogeneous distributed database
- D. Homogeneous stores data at one site while heterogeneous multiple

Answer: A

Q824. What is eventual consistency in distributed systems?

- A. All nodes always have identical data at every single moment
- B. All nodes will converge to the same data state over time
- C. Only the primary node maintains the consistent data state
- D. Data is never consistent across any nodes in the system

Answer: B

Q825. What is a federated database system?

- A. A database system that only supports read-only query operations
- B. An integration of multiple autonomous database systems together
- C. A single database distributed across multiple data centers
- D. A backup system that replicates data to geographically distant

Answer: B

Q826. What is distributed deadlock and how is it detected?

- A. Deadlock involving transactions across multiple sites via graph
- B. Deadlock that only affects read-only transactions across sites
- C. Deadlock that automatically resolves without any intervention
- D. Deadlock that occurs only on a single site in the system

Answer: A

Q827. What is the difference between synchronous and asynchronous replication?

- A. Synchronous only works with two nodes while asynchronous many
- B. Synchronous replicates to all nodes before commit while async not
- C. Asynchronous replicates to all nodes before commit while sync not
- D. Both methods replicate at exactly the same speed and timing

Answer: B

Q828. What is a global schema in a distributed database?

- A. A schema that only exists at the primary coordinator node
- B. A temporary schema created for each distributed query session
- C. A physical schema that defines storage at each individual site
- D. A unified logical view of the entire distributed database

Answer: D

Q829. What is the semi-join operation used for in distributed databases?

- A. Joining all columns from two tables at different remote sites
- B. Splitting a table into horizontal fragments for distribution
- C. Replicating the result of a join across all available sites
- D. Reducing data transfer by sending only relevant join columns

Answer: D

Q830. What is sharding in distributed databases?

- A. Replicating the entire database to every node in the cluster
- B. Encrypting data during transmission between distributed nodes
- C. Compressing data before storing it on distributed server nodes
- D. Partitioning data across nodes where each holds a data subset

Answer: D

Q831. What is the BASE model in NoSQL databases?

- A. Batch Access, Sequential Execution, Cached data approach
- B. Binary Allocation, Sorted Entries, Encrypted data approach
- C. Basically Available, Soft state, Eventually consistent approach
- D. Buffered Access, Static Entries, Compressed data approach

Answer: C

Q832. How does a column-family store differ from a relational table?

- A. Column-family requires a fixed schema like relational tables do
- B. Column-family does not support any form of data querying at all
- C. Column-family groups columns into families with dynamic columns
- D. Column-family stores data in rows like relational tables exactly

Answer: C

Q833. What is the advantage of document databases for hierarchical data?

- A. They require multiple collections for each hierarchy level
- B. They require joins to represent nested data structures
- C. They store nested data naturally within a single document
- D. They flatten all hierarchical data into a single flat table

Answer: C

Q834. What is a graph database best suited for?

- A. Storing large volumes of unstructured text documents only
- B. Storing time-series data from Internet of Things sensors
- C. Processing simple key-value lookups at maximum speed
- D. Modeling and querying highly connected relationship data

Answer: D

Q835. What is eventual consistency in NoSQL systems?

- A. All replicas are always perfectly synchronized at every moment
- B. Consistency is only guaranteed during maintenance windows only
- C. Replicas may temporarily diverge but will converge over time
- D. Data is never consistent across any of the database replicas

Answer: C

Q836. What is data denormalization in NoSQL and why is it used?

- A. Compressing data to reduce storage requirements on disk drives
- B. Removing all redundancy from the data model for efficiency
- C. Encrypting data for security before storing in the database
- D. Duplicating data to avoid expensive join operations at query

Answer: D

Q837. What is the difference between embedded and referenced documents in MongoDB?

- A. Referenced stores related data together while embedded links by ID
- B. Embedded is for read operations while referenced is for write only
- C. Both approaches store data identically with no practical difference
- D. Embedded stores related data together while referenced links by ID

Answer: D

Q838. What is MapReduce in the context of NoSQL databases?

- A. A strategy for backing up NoSQL databases to external storage
- B. A programming model for processing large data sets in parallel
- C. A method for encrypting data stored in NoSQL database systems
- D. A technique for replicating data across multiple server nodes

Answer: B

Q839. What is a wide-column store?

- A. A database that only supports columns with wide text data types
- B. A database that stores each row with exactly the same columns
- C. A database where each row can have different sets of columns
- D. A database where column width is fixed at creation time always

Answer: C

Q840. What is the role of a partition key in NoSQL databases?

- A. It manages user authentication and access control rights
- B. It defines the schema structure for each document record
- C. It determines how data is distributed across cluster nodes
- D. It encrypts data before storing it in the database system

Answer: C

Q841. What is a snowflake schema and how does it differ from a star schema?

- A. Snowflake has no dimension tables while star schema has many
- B. Snowflake normalizes dimension tables while star keeps denormalized
- C. Snowflake has multiple fact tables while star schema has one only
- D. Snowflake is faster for queries while star schema is slower ones

Answer: B

Q842. What is the difference between Extract and Transform phases in ETL?

- A. Extract cleans data while Transform retrieves data from sources
- B. Extract and Transform perform the exact same operations on data
- C. Extract loads data into warehouse while Transform retrieves it
- D. Extract retrieves source data while Transform cleans and converts

Answer: D

Q843. What is a slowly changing dimension (SCD)?

- A. A dimension whose attributes change infrequently over time
- B. A dimension that is updated in real-time with every change
- C. A dimension that never changes after initial data loading
- D. A dimension that is deleted and recreated each data load

Answer: A

Q844. What is the drill-down operation in OLAP?

- A. Moving from detailed data to summarized aggregate level
- B. Sorting data by a specified column in descending order only
- C. Moving from summary data to more detailed granular level
- D. Filtering data to show only records from a specific time

Answer: C

Q845. What is the roll-up operation in OLAP?

- A. Moving from aggregated summary to detailed granular data
- B. Adding new dimension tables to the existing schema design
- C. Deleting old data from the data warehouse storage tables
- D. Moving from detailed data to higher-level aggregated summary

Answer: D

Q846. What is a surrogate key in a data warehouse dimension table?

- A. A natural business key from the source operational system
- B. A composite key made from multiple business attribute columns
- C. A foreign key referencing the primary fact table records
- D. An artificial system-generated key independent of source data

Answer: D

Q847. What is data granularity in a data warehouse?

- A. The frequency at which the data warehouse backup is performed
- B. The level of detail at which data is stored in the warehouse
- C. The total number of dimension tables in the warehouse schema
- D. The level of encryption applied to stored warehouse data

Answer: B

Q848. What is the slice operation in OLAP?

- A. Selecting a single value on one dimension to view data
- B. Aggregating data across all available dimensions at once
- C. Filtering data based on a range of values in a dimension
- D. Rotating the data cube axes to change the perspective view

Answer: A

Q849. What is the dice operation in OLAP?

- A. Rotating the axes of the data cube for different views
- B. Selecting data by filtering on a single dimension value
- C. Aggregating data from the lowest to highest detail level
- D. Selecting a sub-cube by filtering on multiple dimensions

Answer: D

Q850. What is the pivot operation in OLAP?

- A. Filtering data to show only specific dimension values only
- B. Merging two dimension tables into a single combined table
- C. Rotating the data cube axes to get a different perspective
- D. Adding new measures to the existing fact table in warehouse

Answer: C

Q851. Which type of data independence allows changes to the storage structure without affecting the conceptual schema?

- A. Logical data independence
- B. Physical data independence
- C. Schema independence
- D. View independence

Answer: B

Q852. In the three-schema architecture, which schema represents individual user views?

- A. Internal schema
- B. Conceptual schema
- C. External schema
- D. Physical schema

Answer: C

Q853. What is the primary difference between a database and a data warehouse?

- A. A database is for operational data while a data warehouse is for analytical data
- B. A database is larger than a data warehouse
- C. A data warehouse only stores text data
- D. There is no difference between them

Answer: A

Q854. Which DBMS component is responsible for enforcing integrity constraints?

- A. Query optimizer
- B. Integrity constraint manager
- C. Buffer manager
- D. File manager

Answer: B

Q855. What is the purpose of a view in a database system?

- A. To store physical data on disk
- B. To provide a virtual table derived from base tables
- C. To manage network connections
- D. To compile SQL statements

Answer: B

Q856. Which of the following is a characteristic of the relational model?

- A. Data is stored in tree structures
- B. Data is represented as tables with rows and columns
- C. Data must be stored sequentially
- D. Each record must have a parent record

Answer: B

Q857. What role does the data dictionary play in a DBMS?

- A. It stores actual user data
- B. It stores metadata about the database structure
- C. It performs data encryption
- D. It manages user sessions

Answer: B

Q858. Which of the following describes a multi-user DBMS?

- A. A DBMS that runs on a single processor only
- B. A DBMS that allows multiple users to access the database concurrently
- C. A DBMS that stores data in multiple formats
- D. A DBMS limited to read-only access

Answer: B

Q859. What is the difference between intension and extension of a database?

- A. Intension is the schema while extension is the current data
- B. Intension is the data while extension is the schema
- C. Both refer to the same concept
- D. Intension refers to queries while extension refers to indexes

Answer: A

Q860. Which of the following is an example of a DCL command?

- A. CREATE TABLE
- B. SELECT
- C. GRANT
- D. UPDATE

Answer: C

Q861. What is the primary benefit of connection pooling in database architecture?

- A. It encrypts all database connections
- B. It reuses existing connections to reduce overhead of creating new ones
- C. It automatically backs up data
- D. It compresses query results

Answer: B

Q862. In a shared-memory parallel architecture, what is the main bottleneck?

- A. Network bandwidth between nodes
- B. Memory bus contention when many processors access shared memory
- C. Disk storage capacity
- D. Application code complexity

Answer: B

Q863. What is the function of the log manager in DBMS architecture?

- A. It manages user login sessions
- B. It records all database modifications for recovery purposes
- C. It formats query output
- D. It manages network connections

Answer: B

Q864. Which architecture is most suitable for large-scale internet applications?

- A. Single-tier architecture
- B. Two-tier architecture
- C. N-tier or multi-tier architecture
- D. Standalone architecture

Answer: C

Q865. What distinguishes a fat client from a thin client in database architecture?

- A. Fat clients have larger monitors
- B. Fat clients process more application logic locally while thin clients rely on the server
- C. Fat clients use more network bandwidth
- D. There is no real difference

Answer: B

Q866. In DBMS architecture, what does the term 'middleware' refer to?

- A. Hardware between the client and server
- B. Software that connects different applications or system components
- C. The middle table in a database
- D. A type of storage device

Answer: B

Q867. What is the role of the query optimizer in DBMS architecture?

- A. To check SQL syntax
- B. To determine the most efficient execution plan for a query
- C. To store query results
- D. To manage user connections

Answer: B

Q868. What is an embedded database system?

- A. A database that runs within an application process without a separate server
- B. A database stored on embedded hardware only
- C. A database that cannot be queried
- D. A database used only for IoT devices

Answer: A

Q869. In database architecture, what is a stored procedure?

- A. A procedure for physically storing disks
- B. A precompiled collection of SQL statements stored in the database
- C. A method for backing up data
- D. A procedure for installing the DBMS

Answer: B

Q870. What advantage does the three-tier architecture offer over two-tier for security?

- A. It uses stronger encryption algorithms
- B. It adds an intermediate layer that prevents direct client access to the database
- C. It requires fewer passwords
- D. It eliminates the need for firewalls

Answer: B

Q871. What is partial participation in an ER diagram?

- A. All entity instances must participate in the relationship
- B. Some entity instances may not participate in the relationship
- C. The relationship only involves one entity
- D. The entity has no attributes

Answer: B

Q872. How is a weak entity identified in an ER diagram?

- A. By a single rectangle
- B. By a double-bordered rectangle
- C. By a dashed rectangle
- D. By a bold rectangle

Answer: B

Q873. What is a partial key (discriminator) of a weak entity?

- A. A key that uniquely identifies the weak entity globally
- B. An attribute that partially identifies weak entity instances within their owner
- C. A foreign key reference
- D. A candidate key of a strong entity

Answer: B

Q874. In the EER model, what does disjoint specialization mean?

- A. An entity can belong to multiple subclasses simultaneously
- B. An entity can belong to only one subclass at a time
- C. All entities must belong to a subclass
- D. No entity belongs to any subclass

Answer: B

Q875. What is the difference between specialization and generalization in EER?

- A. They are the same process
- B. Specialization is top-down while generalization is bottom-up
- C. Specialization is bottom-up while generalization is top-down
- D. Neither relates to entity hierarchies

Answer: B

Q876. How is a derived attribute represented in an ER diagram?

- A. With a solid oval
- B. With a dashed oval
- C. With a double oval
- D. With an underlined oval

Answer: B

Q877. What is an identifying relationship for a weak entity?

- A. Any relationship the weak entity participates in
- B. The relationship that provides the weak entity with part of its identification
- C. A self-referencing relationship
- D. A many-to-many relationship

Answer: B

Q878. In the ER model, what is a binary relationship?

- A. A relationship involving exactly two entity types
- B. A relationship stored in binary format
- C. A relationship with two attributes
- D. A relationship between an entity and itself

Answer: A

Q879. What does the participation constraint specify in an ER relationship?

- A. The data types of related entities
- B. Whether existence of an entity depends on being related to another entity
- C. The maximum number of entities in a table
- D. The storage format of relationships

Answer: B

Q880. What is an associative entity in ER modeling?

- A. A strong entity with multiple keys
- B. An entity created to resolve a many-to-many relationship
- C. An entity with only derived attributes
- D. An entity without any relationships

Answer: B

Q881. What is a functional dependency in relational databases?

- A. A dependency between two database servers
- B. A constraint where one attribute set determines another
- C. A physical storage dependency
- D. A relationship between database users

Answer: B

Q882. What is the update anomaly in relational databases?

- A. An error that occurs during database installation
- B. Inconsistency caused when the same data must be updated in multiple rows
- C. A problem with database backup procedures
- D. An error in the query optimizer

Answer: B

Q883. What is the insertion anomaly in a relational table?

- A. Inability to insert a new row because required data for another entity is missing
- B. An error in the INSERT SQL statement syntax
- C. A performance issue during batch inserts
- D. A conflict between two simultaneous insertions

Answer: A

Q884. What is the deletion anomaly in relational databases?

- A. A failure to delete data from disk
- B. Unintended loss of data when deleting a row that contains other useful information
- C. A conflict during concurrent deletions
- D. An error in the DELETE SQL syntax

Answer: B

Q885. What is a natural join in the relational model?

- A. A join that only uses primary keys
- B. A join that combines tuples from two relations based on common attributes with equal values
- C. A join that preserves all rows from both tables
- D. A join that creates a Cartesian product

Answer: B

Q886. What constraint does the ON DELETE CASCADE option enforce?

- A. Prevents deletion of any row with a primary key
- B. Automatically deletes referencing rows when the referenced row is deleted
- C. Logs all deletion operations
- D. Requires manual confirmation before deletion

Answer: B

Q887. What is a self-referencing foreign key?

- A. A foreign key that references a column in the same table
- B. A foreign key that references itself recursively with no limit
- C. A key that auto-increments
- D. A key used for database replication

Answer: A

Q888. What is the closed world assumption in the relational model?

- A. The database is closed for modifications after deployment
- B. Only the data present in the database is considered true; absent data is considered false
- C. All tables must be fully populated before queries
- D. The database cannot communicate with external systems

Answer: B

Q889. What is the difference between ON DELETE SET NULL and ON DELETE RESTRICT?

- A. SET NULL deletes the child row while RESTRICT sets it to null
- B. SET NULL sets the foreign key to null while RESTRICT prevents the deletion
- C. Both options behave identically
- D. SET NULL removes the constraint while RESTRICT adds one

Answer: B

Q890. What is a theta join in relational algebra?

- A. A join using only the equality condition
- B. A join that combines relations based on any comparison condition
- C. A join that produces a Cartesian product
- D. A join exclusive to temporal databases

Answer: B

Q891. How does a left outer join differ from an inner join?

- A. Left outer join only returns rows from the left table
- B. Left outer join includes all rows from the left table, padding with nulls where there is no match
- C. Left outer join is faster than inner join
- D. There is no difference

Answer: B

Q892. In relational algebra, can the intersection operation be expressed using other fundamental operations?

- A. No, intersection is a fundamental operation
- B. Yes, $R \cap S = R - (R - S)$
- C. Yes, $R \cap S = R \times S$
- D. Yes, $R \cap S = R \div S$

Answer: B

Q893. What is the result of a full outer join?

- A. Only matching rows from both tables
- B. All rows from both tables with nulls where no match exists
- C. Only the Cartesian product
- D. An empty relation

Answer: B

Q894. What does the grouping (gamma) operation do in extended relational algebra?

- A. Groups tuples by specified attributes and applies aggregate functions
- B. Groups tables into schemas
- C. Groups users by access level
- D. Groups queries for batch processing

Answer: A

Q895. What is a semijoin in relational algebra?

- A. Half of a natural join result
- B. A join that returns tuples from the first relation that have matching tuples in the second
- C. A join that uses only half the attributes
- D. A partial Cartesian product

Answer: B

Q896. Why is relational algebra considered procedural?

- A. Because it uses procedures stored in the database
- B. Because it specifies the sequence of operations to obtain the result
- C. Because it requires procedural programming languages
- D. Because it processes data in procedure calls

Answer: B

Q897. What is relational completeness?

- A. A database that stores all possible data
- B. A query language that can express any query expressible in relational algebra
- C. A relation that has no null values
- D. A database with all constraints defined

Answer: B

Q898. What is the difference between WHERE and HAVING clauses in SQL?

- A. There is no difference
- B. WHERE filters rows before grouping while HAVING filters groups after grouping
- C. HAVING is used before WHERE in query execution
- D. WHERE only works with numeric columns

Answer: B

Q899. What is a correlated subquery in SQL?

- A. A subquery that runs independently of the outer query
- B. A subquery that references columns from the outer query
- C. A query that uses the CORRELATE keyword
- D. A query that joins two subqueries

Answer: B

Q900. What does the COALESCE function do in SQL?

- A. Combines two tables
- B. Returns the first non-null value from a list of expressions
- C. Counts all null values
- D. Converts null to zero

Answer: B

Q901. What is the purpose of the GROUP BY clause in SQL?

- A. To sort the result set
- B. To group rows with the same values in specified columns for aggregate functions
- C. To filter individual rows
- D. To join tables together

Answer: B

Q902. What does the EXISTS keyword do in SQL?

- A. Checks if a table exists in the database
- B. Returns TRUE if the subquery returns one or more rows
- C. Checks if a column exists in a table
- D. Verifies database connection

Answer: B

Q903. What is the purpose of the CASE expression in SQL?

- A. To define a database case study
- B. To provide conditional logic within a SQL query
- C. To change the case of text values
- D. To handle exception cases

Answer: B

Q904. What is an alias in SQL?

- A. A permanent rename of a table
- B. A temporary name given to a table or column for the duration of a query
- C. A backup name for a database
- D. A type of database index

Answer: B

Q905. What is the difference between BEFORE and AFTER triggers?

- A. BEFORE triggers execute before the triggering event while AFTER triggers execute after
- B. BEFORE triggers are faster than AFTER triggers
- C. AFTER triggers cannot modify data
- D. There is no practical difference

Answer: A

Q906. What is a Common Table Expression (CTE) in SQL?

- A. A table shared across multiple databases
- B. A temporary named result set defined within a query using the WITH clause
- C. A permanent table created with CREATE TABLE
- D. A type of materialized view

Answer: B

Q907. What is the difference between a function and a stored procedure in SQL?

- A. Functions must return a value and can be used in expressions while procedures may not return a value
- B. Functions cannot take parameters
- C. Procedures always return a value
- D. There is no difference

Answer: A

Q908. What is dynamic SQL?

- A. SQL that adjusts to different time zones
- B. SQL statements constructed and executed at runtime as strings
- C. SQL that automatically optimizes itself
- D. SQL used only in dynamic websites

Answer: B

Q909. What is the WITH GRANT OPTION used for in SQL privileges?

- A. To grant access for a limited time
- B. To allow the grantee to further grant the same privileges to other users
- C. To grant read-only access
- D. To grant access to system tables

Answer: B

Q910. What is an updatable view in SQL?

- A. Any view that has been recently modified
- B. A view through which INSERT, UPDATE, and DELETE operations can be performed on the base tables
- C. A view that refreshes automatically
- D. A view with an auto-increment column

Answer: B

Q911. What is the purpose of the INSTEAD OF trigger?

- A. To execute instead of a SELECT query
- B. To replace the default DML action on a view with custom logic
- C. To run instead of other triggers
- D. To execute when an error occurs instead of raising the error

Answer: B

Q912. What is a row-level trigger versus a statement-level trigger?

- A. Row-level triggers fire once per affected row while statement-level triggers fire once per SQL statement
- B. Row-level triggers are faster
- C. Statement-level triggers process each row individually
- D. They execute at the same frequency

Answer: A

Q913. What is embedded SQL?

- A. SQL stored inside HTML pages
- B. SQL statements written within a host programming language like C or Java
- C. SQL embedded in PDF documents
- D. SQL that runs only on embedded systems

Answer: B

Q914. What is Boyce-Codd Normal Form (BCNF) and how does it differ from 3NF?

- A. BCNF is the same as 3NF
- B. BCNF is stricter than 3NF, requiring every determinant to be a candidate key
- C. BCNF is weaker than 3NF
- D. BCNF only applies to tables with composite keys

Answer: B

Q915. Can a relation be in 3NF but not in BCNF?

- A. No, 3NF always implies BCNF
- B. Yes, when a non-candidate key attribute determines part of a candidate key
- C. No, BCNF is a weaker form than 3NF
- D. Yes, but only in tables with no candidate keys

Answer: B

Q916. What is Fourth Normal Form (4NF)?

- A. A table in BCNF with no trivial multivalued dependencies
- B. A table with four or fewer columns
- C. A table with exactly four candidate keys
- D. A table that has been normalized four times

Answer: A

Q917. What is the purpose of decomposition in normalization?

- A. To delete unnecessary data
- B. To split a relation into smaller relations to eliminate anomalies while preserving data
- C. To combine multiple tables into one
- D. To encrypt data across multiple tables

Answer: B

Q918. What is a canonical cover (minimal cover) of functional dependencies?

- A. The largest possible set of functional dependencies
- B. A minimal set of functional dependencies equivalent to the original set with no redundancy
- C. A set of dependencies that covers all tables
- D. The first set of dependencies defined during design

Answer: B

Q919. When is denormalization appropriate in database design?

- A. Always, as it improves all aspects of performance
- B. When read performance is critical and the overhead of maintaining redundancy is acceptable
- C. Never, as it violates database theory
- D. Only when the database has fewer than 10 tables

Answer: B

Q920. What is the synthesis approach to normalization?

- A. Combining all tables into one
- B. Building relations from a given set of functional dependencies using a bottom-up approach
- C. Creating tables randomly and then checking normal forms
- D. Synthesizing data from external sources

Answer: B

Q921. What is a prime attribute?

- A. The most important attribute in a table
- B. An attribute that is part of any candidate key
- C. An attribute that stores primary data
- D. The first attribute defined in a table

Answer: B

Q922. What does the decomposition algorithm for BCNF guarantee?

- A. Both lossless join and dependency preservation
- B. Lossless join but not necessarily dependency preservation
- C. Dependency preservation but not lossless join
- D. Neither lossless join nor dependency preservation

Answer: B

Q923. What is a conflict in transaction schedules?

- A. A disagreement between database administrators
- B. When two operations from different transactions access the same data item and at least one is a write
- C. When two transactions have the same name
- D. When the database runs out of storage

Answer: B

Q924. What is conflict serializability?

- A. A schedule where all conflicts are resolved
- B. A schedule that is equivalent to a serial schedule by swapping non-conflicting operations
- C. A schedule with no conflicts
- D. A schedule where conflicts are ignored

Answer: B

Q925. How is conflict serializability tested using a precedence graph?

- A. By counting the number of nodes
- B. A schedule is conflict serializable if and only if its precedence graph has no cycles
- C. By checking if the graph is fully connected
- D. By verifying all nodes have equal edges

Answer: B

Q926. What is a cascading rollback?

- A. Rolling back only the last operation of a transaction
- B. When the rollback of one transaction forces rollbacks of other dependent transactions
- C. Rolling back changes in alphabetical order
- D. A rollback that happens automatically at midnight

Answer: B

Q927. What is the difference between a strict and a cascadeless schedule?

- A. They are the same thing
- B. Cascadeless prevents cascading rollbacks; strict additionally prevents reading uncommitted data items until the writing transaction commits or aborts
- C. Strict schedules allow more concurrency
- D. Cascadeless is stricter than strict schedules

Answer: B

Q928. What is view serializability?

- A. A schedule based on database views
- B. A schedule where each transaction sees the same data values as in some serial schedule
- C. A schedule that only involves view operations
- D. A schedule verified by viewing the database

Answer: B

Q929. What is a blind write in transactions?

- A. A write performed by an unauthorized user
- B. A write operation performed without first reading the data item
- C. A write that produces no output
- D. A write to a hidden column

Answer: B

Q930. What is strict two-phase locking (Strict 2PL)?

- A. A 2PL variant where locks are acquired in strict order
- B. A 2PL variant where all exclusive locks are held until the transaction commits or aborts
- C. A 2PL variant that uses only shared locks
- D. A 2PL variant with no growing phase

Answer: B

Q931. What is rigorous two-phase locking?

- A. A 2PL variant that is more flexible
- B. A 2PL variant where all locks (both shared and exclusive) are held until commit or abort
- C. A 2PL variant that uses no locks
- D. A 2PL variant for read-only transactions

Answer: B

Q932. What is the wait-die deadlock prevention scheme?

- A. Older transactions wait; younger ones are rolled back if they would need to wait
- B. Younger transactions always wait for older ones
- C. Both transactions are rolled back
- D. The database picks a random transaction to abort

Answer: A

Q933. What is the wound-wait deadlock prevention scheme?

- A. Older transactions roll back younger ones; younger transactions wait for older ones
- B. Younger transactions always get priority
- C. Both transactions always wait
- D. Deadlocks are detected after they occur

Answer: A

Q934. How is deadlock detection performed using a wait-for graph?

- A. By counting the number of edges
- B. A deadlock exists if and only if the wait-for graph contains a cycle
- C. By checking if the graph has more than 10 nodes
- D. By measuring the graph's depth

Answer: B

Q935. What is lock granularity in concurrency control?

- A. The size of the lock data structure
- B. The level at which locking is applied: database, table, page, or row level
- C. The number of transactions that can hold a lock
- D. The duration a lock is held

Answer: B

Q936. What is the trade-off between fine-grained and coarse-grained locking?

- A. Fine-grained locking is always better
- B. Fine-grained locking allows more concurrency but has higher lock management overhead
- C. Coarse-grained locking allows more concurrency
- D. There is no trade-off

Answer: B

Q937. What is an intention lock in multi-granularity locking?

- A. A lock that indicates the transaction intends to lock the database
- B. A lock placed on a higher-level node indicating that a finer-grained lock is held on a descendant
- C. A lock that is planned but not yet acquired
- D. A lock used during database design

Answer: B

Q938. What is a livelock in concurrency control?

- A. A lock that keeps the database alive
- B. A situation where a transaction repeatedly retries an operation but never succeeds
- C. A lock on live streaming data
- D. A lock that never expires

Answer: B

Q939. What is the ARIES recovery algorithm?

- A. A simple backup tool
- B. An advanced recovery algorithm using WAL with three phases: analysis, redo, and undo
- C. An encryption algorithm for log files
- D. A database compression algorithm

Answer: B

Q940. In ARIES, what does the analysis phase do?

- A. Analyzes query performance
- B. Scans the log from the last checkpoint to determine dirty pages and active transactions at crash time
- C. Analyzes the database schema
- D. Performs statistical analysis on data

Answer: B

Q941. What is the difference between steal and no-steal buffer management policies?

- A. Steal allows modified pages to be flushed before commit while no-steal does not
- B. Steal refers to taking data from other transactions
- C. No-steal prevents any buffer modifications
- D. They refer to how memory is allocated

Answer: A

Q942. What is the difference between force and no-force buffer management policies?

- A. Force requires all modified pages to be written to disk at commit time while no-force does not
- B. Force means using forced writes for better performance
- C. No-force prevents any writes to disk
- D. They refer to the physical force required to insert disks

Answer: A

Q943. What combination of buffer management policies does ARIES use?

- A. Force, No-steal
- B. No-force, Steal
- C. Force, Steal
- D. No-force, No-steal

Answer: B

Q944. What is a log sequence number (LSN)?

- A. A serial number for database licenses
- B. A unique identifier assigned to each log record to maintain ordering
- C. A number identifying the database sequence
- D. A network sequence number

Answer: B

Q945. What is physiological logging?

- A. Logging based on physical health metrics
- B. A logging approach that identifies pages physically but logs changes logically within each page
- C. A form of logging used only for physical databases
- D. Logging the physical location of data on disk

Answer: B

Q946. What is a compensation log record (CLR) in ARIES?

- A. A log record for financial compensation
- B. A log record written during undo to describe the reversal of an action
- C. A log record that compensates for lost data
- D. A log record for replication

Answer: B

Q947. What is the purpose of a fuzzy checkpoint?

- A. A checkpoint with imprecise data
- B. A checkpoint that does not halt transaction processing while recording the checkpoint information
- C. A checkpoint performed at random intervals
- D. A checkpoint that only saves partial data

Answer: B

Q948. What is a B+ tree and how does it differ from a B-tree?

- A. They are identical structures
- B. In a B+ tree, all data pointers are at the leaf level and leaves are linked, unlike B-trees where data pointers exist at all levels
- C. A B+ tree is always smaller than a B-tree
- D. B+ trees do not support range queries

Answer: B

Q949. What is a hash index?

- A. An index that uses a hash function to map keys directly to their storage location
- B. An index that stores data in alphabetical order
- C. An index used only for text fields
- D. An index that requires no storage space

Answer: A

Q950. What is a covering index?

- A. An index that covers the entire disk
- B. An index that contains all columns needed by a query, eliminating the need to access the base table
- C. An index that covers all tables in the database
- D. An index with a protective encryption layer

Answer: B

Q951. When would a hash index be preferred over a B+ tree index?

- A. For range queries
- B. For exact match (equality) lookups where range queries are not needed
- C. For pattern matching queries
- D. Hash indexes are always preferred

Answer: B

Q952. What is a bitmap index?

- A. An index stored as an image file
- B. An index that uses bit arrays to represent the presence or absence of values for each indexed column
- C. An index on binary data columns
- D. An index that maps bits to bytes

Answer: B

Q953. What is the order (or degree) of a B-tree?

- A. The number of levels in the tree
- B. The maximum number of children a node can have
- C. The number of records in the tree
- D. The total number of nodes

Answer: B

Q954. What is index selectivity?

- A. The ability to select which index to use
- B. The ratio of distinct indexed values to total number of records, indicating how well an index discriminates
- C. How selective the DBA is when creating indexes
- D. The speed of index selection

Answer: B

Q955. What is a non-clustered index?

- A. An unorganized index
- B. An index whose logical order differs from the physical order of the data rows
- C. An index without a B-tree structure
- D. An index that cannot be used for searching

Answer: B

Q956. What is the nested loop join algorithm?

- A. A join that creates nested tables
- B. For each row in the outer table, it scans the entire inner table to find matches
- C. A join that only works with nested queries
- D. A loop that nests multiple joins together

Answer: B

Q957. What is the sort-merge join algorithm?

- A. A join that sorts results after merging
- B. Both tables are sorted on the join attribute and then merged together in a single pass
- C. A join that only works on sorted tables
- D. A combination of sorting and hashing

Answer: B

Q958. What is the hash join algorithm?

- A. A join that uses hash codes for security
- B. One table is hashed into buckets on the join key, then the other table is probed against these buckets
- C. A join that produces hashed output
- D. A join between hash tables

Answer: B

Q959. What is the difference between a rule-based and cost-based query optimizer?

- A. They produce identical plans
- B. Rule-based uses predefined heuristic rules while cost-based uses statistical cost estimates to choose plans
- C. Rule-based is always better
- D. Cost-based does not use any rules

Answer: B

Q960. Why is pushing selections down in a query tree an important optimization?

- A. It makes the query tree shorter
- B. It reduces the size of intermediate results early, decreasing processing costs for subsequent operations
- C. It removes the need for joins
- D. It changes the query semantics for better results

Answer: B

Q961. What is the difference between materialization and pipelining in query execution?

- A. They are the same technique
- B. Materialization stores intermediate results on disk while pipelining passes results directly between operations
- C. Pipelining is always slower
- D. Materialization uses less memory

Answer: B

Q962. What are histograms used for in query optimization?

- A. Displaying query results graphically
- B. Providing detailed data distribution information to improve cardinality estimation accuracy
- C. Showing database usage patterns
- D. Graphing query execution times

Answer: B

Q963. What is cardinality estimation in query optimization?

- A. Estimating the number of tables needed
- B. Predicting the number of rows produced by each operation in the query plan
- C. Estimating the number of users
- D. Calculating the number of indexes required

Answer: B

Q964. What is an index-only scan?

- A. A scan that creates an index
- B. A query execution method that retrieves all needed data from the index without accessing the base table
- C. A scan that only checks if an index exists
- D. A scan of the index metadata

Answer: B

Q965. What is SQL injection and how can it be prevented?

- A. A method to inject data faster; prevented by faster hardware
- B. An attack inserting malicious SQL through user input; prevented by parameterized queries and input validation
- C. An SQL optimization technique; prevented by using indexes
- D. A database backup method; prevented by RAID

Answer: B

Q966. What is mandatory access control (MAC)?

- A. Access control that all users must follow voluntarily
- B. Access control enforced by the system based on security labels and clearance levels, not user discretion
- C. Optional access restrictions
- D. Access control managed by users themselves

Answer: B

Q967. What is discretionary access control (DAC)?

- A. Access control based on time of day
- B. Access control where data owners grant or revoke permissions to other users at their discretion
- C. Access control enforced by the operating system only
- D. Access control that cannot be changed

Answer: B

Q968. What is the Bell-LaPadula model in database security?

- A. A model for database performance tuning
- B. A multilevel security model that enforces no read up and no write down policies
- C. A model for database backup
- D. A relational data model

Answer: B

Q969. What is the difference between encryption at rest and encryption in transit?

- A. They are the same type of encryption
- B. At rest protects stored data on disk while in transit protects data moving across networks
- C. At rest encrypts network data while in transit encrypts stored data
- D. At rest is stronger than in transit

Answer: B

Q970. What is a privilege cascade in database security?

- A. A waterfall display of privileges
- B. When revoking a privilege from a user also revokes it from users who received it through that user
- C. Cascading permissions to all database objects
- D. A list of privileges sorted by priority

Answer: B

Q971. What is column-level security?

- A. Security that applies to all columns equally
- B. Restricting user access to specific columns within a table rather than the entire table
- C. Encryption of column headers
- D. A security column added to every table

Answer: B

Q972. What is row-level security (RLS)?

- A. Security applied to each row of the audit log
- B. Filtering rows based on user attributes so users can only see rows they are authorized to access
- C. Encrypting individual rows
- D. Adding a security row to each table

Answer: B

Q973. What is the two-phase commit (2PC) protocol?

- A. A commitment to finish in two phases of development
- B. A protocol that coordinates distributed transactions by having a prepare phase and a commit/abort phase
- C. A two-step backup process
- D. A protocol for connecting two database servers

Answer: B

Q974. What is the role of the coordinator in the two-phase commit protocol?

- A. To optimize queries across sites
- B. To initiate the commit protocol, collect votes from participants, and make the final commit or abort decision
- C. To balance load across database sites
- D. To manage user authentication

Answer: B

Q975. What is the difference between full replication and partial replication?

- A. Full replication stores more data per site
- B. Full replication copies the entire database to every site while partial replication copies only selected fragments
- C. Full replication is always faster
- D. Partial replication is not possible in distributed databases

Answer: B

Q976. What is a distributed catalog in a distributed database?

- A. A catalog distributed to all users
- B. Metadata about the location, fragmentation, and replication of data across sites
- C. A printed catalog of database features
- D. A central list of all distributed databases

Answer: B

Q977. What is the global deadlock problem in distributed databases?

- A. A deadlock affecting the entire internet
- B. A deadlock involving transactions at multiple sites that cannot be detected by any single site's local deadlock detector
- C. A deadlock in the global catalog
- D. A worldwide database outage

Answer: B

Q978. What is mixed fragmentation in distributed databases?

- A. A mix of correct and corrupted data
- B. A combination of horizontal and vertical fragmentation applied to a relation
- C. Mixing data from different databases
- D. Fragmenting data in a random order

Answer: B

Q979. What is the principle of local autonomy in distributed databases?

- A. Each site must be autonomous from the network
- B. Each site maintains control over its local data and can operate independently
- C. Users can autonomously create databases
- D. Local sites do not need backup

Answer: B

Q980. How does MongoDB store data internally?

- A. In relational tables
- B. In BSON (Binary JSON) documents organized into collections
- C. In key-value pairs only
- D. In column families

Answer: B

Q981. What is sharding in NoSQL databases?

- A. Breaking data into pieces and discarding some
- B. Distributing data across multiple servers by partitioning it based on a shard key
- C. Creating backup copies of data
- D. Encrypting data into shards

Answer: B

Q982. What is the difference between embedded documents and references in MongoDB?

- A. They are the same concept
- B. Embedded documents nest related data within a document while references store relationships using document IDs
- C. Embedded documents are faster to delete
- D. References are used only for primary keys

Answer: B

Q983. What is the Cassandra data model?

- A. A relational model with SQL support
- B. A partitioned row store using partition keys for distribution and clustering keys for sorting within partitions
- C. A pure key-value model
- D. A document-based model

Answer: B

Q984. What is the CAP theorem's significance for NoSQL database design?

- A. It has no relevance to NoSQL databases
- B. It forces NoSQL designers to choose between consistency, availability, and partition tolerance trade-offs
- C. It guarantees all three properties for NoSQL databases
- D. It only applies to SQL databases

Answer: B

Q985. What is a time-series database?

- A. A database that tracks the time of each query
- B. A specialized database optimized for storing and querying time-stamped data points
- C. A database that only works during certain times
- D. A temporary database

Answer: B

Q986. What is the aggregation framework in MongoDB?

- A. A framework for adding data
- B. A pipeline-based data processing framework that transforms and analyzes documents through stages
- C. A framework for managing database clusters
- D. A tool for database administration

Answer: B

Q987. What is the difference between SCD Type 1 and SCD Type 2?

- A. They handle dimension changes identically
- B. Type 1 overwrites old values while Type 2 creates new rows to preserve history
- C. Type 1 preserves history while Type 2 does not
- D. Type 1 is for fact tables while Type 2 is for dimension tables

Answer: B

Q988. What is a surrogate key in a data warehouse?

- A. A natural business key
- B. A system-generated artificial key used as the primary key in dimension tables instead of natural keys
- C. A temporary key used during ETL
- D. A backup of the primary key

Answer: B

Q989. What is the difference between MOLAP, ROLAP, and HOLAP?

- A. They are different programming languages
- B. MOLAP uses multidimensional storage, ROLAP uses relational storage, and HOLAP combines both approaches
- C. They are different database vendors
- D. MOLAP is the fastest and ROLAP is the slowest

Answer: B

Q990. What is a conformed dimension in data warehousing?

- A. A dimension that conforms to naming conventions
- B. A dimension that has the same meaning and content when referenced across multiple fact tables or data marts
- C. A dimension that has been validated
- D. A dimension with confirmed data

Answer: B

Q991. What is a surrogate key and when is it used in relational design?

- A. A key borrowed from another table
- B. A system-generated artificial key used when no natural candidate key is suitable
- C. A temporary key used during testing
- D. A key that replaces foreign keys

Answer: B

Q992. What is the extended projection operation in relational algebra?

- A. A projection that includes all columns
- B. A projection that allows arithmetic expressions and string operations on attributes in the result
- C. A projection that extends the table with new rows
- D. A projection applied to multiple tables

Answer: B

Q993. What is the assignment operation used for in relational algebra?

- A. Assigning users to database roles
- B. Storing the result of a relational algebra expression in a temporary relation variable
- C. Assigning data types to attributes
- D. Assigning primary keys to tables

Answer: B

Q994. What is the purpose of the IN operator in SQL?

- A. To insert new rows
- B. To check if a value matches any value in a specified list or subquery
- C. To index a column
- D. To include all columns in the result

Answer: B

Q995. What is the difference between INNER JOIN and CROSS JOIN in SQL?

- A. They are identical
- B. INNER JOIN matches rows based on a condition while CROSS JOIN produces all combinations of rows from both tables
- C. CROSS JOIN requires a join condition
- D. INNER JOIN produces more rows

Answer: B

Q996. What is the NULL-safe equality operator in SQL?

- A. The = operator handles NULL automatically
- B. The IS NOT DISTINCT FROM operator that treats two NULLs as equal
- C. There is no such operator
- D. The <=> operator that ignores NULL values

Answer: B

Q997. What is a materialized view refresh strategy?

- A. Refreshing the browser to view database changes
- B. The method used to update a materialized view's stored data, either completely (full) or with only changes (incremental)
- C. Clearing the view cache
- D. Restarting the database to refresh views

Answer: B

Q998. What is an extraneous attribute in a functional dependency?

- A. An attribute that is extra in the table
- B. An attribute that can be removed from the left or right side of a dependency without changing the closure
- C. An attribute added by mistake
- D. An attribute from an external database

Answer: B

Q999. What is the non-repeatable read problem?

- A. A read operation that takes too long
- B. When a transaction reads the same data item twice and gets different values because another transaction modified it
- C. A read that cannot be repeated due to errors
- D. A read of non-existent data

Answer: B

Q1000. What is the READ UNCOMMITTED isolation level?

- A. A level where no reads are allowed
- B. The lowest isolation level that allows dirty reads, non-repeatable reads, and phantom reads
- C. A level that prevents all anomalies
- D. A level that only allows uncommitted writes

Answer: B

Q1001. What is the difference between an aborted transaction and a failed transaction?

- A. They are the same state
- B. A failed transaction has encountered an error and cannot proceed, while aborted means it has been rolled back
- C. Aborted transactions can resume while failed ones cannot
- D. Failed transactions are committed while aborted ones are not

Answer: B

Q1002. What is lock compatibility in concurrency control?

- A. Whether locks are compatible with the hardware
- B. The rules determining which lock types can coexist on the same data item, such as shared locks being compatible with each other
- C. Whether locks can be used across databases
- D. Compatibility of lock software versions

Answer: B

Q1003. What is a differential backup and how does it differ from an incremental backup?

- A. They are identical
- B. A differential backup captures all changes since the last full backup while an incremental captures changes since the last backup of any type
- C. Differential is smaller than incremental
- D. Incremental captures more data

Answer: B

Q1004. What is a unique index?

- A. An index on a unique table
- B. An index that enforces uniqueness on the indexed columns, preventing duplicate values
- C. The only index in the database
- D. An index with a unique name

Answer: B

Q1005. What is index fragmentation and how does it affect performance?

- A. Physical breaking of the index file
- B. When index pages become disordered or have excessive free space due to insertions and deletions, degrading scan performance
- C. Splitting an index across databases
- D. Creating too many index levels

Answer: B

Q1006. What is the selectivity of a predicate in query optimization?

- A. How selective the DBA is about indexing
- B. The fraction of rows in a table that satisfy the predicate condition, used to estimate result sizes
- C. The order in which predicates are evaluated
- D. The number of predicates in a query

Answer: B

Q1007. What is a database honeypot?

- A. A storage container for database files
- B. A decoy database system designed to attract and detect unauthorized access attempts
- C. A type of database optimization
- D. A backup database server

Answer: B

Q1008. What is the principle of separation of duties in database security?

- A. Separating the database into different files
- B. Dividing critical security functions among different individuals so no single person has complete control
- C. Separating development and production databases
- D. Keeping security logs separate from data

Answer: B

Q1009. What is a distributed join operation?

- A. Joining multiple distributed systems
- B. A join operation that requires data from multiple sites, involving data transfer across the network
- C. Joining two local tables
- D. A social gathering of database administrators

Answer: B

Q1010. What is the difference between tight coupling and loose coupling in distributed databases?

- A. They refer to how tightly cables connect servers
- B. Tight coupling shares resources like memory and is more integrated, while loose coupling connects independent systems via messages
- C. Tight coupling is always faster
- D. There is no practical difference

Answer: B

Q1011. What is the catalog management challenge in distributed databases?

- A. Managing physical catalogs of database products
- B. Deciding where to store the distributed catalog (centralized, replicated, or partitioned) while balancing lookup speed and update cost
- C. Organizing database documentation
- D. Managing vendor catalogs

Answer: B

Q1012. What is an edge in a graph database?

- A. The boundary of the database
- B. A connection between two nodes representing a relationship, which can have its own properties
- C. The last record in a collection
- D. A database boundary condition

Answer: B

Q1013. What is the replica set in MongoDB?

- A. A set of replicated queries
- B. A group of MongoDB instances that maintain the same data set, providing redundancy and high availability
- C. A collection of duplicate documents
- D. A set of backup scripts

Answer: B

Q1014. What is the MapReduce programming model used in NoSQL?

- A. A model for mapping database locations
- B. A data processing paradigm where Map transforms data into key-value pairs and Reduce aggregates values by key
- C. A model for reducing database size
- D. A navigation model for data maps

Answer: B

Q1015. What is a pivot operation in OLAP?

- A. Rotating the OLAP server hardware
- B. Rotating the data cube to view data from a different perspective by swapping dimensions on axes
- C. Pivoting from one database to another
- D. Changing the database server

Answer: B

Q1016. What is the difference between additive, semi-additive, and non-additive measures?

- A. They refer to the number of measures in a fact table
- B. Additive measures can be summed across all dimensions, semi-additive across some, and non-additive cannot be meaningfully summed
- C. Additive measures add rows while non-additive do not
- D. They are different data types for measures

Answer: B

Q1017. What is a role-playing dimension?

- A. A dimension used in role-based access control
- B. A single physical dimension table referenced multiple times in a fact table with different meanings, like date used as order date and ship date
- C. A dimension for storing user roles
- D. A dimension that changes behavior

Answer: B

Q1018. What is a galaxy schema (fact constellation)?

- A. A schema based on astronomical data
- B. A data warehouse schema containing multiple fact tables that share common dimension tables
- C. A schema with star-shaped dimensions
- D. A schema used only in space databases

Answer: B

Q1019. What is ELT and how does it differ from ETL?

- A. They are the same process
- B. ELT loads raw data into the target first then transforms it there, while ETL transforms data before loading
- C. ELT is faster because it skips extraction
- D. ETL loads data while ELT extracts it

Answer: B

Q1020. What is a dimension hierarchy in a data warehouse?

- A. The importance ranking of dimensions
- B. A logical structure within a dimension that defines levels of aggregation, such as day to month to quarter to year
- C. A hierarchy of dimension tables
- D. The order in which dimensions are created

Answer: B

Databases — Hard

510 questions

Q1021. In the ANSI/SPARC three-schema architecture, which mapping exists between the conceptual and internal schemas?

- A. Application/Internal mapping
- B. External/Conceptual mapping
- C. Conceptual/Internal mapping
- D. External/Internal mapping

Answer: C

Q1022. Which of the following best describes the impedance mismatch problem?

- A. Differences between database model and programming language data types
- B. Mismatch between primary key constraints and foreign key references
- C. Slow network connectivity between the client and the database server
- D. Hardware incompatibility between different types of storage devices

Answer: A

Q1023. Self-describing nature of a database system means:

- A. The system contains metadata along with the data itself
- B. Users do not need to interact with it in any way
- C. The database automatically generates its own queries
- D. Data stored in the system is automatically encrypted

Answer: A

Q1024. Program-data independence is primarily achieved through:

- A. Avoiding the use of all index structures
- B. Using only one single programming language
- C. Data abstraction and catalog separation
- D. Upgrading to faster processing hardware

Answer: C

Q1025. Which of the following is a consequence of violating data independence?

- A. Concurrent transaction handling is improved across the system
- B. Data security and access controls are substantially improved
- C. The system achieves significantly faster query processing speed
- D. Application programs must be modified when data storage changes

Answer: D

Q1026. The database approach provides controlled redundancy by:

- A. Encrypting all data stored in the system
- B. Avoiding the use of keys or constraints
- C. Integrating data and applying normalization
- D. Storing data in multiple separate files

Answer: C

Q1027. Which type of DBMS supports complex data types like spatial and multimedia data?

- A. Flat file based DBMS
- B. Network model DBMS
- C. Hierarchical tree DBMS
- D. Object-Relational DBMS

Answer: D

Q1028. In a multi-user DBMS environment, the concurrency control subsystem is essential to:

- A. Perform scheduled database backup operations to external storage
- B. Ensure correct results when multiple users access data simultaneously
- C. Speed up the processing of single-user queries on a local database
- D. Manage the allocation of available disk space across the system

Answer: B

Q1029. Which of the following database utilities is used to reorganize data for better performance?

- A. Statistical analysis utility
- B. Backup restore utility
- C. Reorganization utility
- D. Data loading utility

Answer: C

Q1030. Data abstraction in DBMS hides complexity at different levels. Which sequence is correct from lowest to highest?

- A. External !' Conceptual !' Internal
- B. Conceptual !' Internal !' External
- C. External !' Internal !' Conceptual
- D. Internal !' Conceptual !' External

Answer: D

Q1031. In a shared-disk parallel architecture, the potential bottleneck is:

- A. The interconnection network to the shared disk
- B. The CPU processing speed of each processor
- C. Query parsing and syntax validation steps
- D. Memory allocation among running transactions

Answer: A

Q1032. The query optimizer in a DBMS uses which approach to find the best execution plan?

- A. Cost-based optimization
- B. Random selection method
- C. User preference settings
- D. Alphabetical table ordering

Answer: A

Q1033. Which type of parallelism involves executing different operations of a query simultaneously?

- A. Inter-query parallelism
- B. None of these types
- C. Intra-query parallelism
- D. Pipeline parallelism

Answer: A

Q1034. In a DBMS architecture, the catalog manager is responsible for:

- A. Managing user connections and authentication sessions
- B. Handling network protocols and data transmission tasks
- C. Processing application logic and business rule validation
- D. Storing and retrieving schema information and statistics

Answer: D

Q1035. The concept of data independence is directly supported by:

- A. Using only a single schema without any additional views
- B. Storing all data exclusively in flat file format on disk
- C. Avoiding the use of views or external schema definitions
- D. The three-level architecture with mappings between levels

Answer: D

Q1036. In a DBMS, the lock manager is part of which subsystem?

- A. Storage management subsystem
- B. Communication layer subsystem
- C. Concurrency control subsystem
- D. Query processing subsystem

Answer: C

Q1037. What is the main advantage of shared-memory parallel architecture?

- A. Lower hardware costs compared to other parallel architectures
- B. Fast communication between processors through shared memory
- C. No synchronization needed between the concurrent processors
- D. No need for an operating system to coordinate the processes

Answer: B

Q1038. The DBMS buffer replacement strategy LRU stands for:

- A. Least Recently Used
- B. Last Record Updated
- C. Latest Read Utility
- D. Least Required Update

Answer: A

Q1039. Which DBMS component is responsible for enforcing authorization and integrity constraints?

- A. Network manager and router module
- B. Query parser and syntax analyzer
- C. Storage manager and file handler
- D. Authorization and integrity manager

Answer: D

Q1040. In a multi-tier architecture, which tier contains the business rules?

- A. Data tier (bottom tier)
- B. Application tier (middle tier)
- C. Client tier (front tier)
- D. Presentation tier (top tier)

Answer: B

Q1041. In the EER model, an overlapping specialization means:

- A. An entity can belong to multiple subclasses
- B. No entity can be specialized at all
- C. An entity belongs to exactly one subclass
- D. Subclasses have no common superclass type

Answer: A

Q1042. A category (union type) in the EER model represents:

- A. A superclass entity that has no subclasses
- B. An entity defined without any attributes
- C. A relationship with no participants at all
- D. A subclass with multiple possible superclasses

Answer: D

Q1043. The concept of aggregation in the ER model is used to:

- A. Treat a relationship set as a higher-level entity
- B. Create weak entities that depend on other types
- C. Delete entities from the conceptual data model
- D. Merge two databases into a single schema model

Answer: A

Q1044. In a ternary relationship among three entities A, B, and C, the cardinality constraint is specified:

- A. Only for the strongest entity participating here
- B. Only for one entity in the ternary relationship
- C. Cardinality does not apply to ternary relationships
- D. For each pair of entities with respect to the third

Answer: D

Q1045. Disjoint constraint (d) on specialization means:

- A. An entity can be a member of at most one subclass
- B. Subclasses can overlap with shared memberships
- C. There is no constraint on subclass membership
- D. An entity must belong to all subclasses at once

Answer: A

Q1046. Attribute inheritance in EER means:

- A. Attributes are shared across all entities uniformly
- B. Subclass entities inherit all attributes of the superclass
- C. Attributes cannot be inherited between entity types
- D. Superclass inherits all attributes from the subclass

Answer: B

Q1047. A lattice in the EER model occurs when:

- A. An entity has no attributes or properties defined on it
- B. A superclass has no subclasses defined in the model at all
- C. A subclass has more than one superclass (multiple inheritance)
- D. All entities in the model are weak entity types entirely

Answer: C

Q1048. The Object-Relational model extends the relational model by adding:

- A. Hierarchical parent-child relationships between records
- B. Network-style pointers linking records across the model
- C. Only more tables and relations without other extensions
- D. Support for complex data types, methods, and inheritance

Answer: D

Q1049. In the Entity-Relationship model, the participation constraint and cardinality ratio together define:

- A. Query optimization parameters used
- B. Physical storage structure of tables
- C. Only the key constraints of entities
- D. Structural constraints of a relationship

Answer: D

Q1050. A recursive relationship is one where:

- A. No entity type participates in the relationship that is defined
- B. The same entity type participates more than once in different roles
- C. Only weak entity types participate in the relationship instance
- D. Two different entity types participate together in one relationship

Answer: B

Q1051. Codd's Rule 0 (Foundation Rule) states:

- A. SQL is the only query language allowed for use within the database system
- B. Indexes are mandatory for every attribute present in all database tables
- C. All data must be stored exclusively in flat files without any table structure
- D. A relational DBMS must manage data entirely through its relational capabilities

Answer: D

Q1052. The information rule (Codd's Rule 1) states that:

- A. Pointers can be used to represent data relationships
- B. All information must be represented as values in tables
- C. Metadata need not be stored in tables or relations
- D. Data can be stored in external files outside the system

Answer: B

Q1053. A relation is in BCNF if:

- A. For every non-trivial functional dependency $X \rightarrow Y$, X is a superkey
- B. All attributes in the relation are prime (part of some key)
- C. It has no multivalued dependencies among any of its attributes
- D. It satisfies only the requirements of first normal form (1NF)

Answer: A

Q1054. The closure of a set of attributes X (X^+) under a set of functional dependencies is:

- A. The set of all foreign keys that reference attribute X
- B. The set of all attributes functionally determined by X
- C. The total number of individual attributes within X
- D. The set of all tables in the database containing X

Answer: B

Q1055. Armstrong's axioms include all EXCEPT:

- A. Augmentation
- B. Reflexivity
- C. Distribution
- D. Transitivity

Answer: C

Q1056. A lossless-join decomposition guarantees that:

- A. The original relation can be reconstructed by joining the decomposed relations
- B. No foreign keys are ever needed between the resulting decomposed relations
- C. Join operations are always faster when applied to the decomposed relations
- D. Some data will always be permanently lost during the decomposition process

Answer: A

Q1057. The dependency preservation property ensures that:

- A. Only primary key dependencies are preserved during the decomposition of tables
- B. All functional dependencies can be checked within individual decomposed relations
- C. No dependencies are maintained after decomposition of the original relation
- D. Dependencies are stored separately from the decomposed relation structures

Answer: B

Q1058. A surrogate key is:

- A. A system-generated unique identifier with no business meaning
- B. A natural key derived directly from the business domain
- C. A composite key made up of multiple business attributes
- D. A foreign key referencing a primary key in another table

Answer: A

Q1059. Codd defined how many rules for a fully relational database?

- A. 12 rules (numbered 0-12, so 13 total)
- B. 10 rules (numbered 1 through 10 only)
- C. 5 rules (covering the basic properties)
- D. 20 rules (for comprehensive coverage)

Answer: A

Q1060. The Closed World Assumption in relational databases means:

- A. Facts not present in the database are considered false
- B. Nothing stored in the database is considered to be true
- C. All possible facts about the domain are stored in it
- D. The database cannot be modified once it has been created

Answer: A

Q1061. A query language is relationally complete if it can express:

- A. All queries expressible in relational calculus
- B. Only join operations between related tables
- C. Only aggregate functions on grouped data
- D. Only selection and projection operations used

Answer: A

Q1062. The aggregate operation in extended relational algebra:

- A. Performs Cartesian products between two unrelated relations
- B. Renames relations and their associated attributes in queries
- C. Selects individual tuples based on a simple filter condition
- D. Computes functions like SUM, AVG, COUNT on groups of tuples

Answer: D

Q1063. Which of the following cannot be expressed in basic relational algebra?

- A. Projection of attributes from relations
- B. Selection with various filter conditions
- C. Recursive queries like transitive closure
- D. Natural join of two compatible relations

Answer: C

Q1064. The concept of safe expressions in relational calculus ensures:

- A. Queries always run faster than other approaches
- B. No NULL values appear in the query results
- C. All tuples in the relation are always selected
- D. The result of a query is a finite set of tuples

Answer: D

Q1065. In tuple relational calculus, the expression $\{t \mid P(t)\}$ means:

- A. Update tuples in the relation where P is a false value
- B. Delete all tuples from the relation where P is true value
- C. The set of all tuples t such that predicate P(t) is true
- D. Insert tuples into the relation that satisfy predicate P

Answer: C

Q1066. The outer union operation:

- A. Only works with relations that have identical schemas and matching attributes
- B. Is the same as regular union and produces exactly identical results every time
- C. Takes the union of relations that are not union-compatible by padding with NULLs
- D. Removes all NULL values from the result set after combining both relations

Answer: C

Q1067. Full outer join preserves:

- A. Only tuples from the left relation with NULL for non-matching rows
- B. Only tuples from the right relation with NULL for non-matching rows
- C. All tuples from both relations, padding with NULLs where necessary
- D. Only tuples that have matching values in both of the joined tables

Answer: C

Q1068. The multiset (bag) version of relational algebra differs from set-based algebra in that:

- A. It removes all duplicates
- B. It allows duplicate tuples
- C. It does not support joins
- D. It only supports selection

Answer: B

Q1069. Which algebraic equivalence rule is used for query optimization?

- A. Removing all joins from the query plan
- B. Pushing selections down the query tree
- C. Eliminating projections from the query
- D. Reversing the order of query operations

Answer: B

Q1070. The assignment operation in relational algebra is used to:

- A. Delete relations permanently from the database schema and storage system
- B. Assign the result of a relational algebra expression to a temporary variable
- C. Modify existing tuples by updating their attribute values in the table
- D. Create new databases and define their schema structure from scratch

Answer: B

Q1071. A correlated subquery is different from a regular subquery because it:

- A. References columns from the outer query and executes once per outer row
- B. Does not reference the outer query and runs completely independently
- C. Is always faster than any other type of subquery in all situations
- D. Executes only once regardless of how many rows the outer query returns

Answer: A

Q1072. The COALESCE function in SQL:

- A. Calculates the average of numeric values
- B. Returns the last value in a specified list
- C. Concatenates strings from multiple columns
- D. Returns the first non-NULL value in a list

Answer: D

Q1073. What is the difference between DELETE and TRUNCATE?

- A. TRUNCATE can have a WHERE clause to filter which specific rows should be removed from tables
- B. DELETE is faster than TRUNCATE because it removes all rows without any logging of the changes
- C. They are identical in function and there is no difference between them in any database system
- D. DELETE can have a WHERE clause and is logged; TRUNCATE removes all rows and is minimally logged

Answer: D

Q1074. The CASE expression in SQL is used for:

- A. Managing concurrent transactions
- B. Dropping indexes from a table
- C. Creating tables in the database
- D. Conditional logic within a query

Answer: D

Q1075. A CROSS JOIN produces:

- A. An empty result set with no rows
- B. The intersection of two data tables
- C. Only rows matching between tables
- D. The Cartesian product of two tables

Answer: D

Q1076. Which SQL statement is used to create an index?

- A. CREATE INDEX
- B. ADD INDEX
- C. INSERT INDEX
- D. MAKE INDEX

Answer: A

Q1077. The EXISTS operator in SQL:

- A. Returns TRUE if the subquery returns at least one row
- B. Checks for NULL values in a specified column only
- C. Returns the count of rows matching in the subquery
- D. Creates a new table based on the subquery results

Answer: A

Q1078. A self-join is:

- A. A cross join of two tables
- B. An outer join between tables
- C. A join of a table with itself
- D. A join without any conditions

Answer: C

Q1079. UNION vs UNION ALL in SQL:

- A. They are identical in behavior and produce exactly the same result set
- B. UNION ALL removes duplicates while UNION keeps all including duplicates
- C. UNION removes duplicates; UNION ALL keeps all rows including duplicates
- D. UNION is faster because it skips the deduplication processing entirely

Answer: C

Q1080. What does the NVL/IFNULL/ISNULL function do?

- A. Replaces NULL with a specified value
- B. Creates NULL constraints on columns
- C. Deletes all rows containing NULLs
- D. Counts NULL values in the column

Answer: A

Q1081. The DENSE_RANK() function differs from RANK() in that:

- A. They are identical in their ranking behavior
- B. DENSE_RANK() only works with string values
- C. DENSE_RANK() skips ranks after ties occur
- D. DENSE_RANK() does not skip ranks after ties

Answer: D

Q1082. Lateral joins (LATERAL keyword) allow:

- A. The right side of the join to reference columns from the left side
- B. Only cross joins to be performed without any join filter condition
- C. Joining tables without specifying any condition between the rows
- D. Only inner joins to be used between the two tables in the query

Answer: A

Q1083. The PARTITION BY clause in window functions:

- A. Creates new tables from partitioned subsets of the original data
- B. Divides the result set into partitions for function calculation
- C. Deletes rows from specific partitions based on filter conditions
- D. Physically partitions a table into separate storage file segments

Answer: B

Q1084. Dynamic SQL refers to:

- A. SQL without variables or parameters of any kind
- B. SQL statements constructed and executed at runtime
- C. SQL that runs only during the compilation phase
- D. Static predefined queries that never change at all

Answer: B

Q1085. The PIVOT operation in SQL:

- A. Joins two tables together
- B. Transforms rows into columns
- C. Transforms columns into rows
- D. Creates an index on data

Answer: B

Q1086. SQL injection can be prevented by:

- A. Using parameterized queries and prepared statements
- B. Storing user passwords in plain text in database
- C. Disabling authentication on the database server
- D. Using dynamic SQL without any input validation

Answer: A

Q1087. A user-defined function (UDF) in SQL differs from a stored procedure in that:

- A. A UDF must return a value and can be used in SELECT statements
- B. A stored procedure can be used directly within SELECT queries
- C. A UDF cannot return any values to the calling SQL statement
- D. They are identical in every way with no functional difference

Answer: A

Q1088. The GROUPING SETS clause in SQL:

- A. Groups all rows into one single group regardless of the criteria
- B. Removes grouping and returns individual rows from the result set
- C. Allows specifying multiple grouping combinations in a single query
- D. Sorts the result set based on the specified column or expression

Answer: C

Q1089. ROLLUP and CUBE in SQL GROUP BY are used for:

- A. Generating subtotals and grand totals across multiple dimensions
- B. Sorting data within each group by a specified column expression
- C. Creating indexes on columns that are used in GROUP BY clauses
- D. Deleting grouped data from the tables based on aggregate values

Answer: A

Q1090. Temporal tables in SQL track:

- A. Network latency and response metrics
- B. Only the current state of stored data
- C. Deleted tables and their old schemas
- D. Historical changes to data over time

Answer: D

Q1091. Fifth Normal Form (5NF) eliminates:

- A. Join dependencies that are not implied by candidate keys
- B. Multivalued dependencies among the stored data attributes
- C. Functional dependencies between attributes in the table
- D. All types of dependencies regardless of their categories

Answer: A

Q1092. A relation is in Domain-Key Normal Form (DKNF) if:

- A. Every constraint is a logical consequence of domain constraints and key constraints
- B. It satisfies only the requirements of first normal form and nothing beyond that
- C. All attributes in the relation must serve as keys for unique identification
- D. It has no constraints of any type defined on its attributes or relationships

Answer: A

Q1093. The synthesis algorithm for 3NF decomposition:

- A. Always produces a BCNF decomposition that preserves all functional dependencies
- B. Uses the canonical cover to create a dependency-preserving, lossless decomposition
- C. Only works for binary relations with exactly two attributes in the schema
- D. Removes all dependencies from the relation during the decomposition procedure

Answer: B

Q1094. It is possible to have a decomposition that is lossless but NOT dependency-preserving in:

- A. BCNF decomposition
- B. 1NF decomposition
- C. Any decomposition
- D. No decomposition

Answer: A

Q1095. The chase algorithm is used to:

- A. Create indexes on frequently queried columns
- B. Optimize storage allocation across disk pages
- C. Test whether a decomposition is lossless
- D. Test whether a query is syntactically correct

Answer: C

Q1096. A trivial functional dependency is one where:

- A. The right side is a subset of the left side
- B. The left side is empty with no attributes
- C. No attributes are involved in the FD
- D. Both sides are identical key attributes

Answer: A

Q1097. An extraneous attribute in an FD $X \twoheadrightarrow Y$ is an attribute that:

- A. Can be removed from X or Y without changing the closure
- B. Cannot be removed under any circumstance from the FD
- C. Must always be present in the functional dependency rule
- D. Is a primary key attribute that cannot ever be removed

Answer: A

Q1098. Normalization theory is based on which mathematical concept?

- A. Functional dependencies and their properties
- B. Probability theory and statistical methods
- C. Set theory only without other foundations
- D. Graph theory only without other foundations

Answer: A

Q1099. When BCNF decomposition is not dependency-preserving, the alternative is:

- A. Go to 4NF which eliminates all multivalued dependencies from the relation
- B. Accept 3NF which guarantees both lossless join and dependency preservation
- C. Denormalize completely to avoid any decomposition of the original table
- D. Use 1NF only which requires just atomic values in every attribute field

Answer: B

Q1100. A fully functionally dependent attribute means:

- A. Removing any attribute from the left side of the FD destroys the dependency
- B. The attribute has no dependencies of any kind on other attributes at all
- C. The attribute is part of the key and contributes to unique identification
- D. The attribute depends on a single column only in the current relation schema

Answer: A

Q1101. View serializability is:

- A. Not related to serializability in any way
- B. Less restrictive than conflict serializability
- C. The same as conflict serializability criteria
- D. More restrictive than conflict serializability

Answer: B

Q1102. A recoverable schedule ensures that:

- A. A transaction commits only after all transactions whose data it read have committed
- B. All transactions must run serially with no concurrent execution allowed at all
- C. Transactions can commit in any order regardless of their data reading dependencies
- D. Rollbacks are never needed because all transactions always complete successfully

Answer: A

Q1103. A cascadeless schedule is one where:

- A. All transactions must abort immediately
- B. Transactions only read committed data
- C. Transactions can read uncommitted data
- D. No reads are allowed in any transaction

Answer: B

Q1104. A strict schedule requires that:

- A. No transaction reads or writes data written by an uncommitted transaction
- B. Transactions can freely read uncommitted data from other transactions
- C. Commits are not required for transactions to be considered complete
- D. Only serial schedules are allowed for all transactions in the system

Answer: A

Q1105. The precedence graph for testing conflict serializability has:

- A. Nodes for SQL statements and edges for their execution ordering
- B. Nodes for transactions and edges for conflicting operation pairs
- C. Nodes for data items and edges for transactions accessing them
- D. No nodes or edges are used in this graph representation at all

Answer: B

Q1106. Testing view serializability is:

- A. NP-complete
- B. Linear time
- C. Logarithmic time
- D. Constant time

Answer: A

Q1107. The READ COMMITTED isolation level prevents:

- A. No anomalies of any kind
- B. All anomalies including phantoms
- C. Only dirty reads alone
- D. Dirty and non-repeatable reads

Answer: C

Q1108. The REPEATABLE READ isolation level prevents:

- A. All anomalies including phantom reads
- B. Dirty reads and non-repeatable reads
- C. Only dirty reads from transactions
- D. No anomalies of any kind at all

Answer: B

Q1109. In a multi-version concurrency control (MVCC) system:

- A. Transactions must be serial with no concurrent execution allowed at all
- B. Multiple versions of data items are maintained for concurrent access
- C. Only one version of each data item exists at any given point in time
- D. Locks are always required for every read and write operation performed

Answer: B

Q1110. A blind write is:

- A. A committed transaction that has been successfully applied
- B. A query optimization technique that reduces disk access
- C. A read operation that is not followed by any write at all
- D. A write operation without reading the current value first

Answer: D

Q1111. Multiple Granularity Locking uses intention locks. An Intention Shared (IS) lock means:

- A. No lower-level locks exist on any of the descendant nodes in the tree
- B. The item cannot be locked by any transaction under any circumstances
- C. A transaction intends to acquire shared locks on finer-granularity items
- D. A transaction has exclusively locked the entire item and all subitems

Answer: C

Q1112. An Intention Exclusive (IX) lock means:

- A. A transaction currently has a shared lock on the item and its subitems
- B. No locks are needed for any data items under this locking protocol
- C. A transaction intends to acquire exclusive locks on finer-granularity items
- D. The transaction is aborted because it cannot acquire the needed lock

Answer: C

Q1113. A Shared Intention Exclusive (SIX) lock means:

- A. No intention locks exist on any node in the hierarchy and all nodes are freely open
- B. The node is shared-locked and the transaction intends to exclusively lock finer items
- C. The node is exclusively locked by the transaction and no other can access it at all
- D. The transaction is complete and all locks have been released from every node here

Answer: B

Q1114. In the Thomas Write Rule (modification to timestamp ordering):

- A. Obsolete writes are ignored instead of causing rollback
- B. All writes cause rollback regardless of timestamp order
- C. Reads are ignored and only writes are processed by it
- D. Timestamps are not used in this modified protocol

Answer: A

Q1115. The phantom problem in locking can be solved by:

- A. Allowing dirty reads between queries
- B. Using only row-level locks in tables
- C. Index locking or predicate locking
- D. Removing all indexes from the tables

Answer: C

Q1116. In multi-version timestamp ordering (MVTO):

- A. Only one version of each data item is maintained
- B. Reads always block writes until they are completed
- C. No timestamps are used for ordering transactions
- D. Each write creates a new version of the data item

Answer: D

Q1117. The wound-wait scheme is a preemptive protocol because:

- A. Younger transactions always preempt older ones holding locks on data items
- B. No rollback ever occurs because all transactions complete without conflict
- C. An older transaction can force a younger one holding a needed lock to roll back
- D. Only reads are preempted while writes always proceed without interruption

Answer: C

Q1118. In tree locking protocol (for B+ tree indexes), a transaction must:

- A. Use no locks at all when accessing data through the index structure
- B. Lock children before parents when traversing the index tree structure
- C. Lock a node before locking its children, and unlock parent before child
- D. Lock all nodes simultaneously before accessing any data in the tree

Answer: C

Q1119. Lock escalation is the process of:

- A. Converting coarse-grained locks into many more fine-grained locks
- B. Converting many fine-grained locks into fewer coarse-grained locks
- C. Adding locks continuously without ever releasing any of them
- D. Removing all locks from the data items in the locked collection

Answer: B

Q1120. Snapshot isolation provides each transaction with:

- A. Only read access to the data without any ability to write modifications
- B. No isolation at all between concurrently executing database transactions
- C. Access to the latest uncommitted data written by other active transactions
- D. A consistent snapshot of the database as of the transaction's start time

Answer: D

Q1121. In ARIES, the analysis phase determines:

- A. The set of dirty pages and active transactions at crash time
- B. Only committed transactions that completed before the crash
- C. The hardware failure cause and the affected disk components
- D. Only the last checkpoint that was recorded in the log file

Answer: A

Q1122. In ARIES, the redo phase:

- A. Replays all logged changes from the earliest dirty page LSN forward
- B. Starts from the end of the log and processes records in reverse
- C. Only replays committed transactions and ignores uncommitted ones
- D. Undoes all transactions that were active at the time of the crash

Answer: A

Q1123. In ARIES, Compensation Log Records (CLRs) are:

- A. Checkpoint records that mark points of database consistency
- B. Log records for committed transactions only in the system
- C. Log records written during undo to record the undo actions
- D. Index maintenance records for B-tree structure updates

Answer: C

Q1124. The steal/no-force buffer management policy means:

- A. Uncommitted pages can be written to disk (steal); committed need not be forced immediately (no-force)
- B. Committed pages are always forced to disk immediately; uncommitted pages are never written to disk
- C. Only committed pages are present on disk; uncommitted data remains exclusively in the memory buffer
- D. All pages are always kept in memory; nothing is written to disk until the entire database shuts down

Answer: A

Q1125. No-steal/force policy would mean:

- A. Uncommitted pages are never written to disk, and committed pages are always forced to disk
- B. Neither undo nor redo is needed with this policy since both cases are eliminated
- C. Uncommitted pages can be written to disk at any time before the transaction commits
- D. Committed pages may not be on disk immediately and can remain only in memory buffer

Answer: A

Q1126. The redoLSN in ARIES is:

- A. The checkpoint LSN that marks the most recent consistency checkpoint
- B. The last LSN that was recorded in the transaction log before the crash
- C. The LSN from which redo must start (earliest recLSN among dirty pages)
- D. The LSN of the last commit record written to the stable storage log

Answer: C

Q1127. Write-ahead logging requires that log records be written to stable storage before:

- A. The transaction starts its initial read operation
- B. Any query is executed against the database tables
- C. The corresponding data page is flushed to disk
- D. The user submits the query to the database system

Answer: C

Q1128. In ARIES, when a transaction is undone, the undo follows:

- A. The commit order of other transactions in the same schedule
- B. The table creation order defined in the database schema file
- C. The lastLSN chain backward through that transaction's log records
- D. A random order without any specific sequence or priority rules

Answer: C

Q1129. Remote backup systems for disaster recovery provide:

- A. A copy of the database at a geographically separate site
- B. Better index structures for improving search performance
- C. Improved normalization of tables to reduce data redundancy
- D. Faster query processing through parallel execution plans

Answer: A

Q1130. The concept of idempotent redo means:

- A. Applying redo multiple times produces the same result as applying it once
- B. Redo can only be applied once and cannot be repeated after that point
- C. Redo changes the result each time it is applied to the database state
- D. Redo is never repeated because recovery always succeeds on first try

Answer: A

Q1131. In a B+ tree, the minimum occupancy of internal nodes (non-root) is:

- A. All pointers must be full
- B. One pointer (child) as minimum
- C. Exactly n pointers at all times
- D. Ceiling of $n/2$ pointers (children)

Answer: D

Q1132. The global depth in extendible hashing indicates:

- A. The total number of records stored across all the buckets
- B. The number of buckets currently allocated in the structure
- C. The number of bits used to determine the directory entry
- D. The number of hash functions used for the key computation

Answer: C

Q1133. Local depth in extendible hashing indicates:

- A. The depth of the B+ tree from the root node to the leaf node level
- B. The total number of collisions that have occurred in the hash table
- C. The number of bits used to determine which bucket a record belongs to
- D. The size of the directory that maps hash values to bucket addresses

Answer: C

Q1134. A hash join in query processing works by:

- A. Partitioning both relations using the same hash function then joining matching partitions
- B. Using an index on the join attribute to look up matching tuples from the inner relation
- C. Sorting both relations on the join attribute first and then merging the sorted result sets
- D. Using nested loops where each tuple of one relation is compared with every tuple of other

Answer: A

Q1135. R-trees are used to index:

- A. Text data and documents only
- B. Temporal data and timestamps
- C. One-dimensional key values only
- D. Spatial (multidimensional) data

Answer: D

Q1136. A function-based index is created on:

- A. A user role that defines access permissions and grants
- B. A database name used to identify the schema information
- C. A table name stored in the database system catalog data
- D. The result of a function applied to one or more columns

Answer: D

Q1137. An inverted index is commonly used in:

- A. Foreign key enforcement
- B. Numerical computations only
- C. Primary key lookups only
- D. Full-text search engines

Answer: D

Q1138. The order of a B+ tree affects performance because:

- A. A higher order means more keys per node, reducing tree height and disk I/O
- B. A higher order always increases the tree height and requires more disk access
- C. Lower order is always better because smaller nodes fit in memory more easily
- D. Order has no effect on performance and does not change disk I/O requirements

Answer: A

Q1139. Bulk loading a B+ tree is more efficient than individual insertions because:

- A. It creates an unbalanced tree that must be rebalanced later
- B. It sorts data first and builds the tree bottom-up, minimizing splits
- C. It skips creating the tree entirely and uses a flat file instead
- D. It uses no sorting and inserts records in their original order

Answer: B

Q1140. A skip list can be used as an alternative to B+ trees for indexing because:

- A. It provides $O(\log n)$ average search, insert, and delete with simpler implementation
- B. It uses less memory than any other data structure used for indexing purposes
- C. It is always faster than B+ trees for every type of database query and operation
- D. It requires no randomization and is fully deterministic in its performance

Answer: A

Q1141. The Selinger optimizer (System R) uses:

- A. Only heuristic rules without any cost estimation being performed
- B. Random plan selection from all possible execution plan choices
- C. Dynamic programming with interesting orders to find optimal plans
- D. No optimization at all with queries running as they are written

Answer: C

Q1142. An interesting sort order in query optimization is:

- A. Any random sort order applied to tuples without considering the query operations
- B. A sort order used only for displaying the final result to the end user query
- C. A sort order that increases cost by requiring additional sorting in later steps
- D. A sort order produced by one operation that is useful for a subsequent operation

Answer: D

Q1143. The number of possible join orders for n relations is:

- A. Grows exponentially (Catalan number)
- B. Always n^2 for n relations
- C. Always n for n given relations total
- D. Exactly $n!$ for n relations total

Answer: A

Q1144. In cost-based optimization, the estimated cost of a nested loop join $R \bowtie S$ is approximately:

- A. Always 1 block access regardless of the sizes of relations R and S used
- B. b_R plus b_S (sum of blocks from both R and S relations) in simple analysis
- C. n_R plus n_S (sum of tuples from both R and S relations) in simple analysis
- D. n_R times b_S plus b_R (b is blocks, n_R is tuples in R) in simple analysis

Answer: D

Q1145. Histogram statistics improve query optimization by:

- A. Increasing query complexity by adding more steps to the optimization process
- B. Providing more accurate selectivity estimates for non-uniform data distributions
- C. Removing the need for indexes by providing direct access to all stored records
- D. Making all estimates uniform regardless of the actual data value distribution

Answer: B

Q1146. A left-deep join tree is preferred by many optimizers because:

- A. It is the only valid tree shape for all joins
- B. It avoids all disk I/O during the join process
- C. It minimizes the number of tables in joins
- D. It allows pipelining of intermediate results

Answer: D

Q1147. Adaptive query processing differs from traditional optimization by:

- A. Only optimizing before execution begins and then running without changes
- B. Adjusting the execution plan during query execution based on runtime statistics
- C. Never changing the plan once it has been selected by the query optimizer
- D. Ignoring all statistics and using a fixed plan for every query submitted

Answer: B

Q1148. The concept of equivalent query expressions means:

- A. Identical query text that is written exactly the same way in SQL
- B. Different relational algebra expressions that produce the same result
- C. Queries that produce different results from the same input relations
- D. Queries executed on different databases with different table schemas

Answer: B

Q1149. Predicate pushdown in distributed queries:

- A. Ignores remote data sources completely and only queries local tables
- B. Sends all data to a central location first before any filtering occurs
- C. Sends filter conditions to remote data sources to reduce data transfer
- D. Removes all predicates from the query and returns all unfiltered data

Answer: C

Q1150. Columnar storage benefits analytical queries because:

- A. No indexes are needed for any query on columnar data store
- B. Row storage is used instead of column-oriented data format
- C. All columns are always read regardless of the query needs
- D. Only the needed columns are read from disk, reducing I/O

Answer: D

Q1151. Statistical database security addresses the problem of:

- A. Managing user roles and their assigned privileges set
- B. Creating backup copies of the database for recovery
- C. Inferring individual data from aggregate query results
- D. Encrypting data at rest on the physical storage media

Answer: C

Q1152. The inference problem in database security occurs when:

- A. No queries are allowed to be executed against the tables
- B. Users deduce restricted information from permitted queries
- C. Users have direct database access without any restrictions
- D. All data is encrypted and cannot be read without the key

Answer: B

Q1153. Label-based security in databases assigns:

- A. Sensitivity labels to data rows and clearance levels to users
- B. Labels only to tables without any row-level classification
- C. No labels to any data regardless of its sensitivity or nature
- D. Random access permissions to all users without any restrictions

Answer: A

Q1154. Homomorphic encryption allows:

- A. Faster decryption of data using optimized algorithms
- B. Only reading encrypted data without any modifications
- C. Computations on encrypted data without decrypting it
- D. Only writing encrypted data without reading it first

Answer: C

Q1155. Database firewall works by:

- A. Encrypting all data stored in the database tables and log files
- B. Only blocking network traffic without inspecting SQL statements
- C. Creating physical firewalls around the database server hardware
- D. Monitoring SQL traffic and blocking unauthorized or suspicious queries

Answer: D

Q1156. The k-anonymity privacy model ensures:

- A. Only k users can access the database at any given time during concurrent operations
- B. Data is stored in k copies across different sites for redundancy and availability
- C. All data is encrypted using a key of length k bits for secure storage and retrieval
- D. Each record is indistinguishable from at least k-1 other records on quasi-identifiers

Answer: D

Q1157. Differential privacy provides:

- A. Full access to all data without any restrictions or privacy controls
- B. Only physical security measures without any mathematical foundations
- C. No privacy guarantees of any kind for the data stored in the database
- D. Mathematical guarantees that query results do not reveal individual data

Answer: D

Q1158. A privilege graph is used to:

- A. Track the flow of granted privileges among users
- B. Design ER diagrams for conceptual data modeling
- C. Optimize query execution plans for better speed
- D. Create indexes on columns for fast data retrieval

Answer: A

Q1159. Cascading revocation in SQL means:

- A. Revoking a privilege also revokes it from all users who received it through the original grant
- B. No revocation occurs and the privilege remains active for all users who currently have it
- C. Only the specified user loses the privilege and no other users are affected by the revocation
- D. All database privileges for every user are revoked regardless of the specific grant involved

Answer: A

Q1160. Column-level encryption differs from TDE in that:

- A. It encrypts specific columns rather than entire database files
- B. It encrypts only at the network level rather than at storage level
- C. It encrypts only table names rather than the actual column data
- D. It is less secure than TDE because it covers fewer data elements

Answer: A

Q1161. The three-phase commit (3PC) protocol improves on 2PC by:

- A. Using only one phase for the commit decision
- B. Removing the prepare phase from the protocol
- C. Adding a pre-commit phase to avoid blocking
- D. Ignoring all failures that occur at any site

Answer: C

Q1162. In distributed databases, the global query optimization must consider:

- A. Only CPU time without considering disk or network access costs
- B. Only disk I/O costs without considering network or CPU overhead
- C. Only local processing costs without considering network overhead
- D. Network communication costs in addition to local processing costs

Answer: D

Q1163. Paxos is a protocol used for:

- A. Achieving consensus in distributed systems
- B. Indexing data for faster retrieval queries
- C. Query optimization in database management
- D. Schema design for database normalization

Answer: A

Q1164. The RAFT consensus algorithm is designed to be:

- A. Faster than all other protocols in every distributed environment
- B. A replacement for SQL in distributed relational database systems
- C. More understandable than Paxos while providing the same guarantees
- D. A type of index for organizing data in distributed hash tables

Answer: C

Q1165. In a federated database system:

- A. No integration is possible between any databases at all
- B. Autonomous databases cooperate to provide integrated access
- C. Only one database exists in the entire distributed system
- D. All databases are identical copies of the same data schema

Answer: B

Q1166. The eventual consistency model guarantees that:

- A. No consistency is provided between any of the replicas in the system
- B. Immediate consistency at all times for every read at every site node
- C. All reads always return the latest written value at every site instantly
- D. All replicas will converge to the same value if no new updates are made

Answer: D

Q1167. The quorum-based protocol requires:

- A. $Q_r = Q_w = 1$ for consistency meaning only one replica needs to respond
- B. Read quorum (Q_r) + Write quorum (Q_w) > total replicas (N) for consistency
- C. No quorum is needed and any single replica can respond to operations
- D. $Q_r + Q_w$ less than N for consistency between read and write quorums

Answer: B

Q1168. Vector clocks in distributed systems are used to:

- A. Capture causal relationships between events at different sites
- B. Measure CPU speed and processing capacity at each network node
- C. Optimize queries by selecting the best execution plan for joins
- D. Tell wall-clock time accurately across all distributed node sites

Answer: A

Q1169. Distributed hash tables (DHTs) are used in:

- A. Peer-to-peer distributed storage systems
- B. Centralized databases on a single machine
- C. Single-machine databases without networks
- D. Only for indexing within one database node

Answer: A

Q1170. The coordinator selection problem in distributed databases can be solved by:

- A. Manual intervention always requiring a human administrator
- B. No algorithm exists for solving the coordinator selection
- C. Random selection only without any formal algorithmic approach
- D. Election algorithms like the Bully algorithm or Ring algorithm

Answer: D

Q1171. The concept of tunable consistency in Cassandra means:

- A. Consistency cannot be changed once the Cassandra cluster has been configured
- B. Only strong consistency is available for all queries in the entire cluster
- C. The consistency level can be configured per query (ONE, QUORUM, ALL, etc.)
- D. Only eventual consistency is possible and no other level can be selected

Answer: C

Q1172. A wide-column store differs from a traditional columnar database in that:

- A. It is identical to a relational database model
- B. All rows must have the same columns defined
- C. Each row can have a different set of columns
- D. It uses only fixed schemas without flexibility

Answer: C

Q1173. CRDTs (Conflict-free Replicated Data Types) are used in distributed NoSQL to:

- A. Automatically resolve conflicts in replicated data without coordination
- B. Compress data to reduce the network transfer size between replica sites
- C. Encrypt data before replicating it across the distributed cluster nodes
- D. Create indexes on replicated data for faster query retrieval operations

Answer: A

Q1174. The gossip protocol in distributed NoSQL systems is used for:

- A. Transaction management for multi-site commit control
- B. Index creation on distributed data across all nodes
- C. Query optimization for distributed join operations
- D. Spreading information about cluster state across nodes

Answer: D

Q1175. NewSQL databases aim to provide:

- A. Only NoSQL features without any ACID guarantee support
- B. No guarantees at all for either consistency or scaling
- C. SQL-like guarantees (ACID) with NoSQL-like scalability
- D. Only traditional SQL features without scalability gains

Answer: C

Q1176. Examples of NewSQL databases include:

- A. Google Spanner, CockroachDB, VoltDB
- B. All of the above are NewSQL types
- C. MongoDB, Redis, Apache Cassandra
- D. MySQL, PostgreSQL, Oracle Database

Answer: A

Q1177. The LSM (Log-Structured Merge) tree used in many NoSQL databases provides:

- A. Only sequential reads without support for random access query patterns
- B. High read throughput only without any optimization for write operations
- C. High write throughput by buffering writes in memory and merging to disk
- D. No performance benefits over traditional B-tree index structures at all

Answer: C

Q1178. Google's Bigtable is:

- A. A traditional relational database with SQL support
- B. A key-value store only without column family data
- C. A distributed, sparse, multi-dimensional sorted map
- D. A graph database using nodes and edge connections

Answer: C

Q1179. Amazon DynamoDB uses:

- A. B+ trees for all data storage and access
- B. Consistent hashing for data partitioning
- C. Only single-node storage without sharding
- D. Relational tables internally for all data

Answer: B

Q1180. Multi-model databases support:

- A. Only relational data organized in standard normalized table format on disk
- B. Only flat files stored sequentially without any structured data model in use
- C. Multiple data models (document, graph, key-value) in a single database engine
- D. Only one data model such as relational without any support for other models

Answer: C

Q1181. A slowly changing dimension (SCD) Type 2 handles changes by:

- A. Overwriting the old value directly in place without keeping any history of changes
- B. Adding a new row with the new value and preserving the old row with history tracking
- C. Deleting the old row entirely and inserting a brand new row with the changed value
- D. Adding a new column to store the updated value alongside the previous column value

Answer: B

Q1182. A degenerate dimension is:

- A. A dimension key in the fact table that has no corresponding dimension table
- B. A dimension that is always empty and contains no rows of data whatsoever
- C. A dimension table with many attributes and a large number of data records
- D. A normalized dimension table that has been split into multiple sub-tables

Answer: A

Q1183. The galaxy schema (fact constellation) contains:

- A. No dimension tables in the schema at all
- B. Only dimension tables without fact tables
- C. Only one fact table with no shared dimensions
- D. Multiple fact tables sharing dimension tables

Answer: D

Q1184. Clustering in data mining groups:

- A. Data by primary key values in ascending sort order
- B. Data based on known categories and predefined classes
- C. Similar data points together without predefined labels
- D. Data randomly without any similarity considerations

Answer: C

Q1185. Decision tree classification in data mining:

- A. Builds a tree structure to classify data based on attribute values
- B. Creates random groupings of data without any logical partitioning
- C. Performs regression only for predicting continuous numeric values
- D. Finds association rules between items in transactional data sets

Answer: A

Q1186. The K-means algorithm is a:

- A. Classification algorithm using labeled training data sets
- B. Association rule algorithm for finding item relationships
- C. Regression algorithm for predicting continuous value data
- D. Clustering algorithm that partitions data into K groups

Answer: D

Q1187. A materialized view in a data warehouse is refreshed:

- A. Only when the database restarts after a scheduled shutdown
- B. Periodically or on demand to reflect source data changes
- C. In real-time always with zero latency from source updates
- D. Never after creation regardless of any source data changes

Answer: B

Q1188. Data lake differs from a data warehouse in that it:

- A. Is always smaller in size than a data warehouse
- B. Requires data transformation before storage occurs
- C. Stores raw, unstructured data in its native format
- D. Only stores structured data in normalized table form

Answer: C

Q1189. The concept of data lineage in data warehousing refers to:

- A. The size of tables measured in rows and storage bytes
- B. Tracking the origin and transformation history of data
- C. The number of queries executed against the data store
- D. The age of the database since its initial creation date

Answer: B

Q1190. Real-time data warehousing aims to:

- A. Minimize the latency between source data changes and warehouse availability
- B. Store data offline in cold storage without any access for active queries
- C. Avoid using ETL processes entirely and load data without transformation
- D. Only process batch data on a periodic schedule without real-time updates

Answer: A

Q1191. In a scenario where multiple applications need different views of employee data, which architecture feature best addresses this?

- A. External schema views tailored to each application need
- B. Conceptual schema duplication across each separate application
- C. Internal schema customization for each application program
- D. Physical storage partitioning for each individual application

Answer: A

Q1192. A university database must track students, courses, and enrollments. Which data independence level is impacted when changing the index structure on the student table?

- A. Physical data independence is impacted directly
- B. Neither level of independence is affected here
- C. Logical data independence is impacted directly
- D. Both logical and physical independence equally

Answer: A

Q1193. When a DBMS enforces referential integrity, what specific problem is being prevented?

- A. Excessive storage consumption from duplicate entries
- B. Unauthorized users gaining access to sensitive records
- C. Slow query performance due to missing index structures
- D. Dangling references to non-existent related records

Answer: D

Q1194. An enterprise migrates from file-based to DBMS. Which challenge is most critical during this transition?

- A. Selecting a new operating system for all server machines
- B. Reducing the total number of end user workstations
- C. Converting existing data and ensuring schema consistency
- D. Upgrading monitor displays for the database administrators

Answer: C

Q1195. A company needs to change the conceptual schema by adding a new attribute to a table without affecting existing applications. This exemplifies:

- A. View-level independence from external schema changes
- B. Instance-level independence from data value updates
- C. Logical data independence from schema modifications
- D. Physical data independence from storage changes

Answer: C

Q1196. In the ANSI/SPARC architecture, which component is responsible for transforming requests between schema levels?

- A. The transaction coordinator
- B. The schema mapping engine
- C. The query optimizer module
- D. The concurrency controller

Answer: B

Q1197. A hospital database requires that patient records cannot be deleted if they have associated billing entries. Which DBMS feature enforces this?

- A. Triggers that log all deletion attempts made
- B. Check constraints on the billing amount field
- C. Referential integrity with restrict delete
- D. Indexing on the patient identifier column

Answer: C

Q1198. When evaluating DBMS overhead versus file systems for a small single-user application, which factor most favors using a file system?

- A. The application must enforce complex integrity constraints
- B. The application needs minimal setup and lower overhead
- C. The application demands sophisticated query optimization
- D. The application requires complex multi-user concurrency

Answer: B

Q1199. A database system uses deferred modification during transaction execution. What does this mean regarding write operations?

- A. Writes are duplicated to two separate disk locations each
- B. Writes go directly to disk before the transaction commits
- C. Writes are encrypted before being stored on the disk drive
- D. Writes are held in buffer until transaction commits first

Answer: D

Q1200. An organization has separate databases for HR, Finance, and Inventory with overlapping employee data. Which integration approach best reduces redundancy?

- A. Creating a centralized database with shared employee records
- B. Allowing each department to independently define employee data
- C. Using flat files to export and import data between departments
- D. Maintaining separate databases with manual data synchronization

Answer: A

Q1201. A company experiences slow queries during peak hours with a two-tier architecture. Which architectural change would best address this?

- A. Increasing the network bandwidth between client and server
- B. Migrating to a three-tier architecture with application server
- C. Replacing the client hardware with more powerful machines
- D. Adding more storage disks to the existing database server

Answer: B

Q1202. In a shared-nothing parallel database, how is data redistribution handled during query joins across partitions?

- A. All data is consolidated to a single node before joining
- B. Joins across partitions are not supported in this setup
- C. Relevant data is shipped between nodes over the network
- D. Data is copied to a shared memory area for join processing

Answer: C

Q1203. A web application handles thousands of concurrent users. Which DBMS architectural feature is most critical for maintaining response times?

- A. Connection pooling and efficient buffer management
- B. Manual memory allocation by the database operator
- C. Verbose logging of every individual query executed
- D. Schema versioning for backward compatibility support

Answer: A

Q1204. When designing a fault-tolerant database architecture, which strategy best ensures continuous availability?

- A. Disabling transaction logging to improve performance speed
- B. Using active-passive replication with automatic failover
- C. Storing all data on a single high-performance server
- D. Relying solely on periodic backup tapes for recovery

Answer: B

Q1205. In a multi-tier architecture, which tier should implement data validation business rules for optimal security?

- A. In the application logic tier on the middle server
- B. Distributed equally across all three architectural tiers
- C. Exclusively in the presentation tier on the client side
- D. Only in the data tier within stored database procedures

Answer: A

Q1206. A distributed database system uses a two-phase commit protocol. What happens if the coordinator fails after sending prepare messages?

- A. The transaction is automatically rolled back at all sites
- B. Each participant independently decides to commit or abort
- C. All participants automatically commit their transactions
- D. Participants remain blocked until the coordinator recovers

Answer: D

Q1207. An organization needs to process analytical queries without affecting transactional performance. Which architectural pattern best achieves this?

- A. Using a single-tier architecture with query prioritization
- B. Implementing separate OLTP and OLAP systems with data feeds
- C. Disabling indexing on the transactional database for speed
- D. Running both workloads on the same database server instance

Answer: B

Q1208. In a cloud-based database architecture, what does the concept of elastic scaling specifically refer to?

- A. Permanently allocating maximum resources for peak loads
- B. Reducing database features to lower computational overhead
- C. Manually adding hardware when usage increases over time
- D. Automatically adjusting resources based on current demand

Answer: D

Q1209. When implementing a microservices architecture, what database pattern avoids tight coupling between services?

- A. Services communicate by directly reading each others data
- B. Each microservice owns its own private database instance
- C. A single transaction spans multiple microservice databases
- D. All microservices share a single centralized database

Answer: B

Q1210. A DBMS uses write-ahead logging in its architecture. What constraint does this impose on buffer management?

- A. Buffer management operates independently of log records
- B. Data pages must be written before any log records at all
- C. Buffer pages can be flushed to disk at any time freely
- D. Log records must be written to disk before data pages

Answer: D

Q1211. A university needs to model a scenario where a student can major in multiple departments, and each department offers multiple majors. Which relationship type best represents this?

- A. Many-to-many between student and department entities
- B. Unary relationship on the student entity type alone
- C. One-to-many from department entity to student entity
- D. One-to-one between student and department entities

Answer: A

Q1212. When converting an ER diagram with a ternary relationship to a relational schema, how many foreign keys does the resulting relation typically contain?

- A. No foreign keys since ternary cannot be directly mapped
- B. Two foreign keys referencing the two main entity types
- C. Three foreign keys referencing each participating entity
- D. One foreign key referencing the primary entity only

Answer: C

Q1213. In the Enhanced ER model, an overlapping specialization constraint means:

- A. Subtypes must collectively include all supertype entities
- B. An entity can belong to at most one subtype at a time
- C. No entity can belong to any of the defined subtype sets
- D. An entity can belong to multiple subtypes simultaneously

Answer: D

Q1214. A hospital system models doctors, patients, and treatments. A doctor treats a patient with a specific treatment. Which ER construct best models this scenario?

- A. An aggregation of doctor-patient with treatment entity
- B. Binary relationship between doctor and patient only
- C. Two separate binary relationships for each pair type
- D. A ternary relationship among all three entity types

Answer: D

Q1215. When modeling a recursive relationship where an employee manages other employees, what special consideration is needed?

- A. The relationship requires two separate entity types always
- B. A weak entity must be created for managed employees here
- C. Role names must distinguish the two participations clearly
- D. Recursive relationships cannot have any cardinality ratios

Answer: C

Q1216. In an EER diagram, the category (union type) construct is used when:

- A. An entity belongs to exactly one of multiple supertype sets
- B. An entity is always independent of all other entity types
- C. An entity has no relationship with any supertype entity set
- D. An entity must belong to all listed supertype entity sets

Answer: A

Q1217. A library system has Books, Authors, and Publishers. A book can have multiple authors, and each author may work with multiple publishers. If we need to track which publisher published which author-specific book, the best modeling approach is:

- A. Two binary relationships with an aggregation construct used
- B. Three binary relationships connecting each pair of entities
- C. A single ternary relationship among Book, Author, Publisher
- D. A single binary relationship between Book and Author only

Answer: C

Q1218. When mapping a total specialization with disjoint constraint to a relational schema, which approach uses the fewest tables?

- A. Create a single table with a type attribute discriminator
- B. Create tables for only the subtypes with inherited attributes
- C. Create tables for supertype and all subtypes separately
- D. Create one table per attribute of the supertype entity type

Answer: B

Q1219. In object-oriented data models, what does the concept of encapsulation specifically provide?

- A. Hiding internal state and exposing only defined methods
- B. Direct access to all internal data attributes of objects
- C. Conversion of objects into flat relational table records
- D. Automatic replication of objects across multiple servers

Answer: A

Q1220. A social network needs to model users who can be friends with each other and also belong to groups. Which combination of ER constructs most accurately captures this?

- A. A ternary relationship among user, friendship, and group
- B. Only binary relationships between user and group entities
- C. A recursive relationship on user plus binary with groups
- D. An aggregation of user entity with a group attribute added

Answer: C

Q1221. Given a relation $R(A, B, C, D)$ with functional dependencies $A \rightarrow B$, $B \rightarrow C$, and $C \rightarrow D$, what is the candidate key?

- A. The attribute set consisting of C only
- B. The attribute set consisting of D only
- C. The attribute set consisting of A only
- D. The attribute set consisting of B only

Answer: C

Q1222. In a relation with attributes $\{A, B, C, D, E\}$ and FDs $\{AB \rightarrow C, C \rightarrow D, D \rightarrow E\}$, what is the closure of $\{A, B\}$?

- A. The closure is the set $\{A, B, D, E\}$ only
- B. The closure is the set $\{A, B, C\}$ only
- C. The closure is the set $\{C, D, E\}$ without A B
- D. The closure is the set $\{A, B, C, D, E\}$

Answer: D

Q1223. A table has a composite primary key (StudentID, CourseID). Which statement about this design is correct?

- A. The combination must be unique but individual parts may repeat
- B. Each attribute alone must uniquely identify every record
- C. Only StudentID needs to be unique across all table records
- D. Neither StudentID nor CourseID can be null independently

Answer: A

Q1224. If a relation has five candidate keys each consisting of two attributes from a set of six, how many superkeys exist at minimum?

- A. Only one superkey exists in this relation total
- B. Exactly five superkeys exist in this relation
- C. At least five but potentially many more superkeys
- D. Exactly ten superkeys exist in this relation

Answer: C

Q1225. A database designer must choose between a natural key (SSN) and a surrogate key (auto-increment ID) for a Person table. Which consideration most favors the surrogate key?

- A. Surrogate keys carry inherent meaning about the entity
- B. Surrogate keys are stable and unaffected by real-world changes
- C. Natural keys provide better query performance in all cases
- D. Natural keys are always shorter than surrogate key values

Answer: B

Q1226. In a relation $R(A, B, C)$ with the functional dependency $AB \rightarrow C$ and no other non-trivial FDs, which normal form is R in?

- A. Second normal form but not third due to transitivity
- B. Third normal form since no transitive dependency exists
- C. Boyce-Codd normal form since AB is a superkey of R
- D. First normal form only since it has partial dependency

Answer: C

Q1227. When designing a relational schema for a many-to-many relationship between Students and Courses with an enrollment date attribute, which approach is correct?

- A. Add enrollment date to the Students table as a new column
- B. Store enrollment date in a separate unrelated lookup table
- C. Create a junction table with enrollment date as attribute
- D. Add enrollment date to the Courses table as a new column

Answer: C

Q1228. If Armstrong axioms are applied to derive $A \rightarrow C$ from $A \rightarrow B$ and $B \rightarrow C$, which specific axiom is used?

- A. Transitivity axiom of functional dependencies
- B. Decomposition axiom of functional dependencies
- C. Reflexivity axiom of functional dependencies
- D. Augmentation axiom of functional dependencies

Answer: A

Q1229. A relation has attributes $\{W, X, Y, Z\}$ with FDs $\{W \rightarrow X, X \rightarrow Y, WY \rightarrow Z\}$. What are all the candidate keys?

- A. The candidate keys are both $\{W\}$ and $\{X\}$ individually
- B. The only candidate key is the single attribute $\{W\}$
- C. The only candidate key is the attribute set $\{W, Y\}$
- D. The only candidate key is the attribute set $\{W, X\}$

Answer: B

Q1230. In a relational database, what specific problem does the phantom read anomaly cause in concurrent transactions?

- A. A transaction reads data that was already committed earlier
- B. A transaction sees new rows that match its previous query range
- C. A transaction cannot read any rows due to an exclusive lock
- D. A transaction overwrites changes made by another transaction

Answer: B

Q1231. Given relations $R(A,B)$ and $S(B,C)$, which expression finds A values associated with all C values in S?

- A. Divide the projection of AB from R by projection of B from S
- B. Divide the projection of AC from R join S by projection of C from S
- C. Divide the natural join of R and S by relation R on A only
- D. Project A from the natural join of R and S directly

Answer: B

Q1232. When optimizing a relational algebra expression, which transformation typically yields the greatest performance improvement?

- A. Replacing all natural joins with Cartesian products instead
- B. Pushing selections down to be applied as early as possible
- C. Performing projections after all joins are completed fully
- D. Delaying all selections until the final result is computed

Answer: B

Q1233. In relational algebra, how can the intersection operation be expressed using fundamental operations?

- A. A minus the result of A minus B gives the intersection
- B. A union B minus the Cartesian product of A and B together
- C. A union B minus A minus B gives the intersection result
- D. The Cartesian product of A and B filtered by equality check

Answer: A

Q1234. A query requires finding employees who work in every department. Which relational algebra operation is most appropriate?

- A. Division of employee-department pairs by all departments
- B. Natural join of employee and department relations directly
- C. Projection of employee names from the department relation
- D. Selection on the Cartesian product with a count condition

Answer: A

Q1235. When converting SELECT DISTINCT A FROM R WHERE B > 5 to relational algebra, the correct expression is:

- A. $R \text{ join } \sigma_{\{B>5\}}(R)$ using natural join with selection first
- B. $\pi_A(\sigma_{\{B>5\}}(R))$ applying selection before projection
- C. $\sigma_{\{B>5\}}(\pi_A(R))$ applying projection before selection
- D. $\pi_A(R) / \sigma_{\{B>5\}}(R)$ using division after projection

Answer: B

Q1236. In query tree optimization, why is it beneficial to perform projections early in the expression tree?

- A. Early projections increase the number of attributes processed
- B. Early projections automatically sort the intermediate results
- C. Early projections eliminate the need for any join operations
- D. Early projections reduce tuple width reducing memory overhead

Answer: D

Q1237. Given $R(A,B,C)$ and $S(C,D,E)$, which expression correctly computes R left outer join S on R.C = S.C?

- A. $(R \text{ join } S) \text{ union } ((R - \pi_{\{A,B,C\}}(R \text{ join } S)) \text{ times } \{(null, null)\})$
- B. $(R \text{ join } S) \text{ union } (\pi_{\{A,B,C\}}(R \text{ join } S) - R)$ with null padding
- C. $R \text{ times } S$ filtered by R.C equals S.C with duplicates removed
- D. $\pi_{\{A,B\}}(R) \text{ join } \sigma_{\{D>0\}}(S)$ with null values for unmatched

Answer: A

Q1238. When two selection conditions are conjunctive (AND), which algebraic equivalence holds?

- A. $\sigma_{\{p \text{ AND } q\}}(R) = \sigma_p(\sigma_q(R))$ using cascaded selections
- B. $\sigma_{\{p \text{ AND } q\}}(R) = \sigma_p(R) \text{ times } \sigma_q(R)$ using product
- C. $\sigma_{\{p \text{ AND } q\}}(R) = \sigma_p(R) \text{ union } \sigma_q(R)$ using union
- D. $\sigma_{\{p \text{ AND } q\}}(R) = \sigma_p(R) - \sigma_q(R)$ using set difference

Answer: A

Q1239. In the context of query optimization, what is the cost of a Cartesian product of relations with m and n tuples respectively?

- A. The cost is proportional to the sum of m plus n tuples
- B. The cost is proportional to the minimum of m or n tuples
- C. The cost is proportional to the maximum of m or n tuples
- D. The cost is proportional to the product of m times n tuples

Answer: D

Q1240. A relational algebra expression uses multiple joins. Which join ordering heuristic typically minimizes intermediate result sizes?

- A. Always join the two smallest intermediate results first
- B. Always join the two largest relations first in the query
- C. Join relations in the order they appear in the SQL query
- D. Join relations in alphabetical order of their table names

Answer: A

Q1241. A query needs to find the second highest salary from an Employee table. Which approach is correct?

- A. `SELECT MIN(salary) FROM Employee WHERE salary > (SELECT MIN(salary) FROM Employee)`
- B. `SELECT DISTINCT salary FROM Employee ORDER BY salary LIMIT 1`
- C. `SELECT MAX(salary) FROM Employee WHERE salary < (SELECT MAX(salary) FROM Employee)`
- D. `SELECT MAX(salary) FROM Employee` without any additional filter

Answer: C

Q1242. What is the difference between UNION and UNION ALL in SQL?

- A. UNION works with different columns while UNION ALL needs same
- B. UNION removes duplicates while UNION ALL keeps all duplicates
- C. UNION keeps duplicates while UNION ALL removes all duplicates
- D. UNION is for joins while UNION ALL is for subquery operations

Answer: B

Q1243. In a correlated subquery, when is the inner query executed?

- A. Once for each row processed by the outer query statement
- B. Only when the outer query returns an empty result set
- C. Once before the outer query begins its execution cycle
- D. Once after the outer query completes all its processing

Answer: A

Q1244. Which SQL construct is best for assigning a rank with gaps to rows based on a column value?

- A. NTILE window function
- B. RANK window function
- C. ROW_NUMBER window function
- D. DENSE_RANK window function

Answer: B

Q1245. A query uses multiple joins. In which order does the SQL engine typically process the clauses?

- A. SELECT, FROM, WHERE, GROUP BY, HAVING, ORDER BY sequence
- B. ORDER BY, SELECT, FROM, WHERE, GROUP BY, HAVING sequence
- C. WHERE, FROM, SELECT, GROUP BY, ORDER BY, HAVING sequence
- D. FROM, WHERE, GROUP BY, HAVING, SELECT, ORDER BY sequence

Answer: D

Q1246. What is the purpose of the COALESCE function in SQL?

- A. It returns the first non-null value from its arguments
- B. It concatenates multiple string values into one result
- C. It converts a value from one data type to another type
- D. It calculates the average of all provided numeric values

Answer: A

Q1247. When using a Common Table Expression (CTE), what advantage does it provide over a regular subquery?

- A. CTEs allow recursive queries that subqueries cannot handle
- B. CTEs bypass all security permissions on underlying base tables
- C. CTEs are always faster than equivalent subquery expressions
- D. CTEs automatically create permanent tables in the database

Answer: A

Q1248. A query needs to find departments where ALL employees earn more than 50000. Which approach is correct?

- A. SELECT dept FROM Employee GROUP BY dept HAVING MAX(salary) > 50000
- B. SELECT dept FROM Employee GROUP BY dept HAVING MIN(salary) > 50000
- C. SELECT dept FROM Employee GROUP BY dept HAVING AVG(salary) > 50000
- D. SELECT dept FROM Employee WHERE salary > 50000 without grouping

Answer: B

Q1249. What is the difference between EXISTS and IN when used with subqueries?

- A. EXISTS only works with correlated subqueries and IN never does
- B. EXISTS and IN always produce different results for same data
- C. EXISTS checks for value matches while IN checks row existence
- D. EXISTS returns true if subquery returns any rows at all

Answer: D

Q1250. In SQL, what does the PIVOT operation accomplish?

- A. It transforms row values into separate column headers
- B. It merges multiple tables into a single combined view
- C. It sorts data by multiple columns in alternating order
- D. It splits a single column into multiple separate tables

Answer: A

Q1251. A recursive CTE is used to traverse an employee hierarchy. What component prevents infinite recursion?

- A. A termination condition in the recursive member WHERE clause
- B. The LIMIT clause applied to the final outer SELECT query
- C. An index on the self-referencing foreign key column used
- D. The DISTINCT keyword in the recursive member query

Answer: A

Q1252. When should a materialized view be refreshed using ON COMMIT versus ON DEMAND?

- A. ON COMMIT for infrequently changed data and ON DEMAND for frequent
- B. ON COMMIT for read-heavy workloads and ON DEMAND for write-heavy
- C. ON COMMIT when data freshness is critical and ON DEMAND when not
- D. ON COMMIT for large tables and ON DEMAND for small table sizes

Answer: C

Q1253. A trigger needs to maintain an audit trail of all salary changes. Which trigger type and timing combination is most appropriate?

- A. BEFORE SELECT trigger on the employee table records
- B. INSTEAD OF DELETE trigger on the employee table only
- C. AFTER UPDATE trigger on the employee salary column
- D. BEFORE INSERT trigger on the employee table only

Answer: C

Q1254. In SQL, what is the difference between RANK(), DENSE_RANK(), and ROW_NUMBER() for tied values?

- A. All three produce identical results for tied values always
- B. ROW_NUMBER skips ranks while RANK and DENSE_RANK assign same to ties
- C. RANK and DENSE_RANK handle ties while ROW_NUMBER assigns unique numbers
- D. RANK skips numbers for ties while DENSE_RANK and ROW_NUMBER do not

Answer: C

Q1255. A stored procedure needs to handle errors gracefully. Which SQL mechanism is most appropriate?

- A. Implementing TRY-CATCH or EXCEPTION handling blocks within
- B. Using nested IF-ELSE statements for every possible error
- C. Wrapping every statement in a separate transaction context
- D. Checking return values of every SQL statement manually each

Answer: A

Q1256. When creating a covering index, what specific advantage does it provide for query performance?

- A. It compresses the data to reduce overall storage requirements
- B. It allows the query to be answered from the index alone
- C. It eliminates the need for primary keys on the covered table
- D. It automatically partitions the table across multiple disks

Answer: B

Q1257. A database has a view defined on a join of three tables. Under which condition is this view updatable?

- A. When all three tables have identical primary key definitions
- B. When the update affects columns from only one base table
- C. When the view includes aggregate functions on joined columns
- D. When the view includes a DISTINCT clause in its definition

Answer: B

Q1258. What is the difference between statement-level and row-level triggers?

- A. Statement-level only works with SELECT while row-level with all DML
- B. Statement-level is faster than row-level in every possible case
- C. Statement-level fires for each row while row-level fires once
- D. Statement-level fires once per statement while row-level per row

Answer: D

Q1259. When using the MERGE statement, what advantage does it provide over separate INSERT, UPDATE, DELETE statements?

- A. MERGE automatically creates indexes on the target table columns
- B. MERGE combines insert, update, and delete in a single atomic step
- C. MERGE is always faster than individual DML statements executed
- D. MERGE bypasses all constraint checking for improved performance

Answer: B

Q1260. A query uses LAG and LEAD window functions. What do these functions specifically provide?

- A. They access data from preceding and following rows respectively
- B. They count the number of rows before and after current row
- C. They calculate running totals across all rows in partition
- D. They determine the first and last values within each partition

Answer: A

Q1261. Given R(A, B, C, D) with FDs {A->B, B->C, C->D}, decompose into BCNF relations.

- A. R1(A,D), R2(B,C) splitting based on attribute positions
- B. R1(A,B,C,D) remains as a single table in BCNF form
- C. R1(A,C), R2(B,D) with cross-referencing foreign key links
- D. R1(A,B), R2(B,C), R3(C,D) each with proper determinants

Answer: D

Q1262. A relation R(Student, Course, Instructor) has FDs {Student,Course->Instructor, Instructor->Course}. What normal form is R in?

- A. R is only in 2NF because of a partial dependency present
- B. R is only in 1NF because of a transitive dependency issue
- C. R is in 3NF but not BCNF due to the second dependency
- D. R is in BCNF since all determinants are candidate keys

Answer: C

Q1263. When decomposing a relation into 3NF using the synthesis algorithm, what is the first step?

- A. Identify all candidate keys of the original relation first
- B. List all attributes and their domains in the relation schema
- C. Compute a minimal cover of the functional dependencies set
- D. Create a separate table for each attribute in the relation

Answer: C

Q1264. A relation has attributes {A, B, C, D, E} with FDs {AB→C, C→D, D→E, E→A}. How many candidate keys does this relation have?

- A. Exactly two candidate keys exist in this relation
- B. Exactly one candidate key exists in this relation
- C. Exactly three candidate keys exist in this relation
- D. Exactly four candidate keys exist in this relation

Answer: D

Q1265. What specific problem can arise when decomposing a relation into BCNF that does not occur with 3NF decomposition?

- A. BCNF decomposition may lose some functional dependencies
- B. BCNF decomposition cannot handle composite primary key sets
- C. BCNF decomposition always results in lossy join operations
- D. BCNF decomposition always creates more tables than needed

Answer: A

Q1266. A table stores Employee(EmpID, Name, DeptID, DeptName, DeptLocation). With FD DeptID→DeptName, DeptLocation, what anomalies exist?

- A. No anomalies exist because the table is fully normalized
- B. Update, insertion, and deletion anomalies all exist here
- C. Only deletion anomalies exist due to DeptID dependency
- D. Only insertion anomalies exist due to the composite key

Answer: B

Q1267. Fifth normal form (5NF) addresses which specific type of dependency?

- A. Join dependencies that cannot be expressed as simpler ones
- B. Inclusion dependencies between attributes of related tables
- C. Functional dependencies between key and non-key attributes
- D. Multivalued dependencies between independent attribute sets

Answer: A

Q1268. In the chase algorithm for testing lossless join, what does the algorithm verify?

- A. Whether the original relation has any candidate key defined
- B. Whether the resulting tables are all in BCNF normal form
- C. Whether the decomposition allows lossless reconstruction
- D. Whether all functional dependencies are preserved in split

Answer: C

Q1269. A relation R(A,B,C) has multivalued dependencies A→→B and A→→C. What is the correct 4NF decomposition?

- A. Keep R(A,B,C) as a single table since it is already in 4NF
- B. Decompose into R1(A,B) and R2(A,C) removing the dependency
- C. Decompose into R1(B,C) and R2(A) with a mapping table added
- D. Decompose into R1(A,B,C) and R2(A) with a foreign key link

Answer: B

Q1270. When is a relation in Domain-Key Normal Form (DKNF)?

- A. When all functional dependencies have superkeys as determinants
- B. When all constraints are consequences of domain and key constraints
- C. When no multivalued dependencies exist in the relation schema
- D. When all attributes are of the same domain type and key type

Answer: B

Q1271. A banking system processes concurrent transfers between accounts. If T1 transfers from A to B and T2 transfers from B to A simultaneously, what specific ACID property prevents incorrect total balances?

- A. Atomicity ensures each transfer fully completes or not at all
- B. All three properties above work together to prevent the issue
- C. Isolation ensures the transfers do not interfere with each
- D. Durability ensures the balances survive server failures

Answer: C

Q1272. In a schedule where T1 reads X, T2 writes X, T1 reads X again, which anomaly occurs and at which isolation level is it first prevented?

- A. Non-repeatable read prevented at Repeatable Read level first
- B. Dirty read anomaly prevented at Read Uncommitted level first
- C. Phantom read anomaly prevented at Serializable level first
- D. No anomaly occurs because both transactions read the data

Answer: A

Q1273. When testing conflict serializability, what does a cycle in the precedence graph indicate?

- A. The schedule needs more transactions to determine its result
- B. The schedule is not conflict serializable and has conflicts
- C. The schedule is conflict serializable and safe to execute
- D. The schedule has no conflicts between any transactions at all

Answer: B

Q1274. A long-running analytical transaction keeps getting aborted due to conflicts with short OLTP transactions. Which strategy best addresses this?

- A. Using snapshot isolation for the analytical transaction query
- B. Increasing the timeout period for the analytical transaction
- C. Converting the analytical query into multiple short transactions
- D. Disabling concurrency control for the analytical transaction

Answer: A

Q1275. In a distributed transaction using the two-phase commit protocol, what happens during the prepare phase?

- A. All participants send their data to the coordinator node
- B. The coordinator unilaterally decides to commit or abort all
- C. All participants commit their local transactions immediately
- D. The coordinator asks participants if they can commit safely

Answer: D

Q1276. A system uses strict two-phase locking. What additional guarantee does this provide over basic two-phase locking?

- A. It prevents deadlocks from occurring between transactions
- B. It holds all exclusive locks until the transaction commits
- C. It allows locks to be released during the growing phase
- D. It eliminates the need for any lock management overhead

Answer: B

Q1277. When using write-ahead logging with UNDO/REDO recovery, what must be true before a dirty page can be flushed to disk?

- A. The page must have been read by at least one other transaction
- B. The buffer pool must be completely full of dirty data pages
- C. The log record for the update must be on stable storage
- D. The transaction must have already committed successfully

Answer: C

Q1278. A system uses multiversion concurrency control (MVCC). What is the primary advantage over lock-based concurrency?

- A. Readers do not block writers and writers do not block readers
- B. MVCC eliminates the need for any transaction isolation level
- C. MVCC uses less memory than lock-based systems in all cases
- D. MVCC guarantees serializability without any additional effort

Answer: A

Q1279. In the ARIES recovery algorithm, what are the three phases of recovery?

- A. Prepare, Commit, and Acknowledge phases for transaction recovery
- B. Scan, Parse, and Execute phases for query recovery management
- C. Analysis, Redo, and Undo phases for crash recovery processing
- D. Backup, Restore, and Verify phases for data recovery operations

Answer: C

Q1280. A transaction T1 reads a set of rows matching a condition, then T2 inserts a new row matching that condition, then T1 re-reads. What anomaly is this and how is it prevented?

- A. Phantom read anomaly prevented by Serializable isolation level
- B. Dirty read anomaly prevented by Read Committed isolation level
- C. Non-repeatable read prevented by Repeatable Read isolation level
- D. Lost update anomaly prevented by Read Uncommitted isolation level

Answer: A

Q1281. In a system using strict two-phase locking, a transaction T1 holds an exclusive lock on item X and T2 requests a shared lock on X. What happens?

- A. T1 is forced to release its exclusive lock for T2 to proceed
- B. T2 is permanently aborted and must restart from the beginning
- C. T2 immediately gets the shared lock alongside the exclusive
- D. T2 must wait until T1 releases the exclusive lock on item X

Answer: D

Q1282. A database uses multigranularity locking with IS, IX, S, SIX, and X modes. Which lock mode allows reading an entire table while intending to write some rows?

- A. SIX (Shared with Intention Exclusive) lock on table
- B. IS (Intention Shared) lock mode on the entire table
- C. S (Shared) lock mode on the entire table resource
- D. IX (Intention Exclusive) lock mode on the table

Answer: A

Q1283. When using timestamp ordering, what happens when a transaction tries to write an item with a read timestamp greater than its own?

- A. The write proceeds normally without any restriction applied
- B. The transaction is rolled back and restarted with new timestamp
- C. The read timestamp is updated to match the write transaction
- D. The item is locked until both transactions complete their work

Answer: B

Q1284. In optimistic concurrency control, what are the three phases of transaction execution?

- A. Prepare, Commit, and Cleanup phases for transaction flow
- B. Read, Validation, and Write phases for each transaction
- C. Parse, Optimize, and Execute phases for query processing
- D. Lock, Execute, and Release phases for each transaction

Answer: B

Q1285. A system detects deadlocks using a wait-for graph. What does a cycle in this graph indicate?

- A. No conflicts exist between the transactions in the cycle
- B. A deadlock exists among the transactions in the cycle now
- C. The transactions in the cycle have all committed already
- D. All transactions in the cycle can proceed without issues

Answer: B

Q1286. When using predicate locking to prevent phantoms, what is the main practical limitation?

- A. Predicate locking requires all transactions to run serially
- B. Predicate locking is computationally expensive to implement
- C. Predicate locking cannot prevent any type of data anomaly
- D. Predicate locks only work with exclusive lock modes applied

Answer: B

Q1287. In snapshot isolation, what specific anomaly can still occur that true serializability prevents?

- A. Phantom reads from newly inserted rows matching conditions
- B. Dirty reads from uncommitted transaction data values
- C. Write skew where two transactions write based on stale reads
- D. Non-repeatable reads from concurrent update operations

Answer: C

Q1288. A database uses tree protocol locking. What constraint does this protocol impose on lock acquisition?

- A. All locks must be acquired simultaneously before any operation
- B. Locks must be acquired in alphabetical order of data items
- C. Locks can only be acquired on leaf nodes of the tree structure
- D. First lock can be on any item but subsequent locks follow tree

Answer: D

Q1289. When comparing wound-wait and wait-die for deadlock prevention, which scheme causes more unnecessary rollbacks and why?

- A. Wait-die causes more because younger transactions die on conflict
- B. Neither scheme causes any unnecessary rollbacks by their design
- C. Wound-wait causes more because it preempts younger transactions
- D. Both schemes cause exactly the same number of rollbacks always

Answer: A

Q1290. In a real-time database system, how should the concurrency control protocol be modified to meet timing constraints?

- A. Use standard two-phase locking with no modifications at all
- B. Disable concurrency control entirely to maximize throughput
- C. Use priority-based protocols that consider transaction deadlines
- D. Lock all data items at the start to prevent any waiting time

Answer: C

Q1291. In the ARIES recovery algorithm, what is the purpose of the analysis phase?

- A. To determine dirty pages and active transactions at crash time
- B. To redo all committed transactions found in the log records
- C. To create a new checkpoint after the recovery process completes
- D. To undo all uncommitted transactions found in the log records

Answer: A

Q1292. During ARIES redo phase, why does it repeat history by redoing all logged actions?

- A. To optimize the recovery process by skipping unnecessary operations
- B. To restore the database to its exact state at the moment of crash
- C. To ensure only committed transaction changes are applied to data
- D. To verify that all log records are correctly formatted and valid

Answer: B

Q1293. In a system using steal/no-force policy, which recovery operations are needed?

- A. Neither UNDO nor REDO operations are needed for the system
- B. Only REDO operations are needed for committed transactions
- C. Both UNDO and REDO operations are needed for full recovery
- D. Only UNDO operations are needed for uncommitted transactions

Answer: C

Q1294. A system crash occurs during a checkpoint. How does the recovery process handle this incomplete checkpoint?

- A. The incomplete checkpoint is treated as if it completed normally
- B. The recovery process ignores the incomplete checkpoint entirely
- C. The database is restored from the most recent full backup only
- D. Recovery uses the previous complete checkpoint as starting point

Answer: D

Q1295. In physiological logging used by ARIES, what does a log record describe?

- A. The logical operation applied to a specific physical page only
- B. The exact byte-level changes made to the physical disk blocks
- C. The complete SQL statement that was executed by the transaction
- D. The schema changes made to the table definition during update

Answer: A

Q1296. A database uses asynchronous replication for disaster recovery. What is the primary risk of this approach?

- A. Network bandwidth usage is higher than synchronous replication
- B. The primary server may experience increased write latency
- C. The replica server cannot handle any read query operations
- D. Some committed transactions may be lost if the primary fails

Answer: D

Q1297. How does the compensation log record (CLR) work in the ARIES algorithm?

- A. CLR creates a backup of the data before any changes are applied
- B. CLR records the original changes made by a committed transaction
- C. CLR logs the undo action and points to the next record to undo
- D. CLR prevents any further modifications to the recovered data items

Answer: C

Q1298. A system needs sub-second recovery time. Which technique best achieves this requirement?

- A. Shadow paging with deferred modification and periodic checkpoints
- B. In-memory database with synchronous replication to a standby
- C. Full database backup every hour with log replay for recovery
- D. Traditional checkpoint-based recovery with large log intervals

Answer: B

Q1299. In a distributed database, what additional complexity does recovery face compared to centralized systems?

- A. Distributed databases do not require any transaction log records
- B. All sites must use identical hardware and software configurations
- C. Recovery in distributed systems is always faster than centralized
- D. Coordinating recovery across multiple sites with network issues

Answer: D

Q1300. When using logical undo during recovery, why might a physical undo be insufficient?

- A. Physical undo is always sufficient for all recovery scenarios
- B. Concurrent transactions may have modified the same page since then
- C. Physical undo cannot read the log records written to stable storage
- D. Logical undo is faster than physical undo in every possible case

Answer: B

Q1301. A B+ tree with order 4 and 1000 records needs at minimum how many levels including the root?

- A. Two levels are needed for this tree configuration setup
- B. Four levels are needed for this tree configuration setup
- C. Five levels are needed for this tree configuration setup
- D. Three levels are needed for this tree configuration setup

Answer: D

Q1302. When should a database administrator choose a hash index over a B+ tree index?

- A. When the indexed column has very low cardinality values only
- B. When equality queries dominate and data size is predictable
- C. When the table is frequently updated with many insert operations
- D. When range queries are the predominant access pattern needed

Answer: B

Q1303. In a B+ tree, what happens during a leaf node split when inserting into a full node?

- A. The leaf splits into two and the middle key propagates upward
- B. The entire tree is rebuilt from scratch with new node sizes
- C. All records in the full node are redistributed to sibling nodes
- D. The insertion is rejected until space becomes available in node

Answer: A

Q1304. A query accesses 5% of records in a table with one million rows. Is using an index beneficial compared to a full table scan?

- A. Index is always beneficial regardless of the percentage accessed
- B. Index is never beneficial when accessing more than 1% of records
- C. Full table scan is better because 5% means too many random reads
- D. It depends on whether the access is sequential or random pattern

Answer: D

Q1305. What is the advantage of using a B+ tree over extendible hashing for concurrent access?

- A. B+ tree never requires any locking for concurrent operations
- B. B+ tree allows finer granularity locking on individual nodes
- C. Extendible hashing cannot support more than one concurrent user
- D. Extendible hashing always requires locking the entire directory

Answer: B

Q1306. In a database with write-heavy workload, what is the impact of having too many indexes on a table?

- A. Write performance degrades because each index must be updated too
- B. There is no impact because indexes only affect read query operations
- C. Write performance improves because indexes speed up all operations
- D. Read performance degrades because indexes conflict with each other

Answer: A

Q1307. What is a partial index and when is it most useful?

- A. An index that is partially stored in memory and partially on disk
- B. An index created during off-peak hours for performance reasons
- C. An index on all rows but only some columns of the data table
- D. An index on a subset of rows matching a predicate condition

Answer: D

Q1308. How does a log-structured merge tree (LSM tree) differ from a B+ tree in write performance?

- A. B+ trees are faster for writes due to their balanced tree structure
- B. Both structures have identical write performance characteristics
- C. LSM trees are faster for writes by using sequential append-only IO
- D. LSM trees are slower for writes due to frequent merge operations

Answer: C

Q1309. A table has a composite index on columns (A, B, C). Which queries can efficiently use this index?

- A. Queries filtering on A, or A and B, or A and B and C together
- B. Only queries filtering on all three columns A, B, and C at once
- C. Queries filtering on any single column A or B or C independently
- D. Only queries filtering on column C alone can use this index

Answer: A

Q1310. What is the write amplification problem in B+ tree indexes and how do LSM trees address it?

- A. LSM trees have worse write amplification than B+ trees in all cases
- B. B+ trees have no write amplification issue at all in practice
- C. Write amplification only occurs in hash-based indexes not in trees
- D. B+ trees cause multiple page writes per update while LSM trees batch

Answer: D

Q1311. When joining three tables R, S, and T, how does the optimizer determine the best join order?

- A. It always joins in the order specified in the SQL FROM clause
- B. It enumerates possible orders and estimates the cost of each
- C. It randomly selects a join order and uses that for execution
- D. It always joins the two smallest tables first without analysis

Answer: B

Q1312. A query has a selection condition with OR predicates on different indexed columns. How does the optimizer typically handle this?

- A. It creates a temporary composite index on both columns at once
- B. It uses index union (bitmap OR) to combine results from indexes
- C. It ignores all indexes and performs a full table scan always
- D. It only uses the index on the first column mentioned in query

Answer: B

Q1313. In a distributed database, what additional factor must the query optimizer consider compared to centralized systems?

- A. The color scheme of the user interface displaying query results
- B. Only the CPU cost of processing each query operation step
- C. Only the memory usage of temporary tables during query execution
- D. The network communication cost of transferring data between sites

Answer: D

Q1314. What is the difference between rule-based and cost-based optimization?

- A. Cost-based does not consider indexes while rule-based always does so
- B. Rule-based uses predefined rules while cost-based uses data statistics
- C. Rule-based is always more accurate than cost-based for complex plans
- D. Rule-based considers data statistics while cost-based uses fixed rules

Answer: B

Q1315. How does the optimizer handle correlated subqueries for performance?

- A. It ignores correlated subqueries and returns approximate results
- B. It always executes correlated subqueries exactly as they are written
- C. It may decorrelate them by transforming into joins when possible
- D. It converts all correlated subqueries into temporary table scans

Answer: C

Q1316. What is adaptive query processing?

- A. Modifying the execution plan during runtime based on actual data
- B. Processing queries without any optimization or planning phase
- C. Using a fixed plan regardless of the data distribution changes
- D. Processing queries only during off-peak hours for efficiency

Answer: A

Q1317. When estimating the cost of a join operation, what is the most critical factor?

- A. The estimated size of the intermediate join result produced
- B. The number of columns selected in the final output projection
- C. The order of columns listed in the SELECT clause of the query
- D. The names of the tables being joined in the SQL statement

Answer: A

Q1318. A query execution plan shows a sort-merge join between two large unsorted tables. Under what condition would a hash join be more efficient?

- A. When both tables have clustered indexes on the join attribute
- B. When the join condition involves a range-based comparison only
- C. When one table fits in memory and an equality join is needed
- D. When both tables are already sorted on the join attribute key

Answer: C

Q1319. What is the purpose of maintaining multiple access paths in the optimizer search space?

- A. To simplify the optimization process by reducing total choices
- B. To find the plan with the lowest cost among various alternatives
- C. To guarantee that indexes are never used for any query plans
- D. To ensure every query produces exactly the same result always

Answer: B

Q1320. How do parameterized queries affect query optimization?

- A. They force the optimizer to always perform a full table scan only
- B. They prevent the optimizer from generating any execution plan
- C. They allow plan caching and reuse across different parameter values
- D. They eliminate the need for any statistics collection on columns

Answer: C

Q1321. An application stores passwords in a database. Which approach provides the strongest security?

- A. Hashing passwords with salt using a slow algorithm like bcrypt
- B. Encrypting passwords with a reversible symmetric algorithm
- C. Encoding passwords with Base64 encoding for basic obfuscation
- D. Storing passwords in plaintext for easy comparison checks

Answer: A

Q1322. A database administrator needs to implement row-level security for a multi-tenant application. Which approach is most effective?

- A. Using views with security predicates based on tenant context
- B. Using application-level filtering without any database controls
- C. Creating separate tables for each tenant in the same database
- D. Giving each tenant a completely separate database instance here

Answer: A

Q1323. What is the risk of a second-order SQL injection compared to a first-order injection?

- A. Second-order only affects SELECT queries while first-order all
- B. Stored malicious input executes later when used in another query
- C. Second-order is less dangerous because the input is validated
- D. There is no practical difference between first and second order

Answer: B

Q1324. When implementing database encryption, what is the key management challenge?

- A. Encryption keys do not need any management after initial creation
- B. Securely storing, rotating, and distributing keys without exposure
- C. Key management only matters for symmetric encryption not asymmetric
- D. Encryption keys are automatically managed by the operating system

Answer: B

Q1325. A healthcare database must comply with HIPAA regulations. Which database security measures are most critical?

- A. Only password protection and regular backups are sufficient here
- B. Access control, encryption, audit trails, and data masking combined
- C. Only data encryption at rest without any other security measures
- D. Only network firewalls without any database-level security controls

Answer: B

Q1326. What is the difference between data-at-rest encryption and data-in-transit encryption?

- A. At-rest is optional while in-transit is always mandatory requirement
- B. At-rest protects network data while in-transit protects stored data
- C. Both protect data at the same stage of the data lifecycle process
- D. At-rest protects stored data while in-transit protects network data

Answer: D

Q1327. In a privilege escalation attack on a database, what specific vulnerability is exploited?

- A. The attacker gains higher privileges than originally authorized
- B. The attacker corrupts data by writing invalid values to tables
- C. The attacker creates excessive connections to overload the server
- D. The attacker exploits weak network protocols to intercept data

Answer: A

Q1328. A database uses fine-grained auditing. How does this differ from standard auditing?

- A. Fine-grained auditing captures specific conditions and column access
- B. Standard auditing provides more detail than fine-grained auditing
- C. Fine-grained auditing is performed only at the application level
- D. Fine-grained auditing only records failed access attempts made

Answer: A

Q1329. What security risk does dynamic SQL pose that static SQL does not?

- A. Dynamic SQL is more vulnerable to injection if input not sanitized
- B. Dynamic SQL is slower and uses more memory than static SQL code
- C. Dynamic SQL bypasses all authentication checks automatically always
- D. Static SQL cannot access any tables while dynamic SQL can access all

Answer: A

Q1330. When implementing database security for a cloud-hosted database, what additional concern must be addressed compared to on-premises?

- A. Cloud databases are inherently more secure than on-premises ones
- B. Shared responsibility model and data sovereignty across regions
- C. Only network security matters for cloud-hosted database instances
- D. Cloud databases do not require any encryption for stored data

Answer: B

Q1331. In a distributed database using two-phase commit, what is the blocking problem and how does three-phase commit address it?

- A. Participants block if coordinator fails; 3PC adds precommit phase
- B. Three-phase commit eliminates all network communication needed
- C. Blocking only occurs during the abort phase of commit protocol
- D. Blocking occurs when a participant crashes before prepare phase

Answer: A

Q1332. A distributed database must choose between strong consistency and high availability during a network partition. According to CAP, which trade-off must be made?

- A. The system must choose between consistency or availability only
- B. The system should shut down until the partition is resolved
- C. The system can maintain both consistency and availability fully
- D. The system must sacrifice both consistency and availability

Answer: A

Q1333. When optimizing a distributed query that joins tables at different sites, what is the primary cost to minimize?

- A. The CPU processing time at each individual site in network
- B. The memory consumption at the coordinator site during join
- C. The network data transfer cost between the distributed sites
- D. The disk storage used for temporary tables at each site node

Answer: C

Q1334. In a multi-master replication setup, how are write conflicts typically resolved?

- A. Using conflict resolution policies like last-writer-wins or merge
- B. By routing all writes to a single designated primary master node
- C. By preventing all writes until conflict is manually resolved
- D. By rejecting all conflicting writes and requiring resubmission

Answer: A

Q1335. What is the Paxos consensus algorithm used for in distributed databases?

- A. Replicating data between nodes using synchronous write operations
- B. Optimizing query execution plans across distributed network nodes
- C. Partitioning tables into fragments for distributed storage placement
- D. Achieving agreement on a value among distributed nodes despite failures

Answer: D

Q1336. In a geo-distributed database, what is the primary challenge of maintaining ACID transactions across data centers?

- A. Network latency between data centers increases transaction time
- B. Geo-distribution eliminates the need for any concurrency control
- C. Data centers always have identical hardware and network configs
- D. Each data center uses different database management software

Answer: A

Q1337. What is the difference between eager and lazy replication in terms of consistency guarantees?

- A. Eager provides eventual consistency while lazy provides strong
- B. Both provide identical consistency guarantees in all situations
- C. Eager has no consistency guarantee while lazy guarantees it all
- D. Eager provides strong consistency while lazy provides eventual

Answer: D

Q1338. A company uses a distributed database with read replicas. How does read-your-writes consistency differ from strong consistency?

- A. Read-your-writes ensures all users see the same data at all
- B. Read-your-writes and strong consistency are the same guarantee
- C. Read-your-writes ensures a user sees their own writes quickly
- D. Read-your-writes provides weaker guarantees than eventual one

Answer: C

Q1339. In a distributed database using vector clocks, what do they provide that physical timestamps cannot?

- A. Automatic conflict resolution without any additional logic now
- B. Higher precision timing than physical clock implementations can
- C. Exact wall-clock time ordering of all events across all nodes
- D. Causal ordering of events without requiring synchronized clocks

Answer: D

Q1340. When designing a distributed database schema, what is the key consideration for choosing between horizontal and vertical fragmentation?

- A. Choose based on query access patterns and data locality needs
- B. Vertical fragmentation always results in less network transfer
- C. The choice has no impact on query performance or utilization
- D. Horizontal fragmentation is always preferred over vertical form

Answer: A

Q1341. When should a system architect choose a document database over a relational database for a new application?

- A. When the application has a fixed well-defined relational schema
- B. When the application requires complex multi-table join operations
- C. When the data model is hierarchical and schema evolves frequently
- D. When strong ACID transactions across entities are the top priority

Answer: C

Q1342. In a distributed NoSQL database using consistent hashing, what happens when a new node is added to the cluster?

- A. All data is completely redistributed across all existing nodes
- B. All existing nodes must be taken offline during the process
- C. The new node starts empty and receives no data automatically
- D. Only a portion of data from adjacent nodes is redistributed

Answer: D

Q1343. What is the gossip protocol used for in distributed NoSQL databases?

- A. Propagating cluster state information between nodes efficiently
- B. Encrypting data during transmission between cluster nodes only
- C. Executing distributed queries across all nodes in the cluster
- D. Rebalancing data partitions when nodes join or leave cluster

Answer: A

Q1344. In a NoSQL database handling a social media feed, why might a fan-out-on-write approach be chosen over fan-out-on-read?

- A. Fan-out-on-write reduces write complexity and storage overhead
- B. Fan-out-on-write requires less storage than fan-out-on-read
- C. Fan-out-on-read eliminates the need for any caching mechanisms
- D. Fan-out-on-write provides faster read times at cost of writes

Answer: D

Q1345. What is the trade-off between strong consistency and partition tolerance in a distributed NoSQL system?

- A. Strong consistency improves availability during network partition
- B. Maintaining strong consistency during partitions requires blocking
- C. Partition tolerance is never needed in modern cloud environments
- D. Strong consistency and partition tolerance have no trade-off

Answer: B

Q1346. In Cassandra data model, what is the significance of the clustering key?

- A. It specifies the replication factor for the data across nodes
- B. It determines the consistency level for read-write operations
- C. It determines which node stores the data in the cluster ring
- D. It defines the sort order of data within a partition on disk

Answer: D

Q1347. When designing a NoSQL schema for a real-time analytics application, what is the primary modeling principle?

- A. Design schema based on entity relationships like relational models
- B. Design schema to maximize normalization for data integrity only
- C. Design schema to minimize data redundancy above all other goals
- D. Design schema based on the query patterns the application needs

Answer: D

Q1348. What is a vector clock conflict resolution limitation in eventually consistent NoSQL systems?

- A. Vector clocks grow in size with the number of participating nodes
- B. Vector clocks cannot detect any concurrent write conflicts at all
- C. Vector clocks only work with key-value stores and no other types
- D. Vector clocks require perfectly synchronized physical time clocks

Answer: A

Q1349. In a multi-model NoSQL database, what advantage does combining document and graph capabilities provide?

- A. It allows modeling both structured data and complex relationships
- B. It reduces storage requirements compared to single-model databases
- C. It guarantees ACID transactions across all data model types used
- D. It eliminates the need for any indexing on the stored data sets

Answer: A

Q1350. A NoSQL database uses quorum-based replication with $N=5$, $W=3$, $R=3$. What consistency guarantee does this provide?

- A. Strong consistency since $W+R$ exceeds N ensuring overlap exists
- B. Eventual consistency with no read-your-writes guarantee at all
- C. No consistency guarantee regardless of the quorum configuration
- D. Consistency only for write operations not for read operations

Answer: A

Q1351. When designing an enterprise data warehouse, what is the key difference between Inmon top-down and Kimball bottom-up approach?

- A. Inmon builds data marts first while Kimball builds warehouse
- B. Both approaches build the data warehouse in exactly same order
- C. Inmon uses star schemas while Kimball uses snowflake schemas
- D. Inmon builds normalized warehouse first while Kimball marts first

Answer: D

Q1352. In a Type 2 slowly changing dimension, how are historical changes tracked?

- A. Changes are not tracked and old values are permanently deleted
- B. The existing record is overwritten with the new attribute values
- C. A new record is added with effective dates preserving old ones
- D. A separate history table stores all previous versions of records

Answer: C

Q1353. What is a degenerate dimension in data warehouse design?

- A. A dimension containing only null values for all its attributes
- B. A dimension key in the fact table without a dimension table
- C. A dimension stored as a separate fully normalized table entity
- D. A dimension that has been deleted from the warehouse schema

Answer: B

Q1354. When designing a fact table, what is the difference between additive, semi-additive, and non-additive measures?

- A. Non-additive measures can be summed across time dimension only
- B. Semi-additive can be summed across all dimensions but additive cannot
- C. Additive measures can be summed across all dimensions while others not
- D. All three types of measures can be summed across any dimension

Answer: C

Q1355. A company implements a real-time data warehouse. What architectural change is needed compared to traditional batch ETL?

- A. Use only flat files instead of database tables for data storage
- B. Implement change data capture and micro-batch or stream processing
- C. Replace ETL with ELT and remove all transformation logic entirely
- D. Remove all indexes from the warehouse for faster write performance

Answer: B

Q1356. In data mining, what is the difference between classification and clustering?

- A. Both methods use identical approaches with labeled training data sets
- B. Clustering uses labeled data while classification uses unlabeled data
- C. Classification uses labeled data while clustering uses unlabeled data
- D. Classification is unsupervised while clustering is supervised learning

Answer: C

Q1357. What is a factless fact table and when is it used in data warehouse design?

- A. A fact table that stores future predictions instead of facts
- B. A fact table with no numeric measures recording only events
- C. A fact table with no dimension keys linking to other tables
- D. A fact table that is temporarily empty during initial loading

Answer: B

Q1358. When implementing ELT instead of ETL, what is the key architectural difference?

- A. ELT does not require any extraction from source systems at all
- B. Data is loaded first then transformed inside the target warehouse
- C. Data is transformed before loading into the target warehouse
- D. ELT skips the loading step entirely and queries sources directly

Answer: B

Q1359. What is a conformed dimension in Kimball dimensional modeling and why is it important?

- A. A dimension that has been fully normalized into sub-dimensions
- B. A dimension shared across multiple data marts for consistency
- C. A dimension containing only surrogate keys without attributes
- D. A dimension used in only one specific data mart exclusively here

Answer: B

Q1360. In association rule mining, what do the metrics support and confidence measure?

- A. Support measures rule accuracy while confidence measures frequency
- B. Support measures query performance while confidence measures quality
- C. Support measures item frequency while confidence measures strength
- D. Support measures data volume while confidence measures processing

Answer: C

Q1361. In a DBMS, what does the concept of impedance mismatch refer to?

- A. Mismatch between different hardware configurations
- B. Mismatch between the data model used by the application and the database
- C. Mismatch between primary and foreign key types
- D. Mismatch between storage media speeds

Answer: B

Q1362. Which of the following correctly describes the closed world assumption in databases?

- A. All data in the database is encrypted by default
- B. Facts not present in the database are considered false
- C. The database cannot accept new data after initial loading
- D. All queries must return results

Answer: B

Q1363. What is the significance of Codd's 12 rules in database systems?

- A. They define network protocol standards
- B. They establish criteria for a system to qualify as a relational DBMS
- C. They specify hardware requirements for databases
- D. They outline data encryption algorithms

Answer: B

Q1364. Which Codd rule states that all information in a relational database must be represented as values in tables?

- A. Rule 0: Foundation Rule
- B. Rule 1: Information Rule
- C. Rule 5: Comprehensive Data Sublanguage Rule
- D. Rule 8: Physical Data Independence

Answer: B

Q1365. In database systems, what does the principle of minimal redundancy aim to achieve?

- A. Maximum storage utilization by duplicating data
- B. Storing each fact in exactly one place to avoid update anomalies
- C. Removing all indexes to save space
- D. Eliminating all foreign key relationships

Answer: B

Q1366. What problem does the lost update anomaly cause in concurrent database access?

- A. Data is permanently deleted from the database
- B. One transaction overwrites changes made by another concurrent transaction
- C. Indexes become corrupted during updates
- D. The database schema is modified unintentionally

Answer: B

Q1367. Which of the following scenarios best illustrates the need for logical data independence?

- A. Changing disk storage from HDD to SSD
- B. Adding a new column to a table without modifying existing applications
- C. Moving the database to a different server
- D. Increasing the buffer pool size

Answer: B

Q1368. What is the phantom read problem in database concurrency?

- A. A transaction reads data that has been permanently deleted
- B. A transaction reads rows that were inserted by another uncommitted transaction
- C. A transaction cannot find previously existing records
- D. The database returns empty result sets randomly

Answer: B

Q1369. In the context of database evolution, what advantage did the relational model provide over the hierarchical model?

- A. Faster physical storage access
- B. Data independence and flexible querying through declarative languages
- C. Built-in support for multimedia data
- D. Automatic data replication

Answer: B

Q1370. What is the dirty read problem in database transactions?

- A. Reading data from a corrupted disk sector
- B. Reading uncommitted data written by another transaction that may roll back
- C. Reading outdated cached data from the buffer
- D. Reading data without proper authentication

Answer: B

Q1371. In a shared-nothing parallel database architecture, how is fault tolerance typically achieved?

- A. By sharing memory across all nodes
- B. Through data replication across multiple independent nodes
- C. By using a single centralized backup server
- D. Through shared disk access

Answer: B

Q1372. What is the primary challenge of implementing a shared-disk architecture?

- A. Limited network bandwidth
- B. Managing concurrent disk access and cache coherence across nodes
- C. Insufficient disk storage
- D. Slow query parsing

Answer: B

Q1373. In the context of DBMS architecture, what is the write-ahead log (WAL) protocol?

- A. All write operations are performed after user confirmation
- B. Log records must be written to stable storage before the corresponding data changes are written to disk
- C. Data is written ahead of any read operation
- D. Logs are written only after transaction completion

Answer: B

Q1374. What is the purpose of the ARIES recovery algorithm in DBMS architecture?

- A. To optimize query execution plans
- B. To provide efficient and correct crash recovery using WAL
- C. To manage user authentication
- D. To allocate memory buffers

Answer: B

Q1375. In a multi-tier architecture, what is the benefit of using an application server as a connection broker?

- A. It eliminates the need for a database entirely
- B. It multiplexes many client connections into fewer database connections, reducing server load
- C. It automatically generates SQL queries
- D. It compresses all data before storage

Answer: B

Q1376. What is the CAP theorem's relevance to distributed database architecture?

- A. It states that a database can achieve consistency, availability, and partition tolerance simultaneously
- B. It proves that a distributed system can guarantee at most two of consistency, availability, and partition tolerance
- C. It defines the maximum number of concurrent users
- D. It specifies the minimum hardware requirements

Answer: B

Q1377. What is the difference between inter-query and intra-query parallelism?

- A. Inter-query parallelism is faster than intra-query parallelism
- B. Inter-query runs different queries in parallel while intra-query parallelizes a single query
- C. Inter-query uses more memory while intra-query uses more disk
- D. They are the same concept with different names

Answer: B

Q1378. In DBMS architecture, what is the purpose of the checkpoint mechanism?

- A. To mark the completion of a user session
- B. To periodically flush modified buffers to disk and record a checkpoint in the log to limit recovery time
- C. To check database integrity constraints
- D. To verify user permissions

Answer: B

Q1379. What architectural pattern does database sharding implement?

- A. Vertical partitioning across different schemas
- B. Horizontal partitioning of data across multiple independent database instances
- C. Replication of entire databases to standby servers
- D. Compression of data into fewer storage nodes

Answer: B

Q1380. In a replicated database architecture, what is the split-brain problem?

- A. A query that is too complex for the optimizer
- B. When network partition causes two nodes to independently accept writes, leading to conflicting states
- C. When a database table is split across two different schemas
- D. When CPU usage is divided between query processing and logging

Answer: B

Q1381. When converting an ER diagram to a relational schema, how is a multivalued attribute handled?

- A. It is stored as a comma-separated list in the parent table
- B. A separate table is created with a foreign key referencing the parent entity
- C. It is ignored during conversion
- D. It becomes a primary key in the parent table

Answer: B

Q1382. In the EER model, what is a shared subclass (multiple inheritance)?

- A. A subclass that inherits attributes from only one superclass
- B. A subclass that is a subset of two or more superclasses
- C. A subclass with no attributes of its own
- D. A superclass with no subclasses

Answer: B

Q1383. When mapping a ternary relationship to a relational schema, how is the resulting table structured?

- A. It contains only the primary key of one entity
- B. It contains the primary keys of all three participating entities as foreign keys
- C. It creates three separate binary relationship tables
- D. It merges all three entities into a single table

Answer: B

Q1384. What is the temporal data model's primary extension to the relational model?

- A. Adding spatial coordinates to every record
- B. Incorporating time dimensions to track data changes over time
- C. Replacing tables with graphs
- D. Removing the need for primary keys

Answer: B

Q1385. In the context of ER modeling, what is the minimum cardinality constraint?

- A. The maximum number of relationships an entity can have
- B. The minimum number of relationship instances each entity must participate in
- C. The minimum number of attributes an entity must have
- D. The minimum number of tables in a database

Answer: B

Q1386. How does the attribute-defined specialization differ from user-defined specialization in EER?

- A. Attribute-defined uses a discriminating attribute while user-defined is determined manually
- B. Attribute-defined is faster to process
- C. User-defined cannot have disjoint constraints
- D. There is no difference between them

Answer: A

Q1387. What happens to a total specialization constraint when converting an EER diagram to a relational schema?

- A. It is automatically enforced by the schema
- B. It cannot be directly enforced in the relational model and requires application-level or trigger enforcement
- C. It is converted to a CHECK constraint
- D. It becomes a foreign key constraint

Answer: B

Q1388. In the object-relational data model, what are user-defined types (UDTs)?

- A. Types that only the database administrator can use
- B. Custom data types created by users that can encapsulate attributes and methods
- C. Temporary types used only during queries
- D. Types that replace all standard SQL types

Answer: B

Q1389. What is the concept of closure in the context of data model operations?

- A. The database cannot accept new connections
- B. The result of any operation on relations is also a relation
- C. All transactions must be closed before new ones begin
- D. The database is shut down properly

Answer: B

Q1390. When mapping overlapping specialization to a relational schema, which approach is most appropriate?

- A. Create a single table for all subclasses with a type column
- B. Create separate tables for each subclass with redundant superclass attributes
- C. Create separate tables for each subclass with foreign keys to the superclass table
- D. Merge all entities into the superclass table only

Answer: C

Q1391. What is Armstrong's axiom of reflexivity in functional dependencies?

- A. If Y is a subset of X, then X functionally determines Y
- B. If X determines Y, then Y determines X
- C. If X determines Y and Y determines Z, then X determines Z
- D. If X determines Y, then XZ determines YZ

Answer: A

Q1392. What is the transitive rule in Armstrong's axioms?

- A. If $X \rightarrow Y$, then $Y \rightarrow X$
- B. If $X \rightarrow Y$ and $Y \rightarrow Z$, then $X \rightarrow Z$
- C. If $X \rightarrow Y$, then $XZ \rightarrow YZ$
- D. If Y is a subset of X, then $X \rightarrow Y$

Answer: B

Q1393. What is the closure of a set of attributes under a set of functional dependencies?

- A. The set of all attributes that can be functionally determined from the given attributes
- B. The removal of all redundant attributes
- C. The set of all tables referencing the attributes
- D. The physical storage allocation for the attributes

Answer: A

Q1394. What is a lossless join decomposition?

- A. A decomposition where no data is physically deleted
- B. A decomposition where the original relation can be perfectly reconstructed by joining the decomposed relations
- C. A decomposition that reduces storage by 50%
- D. A decomposition that creates no foreign keys

Answer: B

Q1395. What is a dependency-preserving decomposition?

- A. A decomposition where all data dependencies are documented
- B. A decomposition where all original functional dependencies can be enforced using the decomposed relations
- C. A decomposition that preserves physical storage format
- D. A decomposition that maintains the same number of rows

Answer: B

Q1396. What is a multivalued dependency in relational theory?

- A. A dependency where an attribute has multiple data types
- B. A dependency where one attribute set determines a set of values independently of other attributes
- C. A dependency between multiple databases
- D. A dependency that only exists in NoSQL databases

Answer: B

Q1397. What is a join dependency in relational database theory?

- A. A constraint that two tables must always be joined in queries
- B. A constraint where a relation can be losslessly decomposed into multiple relations and reconstructed via natural join
- C. A dependency between join operations
- D. A physical storage constraint for joined tables

Answer: B

Q1398. What is a trivial functional dependency?

- A. A dependency that is obvious to the database designer
- B. A dependency $X \rightarrow Y$ where Y is a subset of X
- C. A dependency that involves only primary keys
- D. A dependency that can be derived from other dependencies

Answer: B

Q1399. In relational algebra, what is the difference between division and other operations?

- A. Division is faster than other operations
- B. Division finds tuples in one relation that are associated with all tuples in another relation
- C. Division removes duplicate tuples
- D. Division is used for aggregation

Answer: B

Q1400. What does the augmentation rule in Armstrong's axioms state?

- A. If $X \rightarrow Y$, then $X \rightarrow YZ$ for any Z
- B. If $X \rightarrow Y$, then $XZ \rightarrow YZ$ for any attribute set Z
- C. If $X \rightarrow Y$ and $Y \rightarrow Z$, then $X \rightarrow Z$
- D. If $XY \rightarrow Z$, then $X \rightarrow Z$ and $Y \rightarrow Z$

Answer: B

Q1401. How can the division operation $R \div S$ be expressed using fundamental relational algebra operations?

- A. $R \div S = R - S$
- B. $R \div S = \{ \langle x, y \rangle \in R \mid \forall \langle x, y' \rangle \in S, \langle x, y' \rangle \in R \}$
- C. $R \div S = R \cap S$
- D. $R \div S = \{ \langle x, y \rangle \in R \mid \langle x, y \rangle \in S \}$

Answer: B

Q1402. What is the algebraic equivalence rule for pushing selection through a join?

- A. Selections cannot be pushed through joins
- B. $\sigma_c(R \bowtie S) = \sigma_c(R) \bowtie S$ if condition c involves only attributes of R
- C. $\sigma_c(R \bowtie S) = R \bowtie \sigma_c(S)$ regardless of which relation c references
- D. Selections must always be applied after joins

Answer: B

Q1403. In query optimization, why is pushing projections down the query tree beneficial?

- A. It increases the number of columns processed
- B. It reduces the width of intermediate results, decreasing memory and processing costs
- C. It makes queries more readable
- D. It eliminates the need for selection operations

Answer: B

Q1404. What is the difference between tuple relational calculus and domain relational calculus?

- A. They are identical in expressive power and syntax
- B. Tuple calculus uses tuple variables ranging over relations while domain calculus uses variables ranging over attribute domains
- C. Tuple calculus is more powerful than domain calculus
- D. Domain calculus cannot express joins

Answer: B

Q1405. What does the concept of safe expressions mean in relational calculus?

- A. Expressions that do not cause runtime errors
- B. Expressions guaranteed to produce a finite result set
- C. Expressions that do not modify the database
- D. Expressions that use only primary keys

Answer: B

Q1406. How is the outer union operation different from the standard union?

- A. Outer union requires union compatibility
- B. Outer union can combine relations that are not union compatible by using nulls for missing attributes
- C. Outer union removes all duplicates while union does not
- D. There is no difference

Answer: B

Q1407. What is the commutative property of natural join in relational algebra?

- A. $R \bowtie S$ is not equal to $S \bowtie R$
- B. $R \bowtie S = S \bowtie R$, meaning the order of operands does not affect the result
- C. Natural join can only be applied in one direction
- D. Commutativity only holds for equi-joins

Answer: B

Q1408. In relational algebra, what is the associative property of join?

- A. Joins cannot be reordered
- B. $(R \bowtie S) \bowtie T = R \bowtie (S \bowtie T)$, allowing join reordering for optimization
- C. Only two relations can be joined at a time
- D. Associativity only applies to Cartesian products

Answer: B

Q1409. What is the expressive power relationship between relational algebra and safe relational calculus?

- A. Relational algebra is more powerful than relational calculus
- B. Relational calculus is more powerful than relational algebra
- C. They are equivalent in expressive power for safe expressions
- D. Neither can express joins

Answer: C

Q1410. What is the deferred operation approach in evaluating relational algebra expressions?

- A. All operations are executed immediately
- B. Operations are composed into a query tree and optimized before execution
- C. Operations are deferred until the database is idle
- D. Only selection operations are deferred

Answer: B

Q1411. What is the execution order of clauses in a SQL SELECT statement?

- A. SELECT, FROM, WHERE, GROUP BY, HAVING, ORDER BY
- B. FROM, WHERE, GROUP BY, HAVING, SELECT, ORDER BY
- C. FROM, SELECT, WHERE, GROUP BY, ORDER BY, HAVING
- D. SELECT, WHERE, FROM, HAVING, GROUP BY, ORDER BY

Answer: B

Q1412. What is a recursive Common Table Expression (CTE) used for in SQL?

- A. Creating temporary indexes
- B. Querying hierarchical or recursive data structures like organizational charts
- C. Improving query performance automatically
- D. Creating permanent views

Answer: B

Q1413. What is the difference between RANK() and DENSE_RANK() window functions?

- A. They always produce the same result
- B. RANK() leaves gaps after ties while DENSE_RANK() does not
- C. DENSE_RANK() leaves gaps after ties while RANK() does not
- D. RANK() is faster than DENSE_RANK()

Answer: B

Q1414. What is a lateral join in SQL?

- A. A join that connects tables horizontally
- B. A join where the right-hand subquery can reference columns from the left-hand table
- C. A standard cross join with an alias
- D. A join used exclusively in NoSQL databases

Answer: B

Q1415. What does the SQL MERGE statement do?

- A. Merges two databases into one
- B. Performs insert, update, or delete operations in a single statement based on match conditions
- C. Combines query results like UNION
- D. Merges table schemas

Answer: B

Q1416. What is the difference between a scalar subquery and a table subquery?

- A. They are the same thing
- B. A scalar subquery returns exactly one value while a table subquery returns a result set
- C. A scalar subquery is faster
- D. A table subquery cannot be used in WHERE clauses

Answer: B

Q1417. What is a window function's PARTITION BY clause used for?

- A. To partition a table into separate physical storage
- B. To divide the result set into partitions to which the window function is independently applied
- C. To create table partitions for archiving
- D. To split a query into parallel operations

Answer: B

Q1418. What is the purpose of the ROWS BETWEEN clause in window functions?

- A. To limit the total number of rows returned
- B. To define the window frame specifying which rows relative to the current row are included in the calculation
- C. To filter rows before the window function is applied
- D. To specify the number of rows per page

Answer: B

Q1419. What is the NTH_VALUE() window function used for?

- A. To find the nth database in a server
- B. To retrieve the value of an expression at the nth row within a window frame
- C. To calculate the nth percentile
- D. To insert values at the nth position

Answer: B

Q1420. What is SQL injection and how is it prevented?

- A. A method to inject performance optimizations; prevented by disabling indexes
- B. A security vulnerability where malicious SQL is inserted into queries; prevented using parameterized queries
- C. A technique to insert bulk data; prevented by limiting batch sizes
- D. A method to inject stored procedures; prevented by disabling triggers

Answer: B

Q1421. What is the purpose of EXPLAIN or EXPLAIN PLAN in SQL?

- A. To explain SQL syntax to beginners
- B. To display the execution plan showing how the DBMS will process a query
- C. To explain error messages in detail
- D. To document stored procedures

Answer: B

Q1422. What is a pivot operation in SQL?

- A. Rotating a table physically on screen
- B. Transforming row data into column data for cross-tabulation reporting
- C. Reversing the order of rows
- D. Swapping primary and foreign keys

Answer: B

Q1423. What is the difference between OLTP and OLAP in terms of SQL usage?

- A. OLTP uses SQL while OLAP does not
- B. OLTP uses simple short transactions while OLAP uses complex analytical queries
- C. OLTP is for analytics while OLAP is for transactions
- D. There is no difference

Answer: B

Q1424. What is a temporal table in SQL?

- A. A table that is deleted after each session
- B. A system-versioned table that automatically tracks data changes over time
- C. A table used only for temporary calculations
- D. A table with timestamp columns

Answer: B

Q1425. What are the ACID properties' implications for SQL transaction design?

- A. They require all queries to be written in uppercase
- B. They ensure that transactions are atomic, consistent, isolated, and durable for reliable processing
- C. They define the SQL syntax rules
- D. They specify hardware requirements for databases

Answer: B

Q1426. What is the purpose of the GROUPING SETS clause in SQL?

- A. To create multiple user groups
- B. To specify multiple grouping combinations in a single query instead of multiple UNION queries
- C. To set group permissions
- D. To group database connections

Answer: B

Q1427. What is a deferred constraint in SQL?

- A. A constraint that is never checked
- B. A constraint whose enforcement is postponed until the end of the transaction
- C. A constraint applied only to archived data
- D. A constraint that is always disabled

Answer: B

Q1428. What is the FETCH FIRST or LIMIT clause used for with OFFSET?

- A. To limit the database size
- B. To implement pagination by skipping a number of rows and returning a limited set
- C. To limit the number of database connections
- D. To restrict the number of columns returned

Answer: B

Q1429. What is the difference between CUBE and ROLLUP in SQL?

- A. They produce identical results
- B. ROLLUP generates hierarchical subtotals while CUBE generates subtotals for all combinations of grouping columns
- C. CUBE is faster than ROLLUP
- D. ROLLUP generates more rows than CUBE

Answer: B

Q1430. What is Fifth Normal Form (5NF) or Project-Join Normal Form?

- A. A relation with five or fewer attributes
- B. A relation in 4NF with no non-trivial join dependencies that are not implied by candidate keys
- C. A relation that requires five decompositions
- D. A theoretical form that cannot be achieved in practice

Answer: B

Q1431. What is Domain-Key Normal Form (DKNF)?

- A. A form where all constraints are based on domains and keys only
- B. A form that normalizes domain names
- C. A form that requires key encryption
- D. A form applicable only to key-value databases

Answer: A

Q1432. Why might a BCNF decomposition not preserve all functional dependencies?

- A. BCNF does not use functional dependencies
- B. Some dependencies may span across decomposed relations and cannot be checked without joining them
- C. BCNF only considers primary keys
- D. Dependencies are automatically removed during decomposition

Answer: B

Q1433. What is the minimal cover algorithm used for in normalization?

- A. To find the minimum number of tables
- B. To reduce a set of functional dependencies to an equivalent minimal set
- C. To cover all tables with indexes
- D. To minimize the number of rows in a table

Answer: B

Q1434. In the context of normalization, what is a spurious tuple?

- A. A tuple with all null values
- B. An invalid tuple generated by incorrectly joining decomposed relations
- C. A tuple that violates a domain constraint
- D. A duplicate tuple in a relation

Answer: B

Q1435. How does the 3NF synthesis algorithm ensure dependency preservation?

- A. By keeping all attributes in one table
- B. By creating a separate relation for each functional dependency in the minimal cover
- C. By ignoring dependencies during decomposition
- D. By enforcing all dependencies through triggers

Answer: B

Q1436. What is the relationship between multivalued dependencies and Fourth Normal Form?

- A. 4NF ignores multivalued dependencies
- B. 4NF requires that all non-trivial multivalued dependencies have a superkey on their left side
- C. Multivalued dependencies only exist in 1NF
- D. 4NF eliminates all functional dependencies

Answer: B

Q1437. What is over-normalization and what are its consequences?

- A. Having too many columns in a table
- B. Excessive decomposition that leads to too many joins, degrading query performance
- C. Having too many indexes
- D. Normalizing data that should be kept as text

Answer: B

Q1438. What is the Heath's theorem in normalization?

- A. Every relation can be decomposed into 1NF
- B. If $R(A,B,C)$ has $A \rightarrow B$, then R can be losslessly decomposed into $R_1(A,B)$ and $R_2(A,C)$
- C. All relations must have at least three attributes
- D. Normalization always preserves dependencies

Answer: B

Q1439. What is the difference between the decomposition and synthesis approaches to achieving 3NF?

- A. They always produce identical results
- B. Decomposition starts from a universal relation and decomposes, while synthesis builds relations from functional dependencies
- C. Synthesis is a top-down approach while decomposition is bottom-up
- D. Decomposition is always preferred over synthesis

Answer: B

Q1440. What is the Thomas Write Rule in transaction processing?

- A. All writes must be timestamped
- B. An outdated write can be ignored if a more recent transaction has already written to the same item
- C. Writes must always be performed before reads
- D. Thomas must approve all writes

Answer: B

Q1441. What is the write-ahead logging (WAL) protocol for transaction recovery?

- A. All writes are performed ahead of other operations
- B. Log records must be written to stable storage before the actual database modifications
- C. All log files must be written to tape
- D. Write operations must precede read operations

Answer: B

Q1442. What is the difference between the READ COMMITTED and REPEATABLE READ isolation levels?

- A. They provide the same guarantees
- B. READ COMMITTED prevents dirty reads while REPEATABLE READ also prevents non-repeatable reads
- C. REPEATABLE READ is less restrictive than READ COMMITTED
- D. READ COMMITTED prevents phantom reads while REPEATABLE READ does not

Answer: B

Q1443. What is the SERIALIZABLE isolation level and what anomalies does it prevent?

- A. It only prevents dirty reads
- B. It prevents all anomalies including dirty reads, non-repeatable reads, and phantom reads
- C. It only prevents phantom reads
- D. It provides no isolation guarantees

Answer: B

Q1444. What is a distributed transaction and what protocol manages it?

- A. A transaction that runs on multiple distributed disks; managed by RAID
- B. A transaction spanning multiple database nodes; managed by two-phase commit protocol
- C. A transaction distributed across time zones; managed by clock sync
- D. A transaction split into smaller parts; managed by thread pools

Answer: B

Q1445. What is a compensation transaction?

- A. A transaction that pays for database usage
- B. A transaction that semantically undoes the effects of a previously committed transaction
- C. A transaction that compensates for hardware failures
- D. A backup transaction that runs automatically

Answer: B

Q1446. What is snapshot isolation and how does it differ from true serializability?

- A. They are identical
- B. Snapshot isolation provides each transaction a consistent view at its start but may allow write skew anomalies that serializability prevents
- C. Snapshot isolation is stronger than serializability
- D. Snapshot isolation does not prevent dirty reads

Answer: B

Q1447. What is the write skew anomaly in transaction processing?

- A. When a write operation skews data distribution
- B. When two concurrent transactions read overlapping data and make disjoint updates that together violate a constraint
- C. When disk writes are not aligned properly
- D. When write operations are slower than reads

Answer: B

Q1448. What is the phantom problem in concurrency control and how does it differ from a non-repeatable read?

- A. They are the same problem
- B. Phantoms involve new rows appearing in a re-executed query while non-repeatable reads involve changes to existing rows
- C. Phantoms are caused by hardware issues
- D. Non-repeatable reads involve new rows

Answer: B

Q1449. What is predicate locking and why is it used?

- A. Locking based on column names
- B. Locking based on the condition (predicate) of a query to prevent phantom reads
- C. Locking only primary key columns
- D. Locking rows in a predetermined order

Answer: B

Q1450. What is index locking (next-key locking) as a practical alternative to predicate locking?

- A. Locking the entire index structure
- B. Locking index entries and gaps between them to prevent phantoms without the overhead of predicate locking
- C. Locking only the first index entry
- D. Creating a lock on every index in the database

Answer: B

Q1451. In the multiple-granularity locking protocol, what does an IS (Intention Shared) lock on a table indicate?

- A. The transaction intends to insert new rows
- B. The transaction intends to acquire shared locks on some rows within the table
- C. The table is being shared across databases
- D. The index structure is being modified

Answer: B

Q1452. What is the SIX (Shared with Intention Exclusive) lock mode used for?

- A. Locking six data items simultaneously
- B. A transaction that reads all rows but intends to update only some rows in the locked item
- C. A special lock for six-column tables
- D. A combination of six different lock types

Answer: B

Q1453. What is the validation phase in optimistic concurrency control?

- A. Validating SQL syntax before execution
- B. Checking at commit time whether the transaction has conflicted with other transactions during its execution
- C. Validating user credentials
- D. Checking database schema validity

Answer: B

Q1454. What is the difference between deadlock prevention and deadlock detection?

- A. They are the same approach
- B. Prevention avoids deadlocks proactively using protocols while detection allows deadlocks and resolves them when found
- C. Prevention is always more efficient
- D. Detection prevents deadlocks from occurring

Answer: B

Q1455. In timestamp ordering, what happens when a transaction tries to read a data item written by a younger transaction?

- A. The read proceeds normally
- B. The transaction is rolled back and restarted with a new timestamp
- C. The younger transaction is rolled back
- D. Both transactions are aborted

Answer: B

Q1456. What is the concept of lock escalation in database systems?

- A. Increasing the lock timeout duration
- B. Automatically converting many fine-grained locks to fewer coarse-grained locks to reduce overhead
- C. Escalating lock priorities for important transactions
- D. Moving locks from memory to disk

Answer: B

Q1457. What is the relationship between two-phase locking and serializability?

- A. 2PL does not guarantee serializability
- B. 2PL is a sufficient but not necessary condition for conflict serializability
- C. 2PL is both necessary and sufficient for serializability
- D. 2PL only guarantees view serializability

Answer: B

Q1458. In ARIES recovery, why does the redo phase redo all logged actions including those of uncommitted transactions?

- A. To simplify the algorithm
- B. To restore the exact state of the database at crash time (repeating history) before selectively undoing uncommitted changes
- C. To commit all active transactions
- D. To verify log integrity

Answer: B

Q1459. What is the role of the dirty page table in ARIES recovery?

- A. A list of pages with errors
- B. It tracks modified pages in the buffer pool with their recovery LSN to determine the starting point for redo
- C. A list of pages that need cleaning
- D. A table storing deleted pages

Answer: B

Q1460. What is the concept of idempotent redo in recovery algorithms?

- A. Redo operations that can only be performed once
- B. Redo operations that produce the same result regardless of how many times they are applied
- C. Redo operations that are twice as fast
- D. Redo operations that skip duplicate records

Answer: B

Q1461. What is the difference between physical and logical logging?

- A. Physical logging is done on paper while logical is digital
- B. Physical logging records the actual before/after data values while logical logging records the operation performed
- C. Physical logging is faster
- D. They produce identical log files

Answer: B

Q1462. In a no-steal, force buffer policy, why is recovery simpler but performance worse?

- A. Because no recovery is ever needed
- B. No undo is needed (uncommitted data never on disk) and no redo is needed (committed data always on disk), but forcing writes at commit is expensive
- C. Because the log is not needed
- D. Because all transactions are serialized

Answer: B

Q1463. What is the shadow paging recovery technique?

- A. Creating backup pages in shadow memory
- B. Maintaining a shadow page table pointing to the original pages while current modifications go to new pages
- C. Using dark mode for the database interface
- D. Logging shadows of transactions

Answer: B

Q1464. What is the nested top actions concept in ARIES?

- A. Transactions nested inside other transactions
- B. A mechanism to protect certain actions from being undone during rollback while still being part of the transaction
- C. Top-level actions that contain sub-actions
- D. Actions performed at the top of the log file

Answer: B

Q1465. How does log-based recovery handle distributed transactions after a crash?

- A. Each node recovers independently without coordination
- B. The coordinator consults its log to determine the commit decision and informs participants to redo or undo accordingly
- C. All distributed transactions are automatically aborted
- D. Recovery is not possible for distributed transactions

Answer: B

Q1466. What is the purpose of the transaction table in ARIES?

- A. A table listing all tables in the database
- B. It tracks the state and last LSN of each active transaction for determining which transactions to undo
- C. A table of transaction prices
- D. A lookup table for transaction types

Answer: B

Q1467. What is the difference between hot backup and cold backup?

- A. Hot backup uses heated storage media
- B. Hot backup is taken while the database is running while cold backup requires shutting down the database
- C. Cold backup is faster than hot backup
- D. They produce identical backups

Answer: B

Q1468. What is a function-based index?

- A. An index on a column that stores function names
- B. An index built on the result of a function or expression applied to one or more columns
- C. A function that returns index values
- D. An index that uses functional programming concepts

Answer: B

Q1469. What is the difference between a B+ tree index and a log-structured merge tree (LSM tree)?

- A. They are the same structure with different names
- B. B+ trees optimize for reads with in-place updates while LSM trees optimize for writes using sequential I/O and periodic compaction
- C. LSM trees cannot perform range queries
- D. B+ trees are always superior to LSM trees

Answer: B

Q1470. What is a partial index?

- A. An incomplete index that does not work correctly
- B. An index built on a subset of rows that satisfy a specified condition
- C. An index on half the columns of a table
- D. An index that partially sorts data

Answer: B

Q1471. What is the relationship between index height and query performance in a B+ tree?

- A. Index height has no effect on performance
- B. Each additional level of height adds one more disk I/O to the search path
- C. Taller indexes are always faster
- D. Height only affects insertion performance

Answer: B

Q1472. What is a space-partitioning index like an R-tree used for?

- A. Indexing relational tables only
- B. Indexing multidimensional spatial data such as geographic coordinates and geometric objects
- C. Partitioning disk space
- D. Indexing data across multiple tablespaces

Answer: B

Q1473. What is the write amplification problem in index maintenance?

- A. Writes become louder on disk
- B. A single logical write operation causes multiple physical writes due to index updates and page splits
- C. Writing is amplified to be faster
- D. The write buffer size increases

Answer: B

Q1474. What is an inverted index and where is it commonly used?

- A. An index that sorts data in reverse order
- B. A mapping from content (words or terms) to their locations in documents, commonly used in full-text search engines
- C. An index that inverts primary and foreign keys
- D. A reversed B-tree

Answer: B

Q1475. What is the query optimization search space problem?

- A. Finding the optimal physical location for the database
- B. The number of possible execution plans grows exponentially with the number of relations, making exhaustive search impractical
- C. Finding enough disk space for queries
- D. The search for optimal query syntax

Answer: B

Q1476. What is dynamic programming used for in query optimization?

- A. Dynamically creating database tables
- B. Efficiently finding optimal join orderings by building solutions from smaller subproblems
- C. Programming dynamic SQL statements
- D. Dynamically allocating memory for queries

Answer: B

Q1477. What is the concept of interesting orders in query optimization?

- A. Ordering results in a visually appealing way
- B. Sort orders produced by operations that are useful for later operations like merge joins or ORDER BY
- C. The most interesting queries to optimize
- D. Ordering tables by their creation date

Answer: B

Q1478. What is the difference between a left-deep join tree and a bushy join tree?

- A. Left-deep trees are deeper than bushy trees
- B. Left-deep trees have only base relations as the right input, enabling pipelining, while bushy trees allow any join result as input
- C. Bushy trees are always more efficient
- D. Left-deep trees cannot be optimized

Answer: B

Q1479. What is the impact of correlated subqueries on query performance?

- A. They always improve performance
- B. They can be very expensive because the subquery is re-evaluated for each row of the outer query
- C. They have no impact on performance
- D. They speed up the outer query

Answer: B

Q1480. What is the Semi-Join Reduction technique in distributed query optimization?

- A. Reducing the number of joins by half
- B. Reducing data transfer by sending only the join column values to a remote site to filter before transferring matching rows
- C. Joining only half the table
- D. Reducing the join condition to one column

Answer: B

Q1481. What is the role of a plan cache (query cache) in query optimization?

- A. To cache database backups
- B. To store previously computed execution plans for reuse, avoiding repeated optimization overhead
- C. To cache query results permanently
- D. To store database statistics

Answer: B

Q1482. What is the problem with parameter sniffing in query optimization?

- A. Parameters are stolen by unauthorized users
- B. The optimizer creates a plan based on the first parameter value which may be suboptimal for other values
- C. Parameters are not passed correctly to the query
- D. Sniffing causes network latency

Answer: B

Q1483. What is the Volcano/Iterator model of query execution?

- A. A model based on volcanic data processing
- B. Each operator implements a next() method, pulling one tuple at a time from child operators in a demand-driven fashion
- C. A model for processing data in eruptions
- D. A model that stores data in volcanic rock formations

Answer: B

Q1484. What is an inference attack in database security?

- A. An attack that infers database passwords from usernames
- B. An attack where users deduce confidential data from permitted queries by combining and analyzing non-sensitive results
- C. An attack on the database inference engine
- D. An attack that targets the CPU inference pipeline

Answer: B

Q1485. What is differential privacy in the context of databases?

- A. Using different passwords for different databases
- B. A mathematical framework that adds noise to query results to prevent identification of individuals while maintaining statistical utility
- C. Privacy between different database systems
- D. Different levels of privacy for different users

Answer: B

Q1486. What is the polyinstantiation technique in multilevel secure databases?

- A. Creating multiple instances of the database server
- B. Allowing different versions of the same data at different security levels to prevent inference
- C. Installing the database multiple times
- D. Creating multiple copies of tables for performance

Answer: B

Q1487. What is homomorphic encryption and its relevance to database security?

- A. Encryption that uses similar keys for all data
- B. Encryption that allows computations on encrypted data without decrypting it first
- C. Encryption that only works on homogeneous data
- D. Encryption applied uniformly across all tables

Answer: B

Q1488. What is the Biba model and how does it complement Bell-LaPadula?

- A. The Biba model is identical to Bell-LaPadula
- B. Biba focuses on integrity (no write up, no read down) while Bell-LaPadula focuses on confidentiality
- C. Biba focuses on availability while Bell-LaPadula focuses on integrity
- D. They cannot be used together

Answer: B

Q1489. What is k-anonymity and its limitations in database privacy?

- A. Anonymizing k users from the database
- B. Ensuring each record is indistinguishable from at least k-1 other records on quasi-identifiers, but vulnerable to homogeneity and background knowledge attacks
- C. Making data available to k users only
- D. A key management protocol for k users

Answer: B

Q1490. What is database activity monitoring (DAM)?

- A. Monitoring disk activity for defragmentation
- B. Independent monitoring of database activities in real-time to detect and prevent unauthorized access and policy violations
- C. Monitoring the dam that stores database water cooling
- D. Monitoring CPU activity during queries

Answer: B

Q1491. What is the covert channel problem in multilevel secure databases?

- A. A hidden database administrator account
- B. An unintended communication path that allows information to leak between security levels through indirect means
- C. A covert backup channel
- D. A channel used for database replication

Answer: B

Q1492. What is attribute-based access control (ABAC) and how does it extend RBAC?

- A. Controlling access based on data attributes only
- B. Making access decisions based on attributes of users, resources, actions, and environment, providing finer-grained control than role-based approaches
- C. Adding attributes to roles
- D. A type of encryption based on attributes

Answer: B

Q1493. What is the challenge of key management in database encryption?

- A. Keys are too small to manage
- B. Securely storing, rotating, and distributing encryption keys without exposing them while maintaining access to encrypted data
- C. Keys take too much storage space
- D. There are no challenges with key management

Answer: B

Q1494. What is the three-phase commit (3PC) protocol and what problem does it solve?

- A. A three-step backup process
- B. An extension of 2PC that adds a pre-commit phase to avoid blocking when the coordinator fails
- C. A three-step data migration process
- D. A protocol requiring three coordinators

Answer: B

Q1495. What is the Raft consensus algorithm and how does it compare to Paxos?

- A. Raft is unrelated to distributed systems
- B. Raft achieves similar consensus guarantees as Paxos but is designed to be more understandable with a clear leader-election mechanism
- C. Raft is less reliable than Paxos
- D. Raft works only in non-distributed settings

Answer: B

Q1496. What is eventual consistency in distributed databases?

- A. Consistency that is never achieved
- B. A model where all replicas will converge to the same state eventually if no new updates are made, trading immediate consistency for availability
- C. Consistent performance across all queries
- D. A guarantee that data is always consistent

Answer: B

Q1497. What is vector clock and how is it used in distributed databases?

- A. A synchronized physical clock across all sites
- B. A logical clock mechanism that tracks causal ordering of events across distributed nodes using a vector of counters
- C. A clock that measures query execution time
- D. A hardware clock in the database server

Answer: B

Q1498. What is the CALM theorem and its implication for distributed databases?

- A. It states that all distributed systems must be calm and stable
- B. It states that logically monotonic programs can be made consistent without coordination, guiding coordination-free distributed designs
- C. It defines the maximum load for distributed queries
- D. It specifies memory allocation for distributed systems

Answer: B

Q1499. What is a conflict-free replicated data type (CRDT)?

- A. A data type that prevents all database conflicts
- B. A data structure designed so concurrent updates from different replicas can be merged automatically without conflicts
- C. A data type that only allows read operations
- D. A replicated column type

Answer: B

Q1500. What is the quorum-based protocol in distributed databases?

- A. A vote by database administrators
- B. A protocol requiring read and write operations to access a minimum number of replicas (quorum) to ensure consistency
- C. A protocol for electing a new database coordinator
- D. A minimum number of users needed to access the database

Answer: B

Q1501. What is the network partition handling challenge in distributed databases related to the CAP theorem?

- A. Managing network cable partitions
- B. During a network partition, the system must choose between maintaining consistency (rejecting operations) or availability (allowing operations with potential inconsistency)
- C. Partitioning network traffic for different queries
- D. Managing database partitions across different networks

Answer: B

Q1502. What is Google Spanner's approach to distributed consistency?

- A. It ignores consistency for performance
- B. It uses TrueTime API with synchronized atomic clocks to achieve external consistency (linearizability) across globally distributed data
- C. It uses eventual consistency only
- D. It restricts all data to a single region

Answer: B

Q1503. What is the consistent hashing technique used in NoSQL databases?

- A. Using the same hash function across all operations
- B. A technique that maps both data and nodes to a ring, minimizing data redistribution when nodes are added or removed
- C. Hashing data to ensure consistency
- D. A type of encryption for NoSQL data

Answer: B

Q1504. What is the gossip protocol used for in distributed NoSQL systems?

- A. Spreading rumors about database performance
- B. A protocol where nodes periodically exchange state information with random peers to disseminate cluster membership and health data
- C. A chatting protocol between database users
- D. A protocol for sharing query results

Answer: B

Q1505. What is the LSM (Log-Structured Merge) tree and why do many NoSQL databases use it?

- A. A tree for logging system messages
- B. A write-optimized data structure that batches writes in memory and periodically merges sorted runs to disk
- C. A tree structure for merge operations only
- D. A logging mechanism for SQL databases

Answer: B

Q1506. What is tunable consistency in Cassandra?

- A. Consistency that tunes database performance
- B. The ability to configure the consistency level per query, choosing between ONE, QUORUM, ALL, and other levels
- C. Automatic consistency management
- D. A fixed consistency setting

Answer: B

Q1507. What is the concept of anti-entropy in distributed NoSQL databases?

- A. A measure of data disorder
- B. A mechanism that compares and synchronizes data across replicas to repair inconsistencies
- C. The opposite of data entropy
- D. A technique to compress data

Answer: B

Q1508. What are tombstones in NoSQL databases like Cassandra?

- A. Markers for inactive user accounts
- B. Special markers that indicate deleted data, retained temporarily to ensure deletion propagates across all replicas
- C. Physical markers on disk sectors
- D. A type of index entry

Answer: B

Q1509. What is the difference between CP and AP systems in the context of the CAP theorem?

- A. CP systems are faster than AP systems
- B. CP systems sacrifice availability during partitions to maintain consistency, while AP systems sacrifice consistency to remain available
- C. AP systems always have better performance
- D. CP and AP are identical in practice

Answer: B

Q1510. What is a Bloom filter and how is it used in NoSQL databases?

- A. A filter for cleaning database data
- B. A probabilistic data structure that efficiently tests whether an element might be in a set, used to avoid unnecessary disk reads
- C. A filter for bloom-type queries
- D. A data quality filter

Answer: B

Q1511. What is the Dynamo architecture model that influenced many NoSQL databases?

- A. A model for dynamic query processing
- B. A decentralized, highly available key-value architecture using consistent hashing, vector clocks, and sloppy quorum with hinted handoff
- C. A centralized database management model
- D. A model for database security

Answer: B

Q1512. What is a multi-model database and what advantage does it offer?

- A. A database that models multiple data centers
- B. A database supporting multiple data models (document, graph, key-value) in a single engine, reducing operational complexity
- C. A database with multiple backup models
- D. A database that runs multiple instances

Answer: B

Q1513. What is the Kimball methodology for data warehouse design?

- A. A methodology focused on building a centralized enterprise data warehouse first
- B. A bottom-up approach using dimensional modeling, building data marts first that can be integrated into an enterprise warehouse
- C. A methodology that avoids using star schemas
- D. A methodology that does not use ETL processes

Answer: B

Q1514. What is the Inmon methodology for data warehouse design?

- A. A bottom-up approach starting with data marts
- B. A top-down approach building a normalized enterprise data warehouse first, from which data marts are derived
- C. A methodology that uses only star schemas
- D. A methodology that avoids normalization

Answer: B

Q1515. What is a factless fact table?

- A. A fact table with errors
- B. A fact table that records events or conditions without any numeric measures, capturing only the relationship between dimensions
- C. A table that is not a fact table
- D. A fact table with null values

Answer: B

Q1516. What is a junk dimension in data warehouse design?

- A. A dimension containing useless data
- B. A dimension that combines miscellaneous low-cardinality flags and indicators into a single table
- C. A dimension that should be deleted
- D. A dimension with corrupted data

Answer: B

Q1517. What is a degenerate dimension?

- A. A dimension that has degraded over time
- B. A dimension key in the fact table that has no corresponding dimension table, such as an order number
- C. A dimension with no attributes
- D. A dimension that is no longer used

Answer: B

Q1518. What is change data capture (CDC) and how is it used in data warehousing?

- A. Capturing screenshots of data changes
- B. A technique that identifies and captures only the data that has changed in source systems since the last extraction for efficient ETL
- C. Capturing data from change management systems
- D. A method for tracking schema changes

Answer: B

Q1519. What is a data lake and how does it differ from a data warehouse?

- A. A data lake is a smaller version of a data warehouse
- B. A data lake stores raw, unprocessed data in its native format while a data warehouse stores processed, structured data
- C. They are the same concept
- D. A data lake only stores water-related data

Answer: B

Q1520. What is the concept of grain in dimensional modeling?

- A. The texture of stored data
- B. The level of detail represented by each row in a fact table, defining what one row means
- C. The granularity of storage hardware
- D. The size of each data block

Answer: B

Q1521. What is a periodic snapshot fact table versus a transaction fact table?

- A. They store the same data differently
- B. A transaction fact table records individual events while a periodic snapshot captures cumulative state at regular intervals
- C. Periodic snapshots are always larger
- D. Transaction fact tables contain no measures

Answer: B

Q1522. What is an accumulating snapshot fact table?

- A. A fact table that accumulates errors
- B. A fact table that tracks the lifecycle of a process with multiple date columns representing milestones, updating rows as the process progresses
- C. A fact table that only grows and never changes
- D. A temporary fact table

Answer: B

Q1523. What is the LEAD() window function used for in SQL?

- A. To create a leading index
- B. To access a row at a specified physical offset after the current row within the partition
- C. To lead data migration
- D. To find the first row in a partition

Answer: B

Q1524. What is the saga pattern for long-running transactions?

- A. A storytelling approach to database design
- B. A pattern that splits a long transaction into a sequence of local transactions, each with a compensating transaction for rollback
- C. A type of database narrative log
- D. A sequential access pattern for data

Answer: B

Q1525. What is the strict timestamp ordering protocol?

- A. Ordering transactions by their creation date
- B. A protocol where a transaction that attempts to read or write a data item written by a younger transaction is rolled back and restarted
- C. Ordering log entries by time
- D. Strictly formatting timestamps in the log

Answer: B

Q1526. What is the fill factor of a B+ tree index and why does it matter?

- A. The percentage of the database that is indexed
- B. The percentage of space in each index page that is filled with data, affecting split frequency and space utilization
- C. The total size of the index file
- D. The ratio of index pages to data pages

Answer: B

Q1527. What is a skip list and how is it used as an alternative index structure?

- A. A list that skips every other record
- B. A probabilistic data structure using multiple linked list levels where higher levels skip over elements for $O(\log n)$ average search time
- C. A list of skipped transactions
- D. An index that skips null values

Answer: B

Q1528. What is the cost model used to evaluate whether an index scan or table scan is more efficient?

- A. The cost of purchasing the index software
- B. Comparing estimated I/O costs: index scan cost = index traversal I/Os + data page I/Os versus table scan cost = total data pages
- C. The monetary cost of running the query
- D. The time it takes to create the index

Answer: B

Q1529. What is the Selinger optimization algorithm?

- A. An algorithm for selecting database hardware
- B. A dynamic programming-based algorithm from IBM System R that finds optimal join orderings by considering access paths and interesting orders
- C. A security algorithm for databases
- D. An algorithm for selecting storage structures

Answer: B

Q1530. What is the blocking problem in the two-phase commit protocol?

- A. Network traffic blocking data transfer
- B. If the coordinator fails after sending prepare but before sending the final decision, participants holding locks are blocked indefinitely waiting
- C. Users being blocked from accessing the database
- D. Queries being blocked by long-running transactions

Answer: B

SUBJECT

Data Structures & Algorithms

1,080 Multiple Choice Questions

Data Structures & Algorithms — Easy

360 questions

Q1. What is a data structure?

- A. A specific programming language type
- B. A step-by-step problem procedure
- C. A physical hardware component unit
- D. A way to organize and store data

Answer: D

Q2. What does 'algorithm' mean?

- A. A compiled programming language
- B. A primitive data type definition
- C. A processing hardware component
- D. A step-by-step problem procedure

Answer: D

Q3. Which notation is used to describe the upper bound of an algorithm's time complexity?

- A. Small-o notation
- B. Theta notation
- C. Omega notation
- D. Big-O notation

Answer: D

Q4. What is the time complexity of accessing an element in an array by index?

- A. $O(n)$
- B. $O(\log n)$
- C. $O(1)$
- D. $O(n^2)$

Answer: C

Q5. Which of the following is a linear data structure?

- A. Graph
- B. Array
- C. Heap
- D. Tree

Answer: B

Q6. What is the purpose of pseudocode?

- A. To execute low-level machine instructions
- B. To debug and test running software
- C. To compile a program into bytecode
- D. To describe an algorithm in readable form

Answer: D

Q7. Which term refers to the amount of memory an algorithm uses?

- A. Time complexity
- B. Code complexity
- C. Space complexity
- D. Cyclomatic complexity

Answer: C

Q8. What is $O(1)$ complexity called?

- A. Constant
- B. Quadratic
- C. Logarithmic
- D. Linear

Answer: A

Q9. Which of the following is a non-linear data structure?

- A. Linked List
- B. Stack
- C. Array
- D. Tree

Answer: D

Q10. What is an Abstract Data Type (ADT)?

- A. A concrete implementation of a data structure
- B. A reserved programming language keyword term
- C. A specific type of comparison sorting algorithm
- D. A mathematical model defined by its behavior

Answer: D

Q11. What is a stack?

- A. A LIFO data structure
- B. A FIFO data structure
- C. A random access data structure
- D. A hierarchical data structure

Answer: A

Q12. What is a queue?

- A. A sorted data structure
- B. A FIFO data structure
- C. A LIFO data structure
- D. A tree-based data structure

Answer: B

Q13. Which operation adds an element to the top of a stack?

- A. Push
- B. Dequeue
- C. Enqueue
- D. Pop

Answer: A

Q14. What is a linked list?

- A. Nodes where each node points to the next node
- B. A hierarchical tree-based data structure
- C. Elements stored in contiguous memory locations
- D. A hash-based key-value pair structure

Answer: A

Q15. What is the time complexity of inserting an element at the beginning of a singly linked list?

- A. $O(n^2)$
- B. $O(\log n)$
- C. $O(1)$
- D. $O(n)$

Answer: C

Q16. Which data structure is used to evaluate postfix expressions?

- A. Array
- B. Stack
- C. Queue
- D. Linked List

Answer: B

Q17. What is the main disadvantage of an array?

- A. Random access is not possible
- B. Fixed size in static arrays
- C. It cannot store integers
- D. Elements cannot be sorted

Answer: B

Q18. What does the 'dequeue' operation do?

- A. Adds an element to the front of a queue
- B. Removes an element from the front of a queue
- C. Adds an element to the rear of a queue
- D. Removes an element from the rear of a queue

Answer: B

Q19. In a singly linked list, the last node points to:

- A. The first node
- B. NULL (nothing)
- C. The middle node
- D. Itself (self)

Answer: B

Q20. Which operation removes an element from the top of a stack?

- A. Enqueue
- B. Push
- C. Peek
- D. Pop

Answer: D

Q21. What is a binary tree?

- A. A directed graph with no cycles at all
- B. A tree with exactly 2 total internal nodes
- C. A tree where each node has at most 2 children
- D. A tree where each node has at most 3 children

Answer: C

Q22. What is the root of a tree?

- A. The node with the most children
- B. The topmost node with no parent
- C. The node with no children at all
- D. The leftmost node in the tree

Answer: B

Q23. What is a leaf node?

- A. A node with two children
- B. A node with no children
- C. The root node itself
- D. A node with one child

Answer: B

Q24. What is a Binary Search Tree (BST)?

- A. A binary tree where left child > parent > right child
- B. A tree data structure with exactly two total nodes
- C. A binary tree where left child < parent < right child
- D. A balanced tree structure that is always complete

Answer: C

Q25. What is a graph in data structures?

- A. A chart used for data visualization
- B. A type of contiguous memory array
- C. A sequential linear data structure
- D. A collection of vertices and edges

Answer: D

Q26. What is the maximum number of nodes in a binary tree of height h?

- A. $2^h - 1$
- B. h^2
- C. $2h$
- D. $2^{(h+1)} - 1$

Answer: D

Q27. What is an undirected graph?

- A. A graph where edges have no direction
- B. A graph with no edges at all
- C. A hierarchical tree structure
- D. A graph where edges have direction

Answer: A

Q28. What is the degree of a vertex in an undirected graph?

- A. The assigned weight of the given vertex
- B. The number of edges incident to the vertex
- C. The total number of vertices in the graph
- D. The shortest path length from the vertex

Answer: B

Q29. What is a heap?

- A. A type of directed graph with cycles
- B. A complete binary tree with heap property
- C. A sorted array-based storage structure
- D. A variant of a hash table with chains

Answer: B

Q30. How is a graph typically represented?

- A. Using adjacency matrix or list
- B. Using singly linked lists only
- C. Using stack data structures
- D. Using simple arrays only

Answer: A

Q31. What is linear search?

- A. Searching using a hash table for lookups
- B. Searching only in a sorted array format
- C. Searching each element sequentially from start
- D. Searching by dividing the array in half

Answer: C

Q32. What is the time complexity of linear search in the worst case?

- A. $O(n^2)$
- B. $O(\log n)$
- C. $O(n)$
- D. $O(1)$

Answer: C

Q33. What is the prerequisite for binary search?

- A. The array must be sorted
- B. The array must be of fixed size
- C. The array must contain unique elements
- D. The array must be unsorted

Answer: A

Q34. What is the time complexity of binary search?

- A. $O(\log n)$
- B. $O(n^2)$
- C. $O(n)$
- D. $O(1)$

Answer: A

Q35. In binary search, if the target is greater than the middle element, where do you search next?

- A. Start over again
- B. Middle element only
- C. Left half only
- D. Right half only

Answer: D

Q36. What is the best-case time complexity of linear search?

- A. $O(1)$
- B. $O(n)$
- C. $O(\log n)$
- D. $O(n^2)$

Answer: A

Q37. Which search algorithm does not require the data to be sorted?

- A. Interpolation search
- B. Linear search
- C. Fibonacci search
- D. Binary search

Answer: B

Q38. What does binary search return if the element is not found?

- A. The middle index of the array
- B. The value zero always
- C. The last index of the array
- D. An indicator such as -1 for not found

Answer: D

Q39. How many comparisons does binary search make for an array of 1024 elements in the worst case?

- A. 100
- B. 10
- C. 1024
- D. 512

Answer: B

Q40. What is the space complexity of iterative binary search?

- A. $O(\log n)$
- B. $O(n)$
- C. $O(n^2)$
- D. $O(1)$

Answer: D

Q41. What is bubble sort?

- A. A sort that divides the array in half each time
- B. A sort that selects the minimum element each pass
- C. A sort that uses a heap for element ordering
- D. A sort that swaps adjacent out-of-order elements

Answer: D

Q42. What is the worst-case time complexity of bubble sort?

- A. $O(\log n)$
- B. $O(n \log n)$
- C. $O(n^2)$
- D. $O(n)$

Answer: C

Q43. What is selection sort?

- A. A sort that works only on linked list nodes
- B. A sort that uses binary search for placement
- C. A sort that finds minimum and places it correctly
- D. A sort that selects a random pivot element

Answer: C

Q44. What is insertion sort?

- A. A sort that uses divide and conquer recursion
- B. A sort that inserts each element into correct position
- C. A sort that only works on integer data types
- D. A sort that inserts elements into a heap structure

Answer: B

Q45. Which sorting algorithm is best for nearly sorted data?

- A. Merge sort
- B. Quick sort
- C. Selection sort
- D. Insertion sort

Answer: D

Q46. Is bubble sort a stable sorting algorithm?

- A. Yes, it is stable
- B. Only for integer data
- C. Only for string data
- D. No, it is unstable

Answer: A

Q47. What is the best-case time complexity of insertion sort?

- A. $O(n \log n)$
- B. $O(n^2)$
- C. $O(n)$
- D. $O(1)$

Answer: C

Q48. What is the space complexity of selection sort?

- A. $O(n^2)$
- B. $O(1)$
- C. $O(n)$
- D. $O(\log n)$

Answer: B

Q49. How many passes does bubble sort need in the worst case for n elements?

- A. $n - 1$
- B. $\log n$
- C. n
- D. $n/2$

Answer: A

Q50. Is selection sort a stable sorting algorithm?

- A. Only for numeric data
- B. No, it is unstable
- C. Yes, it is stable
- D. Depends on implementation

Answer: B

Q51. What is a hash function?

- A. A function that encrypts all data
- B. A function that compresses file data
- C. A function that sorts data elements
- D. A function that maps data to fixed-size value

Answer: D

Q52. What is a hash table?

- A. A binary tree data structure
- B. A sorted array structure
- C. A type of linked list structure
- D. A key-value store using hash function

Answer: D

Q53. What is a collision in hashing?

- A. When a key is not found in the table
- B. When two keys hash to the same index
- C. When the hash table is completely full
- D. When the hash function fails to compute

Answer: B

Q54. What is the average time complexity of lookup in a hash table?

- A. $O(n)$
- B. $O(\log n)$
- C. $O(n^2)$
- D. $O(1)$

Answer: D

Q55. What is chaining in hash tables?

- A. Linking multiple hash tables together
- B. Connecting hash functions in sequence
- C. Storing colliding elements in a linked list
- D. A type of hash function computation

Answer: C

Q56. What is the load factor of a hash table?

- A. The ratio of entries to number of slots
- B. The number of hash functions being used
- C. The total number of collisions occurred
- D. The total size of the hash table

Answer: A

Q57. Which of the following is a simple hash function for integers?

- A. $h(k) = k + \text{table\_size}$
- B. $h(k) = k \times \text{table\_size}$
- C. $h(k) = k / \text{table\_size}$
- D. $h(k) = k \bmod \text{table\_size}$

Answer: D

Q58. What is the worst-case time complexity of hash table lookup?

- A. $O(1)$
- B. $O(n^2)$
- C. $O(\log n)$
- D. $O(n)$

Answer: D

Q59. What is the purpose of a hash table?

- A. To store data in sorted ascending order
- B. To compress data for storage saving
- C. To sort data elements efficiently
- D. To provide fast insert, delete, and lookup

Answer: D

Q60. What happens when you insert a key that already exists in a hash table?

- A. The insertion is silently ignored completely
- B. Both old and new values are stored
- C. The old value is typically updated with new
- D. It causes a runtime error exception

Answer: C

Q61. What is in-order traversal of a binary tree?

- A. Right, Root, Left
- B. Left, Right, Root
- C. Left, Root, Right
- D. Root, Left, Right

Answer: C

Q62. What is pre-order traversal of a binary tree?

- A. Right, Left, Root
- B. Root, Left, Right
- C. Left, Right, Root
- D. Left, Root, Right

Answer: B

Q63. What is post-order traversal of a binary tree?

- A. Right, Root, Left
- B. Root, Left, Right
- C. Left, Root, Right
- D. Left, Right, Root

Answer: D

Q64. What data structure is used for level-order (BFS) traversal of a tree?

- A. Array
- B. Queue
- C. Hash table
- D. Stack

Answer: B

Q65. What is the height of a single-node tree?

- A. 1
- B. 0
- C. -1
- D. 2

Answer: B

Q66. How do you find the minimum element in a BST?

- A. Check the root node
- B. Traverse to the rightmost node
- C. Perform in-order traversal
- D. Traverse to the leftmost node

Answer: D

Q67. How do you find the maximum element in a BST?

- A. Check the root node
- B. Traverse to the leftmost node
- C. Traverse to the rightmost node
- D. Use binary search method

Answer: C

Q68. What is the time complexity of searching for an element in a BST of height h?

- A. $O(1)$
- B. $O(n^2)$
- C. $O(h)$
- D. $O(n)$

Answer: C

Q69. What is a binary tree's depth?

- A. The edge count from root to deepest leaf
- B. The total number of nodes in tree
- C. The number of internal non-leaf nodes
- D. The total number of leaf nodes only

Answer: A

Q70. Which traversal of a BST gives elements in sorted order?

- A. Level-order
- B. In-order
- C. Post-order
- D. Pre-order

Answer: B

Q71. What is Breadth-First Search (BFS)?

- A. A search that goes as deep as possible first
- B. A search only applicable to tree structures
- C. A search that uses a stack data structure
- D. A search exploring all neighbors at current depth

Answer: D

Q72. What is Depth-First Search (DFS)?

- A. A search going deep along each branch first
- B. A search only applicable to array structures
- C. A search that uses a queue data structure
- D. A search that explores all neighbors first

Answer: A

Q73. What data structure does BFS use?

- A. Hash table
- B. Heap
- C. Stack
- D. Queue

Answer: D

Q74. What data structure does DFS use?

- A. Array structure
- B. Stack or recursion
- C. Queue structure
- D. Heap structure

Answer: B

Q75. What is the time complexity of BFS on a graph with V vertices and E edges?

- A. $O(V + E)$
- B. $O(E)$
- C. $O(V \times E)$
- D. $O(V)$

Answer: A

Q76. What is the time complexity of DFS on a graph with V vertices and E edges?

- A. $O(V + E)$
- B. $O(V \times E)$
- C. $O(E^2)$
- D. $O(V^2)$

Answer: A

Q77. Can BFS find the shortest path in an unweighted graph?

- A. Yes, BFS can do this
- B. No, BFS cannot do this
- C. Only in tree structures
- D. Only in directed graphs

Answer: A

Q78. What is a connected component in an undirected graph?

- A. A cycle within the undirected graph
- B. A maximal set of vertices all connected by paths
- C. An edge connecting two vertices together
- D. A single isolated vertex in the graph

Answer: B

Q79. What is a cycle in a graph?

- A. An isolated vertex with no edges at all
- B. A path starting and ending at the same vertex
- C. A straight path between two vertices
- D. A tree edge connecting parent to child

Answer: B

Q80. What is a spanning tree of a connected graph?

- A. A subgraph with all vertices and all edges
- B. A complete subgraph with maximum edges
- C. A tree subgraph containing all graph vertices
- D. The shortest path tree from one source

Answer: C

Q81. What is the divide and conquer strategy?

- A. Breaking into subproblems, solving, and combining
- B. Solving all problems using bottom-up iteration
- C. Using a greedy approach at each step
- D. Solving problems by trying all possibilities

Answer: A

Q82. What is a greedy algorithm?

- A. An algorithm making locally optimal choices each step
- B. An algorithm that solves all subproblems first
- C. An algorithm that uses backtracking for search
- D. An algorithm that tries all possible combinations

Answer: A

Q83. Which of the following is an example of divide and conquer?

- A. Merge sort
- B. Linear search
- C. Insertion sort
- D. Bubble sort

Answer: A

Q84. What is dynamic programming?

- A. A type of greedy algorithm with memoization
- B. Solving problems by storing overlapping subproblem results
- C. Programming that changes dynamically at runtime
- D. An algorithm that uses only recursion for solving

Answer: B

Q85. What is backtracking?

- A. Sorting data elements in reverse descending order
- B. Trying solutions and abandoning invalid paths early
- C. Going back to the start of an algorithm run
- D. Reversing the input data before processing

Answer: B

Q86. What is memoization?

- A. Writing memos about code documentation
- B. A type of sorting with memory optimization
- C. Memorizing the algorithm steps by heart
- D. Caching results of function calls for reuse

Answer: D

Q87. Which approach builds solutions from smallest subproblems to larger ones?

- A. Greedy approach
- B. Top-down approach
- C. Bottom-up (tabulation)
- D. Brute force approach

Answer: C

Q88. What is brute force in algorithm design?

- A. An optimized search using heuristics
- B. Trying all possible solutions exhaustively
- C. The most efficient approach available
- D. A divide and conquer optimization

Answer: B

Q89. What two properties must a problem have to be solvable by dynamic programming?

- A. Optimal substructure and overlapping subproblems
- B. Linear structure and constant time operations
- C. Greedy choice property and unique solution
- D. Sorting and searching capabilities

Answer: A

Q90. Which algorithm design paradigm does the activity selection problem use?

- A. Backtracking search approach
- B. Dynamic programming approach
- C. Divide and conquer approach
- D. Greedy algorithm approach

Answer: D

Q91. What is a priority queue?

- A. A type of stack that also supports priority ordering
- B. A queue that processes all elements in strict FIFO order
- C. A sorted array of elements ordered by their value
- D. A structure where higher priority elements are served first

Answer: D

Q92. What data structure typically implements a priority queue?

- A. Stack structure
- B. Heap structure
- C. Array structure
- D. Linked list structure

Answer: B

Q93. What is a trie (prefix tree)?

- A. A tree for storing and retrieving strings by prefix
- B. A graph data structure for string matching
- C. A binary search tree for numeric keys
- D. A type of hash table for string storage

Answer: A

Q94. What is the time complexity of searching for a word of length L in a trie?

- A. $O(\log n)$
- B. $O(L)$
- C. $O(n)$
- D. $O(n \times L)$

Answer: B

Q95. What is a disjoint set (Union-Find)?

- A. A hash set with unique elements
- B. A structure tracking non-overlapping subsets
- C. A set with no elements at all
- D. A sorted set with ordered elements

Answer: B

Q96. What operation does 'find' perform in Union-Find?

- A. Determines which set an element belongs to
- B. Searches for an element in the set
- C. Finds the minimum element in all sets
- D. Finds all elements in a particular set

Answer: A

Q97. What is a min-heap?

- A. A heap where minimum element is at root
- B. A heap where maximum element is at root
- C. A sorted array of elements by value
- D. A balanced BST with min at leftmost

Answer: A

Q98. What is the main advantage of a trie over a hash table for string operations?

- A. Has a simpler implementation overall
- B. Uses less memory usage in practice
- C. Supports prefix-based queries efficiently
- D. Provides faster exact string lookups

Answer: C

Q99. What does 'union' operation do in Union-Find?

- A. Splits a set into two parts
- B. Finds common elements between sets
- C. Merges two disjoint sets into one
- D. Sorts the sets in ascending order

Answer: C

Q100. What is a multiset?

- A. A sorted set with unique elements
- B. A collection allowing duplicate elements
- C. A set with multiple data types
- D. A set of sets nested together

Answer: B

Q101. What is string matching?

- A. Comparing two strings for equality only
- B. Finding occurrences of a pattern in text
- C. Sorting strings in alphabetical order
- D. Concatenating two strings end to end

Answer: B

Q102. What is the brute force string matching time complexity for text of length n and pattern of length m ?

- A. $O(n)$
- B. $O(n + m)$
- C. $O(m)$
- D. $O(n \times m)$

Answer: D

Q103. What is a substring?

- A. A rearrangement of string characters
- B. A contiguous sequence of characters within a string
- C. A string of exactly length 1 character
- D. Any characters from a string in any order

Answer: B

Q104. What is a subsequence of a string?

- A. A sequence derived by deleting characters keeping order
- B. A contiguous part of the original string
- C. A sorted version of the original string
- D. A reversed version of the original string

Answer: A

Q105. What is a palindrome?

- A. A string with only repeated characters throughout
- B. A string with all unique characters
- C. A string with an even number of characters
- D. A string that reads the same forward and backward

Answer: D

Q106. What is an anagram?

- A. A palindrome of another word
- B. A prefix of another given word
- C. A substring of another word
- D. A rearrangement forming another word

Answer: D

Q107. How can you check if two strings are anagrams?

- A. Compare their lengths only for equality
- B. Check if they have the same first character
- C. Reverse one string and compare with the other
- D. Sort both strings and compare, or count frequencies

Answer: D

Q108. What is a prefix of a string?

- A. The middle portion of the string
- B. A reversed version of the string
- C. The last character of the string
- D. A substring starting from the beginning

Answer: D

Q109. What is a suffix of a string?

- A. A substring starting from the beginning
- B. A rotated version of the string
- C. The first character of the string
- D. A substring ending at the last character

Answer: D

Q110. What is string concatenation?

- A. Splitting a string into parts
- B. Reversing a string entirely
- C. Joining two strings end to end
- D. Sorting characters in a string

Answer: C

Q111. What does P stand for in computational complexity?

- A. Problems solvable in polynomial time
- B. Prime number problems class
- C. Probability problems class
- D. Parallel problems class

Answer: A

Q112. What does NP stand for in computational complexity?

- A. New computational Problems class
- B. Non-deterministic Polynomial time
- C. Not Polynomial time class
- D. No Proof available for it

Answer: B

Q113. What is the best time complexity achievable for comparison-based sorting?

- A. $O(\log n)$
- B. $O(n \log n)$
- C. $O(n)$
- D. $O(n^2)$

Answer: B

Q114. Which is more efficient: $O(n)$ or $O(n \log n)$?

- A. Depends on constants
- B. $O(n)$ is more efficient
- C. $O(n \log n)$ is better
- D. They are equivalent

Answer: B

Q115. What is the time complexity of accessing an element in a hash table on average?

- A. $O(n^2)$
- B. $O(\log n)$
- C. $O(1)$
- D. $O(n)$

Answer: C

Q116. What is time-space tradeoff?

- A. Time complexity cannot affect space complexity
- B. Time and space are always equal in measure
- C. Using more time to save space or vice versa
- D. A method to reduce both time and space simultaneously

Answer: C

Q117. Which is faster: $O(\log n)$ or $O(n)$?

- A. They are equivalent
- B. $O(\log n)$ is faster
- C. $O(n)$ is faster
- D. Cannot compare them

Answer: B

Q118. What is the significance of $O(n!)$ complexity?

- A. It is essentially the same as $O(n^2)$ complexity class
- B. It is classified as a polynomial time complexity class
- C. It is very efficient for all possible input sizes given
- D. It is extremely inefficient, growing super-exponentially

Answer: D

Q119. What does 'in-place algorithm' mean?

- A. An algorithm that runs at a specific location
- B. An algorithm using only constant extra space
- C. An algorithm that cannot be relocated in memory
- D. An algorithm that only sorts data in place

Answer: B

Q120. What is the time complexity of finding the maximum element in an unsorted array?

- A. $O(n)$
- B. $O(n^2)$
- C. $O(1)$
- D. $O(\log n)$

Answer: A

Q121. What does FIFO stand for in data structures?

- A. Fast Input Fast Output
- B. First Index First Order
- C. First In First Out
- D. Fixed In Fixed Out

Answer: C

Q122. Which term describes a function calling itself?

- A. Abstraction
- B. Iteration
- C. Recursion
- D. Selection

Answer: C

Q123. What is the base case in a recursive function?

- A. The largest subproblem solved by recursion
- B. The first recursive call made by function
- C. The condition that stops the recursion calls
- D. The condition that starts the recursion loop

Answer: C

Q124. What does LIFO stand for in data structures?

- A. Linear Input Fixed Output
- B. Last Index First Output
- C. Least Index First Order
- D. Last In First Out

Answer: D

Q125. Which complexity class grows the fastest?

- A. Linear
- B. Logarithmic
- C. Exponential
- D. Polynomial

Answer: C

Q126. What is the time complexity of a simple for loop iterating n times?

- A. $O(n)$
- B. $O(1)$
- C. $O(n^2)$
- D. $O(\log n)$

Answer: A

Q127. What does the term 'traversal' mean in data structures?

- A. Copying a structure into another one
- B. Sorting elements in ascending order only
- C. Deleting all elements from a structure
- D. Visiting each element exactly one time

Answer: D

Q128. Which data structure uses the LIFO principle?

- A. Stack
- B. Queue
- C. Graph
- D. Array

Answer: A

Q129. What is a primitive data type?

- A. A complex object-oriented class type
- B. A user-defined composite structure type
- C. A dynamically allocated reference type
- D. A basic built-in type provided by language

Answer: D

Q130. What is the purpose of a flowchart?

- A. To encrypt data for secure transmission
- B. To visually represent an algorithm's steps
- C. To compile source code into binary files
- D. To allocate memory for data structures

Answer: B

Q131. What is the maximum number of elements a stack can pop at once?

- A. Half of them
- B. All elements
- C. Two elements
- D. One element

Answer: D

Q132. Which operation adds an element to the rear of a queue?

- A. Enqueue
- B. Pop
- C. Push
- D. Dequeue

Answer: A

Q133. What is a singly linked list?

- A. Each node contains two separate data field values
- B. Each node points to both next and previous nodes
- C. Each node points only to the next node in list
- D. Each node has a direct link to every other node

Answer: C

Q134. What is stored in each node of a linked list?

- A. Data and a pointer to the next node
- B. Only the memory address of next node
- C. A pointer to both previous and next nodes
- D. Only the data value without any pointer

Answer: A

Q135. Which data structure is used for undo operations in editors?

- A. Queue
- B. Graph
- C. Stack
- D. Array

Answer: C

Q136. What happens when you dequeue from an empty queue?

- A. It inserts a default placeholder element
- B. It returns a null or zero value silently
- C. It causes an underflow error condition
- D. It resets the queue to initial capacity

Answer: C

Q137. What is the index of the first element in a zero-based array?

- A. n
- B. -1
- C. 1
- D. 0

Answer: D

Q138. Which linear structure allows insertion at both ends?

- A. Deque
- B. Array
- C. Queue
- D. Stack

Answer: A

Q139. What is the time complexity of inserting at the head of a linked list?

- A. $O(1)$
- B. $O(n^2)$
- C. $O(n)$
- D. $O(\log n)$

Answer: A

Q140. What is a static array?

- A. An array that automatically sorts its elements
- B. An array with a fixed size set at creation
- C. An array whose size can grow dynamically
- D. An array that stores only string data types

Answer: B

Q141. What is the root node of a tree?

- A. Any leaf node at the bottom of the tree
- B. The node with the most children in tree
- C. The node at the deepest level of tree
- D. The topmost node with no parent at all

Answer: D

Q142. What is a leaf node in a tree?

- A. A node that connects to the root directly
- B. A node that has no children at all
- C. A node that has exactly two child nodes
- D. The first node added to the tree structure

Answer: B

Q143. How many children can a binary tree node have at most?

- A. Four
- B. Three
- C. Two
- D. One

Answer: C

Q144. What is an edge in a graph?

- A. A connection between two vertices
- B. A path visiting all graph vertices
- C. An isolated node with no neighbors
- D. The weight assigned to each vertex

Answer: A

Q145. What is an undirected graph?

- A. A graph where edges have specific direction only
- B. A graph where edges have no specific direction
- C. A graph where all vertices are fully connected
- D. A graph with no edges between any vertices

Answer: B

Q146. What is the height of a single-node tree?

- A. One
- B. Two
- C. Undefined
- D. Zero

Answer: D

Q147. What is a weighted graph?

- A. A graph where all vertices have the same degree
- B. A graph with no cycles between any of its nodes
- C. A graph where every node connects to all others
- D. A graph where edges have assigned numerical values

Answer: D

Q148. What is the degree of a vertex in an undirected graph?

- A. The total weight of all edges from that vertex
- B. The number of edges connected to the vertex
- C. The shortest path from that vertex to root
- D. The number of vertices in the entire graph

Answer: B

Q149. What is a complete binary tree?

- A. A tree where every internal node has no children
- B. A tree with only leaf nodes and no internal nodes
- C. A tree where every node has exactly one child node
- D. A tree where all levels are fully filled except last

Answer: D

Q150. What is an adjacency list used for?

- A. Finding shortest path in weighted graph directly
- B. Storing tree nodes in level-order arrangement
- C. Storing graph structure as lists of neighbors
- D. Sorting vertices in topological ordering only

Answer: C

Q151. What does linear search do?

- A. Uses hashing to find the target element value
- B. Sorts elements first then searches for target
- C. Checks elements sequentially from start to end
- D. Divides the array in half at each comparison

Answer: C

Q152. What is the best-case time complexity of linear search?

- A. $O(n)$
- B. $O(\log n)$
- C. $O(n^2)$
- D. $O(1)$

Answer: D

Q153. What is required for binary search to work correctly?

- A. The array must be sorted beforehand
- B. The array must have an even number of items
- C. The array must contain unique elements only
- D. The array must be stored in a linked list

Answer: A

Q154. What is the worst-case time complexity of binary search?

- A. $O(\log n)$
- B. $O(n)$
- C. $O(n^2)$
- D. $O(1)$

Answer: A

Q155. In binary search, what happens if the target is less than the middle element?

- A. Search continues in the right half only
- B. The search terminates with element not found
- C. Search continues in the left half only
- D. Search restarts from the beginning again

Answer: C

Q156. Which search algorithm does not require sorted data?

- A. Binary search
- B. Fibonacci search
- C. Interpolation search
- D. Linear search

Answer: D

Q157. What is the space complexity of iterative binary search?

- A. $O(n^2)$
- B. $O(n)$
- C. $O(\log n)$
- D. $O(1)$

Answer: D

Q158. What is the worst-case time complexity of linear search?

- A. $O(n^2)$
- B. $O(n)$
- C. $O(\log n)$
- D. $O(1)$

Answer: B

Q159. What value does a search algorithm return if the element is not found?

- A. Typically negative one or a sentinel value
- B. The last index of the array searched
- C. The middle index of the array searched
- D. The first index of the array searched

Answer: A

Q160. How many comparisons does binary search make on 1024 sorted elements at most?

- A. 1024
- B. 10
- C. 100
- D. 512

Answer: B

Q161. What is the basic idea behind bubble sort?

- A. Insert each element into its correct sorted position
- B. Select the minimum element and place it at start
- C. Repeatedly swap adjacent elements if they are wrong
- D. Divide array into halves and merge them sorted

Answer: C

Q162. What is the time complexity of bubble sort in the worst case?

- A. $O(n^2)$
- B. $O(n)$
- C. $O(n \log n)$
- D. $O(\log n)$

Answer: A

Q163. Which sorting algorithm picks the minimum and places it first?

- A. Insertion sort
- B. Merge sort
- C. Selection sort
- D. Quick sort

Answer: C

Q164. Is insertion sort a stable sorting algorithm?

- A. Yes, it is stable
- B. It depends on input
- C. Only for integers
- D. No, it is never stable

Answer: A

Q165. What is the best-case time complexity of insertion sort?

- A. $O(\log n)$
- B. $O(n^2)$
- C. $O(n \log n)$
- D. $O(n)$

Answer: D

Q166. What does the term 'in-place sorting' mean?

- A. Sorting that requires $O(n)$ extra memory space
- B. Sorting performed on an external disk storage device
- C. Sorting that creates a new copy of the input array
- D. Sorting done without significant extra memory use

Answer: D

Q167. Which sort is most efficient for nearly sorted data?

- A. Insertion sort
- B. Heap sort
- C. Merge sort
- D. Selection sort

Answer: A

Q168. What is the worst-case time complexity of selection sort?

- A. $O(n)$
- B. $O(\log n)$
- C. $O(n \log n)$
- D. $O(n^2)$

Answer: D

Q169. What does 'stable sort' mean?

- A. A sort preserving relative order of equal elements
- B. A sort with guaranteed $O(n \log n)$ time complexity
- C. A sort that always completes without crashing at all
- D. A sort that works on all data types without error

Answer: A

Q170. How many passes does bubble sort need for n elements?

- A. Exactly n passes
- B. Exactly one pass
- C. At most n-1 passes
- D. At most $\log n$ passes

Answer: C

Q171. What is a hash function?

- A. A function that sorts data into ordered buckets
- B. A function that compresses data to save disk space
- C. A function that encrypts data for secure storage
- D. A function that maps data to a fixed-size value

Answer: D

Q172. What is a hash table?

- A. A tree structure organized by hash value ordering
- B. A data structure using keys mapped to array indices
- C. A linked list with hash-based pointer connections
- D. A table that stores elements in sorted order only

Answer: B

Q173. What is a collision in hashing?

- A. When a hash table runs out of available memory
- B. When two keys map to the same hash table index
- C. When the hash function produces a negative value
- D. When an element is deleted from the hash table

Answer: B

Q174. What is the average time complexity of hash table lookup?

- A. $O(\log n)$
- B. $O(n^2)$
- C. $O(1)$
- D. $O(n)$

Answer: C

Q175. What is chaining in hash tables?

- A. Linking hash tables together in a sequence chain
- B. Connecting array slots with pointers to next slot
- C. Chaining multiple hash functions for better results
- D. Storing colliding elements in a linked list at each slot

Answer: D

Q176. What is the load factor of a hash table?

- A. The ratio of stored elements to table size
- B. The maximum key value stored in the table
- C. The number of collisions per hash operation
- D. The total memory allocated for the hash table

Answer: A

Q177. What is the division method of hashing?

- A. Dividing the key by two repeatedly until zero
- B. Splitting the key into digits and summing them up
- C. Dividing the table into equal-sized partitions only
- D. Computing hash as key modulo the table size value

Answer: D

Q178. What is an ideal property of a good hash function?

- A. It should distribute keys uniformly across the table
- B. It should always produce the same output value
- C. It should map all keys to a single table slot only
- D. It should produce values in strictly ascending order

Answer: A

Q179. What is direct addressing in hashing?

- A. Using random numbers to determine storage locations
- B. Using a linked list to store all elements together
- C. Using a tree structure to organize hash table entries
- D. Using the key itself as the index into the array

Answer: D

Q180. What data structure does a hash set implement?

- A. A collection of unique elements with fast lookup
- B. A hierarchical tree of categorized elements
- C. An ordered sequence of duplicate elements
- D. A sorted list of key-value paired entries

Answer: A

Q181. What is an in-order traversal of a binary tree?

- A. Visit root, then left subtree, then right subtree
- B. Visit left subtree, then root, then right subtree
- C. Visit root, then right subtree, then left subtree
- D. Visit left subtree, then right subtree, then root

Answer: B

Q182. What does pre-order traversal visit first?

- A. Right child
- B. Left child
- C. Leaf node
- D. Root node

Answer: D

Q183. What does post-order traversal visit last?

- A. Right child
- B. Leaf node
- C. Left child
- D. Root node

Answer: D

Q184. What is level-order traversal?

- A. Visiting nodes level by level from top to bottom
- B. Visiting only nodes at even numbered levels only
- C. Visiting nodes from deepest level to root level
- D. Visiting leaf nodes first then internal nodes last

Answer: A

Q185. Which data structure is used for level-order traversal?

- A. Array
- B. Queue
- C. Heap
- D. Stack

Answer: B

Q186. What is the time complexity of searching in a BST with n nodes?

- A. $O(n^2)$ always
- B. $O(n)$ in worst case
- C. $O(\log n)$ always
- D. $O(1)$ in worst case

Answer: B

Q187. What is the in-order successor of a node in a BST?

- A. The leftmost child of the given node always
- B. The parent node of the given node always
- C. The rightmost node in the entire BST tree
- D. The node with the next larger value in tree

Answer: D

Q188. What does in-order traversal of a BST produce?

- A. Elements in random unsorted order
- B. Elements in level-wise grouped order
- C. Elements in descending sorted order
- D. Elements in ascending sorted order

Answer: D

Q189. What is the minimum value node in a BST?

- A. The leftmost node in the whole tree
- B. The root node of the entire tree
- C. The rightmost node in the whole tree
- D. Any leaf node in the entire tree

Answer: A

Q190. How many traversal orders are standard for binary trees?

- A. Four
- B. Three
- C. Two
- D. Six

Answer: A

Q191. What is Breadth-First Search (BFS)?

- A. Randomly visiting nodes until target is found finally
- B. Searching only the leftmost path in the graph first
- C. Exploring as deep as possible before backtracking
- D. Exploring all neighbors at current depth before deeper

Answer: D

Q192. What is Depth-First Search (DFS)?

- A. Exploring as deep as possible then backtracking up
- B. Visiting nodes in alphabetical order of their labels
- C. Exploring all neighbors before going to next level
- D. Only visiting nodes that are directly connected to root

Answer: A

Q193. Which data structure does BFS use?

- A. Heap
- B. Queue
- C. Array
- D. Stack

Answer: B

Q194. Which data structure does DFS typically use?

- A. Queue
- B. Stack
- C. Heap
- D. Deque

Answer: B

Q195. What is the time complexity of BFS on a graph with V vertices and E edges?

- A. $O(E^2)$
- B. $O(V^2)$
- C. $O(V * E)$
- D. $O(V + E)$

Answer: D

Q196. What is a spanning tree of a graph?

- A. A complete copy of the original graph with all edges
- B. A tree containing only the leaf vertices of the graph
- C. A subgraph containing all vertices with minimum edges
- D. A subgraph that connects only half of the graph vertices

Answer: C

Q197. What is a connected graph?

- A. A graph with no isolated vertices but possible missing edges
- B. A graph where there is a path between every pair of nodes
- C. A graph where all vertices have the same degree exactly
- D. A graph where every pair of vertices has a direct edge

Answer: B

Q198. What is a path in a graph?

- A. A sequence of vertices connected by edges in graph
- B. A cycle that visits every vertex exactly once only
- C. The shortest distance between two vertices in graph
- D. A single edge connecting exactly two graph vertices

Answer: A

Q199. What is a cycle in a graph?

- A. A subgraph with all vertices having degree exactly two
- B. A path that starts and ends at the same vertex
- C. An edge that connects a vertex to itself directly
- D. A path that visits every edge in the graph once

Answer: B

Q200. What is the time complexity of DFS?

- A. $O(V^2)$
- B. $O(E^2)$
- C. $O(V + E)$
- D. $O(V * E)$

Answer: C

Q201. What is the greedy algorithm approach?

- A. Making the locally optimal choice at each step taken
- B. Breaking problems into subproblems and combining them
- C. Trying all possible solutions and picking the best one
- D. Using random choices to approximate the best solution

Answer: A

Q202. What is divide and conquer?

- A. Dividing resources equally among parallel processors
- B. Breaking a problem into subproblems, solving and combining
- C. Conquering the problem by brute force enumeration method
- D. Splitting input into two and sorting each half only

Answer: B

Q203. What is dynamic programming?

- A. Solving problems by storing solutions to subproblems
- B. Programming in a dynamic language like Python code
- C. Creating programs that dynamically allocate all memory
- D. A programming style that changes code at runtime

Answer: A

Q204. What is a brute force algorithm?

- A. An algorithm that uses force-directed graph layouts
- B. An algorithm that tries all possible candidate solutions
- C. An algorithm optimized for maximum efficiency speed
- D. An algorithm that aggressively prunes search space

Answer: B

Q205. What is the main advantage of dynamic programming over recursion?

- A. It does not use any memory for storing computations
- B. It uses less code and is simpler to implement
- C. It always runs in constant time for all input types
- D. It avoids recomputing overlapping subproblems using cache

Answer: D

Q206. Which algorithm technique does binary search use?

- A. Dynamic programming
- B. Greedy approach
- C. Divide and conquer
- D. Backtracking

Answer: C

Q207. What is backtracking?

- A. Going back to the previous version of code for fixes
- B. Backing up data structures before modifying them now
- C. Running the algorithm backwards from output to input
- D. Systematically trying options and undoing bad choices

Answer: D

Q208. What type of problems does the greedy approach work best for?

- A. Problems requiring all possible solutions enumerated
- B. Problems with optimal substructure and greedy property
- C. Problems that require exhaustive brute force searching
- D. Problems with overlapping subproblems needing caching

Answer: B

Q209. What is memoization?

- A. Memorizing the steps of an algorithm for quick recall
- B. Writing memos about algorithm design decisions made
- C. Creating documentation for each function in the code
- D. Caching results of function calls to avoid recomputation

Answer: D

Q210. What is an example of a greedy algorithm?

- A. Making change with fewest coins using largest first
- B. N-Queens problem solved using backtracking exploration
- C. Fibonacci sequence calculation using dynamic programming
- D. Matrix chain multiplication using optimal parenthesization

Answer: A

Q211. What is a priority queue?

- A. A stack that prioritizes the most recently added element
- B. A linked list sorted by element insertion time stamps
- C. A queue that processes elements in random arrival order
- D. A queue where elements are processed by priority order

Answer: D

Q212. What is a heap data structure?

- A. A hash table with elements ordered by insertion time
- B. A sorted array stored in dynamic heap memory allocation
- C. A complete binary tree satisfying the heap ordering rule
- D. A linked list where elements are organized by priority

Answer: C

Q213. What is the difference between a max-heap and a min-heap?

- A. Max-heap stores more elements; min-heap stores fewer items
- B. Max-heap uses arrays; min-heap uses linked list structure
- C. Max-heap has root as maximum; min-heap has root as min
- D. Max-heap has root as minimum; min-heap has root as max

Answer: C

Q214. What is a trie data structure?

- A. A tree for storing strings with shared prefix paths
- B. A generic tree with exactly three children per node
- C. A hash table variant that tries multiple hash functions
- D. A balanced binary tree for efficient numeric searching

Answer: A

Q215. What is the time complexity of inserting into a binary heap?

- A. $O(1)$
- B. $O(n^2)$
- C. $O(n)$
- D. $O(\log n)$

Answer: D

Q216. What is a graph adjacency matrix?

- A. A tree representation of the graph traversal order
- B. A 2D array indicating edges between pairs of vertices
- C. A linked list of all vertices in the graph structure
- D. A hash map from vertex names to edge weight values

Answer: B

Q217. What is a disjoint set?

- A. A set that has been divided into two equal-sized halves
- B. A set that contains overlapping elements with other sets
- C. A collection of non-overlapping sets with no shared items
- D. A sorted set used for binary search on set elements

Answer: C

Q218. What is the extract-max operation in a max-heap?

- A. Adding a new maximum element to the top of the heap
- B. Removing and returning the maximum element from heap
- C. Finding the maximum element without removing it at all
- D. Swapping the maximum and minimum elements in the heap

Answer: B

Q219. How is a heap typically stored in memory?

- A. As a contiguous array with index-based navigation
- B. As a linked list with pointers between nodes
- C. As a hash table with heap position as the key
- D. As a 2D matrix with rows and columns for levels

Answer: A

Q220. What is the purpose of the heapify operation?

- A. Converting an arbitrary array into a valid heap structure
- B. Doubling the size of the heap for more element storage
- C. Removing all elements from the heap data structure now
- D. Sorting all elements of the heap in ascending order output

Answer: A

Q221. What is string matching?

- A. Sorting strings in lexicographic order for indexing
- B. Concatenating multiple strings into one single string
- C. Finding occurrences of a pattern within a longer text
- D. Comparing two strings for exact equality testing

Answer: C

Q222. What is the brute force string matching time complexity?

- A. $O(m)$
- B. $O(n * m)$
- C. $O(n)$
- D. $O(n + m)$

Answer: B

Q223. What is a substring?

- A. A sequence of non-adjacent characters from a string
- B. A reversed copy of the original complete string
- C. A rearrangement of characters in a given string
- D. A contiguous sequence of characters within a string

Answer: D

Q224. What is a palindrome?

- A. A string that contains only unique characters in it
- B. A string that reads the same forwards and backwards
- C. A string where all characters are in alphabetical order
- D. A string with an even number of total characters only

Answer: B

Q225. What is the length of the string 'hello'?

- A. Six
- B. Five
- C. Four
- D. Three

Answer: B

Q226. What is string concatenation?

- A. Splitting a string into individual characters apart
- B. Removing duplicate characters from a given string
- C. Reversing the order of characters within a string
- D. Joining two or more strings end to end together

Answer: D

Q227. What is a prefix of a string?

- A. The last character of the string only always
- B. Any substring starting from the first character
- C. The middle character of the string only always
- D. Any character appearing in the string at all

Answer: B

Q228. What is a suffix of a string?

- A. A reversed copy of the string prefix portion
- B. The string with all vowels removed from it
- C. Any substring ending at the last character point
- D. The first character of the string only always

Answer: C

Q229. What is lexicographic order?

- A. Ordering strings by their total length in characters
- B. Ordering strings by the frequency of their characters
- C. Ordering strings alphabetically like in a dictionary
- D. Ordering strings by their last character value only

Answer: C

Q230. How many substrings does a string of length n have?

- A. $n(n+1)/2$
- B. n^2
- C. $2n$
- D. n

Answer: A

Q231. What is time complexity?

- A. A measure of how runtime grows with input size
- B. The clock speed required to run the algorithm
- C. The actual runtime of a program in seconds only
- D. The time taken to write the algorithm pseudocode

Answer: A

Q232. What is space complexity?

- A. The disk space needed to store source code files
- B. The number of lines of code in the program total
- C. The physical space the computer occupies on a desk
- D. A measure of how memory usage grows with input size

Answer: D

Q233. Which is faster: $O(n)$ or $O(n^2)$?

- A. $O(n)$ is faster for sufficiently large input sizes
- B. $O(n^2)$ is faster for all possible input sizes
- C. Both have identical performance for any input size
- D. It depends entirely on the programming language used

Answer: A

Q234. What does $O(1)$ mean?

- A. The algorithm processes only one element at any time
- B. The algorithm takes exactly one single operation only
- C. The algorithm runs in constant time regardless of input
- D. The algorithm has exactly one line of code to execute

Answer: C

Q235. What is the Big-O of two nested loops each running n times?

- A. $O(n)$
- B. $O(2n)$
- C. $O(n^2)$
- D. $O(\log n)$

Answer: C

Q236. Which complexity class represents logarithmic growth?

- A. $O(\log n)$
- B. $O(n^2)$
- C. $O(n)$
- D. $O(1)$

Answer: A

Q237. What is optimization in algorithm design?

- A. Converting the algorithm to a different language only
- B. Adding more features to the algorithm's functionality
- C. Improving algorithm efficiency in time or space usage
- D. Making the code look cleaner and more readable

Answer: C

Q238. What is the time complexity of a loop that halves n each iteration?

- A. $O(n)$
- B. $O(1)$
- C. $O(n^2)$
- D. $O(\log n)$

Answer: D

Q239. Why do we ignore constants in Big-O notation?

- A. Constants are only relevant in space complexity analysis
- B. Constants don't affect growth rate for large input sizes
- C. Constants are always equal to one in all algorithms
- D. Constants make the notation too hard to write out

Answer: B

Q240. What is the purpose of algorithm analysis?

- A. To create documentation for the algorithm implementation
- B. To predict resource usage and compare algorithm efficiency
- C. To determine the programming language to use for coding
- D. To find bugs and errors in the source code of program

Answer: B

Q241. What does asymptotic analysis focus on?

- A. Algorithm behavior for small inputs only
- B. Algorithm behavior as input size grows very large
- C. Exact running time in milliseconds
- D. The programming language used

Answer: B

Q242. Which of the following describes $O(\log n)$ complexity?

- A. The algorithm runs in constant time
- B. The running time doubles when the input size doubles
- C. The running time increases by a constant when the input size doubles
- D. The running time grows proportionally to the square of the input

Answer: C

Q243. What is the primary difference between a data structure and an algorithm?

- A. Data structures are faster than algorithms
- B. A data structure organizes data while an algorithm defines steps to solve a problem
- C. Algorithms store data while data structures process data
- D. There is no difference between them

Answer: B

Q244. What is the time complexity of performing two consecutive operations, one $O(n)$ and one $O(n^2)$?

- A. $O(n)$
- B. $O(n^3)$
- C. $O(n^2)$
- D. $O(2n^2)$

Answer: C

Q245. Which data structure follows the First-In-First-Out principle?

- A. Stack
- B. Queue
- C. Binary tree
- D. Hash table

Answer: B

Q246. What happens when a recursive function has no base case?

- A. It runs faster
- B. It returns zero
- C. It results in infinite recursion or a stack overflow
- D. It automatically stops after 100 calls

Answer: C

Q247. What is the Big-O complexity of an algorithm that does nothing regardless of input?

- A. $O(n)$
- B. $O(0)$
- C. $O(1)$
- D. $O(\log n)$

Answer: C

Q248. Which notation describes the best-case performance of an algorithm?

- A. Big-O (O)
- B. Big-Omega (Ω)
- C. Big-Theta (Θ)
- D. Little-o (o)

Answer: B

Q249. What is an iterative algorithm?

- A. An algorithm that calls itself
- B. An algorithm that uses loops to repeat steps
- C. An algorithm that solves problems using random choices
- D. An algorithm that cannot be analyzed

Answer: B

Q250. In Big-O notation, what is $O(n)$ commonly called?

- A. Constant time
- B. Logarithmic time
- C. Linear time
- D. Quadratic time

Answer: C

Q251. What is the time complexity of appending an element at the end of a dynamic array (without resizing)?

- A. $O(n)$
- B. $O(\log n)$
- C. $O(1)$
- D. $O(n^2)$

Answer: C

Q252. What is the difference between a stack and a queue?

- A. A stack uses FIFO and a queue uses LIFO
- B. A stack uses LIFO and a queue uses FIFO
- C. Both use FIFO
- D. Both use LIFO

Answer: B

Q253. Why can a doubly linked list be traversed in both directions?

- A. Because it has two data fields per node
- B. Because each node stores pointers to both the next and previous nodes
- C. Because it stores data in pairs
- D. Because it uses two separate arrays

Answer: B

Q254. What is the time complexity of deleting the first element from a singly linked list?

- A. $O(n)$
- B. $O(\log n)$
- C. $O(1)$
- D. $O(n^2)$

Answer: C

Q255. Which operation checks the top element of a stack without removing it?

- A. Pop
- B. Push
- C. Peek
- D. Enqueue

Answer: C

Q256. What is the main advantage of an array over a linked list?

- A. Dynamic size
- B. $O(1)$ random access to any element by index
- C. Easier insertion in the middle
- D. Less memory usage per element

Answer: B

Q257. What is a circular linked list?

- A. A linked list where the last node points back to the first node
- B. A linked list arranged in a circle shape on screen
- C. A linked list with circular data values
- D. A linked list that can only be traversed backwards

Answer: A

Q258. What is the time complexity of accessing the last element in a singly linked list?

- A. $O(1)$
- B. $O(\log n)$
- C. $O(n)$
- D. $O(n^2)$

Answer: C

Q259. Which data structure is commonly used for function call management in programming languages?

- A. Queue
- B. Stack
- C. Linked list
- D. Hash table

Answer: B

Q260. What is a dynamic array?

- A. An array that stores only dynamic data types
- B. An array that automatically resizes when it runs out of capacity
- C. An array allocated at compile time
- D. An array that cannot change its values

Answer: B

Q261. How many edges does a tree with n nodes have?

- A. n
- B. $n + 1$
- C. $n - 1$
- D. $2n$

Answer: C

Q262. What is the difference between a directed and an undirected graph?

- A. Directed graphs have weighted edges
- B. In a directed graph, edges have a direction from one vertex to another
- C. Undirected graphs cannot have cycles
- D. Directed graphs always have fewer edges

Answer: B

Q263. What property does a min-heap maintain between parent and child nodes?

- A. Each parent is greater than its children
- B. Each parent is less than or equal to its children
- C. Each parent equals its children
- D. Children are sorted left to right

Answer: B

Q264. What defines a simple path in a graph?

- A. A path that visits every vertex
- B. A path where no vertex is repeated
- C. A path with exactly one edge
- D. A path that forms a cycle

Answer: B

Q265. What is the in-order traversal order for a binary tree?

- A. Root, Left, Right
- B. Left, Root, Right
- C. Left, Right, Root
- D. Right, Root, Left

Answer: B

Q266. What is an adjacency matrix?

- A. A list of all vertices in a graph
- B. A 2D matrix where entry (i, j) indicates if there is an edge between vertex i and vertex j
- C. A matrix storing the shortest paths
- D. A matrix of node values in a tree

Answer: B

Q267. What is the maximum number of children a node in a binary tree can have?

- A. 1
- B. 2
- C. 3
- D. Unlimited

Answer: B

Q268. What condition must an undirected graph satisfy to be considered connected?

- A. Every vertex must have the same degree
- B. There must be a path between every pair of vertices
- C. It must have no edges
- D. Every vertex must connect to every other vertex directly

Answer: B

Q269. Where is the maximum element located in a max-heap?

- A. At any leaf node
- B. At the last position in the array
- C. At the root
- D. At the middle of the array

Answer: C

Q270. Can a simple graph contain self-loops?

- A. Yes, simple graphs allow self-loops
- B. No, simple graphs do not allow self-loops or multiple edges between the same pair of vertices
- C. Only directed simple graphs allow self-loops
- D. Only weighted simple graphs allow self-loops

Answer: B

Q271. In linear search, where does the search begin?

- A. At the middle of the array
- B. At the end of the array
- C. At the first element of the array
- D. At a random position

Answer: C

Q272. What is the best-case time complexity of binary search?

- A. $O(n)$
- B. $O(\log n)$
- C. $O(1)$
- D. $O(n^2)$

Answer: C

Q273. Can binary search be applied to a linked list?

- A. Yes, and it runs in $O(\log n)$
- B. No, because linked lists do not support efficient random access
- C. Yes, but only on doubly linked lists
- D. Yes, if the list is circular

Answer: B

Q274. What does binary search do at each step?

- A. Checks every element one by one
- B. Divides the search space in half by comparing with the middle element
- C. Swaps elements to sort the array
- D. Hashes the target value

Answer: B

Q275. What is the worst-case number of comparisons for linear search on an array of 100 elements?

- A. 50
- B. 10
- C. 100
- D. 7

Answer: C

Q276. Does linear search require the array to be sorted?

- A. Yes, it always requires sorting
- B. No, linear search works on both sorted and unsorted arrays
- C. Only for integer arrays
- D. Only when searching for the maximum element

Answer: B

Q277. In binary search on a sorted array, if the target equals the middle element, what happens?

- A. The search continues in the left half
- B. The search continues in the right half
- C. The search returns the index of the middle element
- D. The search restarts from the beginning

Answer: C

Q278. How many comparisons does binary search need at most for an array of 16 elements?

- A. 16
- B. 8
- C. 4
- D. 5

Answer: D

Q279. What is the average-case time complexity of linear search?

- A. $O(1)$
- B. $O(n/2)$, which simplifies to $O(n)$
- C. $O(\log n)$
- D. $O(n^2)$

Answer: B

Q280. Which searching algorithm is most suitable when the dataset is very small?

- A. Binary search
- B. Interpolation search
- C. Linear search
- D. Exponential search

Answer: C

Q281. What is the best-case time complexity of bubble sort?

- A. $O(n^2)$
- B. $O(n \log n)$
- C. $O(n)$
- D. $O(1)$

Answer: C

Q282. How does selection sort work?

- A. It inserts each element into its correct position
- B. It repeatedly finds the minimum element and places it at the beginning of the unsorted portion
- C. It divides the array and merges sorted halves
- D. It swaps adjacent elements

Answer: B

Q283. What is the space complexity of bubble sort?

- A. $O(n)$
- B. $O(n^2)$
- C. $O(\log n)$
- D. $O(1)$

Answer: D

Q284. How does insertion sort process elements?

- A. It picks elements randomly and places them in sorted order
- B. It takes each element and inserts it into its correct position in the already sorted portion
- C. It divides the array into halves recursively
- D. It uses a hash function to determine positions

Answer: B

Q285. Which sorting algorithm is similar to how people typically sort playing cards in their hand?

- A. Bubble sort
- B. Selection sort
- C. Insertion sort
- D. Merge sort

Answer: C

Q286. What is the worst-case time complexity of insertion sort?

- A. $O(n)$
- B. $O(n \log n)$
- C. $O(n^2)$
- D. $O(\log n)$

Answer: C

Q287. Is bubble sort an in-place sorting algorithm?

- A. No, it requires $O(n)$ extra space
- B. Yes, it sorts the elements within the original array
- C. No, it creates a copy of the array
- D. It depends on the implementation

Answer: B

Q288. What does the term 'comparison-based sorting' mean?

- A. Sorting that only works on numbers
- B. Sorting that determines order only by comparing pairs of elements
- C. Sorting that compares the array with a sorted version
- D. Sorting that uses bit manipulation

Answer: B

Q289. How many swaps does selection sort perform in the worst case for n elements?

- A. $O(n^2)$
- B. $O(n \log n)$
- C. $O(n)$
- D. $O(1)$

Answer: C

Q290. After the first complete pass of bubble sort on [5, 3, 8, 1, 2], which element is guaranteed to be in its final position?

- A. 5
- B. 1
- C. 8
- D. 2

Answer: C

Q291. What is the primary purpose of a hash function?

- A. To sort data
- B. To map data of arbitrary size to fixed-size values
- C. To encrypt data
- D. To compress data

Answer: B

Q292. What is chaining in the context of hash table collision resolution?

- A. Linking hash tables together
- B. Storing all elements that hash to the same index in a linked list at that index
- C. Creating a chain of hash functions
- D. Connecting the first and last elements of the hash table

Answer: B

Q293. If a hash table has 10 slots and the hash function is $h(k) = k \bmod 10$, where does the key 27 go?

- A. Slot 2
- B. Slot 7
- C. Slot 3
- D. Slot 27

Answer: B

Q294. What is the time complexity of inserting an element into a hash table in the average case?

- A. $O(n)$
- B. $O(\log n)$
- C. $O(1)$
- D. $O(n^2)$

Answer: C

Q295. What is a hash collision?

- A. When the hash table runs out of memory
- B. When two different keys produce the same hash value
- C. When a hash function produces a negative value
- D. When the hash table is empty

Answer: B

Q296. What is the load factor of a hash table with 15 elements and 20 slots?

- A. 0.25
- B. 0.75
- C. 1.33
- D. 15

Answer: B

Q297. What property should a good hash function have?

- A. It should always return the same value
- B. It should distribute keys uniformly across the hash table
- C. It should sort the keys
- D. It should only work with string keys

Answer: B

Q298. Can two different keys have the same hash value?

- A. No, hash functions always produce unique values
- B. Yes, this is called a collision
- C. Only if the keys are equal
- D. Only in broken hash functions

Answer: B

Q299. What is the worst-case time complexity of searching in a hash table?

- A. $O(1)$
- B. $O(\log n)$
- C. $O(n)$
- D. $O(n \log n)$

Answer: C

Q300. What is the difference between a hash map and a hash table?

- A. Hash maps are always faster
- B. A hash map stores key-value pairs while a hash set stores only keys; both use hashing internally
- C. Hash tables cannot handle collisions
- D. There is no difference at all

Answer: B

Q301. What is the time complexity of finding the height of a binary tree?

- A. $O(1)$
- B. $O(\log n)$
- C. $O(n)$
- D. $O(n^2)$

Answer: C

Q302. Where is the in-order successor located if a BST node has a right subtree?

- A. It is the parent of the node
- B. It is the leftmost node in the right subtree
- C. It is the root of the tree
- D. It is the rightmost node in the left subtree

Answer: B

Q303. What is the result of an in-order traversal of a BST?

- A. Elements in reverse order
- B. Elements in random order
- C. Elements in sorted (ascending) order
- D. Elements level by level

Answer: C

Q304. What data structure is used to implement level-order traversal of a tree?

- A. Stack
- B. Queue
- C. Hash table
- D. Priority queue

Answer: B

Q305. How do you count the total number of nodes in a binary tree?

- A. Count only the leaf nodes
- B. Recursively count nodes in left subtree, right subtree, and add 1 for the root
- C. Count only the internal nodes
- D. Use binary search

Answer: B

Q306. What is the time complexity of inserting a new element into a BST?

- A. $O(1)$
- B. $O(n)$ in the worst case
- C. $O(n^2)$
- D. $O(\log n)$ always

Answer: B

Q307. In which traversal order is the root node visited before its subtrees?

- A. Post-order
- B. In-order
- C. Pre-order
- D. Level-order only

Answer: C

Q308. Which direction do you traverse in a BST to find the maximum element?

- A. Always go left until a leaf is reached
- B. Alternate between left and right
- C. Always go right until a node with no right child is reached
- D. Go to the root and stop

Answer: C

Q309. Which tree traversal is commonly used to delete a binary tree safely (deleting children before the parent)?

- A. Pre-order traversal
- B. In-order traversal
- C. Post-order traversal
- D. Level-order traversal

Answer: C

Q310. How many leaf nodes does a full binary tree with n internal nodes have?

- A. n
- B. $n + 1$
- C. $n - 1$
- D. $2n$

Answer: B

Q311. What is the difference between BFS and DFS?

- A. BFS uses a stack and DFS uses a queue
- B. BFS explores neighbors level by level using a queue, while DFS goes deep using a stack
- C. BFS is faster than DFS
- D. DFS finds shortest paths but BFS does not

Answer: B

Q312. What does DFS stand for?

- A. Direct First Search
- B. Depth First Search
- C. Double Find Search
- D. Data Flow System

Answer: B

Q313. How many edges does a spanning tree of a connected graph with V vertices contain?

- A. V
- B. $V - 1$
- C. $V + 1$
- D. $2V$

Answer: B

Q314. Can BFS be used to find the shortest path in an unweighted graph?

- A. No, only Dijkstra's algorithm can find shortest paths
- B. Yes, BFS finds the shortest path in terms of number of edges in an unweighted graph
- C. Only if the graph is a tree
- D. Only if the graph has no cycles

Answer: B

Q315. What is the space complexity of BFS on a graph with V vertices?

- A. $O(1)$
- B. $O(V)$
- C. $O(V^2)$
- D. $O(E)$

Answer: B

Q316. What is an acyclic graph?

- A. A graph with exactly one cycle
- B. A graph that contains no cycles
- C. A graph with all vertices connected
- D. A graph with no edges

Answer: B

Q317. In a weighted graph, what do the edge weights typically represent?

- A. The number of vertices
- B. Costs, distances, or capacities associated with traveling along that edge
- C. The color of the edge
- D. The degree of adjacent vertices

Answer: B

Q318. How is a graph typically represented in code using an adjacency list?

- A. As a 2D array of size $V \times V$
- B. As an array of linked lists (or arrays), where each vertex has a list of its neighbors
- C. As a single sorted array
- D. As a hash table mapping edges to vertices

Answer: B

Q319. What happens if DFS visits a vertex that has already been visited in a directed graph?

- A. It always means a cycle exists
- B. It depends on whether the previously visited vertex is still on the current DFS path
- C. DFS stops immediately
- D. The graph is invalid

Answer: B

Q320. What is the minimum number of edges in a connected undirected graph with V vertices?

- A. V
- B. $V - 1$
- C. $V + 1$
- D. $V / 2$

Answer: B

Q321. What is the divide and conquer approach?

- A. Solving a problem by trying all possible solutions
- B. Breaking a problem into smaller subproblems, solving them, and combining the solutions
- C. Using a greedy choice at each step
- D. Storing solutions to avoid recomputation

Answer: B

Q322. What are the two key properties required for a dynamic programming solution?

- A. Greedy choice and local optima
- B. Optimal substructure and overlapping subproblems
- C. Random choices and large inputs
- D. Divide and conquer and backtracking

Answer: B

Q323. What is an example of a problem solved by the greedy approach?

- A. 0/1 Knapsack problem
- B. Making change with the fewest coins (when denominations are canonical)
- C. Traveling Salesman Problem
- D. Matrix chain multiplication

Answer: B

Q324. What is the key difference between greedy algorithms and dynamic programming?

- A. Greedy algorithms are always slower
- B. Greedy algorithms make locally optimal choices without reconsidering, while DP considers all subproblem solutions
- C. Dynamic programming never gives optimal solutions
- D. They are the same approach

Answer: B

Q325. What is the bottom-up approach in dynamic programming?

- A. Solving the problem using recursion from the top
- B. Solving smaller subproblems first and building up to the original problem using a table
- C. Randomly solving subproblems
- D. Starting from the solution and working backwards

Answer: B

Q326. How does backtracking reduce the search space compared to brute force?

- A. It sorts the candidates first
- B. It abandons partial solutions as soon as they are determined to be invalid, avoiding exploration of dead-end branches
- C. It uses hashing to skip solutions
- D. It randomly selects solutions to test

Answer: B

Q327. What is the time complexity of computing Fibonacci numbers using memoization?

- A. $O(2^n)$
- B. $O(n^2)$
- C. $O(n)$
- D. $O(\log n)$

Answer: C

Q328. Which algorithm design paradigm does merge sort belong to?

- A. Greedy
- B. Dynamic programming
- C. Divide and conquer
- D. Backtracking

Answer: C

Q329. What is the purpose of the top-down approach in dynamic programming?

- A. To sort elements from top to bottom
- B. To solve the original problem recursively and cache results to avoid redundant computation
- C. To iterate through the array from beginning to end
- D. To find the maximum element first

Answer: B

Q330. What is a brute force approach to problem solving?

- A. An optimized algorithm
- B. An approach that tries every possible solution and checks which one is correct
- C. An approach that uses heuristics
- D. An approach based on random sampling

Answer: B

Q331. What is the time complexity of inserting a word of length L into a trie?

- A. $O(1)$
- B. $O(L)$
- C. $O(n)$
- D. $O(L * n)$

Answer: B

Q332. What is the advantage of a priority queue over a sorted array for maintaining a dynamic set?

- A. Priority queues use less memory
- B. Priority queues support $O(\log n)$ insertion while sorted arrays require $O(n)$ for insertion
- C. Priority queues allow $O(1)$ random access
- D. Sorted arrays cannot store integers

Answer: B

Q333. What does the find operation do in a Union-Find data structure?

- A. Finds the minimum element in a set
- B. Returns the representative (root) of the set containing a given element
- C. Finds the union of two sets
- D. Searches for an element in a sorted array

Answer: B

Q334. What is a min-heap's property regarding the root element?

- A. The root is the largest element
- B. The root is the median element
- C. The root is the smallest element
- D. The root has no special property

Answer: C

Q335. What is the time complexity of the extract-min operation in a binary min-heap?

- A. $O(1)$
- B. $O(\log n)$
- C. $O(n)$
- D. $O(n \log n)$

Answer: B

Q336. What is a trie primarily used for?

- A. Sorting numbers
- B. Efficient storage and retrieval of strings, supporting prefix-based operations
- C. Graph traversal
- D. Matrix operations

Answer: B

Q337. What is the union operation in Union-Find?

- A. Finding all elements in a set
- B. Merging two disjoint sets into one
- C. Splitting a set into two parts
- D. Sorting elements within a set

Answer: B

Q338. For a node at index i in a binary heap stored in an array (0-indexed), what is the index of its left child?

- A. $i + 1$
- B. $2i$
- C. $2i + 1$
- D. $i / 2$

Answer: C

Q339. What is a multimap data structure?

- A. A map that can store multiple values for the same key
- B. A map that uses multiple hash functions
- C. A map with multiple levels
- D. A map that can only store maps

Answer: A

Q340. What is the delete operation in a binary heap called?

- A. Remove
- B. Extract
- C. Pop
- D. Decrease-key followed by extract-min

Answer: D

Q341. What is the time complexity of comparing two strings of length n for equality?

- A. $O(1)$
- B. $O(n)$
- C. $O(n^2)$
- D. $O(\log n)$

Answer: B

Q342. Is 'ace' a valid subsequence of 'abcde'?

- A. No, because the characters are not contiguous
- B. Yes, because the characters appear in the same relative order in the original string
- C. Only if we rearrange the original string
- D. No, because 'ace' is not a word

Answer: B

Q343. How many total substrings does a string of length n have (including the empty string)?

- A. n
- B. n^2
- C. $n(n+1)/2 + 1$
- D. 2^n

Answer: C

Q344. In lexicographic order, which string comes first: 'apple' or 'application'?

- A. 'application' because it is longer
- B. 'apple' because at the first differing position, 'e' < 'i'
- C. They are considered equal
- D. 'apple' because shorter strings always come first

Answer: B

Q345. What is the time complexity of reversing a string of length n?

- A. $O(1)$
- B. $O(\log n)$
- C. $O(n)$
- D. $O(n^2)$

Answer: C

Q346. What is the time complexity of the naive (brute force) pattern matching algorithm?

- A. $O(n + m)$
- B. $O(n * m)$ in the worst case
- C. $O(n)$
- D. $O(m)$

Answer: B

Q347. How can you check if a string is a palindrome?

- A. Sort the string and check if it's the same
- B. Compare the string with its reverse
- C. Check if all characters are the same
- D. Count the number of vowels

Answer: B

Q348. What is the difference between a substring and a subsequence?

- A. They are the same thing
- B. A substring is contiguous characters; a subsequence can skip characters but maintains order
- C. A subsequence is always longer
- D. A substring can skip characters

Answer: B

Q349. What does the term 'pattern matching' mean in the context of strings?

- A. Creating patterns from strings
- B. Finding occurrences of a pattern string within a text string
- C. Matching two strings for equality
- D. Sorting strings by pattern

Answer: B

Q350. What is string immutability?

- A. Strings that can be modified in place
- B. Strings that cannot be changed after creation; any modification creates a new string
- C. Strings stored in read-only memory
- D. Strings with fixed length

Answer: B

Q351. What is the Big-O complexity of an algorithm with running time $3n + 5$?

- A. $O(5)$
- B. $O(3n)$
- C. $O(n)$
- D. $O(n + 5)$

Answer: C

Q352. Which grows faster: $O(n^2)$ or $O(2^n)$?

- A. $O(n^2)$ grows faster
- B. They grow at the same rate
- C. $O(2^n)$ grows faster
- D. It depends on n

Answer: C

Q353. What is the time complexity of a loop that runs from 1 to n and increments by 2 each time?

- A. $O(n^2)$
- B. $O(n)$
- C. $O(n/2)$, which is $O(n)$
- D. $O(\log n)$

Answer: C

Q354. What does it mean to optimize an algorithm?

- A. To make the code longer
- B. To improve the algorithm's time or space efficiency
- C. To add more features
- D. To translate it to a faster programming language

Answer: B

Q355. What is the time complexity of checking if a number n is even or odd?

- A. $O(n)$
- B. $O(\log n)$
- C. $O(1)$
- D. $O(n^2)$

Answer: C

Q356. Why is $O(1)$ the most efficient time complexity?

- A. Because it means the algorithm does nothing
- B. Because the execution time does not increase with input size
- C. Because it uses no memory
- D. Because it always runs in exactly 1 millisecond

Answer: B

Q357. What is the space complexity of an algorithm that uses only a fixed number of variables?

- A. $O(n)$
- B. $O(1)$
- C. $O(\log n)$
- D. $O(n^2)$

Answer: B

Q358. What does it mean when we say an algorithm runs in 'polynomial time'?

- A. It always finishes in under a second
- B. Its running time is bounded by n^k for some constant k
- C. It uses polynomial equations
- D. It can only process polynomials

Answer: B

Q359. If an algorithm has two independent steps with complexities $O(n)$ and $O(m)$, what is the overall complexity?

- A. $O(n * m)$
- B. $O(n + m)$
- C. $O(\max(n, m))$
- D. $O(n^2)$

Answer: B

Q360. Why do we use Big-O notation instead of measuring exact execution time?

- A. Because exact time is always the same
- B. Because Big-O describes how the algorithm scales with input size, independent of hardware and implementation
- C. Because Big-O is more precise
- D. Because exact time cannot be measured

Answer: B

Data Structures & Algorithms — Medium

360 questions

Q361. What is the time complexity of the recurrence $T(n) = 2T(n/2) + n$?

- A. $O(n^2)$
- B. $O(\log n)$
- C. $O(n)$
- D. $O(n \log n)$

Answer: D

Q362. Which case of the Master Theorem applies when $f(n) = 9, \log n$?

- A. Case 3 applies
- B. No case applies
- C. Case 2 applies
- D. Case 1 applies

Answer: C

Q363. What does Omega (Ω) represent?

- A. Exact tight bound complexity
- B. Worst-case upper bound complexity
- C. Best-case or lower bound complexity
- D. Average-case expected complexity

Answer: C

Q364. If an algorithm has complexities $O(n^2)$ in the worst case and $\Omega(n)$ in the best case, which is true?

- A. The algorithm always runs in $O(n)$ time
- B. The algorithm always runs in $O(n^2)$ time
- C. The algorithm runs in $O(n \log n)$ always
- D. Performance varies between $\Omega(n)$ and $O(n^2)$

Answer: D

Q365. What is amortized analysis?

- A. Analyzing the worst case of each operation individually
- B. Averaging the cost of operations over a sequence
- C. Analyzing only the best case of each operation
- D. Computing the cost using a sorting-based analysis

Answer: B

Q366. Which growth rate is faster: $O(2^n)$ or $O(n^3)$?

- A. They are the same
- B. $O(n^3)$ is faster
- C. $O(2^n)$ is faster
- D. Depends on input size

Answer: C

Q367. What is the time complexity of a nested loop where both loops run n times?

- A. $O(n \log n)$
- B. $O(n^2)$
- C. $O(2n)$
- D. $O(n)$

Answer: B

Q368. In asymptotic analysis, which of these is the correct ordering from slowest to fastest growth?

- A. $O(1) < O(\log n) < O(n) < O(n^2)$
- B. $O(1) < O(n) < O(\log n) < O(n^2)$
- C. $O(n) < O(\log n) < O(1) < O(n^2)$
- D. $O(\log n) < O(1) < O(n) < O(n^2)$

Answer: A

Q369. What is the recurrence relation for binary search?

- A. $T(n) = 2T(n/2) + 1$
- B. $T(n) = T(n-1) + 1$
- C. $T(n) = T(n/2) + n$
- D. $T(n) = T(n/2) + 1$

Answer: D

Q370. What does Theta (Θ) represent?

- A. Only upper bound estimate
- B. Tight bound (both upper and lower)
- C. Only lower bound estimate
- D. Average case analysis only

Answer: B

Q371. What is a circular queue?

- A. A queue implemented using a linked list
- B. A queue where the rear connects back to front
- C. A double-ended queue with two pointers
- D. A queue that automatically sorts elements

Answer: B

Q372. What is a doubly linked list?

- A. A list with pointers to both next and previous
- B. A list that stores two data values per node
- C. A list structure with two separate heads
- D. A circular list with bidirectional traversal

Answer: A

Q373. What is a deque (double-ended queue)?

- A. A stack implemented as a queue structure
- B. A priority queue with weighted elements
- C. A queue allowing deletion from both ends only
- D. A structure allowing insert and delete at both ends

Answer: D

Q374. What is the time complexity of deleting a node from the middle of a singly linked list (given a pointer to the previous node)?

- A. $O(1)$
- B. $O(\log n)$
- C. $O(n)$
- D. $O(n^2)$

Answer: A

Q375. How can you implement a queue using two stacks?

- A. It is not possible to implement using stacks
- B. Use both stacks for enqueue operations only
- C. Use one stack for enqueue and another for dequeue
- D. Merge both stacks into a single combined structure

Answer: C

Q376. What is the advantage of a linked list over an array?

- A. Better CPU cache performance overall
- B. Faster random access by index
- C. Dynamic size and efficient insertion
- D. Less memory usage per element

Answer: C

Q377. What is a sentinel node in a linked list?

- A. A dummy node to simplify boundary conditions
- B. The node containing the maximum data value
- C. A node that stores the total list size
- D. The last node in the list structure

Answer: A

Q378. What is the condition for a circular queue being full (array implementation with size N)?

- A. $\text{front} == 0 \ \&\& \ \text{rear} == N - 1$
- B. $(\text{rear} + 1) \% N == \text{front}$
- C. $\text{rear} == N - 1$
- D. $\text{front} == \text{rear}$

Answer: B

Q379. What is the time complexity of searching for an element in an unsorted linked list?

- A. $O(n^2)$
- B. $O(\log n)$
- C. $O(1)$
- D. $O(n)$

Answer: D

Q380. Which application commonly uses a stack?

- A. Breadth-first graph search traversal
- B. Level-order tree traversal processing
- C. Function call management (call stack)
- D. Round-robin job scheduling system

Answer: C

Q381. What is the time complexity of searching in a balanced BST?

- A. $O(\log n)$
- B. $O(n \log n)$
- C. $O(1)$
- D. $O(n)$

Answer: A

Q382. What is the worst-case time complexity of search in an unbalanced BST?

- A. $O(\log n)$
- B. $O(n^2)$
- C. $O(1)$
- D. $O(n)$

Answer: D

Q383. What is the space complexity of an adjacency matrix for a graph with V vertices?

- A. $O(V)$
- B. $O(V + E)$
- C. $O(E)$
- D. $O(V^2)$

Answer: D

Q384. What is a complete binary tree?

- A. A tree where every node has exactly 2 children always
- B. A tree where all leaf nodes are at the same level
- C. A tree with the maximum possible depth for its nodes
- D. A tree with all levels full except possibly the last level

Answer: D

Q385. In a max-heap, where is the maximum element?

- A. At the root
- B. At the last level
- C. At a leaf node
- D. It could be anywhere

Answer: A

Q386. What is the space complexity of an adjacency list for a graph with V vertices and E edges?

- A. $O(E)$
- B. $O(V^2)$
- C. $O(V + E)$
- D. $O(V)$

Answer: C

Q387. What is the in-order traversal of a BST with values 4, 2, 6, 1, 3, 5, 7 (root=4)?

- A. 1, 2, 3, 4, 5, 6, 7
- B. 4, 2, 1, 3, 6, 5, 7
- C. 4, 2, 6, 1, 3, 5, 7
- D. 1, 3, 2, 5, 7, 6, 4

Answer: A

Q388. What type of graph has edges with associated weights?

- A. Weighted graph
- B. Directed graph
- C. Bipartite graph
- D. Complete graph

Answer: A

Q389. What is a full binary tree?

- A. A tree with the maximum number of nodes
- B. A tree where every level is completely filled
- C. A tree where every node has 0 or 2 children
- D. A tree where all leaves are at same depth

Answer: C

Q390. What is the time complexity of inserting an element into a max-heap?

- A. $O(n)$
- B. $O(1)$
- C. $O(\log n)$
- D. $O(n \log n)$

Answer: C

Q391. What is interpolation search?

- A. A recursive linear search with memoization cache
- B. A search that always checks the middle element
- C. A search that uses a hash table for fast lookups
- D. A search that estimates position based on value distribution

Answer: D

Q392. What is the average time complexity of interpolation search on uniformly distributed data?

- A. $O(\log n)$
- B. $O(1)$
- C. $O(n)$
- D. $O(\log \log n)$

Answer: D

Q393. What is jump search?

- A. A search that uses recursion to jump elements
- B. A search that jumps by blocks then does linear scan
- C. A search that jumps to random positions in array
- D. A search that skips every other element each pass

Answer: B

Q394. What is the optimal block size for jump search on an array of n elements?

- A. $\log n$
- B. $\sqrt{n}$
- C. n
- D. $n/2$

Answer: B

Q395. What is exponential search?

- A. A search with exponential time complexity overall
- B. A search designed for exponentially growing datasets
- C. A search that doubles the index then uses binary search
- D. A search that uses exponential mathematical functions

Answer: C

Q396. What is the time complexity of exponential search?

- A. $O(n)$
- B. $O(\sqrt{n})$
- C. $O(n \log n)$
- D. $O(\log n)$

Answer: D

Q397. What is ternary search?

- A. A three-pass sequential linear search method
- B. A search specifically for the third element only
- C. A search dividing array into three parts each step
- D. A search that checks three elements at a time

Answer: C

Q398. Why is binary search preferred over ternary search in practice?

- A. Binary search uses less memory per recursive call
- B. Binary search is always faster in every case
- C. Ternary search does not work on sorted arrays at all
- D. Ternary search makes more comparisons per step overall

Answer: D

Q399. What is the space complexity of recursive binary search?

- A. $O(n)$
- B. $O(n^2)$
- C. $O(\log n)$
- D. $O(1)$

Answer: C

Q400. In which scenario is linear search more efficient than binary search?

- A. When the array has millions of elements
- B. Linear search is never more efficient
- C. Large sorted arrays with many elements
- D. When the array is unsorted and small

Answer: D

Q401. What is the time complexity of merge sort?

- A. $O(n)$
- B. $O(\log n)$
- C. $O(n \log n)$
- D. $O(n^2)$

Answer: C

Q402. What is the space complexity of merge sort?

- A. $O(n)$
- B. $O(n^2)$
- C. $O(1)$
- D. $O(\log n)$

Answer: A

Q403. What is the average-case time complexity of quick sort?

- A. $O(\log n)$
- B. $O(n \log n)$
- C. $O(n^2)$
- D. $O(n)$

Answer: B

Q404. What is the worst-case time complexity of quick sort?

- A. $O(n)$
- B. $O(\log n)$
- C. $O(n^2)$
- D. $O(n \log n)$

Answer: C

Q405. What is heap sort's time complexity?

- A. $O(\log n)$
- B. $O(n)$
- C. $O(n \log n)$
- D. $O(n^2)$

Answer: C

Q406. Is merge sort a stable sorting algorithm?

- A. Only for linked lists
- B. Yes, it is stable
- C. Only for array inputs
- D. No, it is unstable

Answer: B

Q407. What is the divide-and-conquer strategy in merge sort?

- A. Divide the array into three equal parts
- B. Divide in half, sort each half, then merge
- C. Divide by selecting a random pivot element
- D. Divide into sorted and unsorted portions

Answer: B

Q408. What is the role of the pivot in quick sort?

- A. It is always the first element chosen
- B. It is always the exact median value
- C. It partitions into smaller and larger groups
- D. It merges two sorted subarrays together

Answer: C

Q409. What is counting sort?

- A. A comparison-based sorting method
- B. A non-comparison sort counting occurrences
- C. A sort that uses a heap structure
- D. A divide-and-conquer sorting approach

Answer: B

Q410. What is the space complexity of heap sort?

- A. $O(n \log n)$
- B. $O(n)$
- C. $O(\log n)$
- D. $O(1)$

Answer: D

Q411. What is open addressing in hash tables?

- A. Using multiple hash tables for distribution
- B. Finding another open slot when collision occurs
- C. Doubling the table size on every collision
- D. Using linked lists for collision resolution

Answer: B

Q412. What is linear probing?

- A. Using a linear hash function for keys
- B. Checking next consecutive slots on collision
- C. Using a linked list for collision handling
- D. Probing with a linear mathematical equation

Answer: B

Q413. What is the primary clustering problem in linear probing?

- A. Keys are stored in sorted order clusters
- B. Consecutive occupied slots form long clusters
- C. Hash values clustering near zero index
- D. Multiple hash functions produce same values

Answer: B

Q414. What is quadratic probing?

- A. Using four different hash functions total
- B. Squaring the hash value before modding
- C. Using a quadratic hash function for keys
- D. Probing at positions $h+1^2$, $h+2^2$, $h+3^2$, etc.

Answer: D

Q415. What is double hashing?

- A. Using two hash functions for probe step size
- B. Hashing the hash value itself recursively
- C. Hashing a value twice in succession
- D. Using two separate hash tables for storage

Answer: A

Q416. Why should the hash table size be a prime number?

- A. It is required by all hash functions
- B. It distributes hash values more uniformly
- C. It speeds up the hash value computation
- D. It makes the table smaller in memory

Answer: B

Q417. What is rehashing?

- A. Hashing the same key multiple times in a row
- B. Removing all elements and starting fresh over
- C. Changing the hash function to a new one
- D. Creating a larger table and reinserting all elements

Answer: D

Q418. What is the multiplication method for hashing?

- A. $h(k) = k \bmod (m \times 2)$
- B. $h(k) = k^2 \bmod m$
- C. $h(k) = \text{floor}(m \times (k \times A \bmod 1))$ where $0 < A < 1$
- D. $h(k) = k \times m$

Answer: C

Q419. What is the advantage of chaining over open addressing?

- A. Chaining is always faster for all operations
- B. Chaining uses less memory per element overall
- C. Chaining handles high load factors and deletion better
- D. Chaining has better CPU cache performance overall

Answer: C

Q420. What is a hash set?

- A. A multiset allowing duplicate elements
- B. A set using hashing for fast unique lookups
- C. A sorted set with ordered elements
- D. A set implemented using sorted arrays

Answer: B

Q421. What is an AVL tree?

- A. A tree with at most 3 children per node always
- B. A self-balancing BST with height difference at most 1
- C. A tree data structure used only for searching
- D. A binary tree with no balancing guarantees

Answer: B

Q422. What are the rotation operations in an AVL tree?

- A. Merge and split operations on subtrees
- B. Insert and delete operations on tree nodes
- C. Push and pop operations on internal nodes
- D. Left, right, left-right, and right-left rotations

Answer: D

Q423. When is a left rotation performed in an AVL tree?

- A. When the tree is already perfectly balanced
- B. When the right subtree is too heavy (RR case)
- C. When the left subtree is too heavy overall
- D. After every single insertion operation occurs

Answer: B

Q424. How do you delete a node with two children in a BST?

- A. Simply remove the node directly from tree
- B. Swap with the root node then remove it
- C. Replace with in-order successor then delete it
- D. Set the node value to null and leave it

Answer: C

Q425. What is the time complexity of insertion in a Red-Black tree?

- A. $O(1)$
- B. $O(n)$
- C. $O(n \log n)$
- D. $O(\log n)$

Answer: D

Q426. What properties must a Red-Black tree satisfy?

- A. All internal and external nodes must always be colored red only
- B. Root is black, red nodes have black children, equal black-height on all paths
- C. Height is always kept at exactly $\log n$ for any n total nodes
- D. All leaf nodes are always colored red throughout the entire tree

Answer: B

Q427. What is the Lowest Common Ancestor (LCA) of two nodes in a tree?

- A. The deepest node that is ancestor of both
- B. The node with the smallest key value
- C. The root node of the entire tree
- D. The immediate parent of both given nodes

Answer: A

Q428. What is Morris traversal?

- A. A traversal technique for graphs only
- B. An in-order traversal using $O(1)$ extra space
- C. A traversal using an explicit stack structure
- D. A breadth-first traversal using a queue

Answer: B

Q429. How can you check if a binary tree is a valid BST?

- A. Check if each node has at most 2 children
- B. Count the number of nodes in the tree
- C. Verify in-order traversal is strictly increasing
- D. Check if the root is the minimum value node

Answer: C

Q430. What is a segment tree used for?

- A. Storing sorted segments of linked lists
- B. Efficient range queries and point updates on arrays
- C. Segmenting a network into subnetworks
- D. Storing geometric line segments only

Answer: B

Q431. What is Dijkstra's algorithm used for?

- A. Finding all cycles in a given graph
- B. Finding shortest paths with non-negative weights
- C. Sorting all vertices by their degree
- D. Finding connected components in a graph

Answer: B

Q432. What is the time complexity of Dijkstra's algorithm using a binary heap?

- A. $O(V \times E)$
- B. $O(V + E)$
- C. $O((V + E) \log V)$
- D. $O(V^2)$

Answer: C

Q433. What is Kruskal's algorithm used for?

- A. Shortest path computation
- B. Topological sorting of vertices
- C. Cycle detection in graphs
- D. Finding Minimum Spanning Tree

Answer: D

Q434. What is Prim's algorithm used for?

- A. Shortest path computation
- B. Finding Minimum Spanning Tree
- C. Graph coloring assignment
- D. Topological sort ordering

Answer: B

Q435. What is topological sorting?

- A. Linear ordering of DAG vertices respecting edge directions
- B. Sorting edges by their assigned weight values
- C. Sorting all vertices by their numeric label values
- D. Sorting all vertices by their in-degree values

Answer: A

Q436. What type of graph is required for topological sorting?

- A. An undirected graph type
- B. A complete graph type
- C. A weighted graph type
- D. Directed Acyclic Graph (DAG)

Answer: D

Q437. What is the Bellman-Ford algorithm used for?

- A. Finding connected components in undirected graphs
- B. Finding the Minimum Spanning Tree of a graph
- C. Finding shortest paths even with negative weights
- D. Topological sorting of directed acyclic graphs

Answer: C

Q438. What is the time complexity of the Bellman-Ford algorithm?

- A. $O(V \times E)$
- B. $O(V + E)$
- C. $O(V^2)$
- D. $O(E \log V)$

Answer: A

Q439. How can you detect a cycle in a directed graph?

- A. By counting edges in the graph
- B. Using DFS and checking for back edges
- C. Using BFS traversal only
- D. By finding the shortest path first

Answer: B

Q440. What is the Union-Find (Disjoint Set Union) data structure used for?

- A. Finding shortest paths between graph vertices
- B. Storing key-value pairs with fast lookup
- C. Tracking set membership and merging sets efficiently
- D. Sorting elements into ordered groups

Answer: C

Q441. What is the time complexity of the dynamic programming solution for the 0/1 knapsack problem with n items and capacity W ?

- A. $O(n \log n)$
- B. $O(2^n)$
- C. $O(n \times W)$
- D. $O(n)$

Answer: C

Q442. What is the Longest Common Subsequence (LCS) problem?

- A. Finding the longest subsequence present in both sequences
- B. Finding the longest increasing subsequence in an array
- C. Finding the longest common substring between two strings
- D. Finding the longest common prefix of two strings

Answer: A

Q443. What is the greedy choice property?

- A. Always choosing the most expensive option first
- B. A locally optimal choice leads to global optimum
- C. A property that all algorithms must have
- D. A property unique to dynamic programming only

Answer: B

Q444. What is the coin change problem (minimum coins)?

- A. Finding minimum coins to make a target amount
- B. Sorting coins by their face value order
- C. Counting total coins in a collection
- D. Finding the maximum number of coins used

Answer: A

Q445. What is the time complexity of the recursive Fibonacci without memoization?

- A. $O(n)$
- B. $O(\log n)$
- C. $O(n^2)$
- D. $O(2^n)$

Answer: D

Q446. How does memoization improve the Fibonacci computation?

- A. It reduces time from $O(2^n)$ to $O(n)$ by caching
- B. It makes the computation run in $O(\log n)$ time
- C. It does not improve it at all
- D. It makes the computation run in $O(1)$ time

Answer: A

Q447. What is the N-Queens problem?

- A. A graph traversal problem with N vertices
- B. Sorting N queens by their rank and position
- C. Finding N queens in a given dataset collection
- D. Placing N queens so no two attack each other

Answer: D

Q448. What is the difference between top-down and bottom-up dynamic programming?

- A. Top-down is always faster in practice than the bottom-up approach overall
- B. Top-down uses recursion with memoization; bottom-up uses tabulation
- C. Bottom-up always uses significantly more memory than the top-down approach
- D. They are exactly the same fundamental approach overall

Answer: B

Q449. What is Huffman coding an example of?

- A. Greedy algorithm
- B. Backtracking
- C. Dynamic programming
- D. Divide and conquer

Answer: A

Q450. What is the Longest Increasing Subsequence (LIS) problem's time complexity with DP?

- A. $O(n)$
- B. $O(2^n)$
- C. $O(n \log n)$
- D. $O(n^2)$

Answer: D

Q451. What is path compression in Union-Find?

- A. Making every node in find path point to root
- B. Compressing the storage used by the structure
- C. Shortening paths between nodes in a graph
- D. Compressing the data stored in each node

Answer: A

Q452. What is union by rank in Union-Find?

- A. Attaching shorter tree under taller tree during union
- B. Sorting all elements by their assigned rank
- C. Ranking all elements before performing union
- D. Using ranked hash functions for set operations

Answer: A

Q453. What is the nearly constant amortized time complexity of Union-Find with path compression and union by rank?

- A. $O(1)$ constant time
- B. $O(\log \log n)$ iterated log
- C. $O(\log n)$ logarithmic
- D. $O(\alpha(n))$ inverse Ackermann

Answer: D

Q454. What is a Fibonacci heap?

- A. A heap with better amortized decrease-key and insert
- B. A Fibonacci sequence number generator structure
- C. A binary heap variant with extra child pointers
- D. A heap that stores Fibonacci numbers only

Answer: A

Q455. What is a suffix array?

- A. An array of string prefixes
- B. A compressed trie structure
- C. A sorted array of all string suffixes
- D. An array suffix pointer

Answer: C

Q456. What is a sparse table?

- A. A table with many empty cells throughout
- B. A compressed hash table for sparse data
- C. A sparse matrix representation structure
- D. A structure for $O(1)$ static range minimum queries

Answer: D

Q457. What is a k-d tree?

- A. A space-partitioning tree for k-dimensional data
- B. A tree with k children per node always
- C. A tree of exactly depth k levels
- D. A k-ary heap data structure

Answer: A

Q458. What is a circular buffer?

- A. A buffer shaped like a circle visually
- B. A buffer for circular linked list nodes
- C. A double-ended buffer with two pointers
- D. A fixed-size buffer that wraps around when full

Answer: D

Q459. What is an interval tree?

- A. A tree for querying overlapping intervals efficiently
- B. A tree storing time intervals as keys
- C. A segment tree with interval-based updates
- D. A tree with regular spacing between nodes

Answer: A

Q460. What is the time complexity of inserting into a trie for a word of length L?

- A. $O(n)$
- B. $O(\log n)$
- C. $O(L)$
- D. $O(1)$

Answer: C

Q461. What is the KMP (Knuth-Morris-Pratt) algorithm?

- A. A string compression algorithm for reducing text size
- B. A hash-based search method for substring matching tasks
- C. A pattern matcher using failure function to skip redundancy
- D. A comparison-based sorting algorithm for string arrays

Answer: C

Q462. What is the time complexity of the KMP algorithm?

- A. $O(n + m)$
- B. $O(m \log n)$
- C. $O(n \times m)$
- D. $O(n^2)$

Answer: A

Q463. What is the Rabin-Karp algorithm?

- A. A divide and conquer string matching algorithm
- B. A string matching algorithm using rolling hashes
- C. A string sorting algorithm by character values
- D. A tree-based matching algorithm using tries

Answer: B

Q464. What is a rolling hash in Rabin-Karp?

- A. A hash that updates efficiently as the window slides
- B. A multi-level hash with cascading computation
- C. A hash that changes randomly each call
- D. A hash stored in a circular buffer structure

Answer: A

Q465. What is the longest palindromic substring problem?

- A. Finding the longest string in a collection
- B. Reversing a string to find palindromes
- C. Finding the longest contiguous palindrome substring
- D. Finding all palindromes in a given string

Answer: C

Q466. What is the time complexity of finding the longest palindromic substring using dynamic programming?

- A. $O(n^3)$
- B. $O(n)$
- C. $O(n^2)$
- D. $O(n \log n)$

Answer: C

Q467. What is string hashing used for?

- A. Efficiently comparing strings via hash values
- B. Encrypting strings for secure storage
- C. Compressing strings to save disk space
- D. Sorting strings into alphabetical order

Answer: A

Q468. What is the Boyer-Moore algorithm?

- A. A pattern matcher using bad character and good suffix rules
- B. A brute force matching algorithm for strings
- C. A string sorting algorithm by character frequency
- D. A compression algorithm for text data files

Answer: A

Q469. What is the Z-array of a string?

- A. An array where $Z[i]$ is longest prefix match starting at i
- B. The reverse of the original string as array
- C. An array of character frequencies for the string
- D. An array of zeros for initialization

Answer: A

Q470. How is the Z-algorithm used for pattern matching?

- A. By hashing the pattern and comparing values
- B. By using a suffix tree of the text string
- C. By sorting the text characters first
- D. By concatenating pattern and text then computing Z-array

Answer: D

Q471. What is an NP-complete problem?

- A. A problem that cannot be solved at all
- B. A problem with no known algorithm existing
- C. A problem in NP as hard as any NP problem
- D. A problem that runs in constant time always

Answer: C

Q472. What is an NP-hard problem?

- A. A problem with a known polynomial time solution
- B. A problem that is always in the class P
- C. A problem that is easy to solve efficiently
- D. A problem at least as hard as NP-complete problems

Answer: D

Q473. What is memoization's space-time tradeoff?

- A. It uses no extra space at all
- B. It uses extra space to avoid redundant computation
- C. It increases time to save storage space
- D. It reduces both space and time simultaneously

Answer: B

Q474. What is the relationship between P and NP?

- A. P and NP are completely disjoint sets
- B. $P = NP$ has been proven true conclusively
- C. $NP \neq P$ has been proven true conclusively
- D. $P \neq NP$, but whether $P = NP$ is open

Answer: D

Q475. What is a polynomial reduction?

- A. Reducing the degree of a polynomial expression
- B. Transforming one problem into another in polynomial time
- C. Reducing the time complexity of an algorithm
- D. Simplifying an algorithm's code implementation

Answer: B

Q476. What is tail recursion optimization?

- A. Adding a tail data structure to the recursion stack
- B. Removing all recursive calls from the entire program
- C. Making recursive call the last operation for iteration optimization
- D. Making the recursion run in reverse backward order

Answer: C

Q477. What is loop unrolling?

- A. Removing all loops from the source code base entirely
- B. Making loops run in the reverse backward execution order
- C. Executing multiple iterations per loop cycle to reduce overhead
- D. Converting all iterative loops to recursion-based calls

Answer: C

Q478. What is cache-friendly code?

- A. Code with no memory access at all
- B. Code that runs entirely in cache memory
- C. Code that uses caching library APIs
- D. Code that maximizes CPU cache hit rates

Answer: D

Q479. What is the difference between best case, worst case, and average case complexity?

- A. Best is minimum time, worst is maximum, average is expected
- B. Worst case is purely theoretical and never practical
- C. They are always the same value for any algorithm
- D. Best case only applies to small input sizes

Answer: A

Q480. Why is $O(n \log n)$ considered efficient for sorting?

- A. It works only for small input sizes efficiently
- B. It uses the least memory of all algorithms
- C. It matches the comparison-based sorting lower bound
- D. It is the fastest possible for all sorting

Answer: C

Q481. What is the time complexity of $T(n) = T(n-1) + n$?

- A. $O(n)$
- B. $O(2^n)$
- C. $O(n \log n)$
- D. $O(n^2)$

Answer: D

Q482. Which notation gives both upper and lower asymptotic bounds?

- A. Small-o notation
- B. Big-O notation
- C. Theta notation
- D. Omega notation

Answer: C

Q483. What is the space complexity of a recursive function with depth n ?

- A. $O(n)$ due to call stack frames
- B. $O(1)$ constant space always
- C. $O(n^2)$ due to nested recursion
- D. $O(\log n)$ due to stack halving

Answer: A

Q484. Which recurrence relation represents binary search?

- A. $T(n) = 2T(n/2) + n$
- B. $T(n) = 2T(n-1) + 1$
- C. $T(n) = T(n-1) + 1$
- D. $T(n) = T(n/2) + 1$

Answer: D

Q485. What is amortized analysis used for?

- A. Measuring best-case time for a single run
- B. Finding worst case of each single operation
- C. Calculating space used by a single function
- D. Computing the average cost over many operations

Answer: D

Q486. What is tail recursion?

- A. Recursion where the call is the last operation
- B. Recursion that uses an explicit stack structure
- C. Recursion with no defined base case at all
- D. Recursion with multiple simultaneous recursive calls

Answer: A

Q487. Which complexity is better for large inputs: $O(n \log n)$ or $O(n^2)$?

- A. $O(n^2)$ is always better for large inputs
- B. $O(n \log n)$ is better for large inputs
- C. Neither is better; they alternate in speed
- D. Both have identical growth for large inputs

Answer: B

Q488. What is the difference between an algorithm and a heuristic?

- A. Algorithms guarantee solutions; heuristics may not
- B. Algorithms use hardware; heuristics use software
- C. Heuristics guarantee correctness; algorithms do not
- D. Heuristics are faster; algorithms are always slower

Answer: A

Q489. What does the substitution method prove in recurrence solving?

- A. It proves by guessing and verifying via induction
- B. It proves by converting recurrence to iteration
- C. It proves by drawing a complete recursion tree
- D. It proves by applying the Master Theorem formula

Answer: A

Q490. What is the recurrence for merge sort?

- A. $T(n) = 2T(n/2) + n$
- B. $T(n) = T(n/2) + n$
- C. $T(n) = 2T(n-1) + 1$
- D. $T(n) = T(n-1) + n$

Answer: A

Q491. What is the advantage of a circular queue over a linear queue?

- A. It stores more elements in the same memory space
- B. It allows random access to any element directly
- C. It efficiently reuses vacated front space in array
- D. It automatically sorts elements during insertion

Answer: C

Q492. How does a doubly linked list differ from a singly linked list?

- A. Each node has pointers to both next and previous
- B. It stores two data fields in each single node
- C. It uses half the memory of a singly linked list
- D. It can only be traversed in the forward direction

Answer: A

Q493. What is the time complexity of searching in an unsorted linked list?

- A. $O(1)$
- B. $O(n \log n)$
- C. $O(\log n)$
- D. $O(n)$

Answer: D

Q494. What is a sentinel node in a linked list?

- A. A dummy node used to simplify boundary operations
- B. A node that stores the list's total length value
- C. A special node that stores metadata about types
- D. A node that points to the middle of the list

Answer: A

Q495. How is a stack implemented using two queues?

- A. Push to both queues; pop alternates between the two
- B. Push to Q2 always; pop alternates and merges both queues
- C. Push and pop both operate on Q1; Q2 stores overflow
- D. Push to Q1 always; for pop, move Q1 to Q2 and dequeue

Answer: D

Q496. What is the time complexity of inserting at a specific position in an array?

- A. $O(n)$
- B. $O(n^2)$
- C. $O(\log n)$
- D. $O(1)$

Answer: A

Q497. What problem does a circular linked list solve?

- A. It automatically maintains elements in sorted order
- B. It reduces memory usage by half compared to singly
- C. It allows continuous traversal without null termination
- D. It enables $O(1)$ search for any element in list

Answer: C

Q498. Which data structure is best for implementing a browser's back button?

- A. Deque
- B. Array
- C. Stack
- D. Queue

Answer: C

Q499. What is the disadvantage of an array-based stack?

- A. It requires $O(n)$ time for every push operation
- B. It cannot store primitive data type values at all
- C. It has a fixed capacity that can cause overflow
- D. It does not support push and pop operations well

Answer: C

Q500. What is the amortized cost of appending to a dynamic array?

- A. $O(n^2)$
- B. $O(\log n)$
- C. $O(1)$
- D. $O(n)$

Answer: C

Q501. What is the maximum number of nodes in a binary tree of height h ?

- A. $2^{(h+1)} - 1$
- B. $h^2 + 1$
- C. 2^h
- D. $2h + 1$

Answer: A

Q502. What is the difference between adjacency list and adjacency matrix?

- A. Lists cannot represent weighted graphs but matrices can
- B. Lists use $O(V+E)$ space; matrices use $O(V^2)$ space
- C. Both always use exactly the same amount of space
- D. Lists use $O(V^2)$ space; matrices use $O(V+E)$ space

Answer: B

Q503. What is a binary search tree (BST) property?

- A. Nodes are arranged in ascending order left to right
- B. Left child is less; right child is greater than parent
- C. All nodes have exactly two children always
- D. Every level must be completely filled with nodes

Answer: B

Q504. What is a full binary tree?

- A. A tree where each node has exactly one child node
- B. A tree where all leaf nodes are at the same level
- C. A tree where every node has zero or two children
- D. A tree with the maximum possible number of nodes

Answer: C

Q505. How is a graph different from a tree?

- A. Trees have weighted edges; graphs have unweighted ones
- B. Graphs must be connected; trees can be disconnected
- C. Trees can have cycles; graphs cannot have any cycles
- D. Graphs can have cycles; trees are always acyclic

Answer: D

Q506. What is a self-loop in a graph?

- A. An edge shared between two separate graph components
- B. A path that returns to the starting vertex via others
- C. A cycle that passes through every vertex exactly once
- D. An edge that connects a vertex to itself directly

Answer: D

Q507. What is the space complexity of an adjacency matrix for V vertices?

- A. $O(V)$
- B. $O(V + E)$
- C. $O(E)$
- D. $O(V^2)$

Answer: D

Q508. What is a balanced binary tree?

- A. A tree where left and right subtree heights differ by at most one
- B. A tree where all nodes have exactly two children always
- C. A tree where all leaf nodes store the same data value
- D. A tree that contains equal number of left and right nodes

Answer: A

Q509. What is a directed acyclic graph (DAG)?

- A. An undirected graph with at least one cycle present
- B. A tree with multiple root nodes and directed edges
- C. A graph where all edges are bidirectional and acyclic
- D. A directed graph with no cycles between any vertices

Answer: D

Q510. What is the depth of a node in a tree?

- A. The maximum number of nodes at that particular level
- B. The total number of nodes in the entire tree structure
- C. The number of edges from the root to that specific node
- D. The number of children that the node currently has

Answer: C

Q511. How does interpolation search improve over binary search?

- A. It estimates position based on value distribution
- B. It uses extra memory to cache previous search results
- C. It works on unsorted arrays unlike binary search
- D. It always guarantees $O(1)$ time for any input data

Answer: A

Q512. What is the average-case time complexity of interpolation search on uniform data?

- A. $O(\log n)$
- B. $O(\log \log n)$
- C. $O(\sqrt{n})$
- D. $O(n)$

Answer: B

Q513. What is exponential search used for?

- A. Finding range for binary search on unbounded arrays
- B. Searching in a hash table with open addressing
- C. Finding an element in an unsorted linked list
- D. Searching in a balanced binary search tree structure

Answer: A

Q514. What is the time complexity of searching in a balanced BST?

- A. $O(n)$
- B. $O(\log n)$
- C. $O(n^2)$
- D. $O(1)$

Answer: B

Q515. What is ternary search?

- A. A search that compares three elements at each single step
- B. A search that uses three different search algorithms together
- C. A search that requires three sorted arrays simultaneously
- D. A search that divides the search space into three equal parts

Answer: D

Q516. What is jump search and what is its optimal block size?

- A. Search every other element; block size is always two
- B. Jump to random positions; block size is chosen randomly
- C. Skip logarithmic blocks; block size is always $\log n$
- D. Skip fixed blocks then linear search; block size is $\sqrt{n}$

Answer: D

Q517. What is the space complexity of recursive binary search?

- A. $O(1)$
- B. $O(n \log n)$
- C. $O(\log n)$
- D. $O(n)$

Answer: C

Q518. How does Fibonacci search determine the split point?

- A. It always splits the array exactly at the golden ratio
- B. It uses Fibonacci numbers to divide the search space
- C. It uses random Fibonacci-like sequences for splitting
- D. It hashes elements using Fibonacci number as the key

Answer: B

Q519. What is the key advantage of binary search over linear search?

- A. It can find multiple elements in a single search
- B. It has $O(\log n)$ time versus $O(n)$ on sorted data
- C. It uses less memory than linear search algorithm
- D. It works on any type of data without restrictions

Answer: B

Q520. What happens if binary search is applied to an unsorted array?

- A. It will automatically sort the array first then search
- B. It will always return the correct result anyway
- C. It will convert to linear search as a fallback mode
- D. It may skip the target and return incorrect result

Answer: D

Q521. What is the time complexity of merge sort in all cases?

- A. $O(n)$
- B. $O(n \log n)$
- C. $O(n^2)$
- D. $O(\log n)$

Answer: B

Q522. What is the worst-case pivot selection in quick sort?

- A. Randomly selecting the pivot element each time
- B. Always picking the median as pivot element
- C. Always picking smallest or largest as pivot element
- D. Picking the middle index element as the pivot

Answer: C

Q523. What is the space complexity of merge sort?

- A. $O(n)$
- B. $O(\log n)$
- C. $O(n^2)$
- D. $O(1)$

Answer: A

Q524. How does counting sort achieve $O(n+k)$ time complexity?

- A. It counts occurrences of each value and computes positions
- B. It divides elements into buckets based on digit values
- C. It recursively splits and merges the input array halves
- D. It compares pairs of elements to find their order

Answer: A

Q525. Why is quick sort preferred over merge sort in practice?

- A. Quick sort has better cache performance and low overhead
- B. Quick sort uses more memory which speeds up sorting
- C. Quick sort has better worst-case time complexity
- D. Quick sort is always stable unlike merge sort algorithm

Answer: A

Q526. What is the time complexity of heap sort?

- A. $O(\log n)$
- B. $O(n \log n)$
- C. $O(n^2)$
- D. $O(n)$

Answer: B

Q527. What is radix sort and when is it used?

- A. A comparison sort that divides data by value ranges
- B. A recursive sort that always selects median as pivot
- C. A non-comparison sort processing digits from least significant
- D. A sorting method that uses hashing for element placement

Answer: C

Q528. What is the key operation in the partition step of quick sort?

- A. Merging two sorted halves into a single sorted array
- B. Finding the minimum element in the unsorted portion
- C. Rearranging elements around a chosen pivot value point
- D. Comparing every pair of elements and swapping if needed

Answer: C

Q529. Is selection sort a stable algorithm?

- A. It is stable only when elements are all distinct
- B. It is stable only for arrays under a certain size
- C. No, it is not stable due to long-range swapping
- D. Yes, it is always stable for all possible inputs

Answer: C

Q530. What is shell sort?

- A. A generalized insertion sort using decreasing gap sizes
- B. A comparison sort based on the shell command utility
- C. A variant of merge sort using shell-like structures
- D. A sorting algorithm that uses a shell data structure

Answer: A

Q531. What is open addressing in hash tables?

- A. Storing colliding elements in a separate overflow area
- B. Using open-source hash functions for all key types
- C. Finding alternative slots within the same table for collisions
- D. Leaving table slots open and unoccupied permanently

Answer: C

Q532. What is linear probing?

- A. Checking slots at quadratically increasing intervals
- B. Checking sequential slots one by one after collision
- C. Using a second hash function to find alternate slots
- D. Probing in a linear linked list attached to each slot

Answer: B

Q533. What is the primary clustering problem in linear probing?

- A. Occupied slots form long contiguous runs over time
- B. Clusters form in the linked lists at each slot
- C. All keys hash to the same single slot in table
- D. Keys cluster at the end of the hash table only

Answer: A

Q534. How does quadratic probing reduce clustering?

- A. It uses a secondary hash function for probing step
- B. It probes at quadratically increasing intervals from hash
- C. It uses random slot selection for each collision
- D. It sorts the table after each insertion operation

Answer: B

Q535. What is double hashing?

- A. Using the same hash function twice on the same key
- B. Applying one hash function to the result of another
- C. Hashing both the key and value separately for storage
- D. Using two different hash functions for probe sequence

Answer: D

Q536. When should a hash table be resized?

- A. When the first collision occurs in the table structure
- B. When all slots in the table are completely occupied
- C. When the load factor exceeds a certain threshold value
- D. When the hash function starts producing negative values

Answer: C

Q537. What is the multiplication method of hashing?

- A. Multiplying the key by the table size directly
- B. Multiplying the hash by the load factor for index
- C. Multiplying key by constant, extracting fractional part
- D. Multiplying all keys together to get a single hash

Answer: C

Q538. What is rehashing?

- A. Removing all elements and starting with an empty table
- B. Rebuilding the hash table with a new size or hash function
- C. Reversing the hash function to retrieve original keys
- D. Recomputing hash values using the same table and function

Answer: B

Q539. What is the birthday paradox's relevance to hashing?

- A. It explains why hash functions must use prime table sizes
- B. It shows collisions are more likely than intuition suggests
- C. It demonstrates that larger tables always prevent collisions
- D. It proves that hash tables never experience collisions

Answer: B

Q540. What is the difference between a hash map and a hash set?

- A. Hash maps are faster; hash sets use more memory space
- B. Hash maps store key-value pairs; hash sets store keys only
- C. Hash maps use chaining; hash sets use open addressing
- D. Hash sets allow duplicates; hash maps do not allow them

Answer: B

Q541. What is an AVL tree?

- A. A multi-way tree used primarily for database indexing
- B. A self-balancing BST with height difference at most one
- C. A BST where all leaves are at the exact same level
- D. A tree where each node stores at most two keys only

Answer: B

Q542. What rotation is performed for a left-left imbalance in an AVL tree?

- A. Left rotation
- B. Right-left rotation
- C. Right rotation
- D. Left-right rotation

Answer: C

Q543. What is a red-black tree?

- A. A binary tree that alternates red and black at each level
- B. A BST where nodes are colored red or black with rules
- C. A tree using red and black to indicate data priority
- D. A tree where red nodes are leaves and black are internal

Answer: B

Q544. What is the maximum height of an AVL tree with n nodes?

- A. $O(n^2)$
- B. $O(n)$
- C. $O(\log n)$
- D. $O(\sqrt{n})$

Answer: C

Q545. How does deletion work in a BST when the node has two children?

- A. Simply remove the node and leave children as orphan nodes
- B. Delete both children and reconnect their subtrees together
- C. Swap with the root node and then delete the root node
- D. Replace with in-order successor or predecessor, then delete it

Answer: D

Q546. What is the purpose of a B-tree?

- A. Minimizing disk I/O for database and file system access
- B. Compressing tree data for network data transmission
- C. Storing data in a binary format for fast processing
- D. Balancing workload across multiple processor threads

Answer: A

Q547. What is a segment tree used for?

- A. Segmenting data into equal-sized portions for sorting
- B. Creating visual segments for graphical tree rendering
- C. Answering range queries and updates on an array efficiently
- D. Storing segments of text for string matching algorithms

Answer: C

Q548. What is a binary heap?

- A. A binary tree where all nodes have the exact same value
- B. A sorted binary tree with elements in ascending order
- C. A complete binary tree satisfying the heap property order
- D. A binary tree where leaves store minimum values always

Answer: C

Q549. What is the time complexity of building a heap from an array?

- A. $O(n \log n)$
- B. $O(n)$
- C. $O(n^2)$
- D. $O(\log n)$

Answer: B

Q550. How is a node inserted into an AVL tree?

- A. Insert as in BST, then rotate to fix any imbalance
- B. Always insert at the root and push other nodes down
- C. Insert at the deepest level regardless of key value
- D. Insert at a random position and then sort the tree

Answer: A

Q551. What does Dijkstra's algorithm find?

- A. All possible paths between two specific graph vertices
- B. Shortest paths from a single source to all vertices
- C. Minimum spanning tree of a weighted graph structure
- D. Maximum flow through a network with capacity constraints

Answer: B

Q552. What is the time complexity of Dijkstra's algorithm with a binary heap?

- A. $O(V + E)$
- B. $O(E * V)$
- C. $O((V+E) \log V)$
- D. $O(V^2)$

Answer: C

Q553. What does Kruskal's algorithm compute?

- A. Minimum spanning tree using edge sorting approach
- B. Shortest path between two specific vertices only
- C. All connected components of an undirected graph
- D. Topological ordering of a directed acyclic graph

Answer: A

Q554. What does Prim's algorithm compute?

- A. Minimum spanning tree by growing from a vertex
- B. Shortest paths from all sources to all destinations
- C. Strongly connected components of a directed graph
- D. Maximum matching in a bipartite graph structure

Answer: A

Q555. What is topological sorting?

- A. Arranging vertices in a circle for visual graph display
- B. Sorting edges by weight for minimum spanning tree building
- C. Sorting graph vertices by their degree in ascending order
- D. Linear ordering of vertices so each edge goes forward

Answer: D

Q556. Can topological sort be performed on a graph with cycles?

- A. Yes, but the result may not be unique or deterministic
- B. No, topological sort requires a directed acyclic graph
- C. No, but it can be performed on undirected cyclic graph
- D. Yes, cycles are handled by ignoring back edges only

Answer: B

Q557. What limitation does Dijkstra's algorithm have?

- A. It only works on undirected graphs not directed ones
- B. It requires the graph to be a complete graph always
- C. It cannot handle graphs with more than 100 vertices
- D. It fails on graphs with negative weight edge values

Answer: D

Q558. What is the Union-Find data structure used for?

- A. Storing vertex and edge data in compressed format
- B. Searching for elements in a sorted graph structure
- C. Finding unions and intersections of sorted arrays
- D. Tracking connected components and detecting cycles

Answer: D

Q559. What is the Bellman-Ford algorithm used for?

- A. Finding maximum flow in a network with capacities
- B. Finding minimum spanning trees in undirected graphs
- C. Finding shortest paths even with negative edge weights
- D. Finding all bridges and articulation points in graphs

Answer: C

Q560. What is the time complexity of the Floyd-Warshall algorithm?

- A. $O(E^2)$
- B. $O(V^3)$
- C. $O(V^2)$
- D. $O(V * E)$

Answer: B

Q561. What is the difference between memoization and tabulation?

- A. Memoization uses tables; tabulation uses recursive calls
- B. Both are identical approaches with different naming only
- C. Memoization is bottom-up; tabulation is top-down approach
- D. Memoization is top-down; tabulation is bottom-up approach

Answer: D

Q562. What are the two key properties for dynamic programming applicability?

- A. Optimal substructure and overlapping subproblems present
- B. Greedy choice property and matroid structure correctness
- C. Logarithmic depth and polynomial branching factor only
- D. Divide and conquer with non-overlapping subproblem cases

Answer: A

Q563. What is the time complexity of the 0/1 knapsack problem using DP?

- A. $O(n * W)$
- B. $O(n \log n)$
- C. $O(n^2)$
- D. $O(2^n)$

Answer: A

Q564. What is the longest common subsequence (LCS) problem?

- A. Finding the longest substring that appears in one string
- B. Finding the longest path between two nodes in a graph
- C. Finding the longest sorted subsequence in one sequence
- D. Finding the longest subsequence common to two sequences

Answer: D

Q565. What is the activity selection problem?

- A. Selecting activities based on their priority level values
- B. Selecting maximum non-overlapping activities by finish time
- C. Selecting activities that require the minimum total cost
- D. Selecting the activity with the longest duration possible

Answer: B

Q566. What is the rod cutting problem?

- A. Finding the minimum number of cuts to divide a rod evenly
- B. Cutting a rod into equal pieces for uniform distribution
- C. Cutting a rod into pieces to maximize total sale revenue
- D. Cutting rods from a sheet to minimize material waste cost

Answer: C

Q567. How does branch and bound differ from backtracking?

- A. Branch and bound uses bounds to prune unpromising branches
- B. Backtracking uses bounds; branch and bound does not use any
- C. Branch and bound only works on graph problems specifically
- D. They are identical algorithms with different names only

Answer: A

Q568. What is the coin change problem using dynamic programming?

- A. Finding minimum coins needed to make a given amount value
- B. Sorting coins by value for efficient change distribution
- C. Finding the total value of all coins combined together
- D. Counting the total number of coins in a given collection

Answer: A

Q569. What is the fractional knapsack problem?

- A. A problem of dividing a knapsack into fractional parts
- B. A knapsack problem where items must be taken completely
- C. A knapsack problem where items can be taken partially
- D. A problem of distributing items into fractional groups

Answer: C

Q570. What is the time complexity of the LCS problem using DP?

- A. $O(m + n)$
- B. $O(m * n)$
- C. $O(2^{(m+n)})$
- D. $O(m^2 + n^2)$

Answer: B

Q571. What is a Fibonacci heap?

- A. A min-heap that stores only Fibonacci numbers as values
- B. A binary heap where the size follows Fibonacci sequence
- C. A heap with amortized $O(1)$ insert and decrease-key time
- D. A heap based on Fibonacci numbers for index calculation

Answer: C

Q572. What is a suffix array?

- A. An array storing the frequency of each character in string
- B. An array of all prefixes of a string sorted in order
- C. An array of all suffixes of a string sorted lexicographically
- D. An array of indices where a specific pattern is found

Answer: C

Q573. What is path compression in Union-Find?

- A. Compressing file paths for efficient storage in memory
- B. Shortening the physical path between two graph vertices
- C. Making each node point directly to the root during find
- D. Compressing data along the path in a tree structure

Answer: C

Q574. What is union by rank in Union-Find?

- A. Attaching the shorter tree under the root of taller tree
- B. Ranking all elements before performing any union operation
- C. Unioning sets based on alphabetical rank of elements
- D. Sorting elements by rank before adding to disjoint sets

Answer: A

Q575. What is a trie's space complexity for n strings of average length m?

- A. $O(m)$
- B. $O(n^2)$
- C. $O(n * m)$
- D. $O(n)$

Answer: C

Q576. What is an interval tree?

- A. A binary tree with nodes at regular interval positions
- B. A tree that stores intervals for overlap and point queries
- C. A tree where node values differ by a constant interval
- D. A tree that stores time intervals for scheduling queries

Answer: B

Q577. What is the time complexity of decrease-key in a Fibonacci heap?

- A. $O(n)$
- B. $O(n \log n)$
- C. $O(1)$ amortized
- D. $O(\log n)$

Answer: C

Q578. What is a k-d tree used for?

- A. Dividing a tree into k balanced subtrees for search
- B. Storing k different data types in a single tree node
- C. Organizing points in k-dimensional space for queries
- D. Creating k copies of a binary tree for redundancy

Answer: C

Q579. What is the advantage of a balanced BST over a hash table?

- A. It provides faster average-case lookup than hashing does
- B. It has simpler implementation than hash table structures
- C. It maintains elements in sorted order for range queries
- D. It uses less memory than a hash table for same elements

Answer: C

Q580. What is a binomial heap?

- A. A binary heap with binomial distribution of element values
- B. A heap that performs binomial search for extract min
- C. A collection of binomial trees satisfying heap property
- D. A heap using binomial coefficients for priority values

Answer: C

Q581. What is the KMP algorithm?

- A. A string sorting algorithm using key-median partitioning
- B. A pattern matching algorithm using a prefix failure table
- C. A string compression algorithm using key-mapped patterns
- D. A substring search using kernel-mapped positioning index

Answer: B

Q582. What is the time complexity of the KMP algorithm?

- A. $O(n^2)$
- B. $O(m^2)$
- C. $O(n * m)$
- D. $O(n + m)$

Answer: D

Q583. What is the Rabin-Karp algorithm?

- A. A text compression algorithm using repeated substrings
- B. A string sorting algorithm based on radix computation
- C. A pattern matching algorithm using rolling hash function
- D. A string encryption algorithm using the Karp cipher key

Answer: C

Q584. What is the failure function in KMP?

- A. A table mapping each character to its failure position
- B. A table storing longest proper prefix that is also suffix
- C. A function that detects when the algorithm has failed
- D. A function that counts the number of failed comparisons

Answer: B

Q585. What is a rolling hash?

- A. A hash that changes completely for each new substring
- B. A hash that rotates character positions in the input string
- C. A hash computed by rolling through all possible hash functions
- D. A hash updated incrementally by adding and removing characters

Answer: D

Q586. What is the longest palindromic substring problem?

- A. Finding the longest string that contains a palindrome inside
- B. Finding the longest contiguous palindrome within a string
- C. Generating the longest possible palindrome from characters
- D. Finding all palindromes in a string and returning longest

Answer: B

Q587. What is the Boyer-Moore algorithm's key idea?

- A. Comparing only the first and last characters of pattern
- B. Comparing every other character for faster matching speed
- C. Comparing pattern from right to left with shift heuristics
- D. Comparing characters from left to right systematically

Answer: C

Q588. What is the Z-algorithm used for in string processing?

- A. Sorting strings in reverse lexicographic Z to A order
- B. Compressing strings using Z-encoding for smaller storage
- C. Encrypting strings using the Z-cipher transformation
- D. Computing Z-array for pattern matching in linear time

Answer: D

Q589. What is string hashing used for?

- A. Sorting strings in hash-based order for faster access
- B. Compressing strings to reduce storage space requirements
- C. Efficiently comparing strings and finding pattern matches
- D. Encrypting strings for secure network communication only

Answer: C

Q590. What is the time complexity of checking if a string is a palindrome?

- A. $O(n \log n)$
- B. $O(1)$
- C. $O(n^2)$
- D. $O(n)$

Answer: D

Q591. What is the difference between best, average, and worst case?

- A. They describe input scenarios giving minimum, expected, and maximum cost
- B. They represent different programming languages and their speeds
- C. They measure performance for small, medium, and large inputs
- D. They describe memory, disk, and network usage respectively

Answer: A

Q592. What is the time complexity of an algorithm with three sequential loops of $O(n)$?

- A. $O(3n)$
- B. $O(n)$
- C. $O(n^2)$
- D. $O(n^3)$

Answer: B

Q593. What is the space-time tradeoff?

- A. Trading code readability for faster execution time always
- B. Trading algorithm correctness for improved speed results
- C. Trading disk space for network bandwidth in algorithms
- D. Trading space complexity for time complexity or vice versa

Answer: D

Q594. What is the complexity of matrix multiplication of two $n \times n$ matrices?

- A. $O(n^2)$
- B. $O(n^3)$
- C. $O(n \log n)$
- D. $O(n)$

Answer: B

Q595. What is loop unrolling as an optimization technique?

- A. Converting recursive loops into iterative ones for speed
- B. Removing all loops from the algorithm entirely always
- C. Unrolling nested loops into a single flat loop structure
- D. Reducing loop overhead by executing multiple iterations per cycle

Answer: D

Q596. What is memoization's effect on the Fibonacci sequence computation?

- A. It changes complexity from $O(n^2)$ to $O(n \log n)$ time
- B. It has no effect on the time complexity of Fibonacci
- C. It changes complexity from $O(n)$ to $O(1)$ constant time
- D. It changes complexity from $O(2^n)$ to $O(n)$ linear time

Answer: D

Q597. What is the significance of polynomial vs exponential complexity?

- A. Polynomial algorithms are considered efficient; exponential are not
- B. Exponential algorithms are always faster than polynomial ones
- C. Both are equally efficient for all practical input sizes
- D. Polynomial algorithms are impractical; exponential are efficient

Answer: A

Q598. What is cache-friendly code?

- A. Code that disables caching for more predictable performance
- B. Code that accesses memory in patterns that maximize cache hits
- C. Code that uses a custom cache data structure for lookups
- D. Code that stores all data in the CPU cache permanently

Answer: B

Q599. What is the complexity class P?

- A. Problems solvable in polynomial time by deterministic machine
- B. Problems that have no known solution in any time bound
- C. Problems solvable in polynomial time by nondeterministic machine
- D. Problems solvable only in exponential time by any machine

Answer: A

Q600. What is tail call optimization?

- A. Removing the last call in a function for faster returns
- B. Optimizing function calls at the tail of a linked list
- C. Optimizing the last element of an array for faster access
- D. Reusing the current stack frame for a tail recursive call

Answer: D

Q601. What is the time complexity of $T(n) = 4T(n/2) + n$ using the Master Theorem?

- A. $O(n)$
- B. $O(n \log n)$
- C. $O(n^2)$
- D. $O(n^2 \log n)$

Answer: C

Q602. What is the difference between recursion and iteration in terms of space usage?

- A. Recursion always uses less space
- B. Iteration typically uses less space because it doesn't require a call stack
- C. They always use the same amount of space
- D. Iteration always uses more space due to loop variables

Answer: B

Q603. What is the solution to the recurrence $T(n) = T(n/2) + O(1)$?

- A. $O(n)$
- B. $O(n \log n)$
- C. $O(\log n)$
- D. $O(1)$

Answer: C

Q604. If $f(n) = 5n^3 + 2n^2 + 100n + 42$, what is the Big-O classification?

- A. $O(n^2)$
- B. $O(5n^3)$
- C. $O(n^3)$
- D. $O(n^4)$

Answer: C

Q605. What does it mean if an algorithm is said to run in polynomial time?

- A. Its running time is $O(2^n)$
- B. Its running time can be expressed as $O(n^k)$ for some constant k
- C. Its running time is always $O(n)$
- D. It cannot be analyzed using Big-O notation

Answer: B

Q606. What is the recursion tree method used for?

- A. Sorting elements in a tree
- B. Visualizing the work done at each level of a recurrence to find the total cost
- C. Converting iterative code to recursive code
- D. Building binary search trees

Answer: B

Q607. Which of the following is NOT a valid step in the substitution method for solving recurrences?

- A. Guess the form of the solution
- B. Use mathematical induction to prove the guess
- C. Assume the solution holds for smaller values
- D. Differentiate the recurrence with respect to n

Answer: D

Q608. What is the space complexity of an algorithm that creates a new array of size n within a loop that runs n times?

- A. $O(n)$
- B. $O(n^2)$
- C. $O(\log n)$
- D. $O(1)$

Answer: A

Q609. What does little-o notation $o(f(n))$ signify?

- A. The growth rate is exactly $f(n)$
- B. The growth rate is strictly less than $f(n)$ asymptotically
- C. The growth rate is at least $f(n)$
- D. The growth rate is strictly greater than $f(n)$

Answer: B

Q610. How does the accounting method of amortized analysis work?

- A. It counts the total number of function calls
- B. It assigns different charges to operations so cheap operations subsidize expensive ones
- C. It measures the actual runtime on a specific machine
- D. It ignores expensive operations in the analysis

Answer: B

Q611. How can you implement a min-queue that supports enqueue, dequeue, and getMin all in amortized $O(1)$?

- A. Use a single array with sorting
- B. Use two stacks, each augmented to track minimums
- C. Use a linked list sorted by value
- D. Use a binary search tree

Answer: B

Q612. What is the time complexity of inserting an element at a specific position i in an array of size n ?

- A. $O(1)$
- B. $O(i)$
- C. $O(n - i)$
- D. $O(n)$

Answer: D

Q613. What is the advantage of using a circular buffer over a linear array for a queue?

- A. Circular buffers use less memory
- B. They avoid the need to shift elements after dequeue operations
- C. They support random access faster
- D. They can store more elements

Answer: B

Q614. What happens when a dynamic array doubles its size upon reaching capacity?

- A. Elements are sorted into the new array
- B. A new array of double the size is allocated and all elements are copied over
- C. The old array is extended in place without copying
- D. Only the new elements are placed in the new array

Answer: B

Q615. How can you use a stack to check if parentheses in an expression are balanced?

- A. Count opening and closing parentheses
- B. Push opening brackets onto the stack and pop when a matching closing bracket is found
- C. Sort the brackets and compare halves
- D. Use two stacks, one for each type

Answer: B

Q616. What is the time complexity of reversing a singly linked list using iteration?

- A. $O(n^2)$
- B. $O(n \log n)$
- C. $O(1)$
- D. $O(n)$

Answer: D

Q617. What is a priority queue and how does it differ from a regular queue?

- A. A priority queue is the same as a regular queue but faster
- B. Elements are dequeued based on priority rather than insertion order
- C. A priority queue only stores integers
- D. A priority queue reverses the order of insertion

Answer: B

Q618. What is the space overhead of a singly linked list compared to an array for storing n integers?

- A. No extra space
- B. $O(1)$ extra space
- C. $O(n)$ extra space for storing next pointers
- D. $O(n^2)$ extra space

Answer: C

Q619. How does a deque (double-ended queue) differ from a standard queue?

- A. A deque is slower than a queue
- B. A deque allows insertion and deletion at both the front and rear
- C. A deque only allows deletion from both ends
- D. A deque can only store two elements

Answer: B

Q620. What is the time complexity of finding the middle element of a singly linked list in one pass?

- A. $O(n^2)$
- B. $O(n)$
- C. $O(\log n)$
- D. $O(1)$

Answer: B

Q621. What is the time complexity of inserting an element into a binary search tree in the average case?

- A. $O(1)$
- B. $O(\log n)$
- C. $O(n)$
- D. $O(n \log n)$

Answer: B

Q622. What is the relationship between the number of edges and vertices in a complete graph K_n ?

- A. n edges
- B. $n(n-1)/2$ edges
- C. n^2 edges
- D. $2n$ edges

Answer: B

Q623. What is the advantage of an adjacency list representation over an adjacency matrix for sparse graphs?

- A. Adjacency lists use $O(V^2)$ space regardless of edges
- B. Adjacency lists use $O(V + E)$ space, which is much less than $O(V^2)$ for sparse graphs
- C. Adjacency lists allow $O(1)$ edge lookup
- D. Adjacency lists are always slower

Answer: B

Q624. What is the time complexity of deleting the minimum element from a min-heap with n elements?

- A. $O(1)$
- B. $O(\log n)$
- C. $O(n)$
- D. $O(n \log n)$

Answer: B

Q625. What is a perfect binary tree?

- A. A tree where every node has exactly one child
- B. A tree where all internal nodes have two children and all leaves are at the same depth
- C. A tree where nodes are sorted
- D. A tree with the minimum possible height for n nodes

Answer: B

Q626. What is the difference between a tree and a graph?

- A. Trees can have cycles but graphs cannot
- B. A tree is a connected acyclic graph, while a graph can have cycles and disconnected components
- C. Graphs always have fewer edges than trees
- D. Trees can only store numbers

Answer: B

Q627. How is a binary heap typically stored in memory?

- A. As a linked list of nodes
- B. As a contiguous array using index formulas for parent and children
- C. As a hash table
- D. As a 2D matrix

Answer: B

Q628. What is the in-degree of a vertex in a directed graph?

- A. The number of edges going out from the vertex
- B. The total number of edges in the graph
- C. The number of edges coming into the vertex
- D. The number of self-loops at the vertex

Answer: C

Q629. What is the sum of degrees of all vertices in an undirected graph with E edges?

- A. E
- B. $2E$
- C. E^2
- D. $V + E$

Answer: B

Q630. What is a degenerate (skewed) binary tree?

- A. A tree where every node has exactly two children
- B. A tree where every node has at most one child, resembling a linked list
- C. A tree with all leaves at the same level
- D. A tree stored using an adjacency matrix

Answer: B

Q631. What is the time complexity of jump search on a sorted array of n elements?

- A. $O(n)$
- B. $O(\log n)$
- C. $O(\sqrt{n})$
- D. $O(n \log n)$

Answer: C

Q632. How does exponential search work?

- A. It searches all elements exponentially faster
- B. It finds the range where the target lies by doubling the index, then applies binary search within that range
- C. It multiplies each element by 2
- D. It sorts the array first, then searches linearly

Answer: B

Q633. What is the worst-case time complexity of ternary search?

- A. $O(n)$
- B. $O(\log_3 n)$, which is $O(\log n)$
- C. $O(\sqrt{n})$
- D. $O(n \log n)$

Answer: B

Q634. What is the key condition that makes interpolation search perform better than binary search?

- A. The array must be in descending order
- B. The data must be uniformly distributed
- C. The array size must be a power of 2
- D. The array must contain only integers

Answer: B

Q635. How can binary search be modified to find the leftmost (first) occurrence of a duplicate element?

- A. Return the first match found
- B. When the target is found, continue searching in the left half instead of returning immediately
- C. Sort the array again after finding the element
- D. Use linear search after finding any occurrence

Answer: B

Q636. What is the Fibonacci search technique?

- A. Searching for Fibonacci numbers in an array
- B. A search method that divides the array using Fibonacci numbers instead of halving
- C. A recursive search that calls itself twice
- D. A hash-based search technique

Answer: B

Q637. What is the time complexity of searching in an unsorted hash table versus a sorted array using binary search?

- A. Hash table: $O(n)$, Binary search: $O(\log n)$
- B. Hash table: $O(1)$ average, Binary search: $O(\log n)$
- C. Both are $O(\log n)$
- D. Both are $O(1)$

Answer: B

Q638. What is a search tree and what operation does it optimize?

- A. A tree used for sorting elements
- B. A tree data structure that organizes data to optimize search operations to $O(\log n)$
- C. A tree that stores only search queries
- D. A tree with a fixed number of nodes

Answer: B

Q639. How does binary search on answer (parametric search) work?

- A. It searches for the answer in a sorted array of answers
- B. It binary searches over the answer space using a feasibility check function
- C. It sorts the answers before searching
- D. It applies binary search twice

Answer: B

Q640. What is the primary advantage of exponential search over binary search for searching in an unbounded or infinite list?

- A. It uses less memory
- B. It does not require knowing the array size beforehand
- C. It is always faster
- D. It does not require the array to be sorted

Answer: B

Q641. What is the recurrence relation for the time complexity of merge sort?

- A. $T(n) = T(n-1) + n$
- B. $T(n) = 2T(n/2) + O(n)$
- C. $T(n) = T(n/2) + O(1)$
- D. $T(n) = 4T(n/2) + O(n)$

Answer: B

Q642. Why is quicksort generally faster than merge sort in practice despite having the same average complexity?

- A. Quicksort has a lower worst-case complexity
- B. Quicksort has better cache performance and smaller constant factors due to in-place operations
- C. Quicksort uses less stack space
- D. Quicksort is always stable

Answer: B

Q643. What pivot selection strategy helps avoid quicksort's worst case?

- A. Always picking the first element
- B. Always picking the last element
- C. Choosing the median of the first, middle, and last elements
- D. Picking the largest element

Answer: C

Q644. What is counting sort's key limitation?

- A. It cannot sort integers
- B. It requires $O(k)$ extra space where k is the range of input values, making it impractical for large ranges
- C. It is always slower than quicksort
- D. It cannot sort arrays larger than 100 elements

Answer: B

Q645. How does radix sort achieve linear time complexity?

- A. By using comparison-based sorting internally
- B. By sorting digit by digit from least significant to most significant using a stable sort like counting sort
- C. By randomly partitioning the array
- D. By using a binary search tree

Answer: B

Q646. What is the time complexity of heap sort and is it stable?

- A. $O(n^2)$ and stable
- B. $O(n \log n)$ and stable
- C. $O(n \log n)$ and not stable
- D. $O(n)$ and stable

Answer: C

Q647. What is shell sort and how does it improve insertion sort?

- A. It sorts only the outer shell of the array
- B. It insertion-sorts elements that are far apart first, reducing the gap progressively
- C. It uses a shell data structure
- D. It sorts the array into a shell shape

Answer: B

Q648. Is merge sort preferred over quicksort for sorting linked lists? Why?

- A. No, quicksort is always preferred
- B. Yes, because merge sort does not require random access and linked list merge is $O(1)$ extra space
- C. No, because merge sort is slower on linked lists
- D. Yes, but only for doubly linked lists

Answer: B

Q649. What is the purpose of the partition step in quicksort?

- A. To merge two sorted subarrays
- B. To place the pivot in its correct sorted position with smaller elements before it and larger after
- C. To find the median of the array
- D. To count the number of elements

Answer: B

Q650. What is the difference between LSD (Least Significant Digit) and MSD (Most Significant Digit) radix sort?

- A. LSD processes from the least significant digit and requires a stable sort; MSD processes from the most significant digit and can use recursion
- B. They are identical algorithms
- C. LSD is faster for all inputs
- D. MSD cannot sort strings

Answer: A

Q651. How does quadratic probing resolve collisions?

- A. By probing at positions $h(k)+1$, $h(k)+2$, $h(k)+3$, ...
- B. By probing at positions $h(k)+1^2$, $h(k)+2^2$, $h(k)+3^2$, ...
- C. By using a second hash function
- D. By rehashing the entire table

Answer: B

Q652. What is secondary clustering in hash tables?

- A. When all elements hash to the first slot
- B. When keys that hash to the same slot follow the same probe sequence
- C. When the hash table is divided into two clusters
- D. When the load factor exceeds 2

Answer: B

Q653. What is the expected number of probes for a successful search in a hash table with open addressing at load factor α ?

- A. $1/(1-\alpha)$
- B. $(1/\alpha) \ln(1/(1-\alpha^2))$
- C. $1/\alpha$
- D. $\log(1/(1-\alpha))$

Answer: B

Q654. Why is it recommended to use a prime number for hash table size when using the division method?

- A. Prime numbers are faster to compute
- B. It helps distribute keys more uniformly by reducing patterns from common factors
- C. Prime numbers use less memory
- D. It is not recommended; powers of 2 are always better

Answer: B

Q655. What is the typical threshold load factor at which a hash table should be resized?

- A. 0.25
- B. 0.5
- C. 0.75
- D. 1.0

Answer: C

Q656. What is the time complexity of rehashing a hash table with n elements?

- A. $O(1)$
- B. $O(\log n)$
- C. $O(n)$
- D. $O(n^2)$

Answer: C

Q657. How does double hashing avoid clustering?

- A. It uses two hash tables
- B. It uses a second hash function to determine the probe step size, making each key's probe sequence different
- C. It hashes the value twice before inserting
- D. It doubles the table size on each collision

Answer: B

Q658. What is the advantage of open addressing over chaining?

- A. Open addressing is always faster
- B. Open addressing has better cache performance because all data is stored in the table array itself
- C. Open addressing handles higher load factors better
- D. Open addressing never needs resizing

Answer: B

Q659. After approximately how many insertions into a hash table with m slots does the probability of at least one collision exceed 50%?

- A. m insertions
- B. $m/2$ insertions
- C. About $\sqrt{2m}$ insertions
- D. About $\log(m)$ insertions

Answer: C

Q660. What is the difference between separate chaining and open addressing?

- A. Separate chaining uses linked lists at each slot while open addressing probes for the next empty slot
- B. They are the same technique with different names
- C. Open addressing uses linked lists while chaining uses arrays
- D. Separate chaining never has collisions

Answer: A

Q661. What is the balance factor in an AVL tree?

- A. The number of nodes in the tree
- B. The difference in heights of the left and right subtrees of a node
- C. The total height of the tree
- D. The number of rotations performed

Answer: B

Q662. When is a right-left (RL) rotation performed in an AVL tree?

- A. When a node is inserted in the left subtree of the left child
- B. When a node is inserted in the right subtree of the left child
- C. When a node is inserted in the left subtree of the right child
- D. When a node is inserted in the right subtree of the right child

Answer: C

Q663. How is a node with two children deleted from a BST?

- A. Simply remove the node
- B. Replace it with its in-order successor or predecessor, then delete that replacement node
- C. Swap it with the root and remove the root
- D. Delete both children first, then remove the node

Answer: B

Q664. What is Morris traversal and what is its space complexity?

- A. A traversal using two stacks with $O(n)$ space
- B. An in-order traversal using threading to achieve $O(1)$ space without a stack or recursion
- C. A traversal that visits nodes randomly in $O(n)$ space
- D. A BFS traversal using $O(n)$ space

Answer: B

Q665. What is a Red-Black tree's maximum height in terms of the number of nodes n ?

- A. Exactly $\log_2(n)$
- B. At most $2 * \log_2(n + 1)$
- C. Exactly n
- D. At most $n/2$

Answer: B

Q666. What is a segment tree's time complexity for a range sum query after $O(n)$ construction?

- A. $O(n)$
- B. $O(\log n)$
- C. $O(1)$
- D. $O(" n)$

Answer: B

Q667. How can you verify if a given binary tree is a valid BST?

- A. Check if each node is greater than its children
- B. Perform in-order traversal and verify the result is strictly increasing
- C. Check if the root is the median value
- D. Count the number of nodes and compare with height

Answer: B

Q668. Why is bottom-up heap construction $O(n)$ rather than $O(n \log n)$?

- A. Because it skips half the elements
- B. Because most nodes are near the leaves and require very few sift-down operations, forming a convergent series
- C. Because it uses insertion sort internally
- D. Because it only processes the root node

Answer: B

Q669. What is a B-tree primarily used for?

- A. In-memory sorting
- B. Maintaining sorted data in database and file systems with minimal disk reads
- C. Network routing
- D. String pattern matching

Answer: B

Q670. What is the LCA (Lowest Common Ancestor) of two nodes in a binary tree?

- A. The root of the tree
- B. The deepest node that is an ancestor of both given nodes
- C. The parent of the deeper node
- D. The node at the midpoint between the two nodes

Answer: B

Q671. What is the time complexity of Dijkstra's algorithm using a Fibonacci heap?

- A. $O(V^2)$
- B. $O(V \log V + E)$
- C. $O(E \log V)$
- D. $O(V * E)$

Answer: B

Q672. How does Kruskal's algorithm find a minimum spanning tree?

- A. By starting from a vertex and adding the cheapest edge to a neighbor
- B. By sorting all edges by weight and adding them if they don't create a cycle, using Union-Find
- C. By finding the shortest path from source to all vertices
- D. By removing the most expensive edges

Answer: B

Q673. What is topological sorting and on which type of graph is it performed?

- A. Sorting vertices by degree on any graph
- B. A linear ordering of vertices such that for every directed edge (u, v) , u comes before v ; only on DAGs
- C. Sorting edges by weight in an undirected graph
- D. Ordering vertices alphabetically

Answer: B

Q674. How does the Bellman-Ford algorithm differ from Dijkstra's algorithm?

- A. Bellman-Ford is always faster
- B. Bellman-Ford can handle negative edge weights and detect negative cycles, while Dijkstra cannot
- C. Dijkstra handles negative weights but Bellman-Ford does not
- D. They produce different results on the same graph

Answer: B

Q675. What is the time complexity of Prim's algorithm using a binary heap?

- A. $O(V^2)$
- B. $O(E \log V)$
- C. $O(V * E)$
- D. $O(V + E)$

Answer: B

Q676. How can you detect a cycle in an undirected graph using DFS?

- A. Check if any vertex has degree > 2
- B. If DFS visits an already-visited vertex that is not the parent of the current vertex, a cycle exists
- C. Count the number of edges; a cycle exists if $E > V$
- D. Use topological sort

Answer: B

Q677. What does the Union-Find path compression optimization do?

- A. It compresses the data stored at each node
- B. It flattens the tree by making every node point directly to the root during find operations
- C. It removes unused paths from the graph
- D. It compresses the graph into fewer vertices

Answer: B

Q678. What is the Floyd-Warshall algorithm used for?

- A. Finding the minimum spanning tree
- B. Finding shortest paths between all pairs of vertices
- C. Topological sorting
- D. Detecting bipartite graphs

Answer: B

Q679. What is the time complexity of finding connected components in an undirected graph?

- A. $O(V^2)$
- B. $O(V + E)$
- C. $O(V * E)$
- D. $O(E^2)$

Answer: B

Q680. What is a back edge in DFS and what does it indicate?

- A. An edge going from a node to its parent
- B. An edge from a node to an ancestor in the DFS tree, indicating a cycle in directed graphs
- C. An edge that was removed during DFS
- D. An edge connecting two different components

Answer: B

Q681. What is the time complexity of the dynamic programming solution for the Longest Common Subsequence?

- A. $O(n)$
- B. $O(n * m)$ where n and m are the lengths of the two sequences
- C. $O(n^3)$
- D. $O(2^n)$

Answer: B

Q682. Why does the greedy approach fail for the 0/1 Knapsack problem but work for the fractional Knapsack?

- A. The greedy approach is always incorrect
- B. In 0/1 Knapsack, items cannot be divided, so a locally optimal choice may block a better global solution
- C. The fractional Knapsack has fewer items
- D. Both can be solved greedily

Answer: B

Q683. How does the N-Queens backtracking solution work?

- A. It places all queens randomly and checks if the placement is valid
- B. It places queens one row at a time, backtracking when a placement conflicts with existing queens
- C. It sorts the board positions and assigns queens greedily
- D. It uses dynamic programming to find all placements

Answer: B

Q684. What is the Longest Increasing Subsequence (LIS) problem's optimal time complexity?

- A. $O(n^2)$
- B. $O(n \log n)$
- C. $O(n)$
- D. $O(2^n)$

Answer: B

Q685. What is Huffman coding and which algorithm design paradigm does it use?

- A. A sorting algorithm using divide and conquer
- B. An optimal prefix-free encoding algorithm using the greedy approach
- C. A string matching algorithm using dynamic programming
- D. A graph algorithm using backtracking

Answer: B

Q686. What is the difference between the 0/1 Knapsack and the fractional Knapsack problem?

- A. 0/1 Knapsack allows fractions while fractional does not
- B. In 0/1 Knapsack items are taken whole or not at all; in fractional Knapsack items can be divided
- C. They are the same problem
- D. Fractional Knapsack is always harder

Answer: B

Q687. What is the rod cutting problem in dynamic programming?

- A. Cutting a graph into components
- B. Finding the maximum revenue by cutting a rod of length n into pieces with given prices for each length
- C. Dividing an array into equal parts
- D. Cutting a string at specific positions

Answer: B

Q688. What is branch and bound and how does it differ from backtracking?

- A. They are the same algorithm
- B. Branch and bound uses bounds (upper/lower) to prune branches that cannot yield better solutions than the current best
- C. Branch and bound only works on trees
- D. Backtracking uses bounds while branch and bound does not

Answer: B

Q689. What is the space optimization technique for the 0/1 Knapsack DP solution?

- A. Use a 3D array instead of 2D
- B. Use only two 1D arrays (or one array traversed in reverse) since each row only depends on the previous row
- C. Use a hash table
- D. Use a linked list

Answer: B

Q690. What is the activity selection problem and which approach solves it optimally?

- A. Selecting the minimum number of activities; solved by DP
- B. Selecting the maximum number of non-overlapping activities; solved by a greedy algorithm that sorts by finish time
- C. Selecting activities with maximum total weight; solved by brute force
- D. Assigning activities to workers; solved by backtracking

Answer: B

Q691. What is the amortized time complexity of Union-Find operations with both path compression and union by rank?

- A. $O(1)$
- B. $O(\log n)$
- C. $O(\alpha(n))$ where α is the inverse Ackermann function
- D. $O(n)$

Answer: C

Q692. What is a Fibonacci heap's key advantage over a binary heap?

- A. Simpler implementation
- B. $O(1)$ amortized decrease-key operation instead of $O(\log n)$
- C. Less memory usage
- D. Faster extract-min

Answer: B

Q693. What is a suffix array and what problem does it help solve?

- A. An array of suffix sums for range queries
- B. A sorted array of all suffixes of a string, enabling efficient pattern matching and string analysis
- C. An array storing the suffixes of Fibonacci numbers
- D. A compressed version of a string

Answer: B

Q694. What is a sparse table used for?

- A. Storing sparse matrices
- B. Answering static range minimum/maximum queries in $O(1)$ after $O(n \log n)$ preprocessing
- C. Building sparse graphs
- D. Compressing data

Answer: B

Q695. How does a k-d tree determine which dimension to split on at each level?

- A. It always splits on the first dimension
- B. It cycles through the dimensions at each level (depth mod k)
- C. It randomly selects a dimension
- D. It splits on the dimension with the most variance

Answer: B

Q696. What is the time complexity of merge operation in a binomial heap?

- A. $O(n)$
- B. $O(\log n)$
- C. $O(1)$
- D. $O(n \log n)$

Answer: B

Q697. What is an interval tree used for?

- A. Storing time intervals for scheduling
- B. Finding all intervals that overlap with a given query interval efficiently
- C. Measuring time intervals
- D. Sorting intervals by length

Answer: B

Q698. What is the space complexity of a trie storing n strings with average length m over an alphabet of size $< 2^b$

- A. $O(n)$
- B. $O(n * m * < 2^b)$
- C. $O(n * m)$
- D. $O(< 2^b)$

Answer: B

Q699. What is a balanced BST's advantage over a hash table for ordered operations?

- A. Balanced BSTs have $O(1)$ lookup
- B. Balanced BSTs support ordered operations like range queries, in-order traversal, and predecessor/successor in $O(\log n)$
- C. Hash tables cannot store strings
- D. Balanced BSTs use less memory

Answer: B

Q700. What is a circular buffer and where is it commonly used?

- A. A buffer that stores data in a circle shape on disk
- B. A fixed-size buffer that wraps around, commonly used in streaming, producer-consumer patterns, and I/O buffering
- C. A buffer that can grow infinitely
- D. A buffer used only for circular linked lists

Answer: B

Q701. What is the failure function (partial match table) in the KMP algorithm?

- A. A table indicating which characters to skip in the text
- B. A table storing the length of the longest proper prefix that is also a suffix for each prefix of the pattern
- C. A table of hash values for the pattern
- D. A table mapping characters to their positions

Answer: B

Q702. How does the Rabin-Karp algorithm handle spurious matches?

- A. It ignores them completely
- B. It performs a character-by-character comparison when hash values match to verify the actual match
- C. It uses a different hash function each time
- D. It rehashes the entire text

Answer: B

Q703. What is the Z-array for a string S?

- A. An array of zeros
- B. $Z[i]$ = length of the longest substring starting at i that matches a prefix of S
- C. An array sorting characters alphabetically
- D. The reverse of the string stored as an array

Answer: B

Q704. What is the key idea behind the Boyer-Moore algorithm?

- A. Searching from left to right in the pattern
- B. Comparing the pattern from right to left and using bad character and good suffix heuristics to skip positions
- C. Using hash values for comparison
- D. Building a trie of all substrings

Answer: B

Q705. What is the time complexity of computing the longest palindromic substring using dynamic programming?

- A. $O(n)$
- B. $O(n \log n)$
- C. $O(n^2)$
- D. $O(n^3)$

Answer: C

Q706. What is a rolling hash and how does it enable efficient string matching?

- A. A hash that rolls dice for randomness
- B. A hash function that can compute the hash of the next window in $O(1)$ by removing the leading character and adding the trailing character
- C. A hash that is computed once and never changes
- D. A hash function specific to number sequences

Answer: B

Q707. How does the KMP algorithm achieve $O(n + m)$ time complexity?

- A. By sorting the text first
- B. By never re-comparing characters that have already been matched, using the failure function to resume matching
- C. By using hash values instead of character comparison
- D. By dividing the text into blocks

Answer: B

Q708. What is the time complexity of checking if one string is a rotation of another?

- A. $O(n^2)$
- B. $O(n)$ by concatenating the string with itself and searching for the other string
- C. $O(n \log n)$
- D. $O(2^n)$

Answer: B

Q709. In polynomial string hashing, what values are typically chosen for the base p and modulus M ?

- A. $p = 1$ and $M = 10$
- B. p greater than the alphabet size and M a large prime to minimize collisions
- C. Both should be powers of 2
- D. $p = 0$ and $M =$ the string length

Answer: B

Q710. What is the best-case time complexity of the Boyer-Moore algorithm for pattern matching?

- A. $O(n * m)$
- B. $O(n + m)$
- C. $O(n / m)$
- D. $O(n)$

Answer: C

Q711. What is the relationship between the complexity classes P and NP ?

- A. P is a strict subset of NP and $P = NP$ is proven
- B. $P \neq NP$ (P is contained in NP), but whether $P = NP$ is unknown
- C. $NP \subset P$ (NP is a strict subset of P)
- D. P and NP are disjoint sets

Answer: B

Q712. What is the time complexity of Strassen's matrix multiplication algorithm?

- A. $O(n^3)$
- B. $O(n^2)$
- C. $O(n^{(\log_2 7)})$ or $O(n^{2.807})$
- D. $O(n^2 \log n)$

Answer: C

Q713. Which of the following is a well-known NP-complete problem?

- A. Binary search
- B. Sorting an array
- C. The Boolean Satisfiability Problem (SAT)
- D. Finding the minimum in an array

Answer: C

Q714. What is the difference between time complexity and space complexity?

- A. They always have the same value
- B. Time complexity measures the number of operations as a function of input size; space complexity measures the memory used
- C. Space complexity is always larger
- D. Time complexity measures clock time in seconds

Answer: B

Q715. What is the significance of the $O(n \log n)$ barrier for comparison-based sorting?

- A. It can be beaten by any sorting algorithm
- B. No comparison-based sorting algorithm can have a worst case better than $O(n \log n)$
- C. $O(n \log n)$ is the average case but $O(n)$ is achievable in the worst case
- D. Only merge sort achieves this bound

Answer: B

Q716. How does loop unrolling improve performance at the hardware level?

- A. By reducing the total number of operations
- B. By reducing branch overhead and enabling better instruction pipelining and cache utilization
- C. By converting the loop to recursion
- D. By eliminating the need for variables

Answer: B

Q717. What is the time-space tradeoff in memoization?

- A. Memoization uses more time to save space
- B. Memoization uses extra space to store computed results, reducing time by avoiding redundant computation
- C. Memoization has no tradeoff
- D. Memoization reduces both time and space

Answer: B

Q718. What is the complexity of adding two n-digit numbers?

- A. $O(1)$
- B. $O(n)$
- C. $O(n^2)$
- D. $O(\log n)$

Answer: B

Q719. What is cache-oblivious algorithm design?

- A. Algorithms that don't use any cache
- B. Algorithms designed to efficiently use the memory hierarchy without knowing cache parameters
- C. Algorithms that ignore memory constraints
- D. Algorithms that clear the cache before running

Answer: B

Q720. What is the difference between worst-case and average-case complexity?

- A. They are always the same
- B. Worst-case gives the maximum time over all inputs of size n ; average-case gives the expected time over a distribution of inputs
- C. Average-case is always worse than worst-case
- D. Worst-case only considers the best input

Answer: B

Data Structures & Algorithms — Hard

360 questions

Q721. What is the time complexity of the recurrence $T(n) = 3T(n/4) + n \log n$?

- A. $O(n^{(\log_4(3))})$
- B. $O(n^2)$
- C. $O(n \log^2 n)$
- D. $O(n \log n)$

Answer: D

Q722. Which of the following recurrences cannot be solved by the Master Theorem?

- A. $T(n) = T(n-1) + n$
- B. $T(n) = 2T(n/2) + n \log n$
- C. $T(n) = 2T(n/2) + n$
- D. $T(n) = 4T(n/2) + n^2$

Answer: A

Q723. What is the amortized cost of insertion in a dynamic array that doubles its size when full?

- A. $O(1)$
- B. $O(n^2)$
- C. $O(\log n)$
- D. $O(n)$

Answer: A

Q724. Using the substitution method, what is the solution to $T(n) = T(n-1) + n$ with $T(1) = 1$?

- A. $O(2^n)$
- B. $O(n \log n)$
- C. $O(n)$
- D. $O(n^2)$

Answer: D

Q725. What is the lower bound for any comparison-based sorting algorithm?

- A. $O(n \log n)$
- B. $O(n^2)$
- C. $O(\log n)$
- D. $O(n)$

Answer: A

Q726. In the potential method of amortized analysis, the amortized cost is defined as:

- A. Actual cost $\times$ potential function
- B. Actual cost + change in potential
- C. Actual cost - change in potential
- D. Actual cost / potential function

Answer: B

Q727. What complexity class does the problem of finding the median of an unsorted array belong to?

- A. $O(n^2)$ minimum brute force
- B. $O(n \log n)$ using sort first
- C. $O(n)$ using deterministic selection
- D. $O(\log n)$ using binary search

Answer: C

Q728. What is the space complexity of a recursive algorithm with depth d and $O(1)$ space per call?

- A. $O(d)$
- B. $O(2^d)$
- C. $O(d^2)$
- D. $O(1)$

Answer: A

Q729. Which of the following is true about the relationship between time and space complexity?

- A. Space complexity can exceed time complexity
- B. Space complexity cannot exceed time complexity
- C. Time complexity is always greater than space
- D. Time and space complexity are always equal

Answer: B

Q730. What is the time complexity of $T(n) = 2T(n/2) + n/\log n$?

- A. $O(n)$
- B. $O(n \log n)$
- C. $O(n \log \log n)$
- D. $O(n (\log n)^2)$

Answer: C

Q731. What is the amortized time complexity of enqueue and dequeue operations when implementing a queue with two stacks?

- A. $O(1)$ amortized for both
- B. $O(1)$ enqueue, $O(n)$ dequeue
- C. $O(n)$ enqueue, $O(1)$ dequeue
- D. $O(n)$ for both operations

Answer: A

Q732. How can you detect a cycle in a linked list efficiently?

- A. By sorting the list nodes first in order
- B. By converting the list into an array form
- C. Using two nested loops over all nodes
- D. Using Floyd's tortoise and hare algorithm

Answer: D

Q733. What is the time complexity of reversing a singly linked list iteratively?

- A. $O(1)$
- B. $O(n^2)$
- C. $O(n)$
- D. $O(n \log n)$

Answer: C

Q734. In a skip list, what is the expected time complexity of search?

- A. $O(n)$
- B. $O(1)$
- C. $O(\log n)$
- D. $O(n \log n)$

Answer: C

Q735. What is an XOR linked list?

- A. A circular linked list variant with XOR pointers
- B. A linked list that stores only odd-valued elements
- C. A linked list using XOR for data encryption purposes
- D. A doubly linked list using XOR of addresses to save memory

Answer: D

Q736. How do you find the starting node of a cycle in a linked list using Floyd's algorithm?

- A. Use a hash table to store all visited node addresses
- B. Count all nodes in the list and compute the start
- C. Reset one pointer to head and move both one step
- D. Reverse the entire list and then detect the cycle

Answer: C

Q737. What is the minimum number of stacks needed to implement a deque?

- A. 1
- B. 3
- C. 2
- D. 4

Answer: C

Q738. What is the space overhead per node in a doubly linked list compared to a singly linked list?

- A. One extra data field
- B. No extra overhead
- C. Two extra pointers
- D. One extra pointer

Answer: D

Q739. What is the time complexity of merging two sorted linked lists into one sorted list?

- A. $O(n \times m)$
- B. $O(n + m)$
- C. $O(n \log m)$
- D. $O(n^2)$

Answer: B

Q740. What is a self-organizing list?

- A. A list that reorders based on access patterns
- B. A list that balances itself like a tree structure
- C. A list that maintains sorted order automatically
- D. A list that automatically deletes duplicate items

Answer: A

Q741. What is the time complexity of building a heap from an unsorted array of n elements?

- A. $O(n \log n)$
- B. $O(n)$
- C. $O(\log n)$
- D. $O(n^2)$

Answer: B

Q742. What is a threaded binary tree?

- A. A tree where null pointers become in-order links
- B. A binary tree with multiple root nodes
- C. A tree with thread-safe atomic operations
- D. A tree used in multithreading concurrency

Answer: A

Q743. What is the minimum number of nodes in an AVL tree of height h ?

- A. $N(h) = N(h-1) + N(h-2) + 1$
- B. $2^{(h+1)} - 1$
- C. 2^h
- D. $h + 1$

Answer: A

Q744. In a directed graph with V vertices, what is the maximum number of edges?

- A. V^2
- B. V
- C. $V(V-1)$
- D. $V(V-1)/2$

Answer: C

Q745. What is a B-tree of order m ?

- A. A tree data structure with m distinct roots
- B. A self-balancing tree with at most m children
- C. A binary tree with exactly m total levels
- D. A graph with m connected components total

Answer: B

Q746. What is a Red-Black tree's maximum height for n nodes?

- A. $n/2$
- B. n
- C. $\log n$
- D. $2 \log(n+1)$

Answer: D

Q747. What is a bipartite graph?

- A. A graph divisible into two disjoint vertex sets
- B. A fully connected complete graph structure
- C. A graph with exactly two vertices in total
- D. A graph with exactly two connected components

Answer: A

Q748. What is the time complexity of deleting the root from a max-heap of n elements?

- A. $O(n \log n)$
- B. $O(n)$
- C. $O(1)$
- D. $O(\log n)$

Answer: D

Q749. How many distinct BSTs can be formed with n distinct keys?

- A. 2^n (powers of two)
- B. n^2 (square of n)
- C. $C(n)$ the Catalan number
- D. $n!$ (factorial of n)

Answer: C

Q750. What is the difference between a strongly connected and weakly connected directed graph?

- A. Strongly connected means directed paths exist between all pairs; weakly only without directions
- B. There is no meaningful difference between strongly and weakly connected directed graphs
- C. Strongly connected graphs always contain more edges than weakly connected directed graphs
- D. Weakly connected graphs always have more vertices and edges than strongly connected ones

Answer: A

Q751. What is the worst-case time complexity of interpolation search?

- A. $O(n)$
- B. $O(\log n)$
- C. $O(\log \log n)$
- D. $O(\sqrt{n})$

Answer: A

Q752. What is Fibonacci search?

- A. A search on Fibonacci heap data structures
- B. A recursive search with Fibonacci-like recurrence
- C. Searching for Fibonacci numbers in an array
- D. A search using Fibonacci numbers to divide array

Answer: D

Q753. What is the time complexity of searching in a balanced binary search tree?

- A. $O(n \log n)$
- B. $O(n)$
- C. $O(\log n)$
- D. $O(1)$

Answer: C

Q754. How does binary search handle finding the first occurrence of a duplicate element in a sorted array?

- A. Standard binary search always finds the first occurrence automatically
- B. Binary search cannot handle duplicate elements in sorted arrays
- C. By using linear search after binary search locates any occurrence
- D. By modifying binary search to continue searching left after finding it

Answer: D

Q755. What is the order-statistic problem?

- A. Counting elements in a range
- B. Finding the median of sorted data
- C. Finding the k-th smallest element
- D. Sorting all elements into order

Answer: C

Q756. What is the time complexity of the Quickselect algorithm in the average case?

- A. $O(n \log n)$
- B. $O(n^2)$
- C. $O(\log n)$
- D. $O(n)$

Answer: D

Q757. How can you search in a sorted and rotated array?

- A. Interpolation search handles rotation natively
- B. Only linear search works for this case
- C. Modified binary search identifies the sorted half
- D. Ternary search is the only viable approach

Answer: C

Q758. What is the time complexity of searching in a skip list?

- A. $O(\log n)$ expected
- B. $O(1)$ constant
- C. $O(n)$ linear
- D. $O(n \log n)$ total

Answer: A

Q759. What is two-pointer technique in searching?

- A. Using two separate arrays for searching data
- B. Searching for two distinct elements at once
- C. Binary search implemented with two mid points
- D. Using two pointers from different positions to converge

Answer: D

Q760. What is the time complexity of finding the median of two sorted arrays of sizes m and n ?

- A. $O((m + n) \log(m + n))$
- B. $O(m \times n)$
- C. $O(m + n)$
- D. $O(\log(\min(m, n)))$

Answer: D

Q761. What is radix sort's time complexity for n integers with d digits?

- A. $O(n \times d^2)$
- B. $O(d \times (n + k))$ where k is the base
- C. $O(n^2)$
- D. $O(n \log n)$

Answer: B

Q762. What is the worst-case time complexity of randomized quick sort?

- A. $O(n)$
- B. $O(n^2)$
- C. $O(n \log^2 n)$
- D. $O(n \log n)$

Answer: B

Q763. What is bucket sort?

- A. A sort that only works on string data
- B. A sort distributing elements into buckets then merging
- C. A merge sort variant with fixed partitions
- D. A comparison sort using memory buckets

Answer: B

Q764. When is bucket sort most efficient?

- A. When data is in reverse sorted order
- B. When the array is extremely large in size
- C. When input is uniformly distributed over a range
- D. When data has many duplicate values present

Answer: C

Q765. What is introsort?

- A. A merge sort variant with optimizations
- B. A pure insertion sort variant
- C. A hybrid starting with quicksort switching to heapsort
- D. A non-comparison sort using radix method

Answer: C

Q766. What is TimSort?

- A. A real-time sorting algorithm for streaming data
- B. A sort that runs in constant time for all input
- C. A sort named after its time complexity class
- D. A hybrid stable sort using merge and insertion sort

Answer: D

Q767. Is quick sort a stable sorting algorithm?

- A. Only with Lomuto partition scheme
- B. Only when using random pivots
- C. Yes, always stable
- D. No, in its standard implementation

Answer: D

Q768. What is the time complexity of sorting n strings of maximum length s using radix sort?

- A. $O(s \log n)$
- B. $O(n^2 \times s)$
- C. $O(n \times s)$
- D. $O(n \log n)$

Answer: C

Q769. What is the Dutch National Flag problem in the context of sorting?

- A. Coloring a graph with three distinct colors
- B. Sorting elements of three different colors
- C. Sorting using three different algorithms
- D. Three-way partitioning around a pivot value

Answer: D

Q770. What is external sorting?

- A. Sorting using network-based remote requests
- B. Sorting data too large for memory using disk
- C. Sorting using external third-party libraries
- D. Sorting on external hardware devices only

Answer: B

Q771. What is universal hashing?

- A. Hashing all elements to the same single bucket
- B. Using the same hash function across all applications
- C. Randomly selecting a hash function from a family of functions
- D. A single hash function that works for all data types

Answer: C

Q772. What is perfect hashing?

- A. A hash function that always distributes keys evenly
- B. A hash function with no collisions for any input set
- C. A hashing technique that requires no table at all
- D. A two-level scheme guaranteeing $O(1)$ worst-case lookup

Answer: D

Q773. What is a Bloom filter?

- A. A type of hash table with chaining
- B. A comparison-based sorting algorithm
- C. A filter for removing duplicate elements
- D. A probabilistic structure testing set membership

Answer: D

Q774. What is cuckoo hashing?

- A. A technique using random hash function selection
- B. A technique using two hash functions with displacement
- C. A hashing technique inspired by bird nesting
- D. A type of perfect hashing for static keys

Answer: B

Q775. What is consistent hashing?

- A. Using the same hash function consistently always
- B. A collision-free hashing technique for all inputs
- C. A hashing that always produces the same output
- D. A scheme remapping only a fraction of keys on change

Answer: D

Q776. What is the expected maximum chain length in a hash table with n elements and n slots using chaining?

- A. $O(\log n / \log \log n)$
- B. $O(n)$
- C. $O(1)$
- D. $O(n)$

Answer: A

Q777. What is Robin Hood hashing?

- A. A two-level hashing scheme with perfect balance
- B. A charitable hashing distribution across buckets
- C. Stealing hash values from other tables entirely
- D. Open addressing where longer probes displace shorter ones

Answer: D

Q778. What is the purpose of a cryptographic hash function?

- A. To speed up hash table lookups efficiently
- B. To compress files to smaller size format
- C. To produce collision-resistant irreversible fixed output
- D. To sort data in a deterministic order

Answer: C

Q779. What is hopscotch hashing?

- A. An open addressing scheme with bounded neighborhoods
- B. A random displacement hashing for load balance
- C. A hashing game-based approach to distribution
- D. A chaining variant using skip list buckets

Answer: A

Q780. What is the false positive probability of a Bloom filter with m bits, k hash functions, and n inserted elements?

- A. $(1 - e^{-(kn/m)})^k$ approx
- B. n / m ratio
- C. $1 / 2^k$ value
- D. k / m ratio

Answer: A

Q781. What is a splay tree?

- A. A B-tree variant used in database indexing
- B. A self-adjusting BST that moves accessed nodes to root
- C. A tree that spreads elements across multiple arrays
- D. A tree that is always perfectly balanced

Answer: B

Q782. What is the time complexity of range query and update in a segment tree?

- A. $O(n)$ for both operations
- B. $O(n)$ query, $O(1)$ update
- C. $O(1)$ query, $O(n)$ update
- D. $O(\log n)$ for both operations

Answer: D

Q783. What is a Fenwick tree (Binary Indexed Tree)?

- A. A balanced BST variant with extra pointers
- B. A structure for $O(\log n)$ prefix sums and updates
- C. A tree used in networking protocols
- D. A binary tree indexed by key values

Answer: B

Q784. What is lazy propagation in a segment tree?

- A. Deferring updates to children until they are queried
- B. Lazily inserting elements one at a time
- C. Skipping unnecessary queries for performance
- D. Delaying the tree construction until needed

Answer: A

Q785. What is a treap?

- A. A BST-heap hybrid with keys and random priorities
- B. A tree-shaped trap data structure
- C. A triple-ended queue data structure
- D. A tree with array-like access properties

Answer: A

Q786. What is the Euler tour technique for trees?

- A. Linearizing a tree by recording DFS visit times
- B. A graph coloring technique for tree nodes
- C. Touring all leaf nodes in level order
- D. Finding Euler paths in tree structures

Answer: A

Q787. What is Heavy-Light Decomposition (HLD) of a tree?

- A. Decomposing into chains for efficient path queries
- B. Separating heavy and light weight nodes
- C. Removing heavy edges from the tree structure
- D. Balancing an unbalanced tree by restructuring

Answer: A

Q788. What is the time complexity of LCA using binary lifting (sparse table)?

- A. $O(n^2)$ per query with no preprocessing needed
- B. $O(n)$ per query with no preprocessing
- C. $O(1)$ per query after $O(n^2)$ preprocessing
- D. $O(\log n)$ per query after $O(n \log n)$ preprocessing

Answer: D

Q789. What is a persistent segment tree?

- A. A segment tree preserving all versions via path copying
- B. A segment tree that never changes after construction
- C. A segment tree stored on disk permanently
- D. A compressed segment tree with reduced memory use

Answer: A

Q790. How does a B+ tree differ from a B-tree?

- A. B+ tree is not a balanced tree structure
- B. B+ tree has more levels than a B-tree
- C. B+ tree stores all data at leaves with linked leaf nodes
- D. B+ tree has fewer keys per node overall

Answer: C

Q791. What is the Floyd-Warshall algorithm?

- A. An all-pairs shortest path DP algorithm
- B. A minimum spanning tree algorithm only
- C. A single-source shortest path algorithm
- D. A graph coloring assignment algorithm

Answer: A

Q792. What is the time complexity of Floyd-Warshall algorithm?

- A. $O(V^3)$
- B. $O(V^2)$
- C. $O(V \times E)$
- D. $O(V + E)$

Answer: A

Q793. What are Strongly Connected Components (SCCs) in a directed graph?

- A. Maximal subgraphs with mutual reachability
- B. Components with no cycles at all
- C. Components with the most vertices total
- D. Components with the most edges overall

Answer: A

Q794. Which algorithms can find Strongly Connected Components?

- A. Dijkstra's shortest path algorithm
- B. Prim's spanning tree algorithm
- C. Kosaraju's and Tarjan's algorithms
- D. Floyd-Warshall all-pairs algorithm

Answer: C

Q795. What is an articulation point (cut vertex) in a graph?

- A. A vertex with no edges connected to it
- B. The vertex with the smallest key value
- C. A vertex whose removal disconnects the graph
- D. A vertex with the maximum degree overall

Answer: C

Q796. What is a bridge in a graph?

- A. An edge with the maximum weight assigned
- B. An edge whose removal disconnects the graph
- C. The longest edge in the graph by weight
- D. An edge connecting two separate components

Answer: B

Q797. What is the maximum flow problem?

- A. Finding the maximum number of vertices
- B. Finding maximum flow from source to sink
- C. Finding the maximum edge weight in graph
- D. Sorting edges by their flow capacity

Answer: B

Q798. What is the Ford-Fulkerson method for maximum flow?

- A. An iterative method finding augmenting paths until none remain
- B. A divide and conquer approach for solving network flows
- C. A greedy algorithm for computing the minimum cut value
- D. A dynamic programming approach for computing max flow

Answer: A

Q799. What does the Max-Flow Min-Cut theorem state?

- A. Maximum flow equals the number of all paths
- B. Maximum flow equals minimum edge weight
- C. Maximum flow equals the minimum cut capacity
- D. Minimum cut equals the total number of edges

Answer: C

Q800. What is the time complexity of Dijkstra's algorithm with a Fibonacci heap?

- A. $O(V \times E)$
- B. $O((V + E) \log V)$
- C. $O(V \log V + E)$
- D. $O(V^2)$

Answer: C

Q801. What is the matrix chain multiplication problem's time complexity using DP?

- A. $O(n^3)$
- B. $O(n!)$
- C. $O(2^n)$
- D. $O(n^2)$

Answer: A

Q802. What is the Traveling Salesman Problem (TSP)?

- A. Finding the shortest path between two specific cities
- B. Finding the fastest route avoiding all toll roads
- C. Finding minimum cost Hamiltonian cycle visiting all cities
- D. A minimum spanning tree problem for city networks

Answer: C

Q803. What is the time complexity of solving TSP using dynamic programming (Held-Karp)?

- A. $O(n^3)$
- B. $O(n^2 \times 2^n)$
- C. $O(n!)$
- D. $O(2^n)$

Answer: B

Q804. What is branch and bound?

- A. A hashing technique for collision resolution
- B. A type of comparison-based sorting method
- C. An optimization technique pruning suboptimal branches
- D. A graph traversal method for shortest paths

Answer: C

Q805. What is the edit distance (Levenshtein distance) problem?

- A. The minimum single-character edits between two strings
- B. The difference in the lengths of two strings
- C. The number of common characters in two strings
- D. The distance between two arrays in memory

Answer: A

Q806. What is the difference between the 0/1 knapsack and fractional knapsack?

- A. 0/1 uses greedy approach; fractional uses DP approach
- B. Fractional knapsack is computationally harder overall
- C. They are exactly the same problem formulation
- D. 0/1 requires whole items (DP); fractional allows fractions (greedy)

Answer: D

Q807. What is a randomized algorithm?

- A. An algorithm with random bugs in code
- B. An algorithm using random numbers in its decisions
- C. An algorithm that shuffles input before sorting
- D. An algorithm that produces random wrong output

Answer: B

Q808. What is an approximation algorithm?

- A. An algorithm that gives approximate time complexity
- B. An algorithm finding near-optimal solutions in polynomial time
- C. An imprecise algorithm with unreliable output
- D. An algorithm that estimates the input data size

Answer: B

Q809. What is the subset sum problem?

- A. Finding the sum of all array elements
- B. Summing consecutive elements in an array
- C. Finding the maximum valued subset overall
- D. Determining if a subset sums to a target value

Answer: D

Q810. What is the meet-in-the-middle technique?

- A. A two-pointer approach for sorted arrays
- B. A median-finding technique using partitioning
- C. Dividing the array in half for sorting
- D. Splitting input in half and combining to reduce complexity

Answer: D

Q811. What is a suffix tree?

- A. A tree that stores all string prefixes
- B. A compressed trie of all suffixes for $O(m)$ matching
- C. An array of suffix pointers with sorting
- D. A binary tree of sorted suffix indices

Answer: B

Q812. What is a van Emde Boas tree?

- A. A tree with $O(\log \log U)$ operations for integer keys
- B. A tree with variable branching factor per level
- C. A cache-oblivious tree for external memory
- D. A self-balancing BST with color properties

Answer: A

Q813. What is a rope data structure?

- A. A linked list variant for characters
- B. A hash map for string key-value pairs
- C. A graph connecting strings with edges
- D. A balanced binary tree for efficient string operations

Answer: D

Q814. What is a count-min sketch?

- A. A data structure for exact element counting
- B. A compressed hash table with min values
- C. A min-heap variant tracking element counts
- D. A probabilistic structure for approximate frequency counts

Answer: D

Q815. What is a link-cut tree?

- A. A tree where links can be physically cut
- B. A doubly linked tree with removable nodes
- C. A dynamic forest supporting link, cut, and path queries
- D. A tree with removable edges but fixed vertices

Answer: C

Q816. What is a wavelet tree?

- A. A tree for signal processing operations
- B. A frequency-based heap for priority ordering
- C. A Fourier transform data structure for arrays
- D. A tree for rank, select, and range queries on sequences

Answer: D

Q817. What is a persistent data structure?

- A. A data structure saved to disk permanently
- B. A data structure that cannot ever be deleted
- C. A data structure that persists across program runs
- D. A structure preserving all previous versions on modification

Answer: D

Q818. What is a binomial heap?

- A. A heap with binomial probability distribution
- B. A heap used in statistical analysis methods
- C. A binary heap with extra pointer fields added
- D. A collection of binomial trees with $O(\log n)$ merge

Answer: D

Q819. What is an order-statistic tree?

- A. An augmented BST supporting $O(\log n)$ rank and select
- B. A tree for statistical analysis of data sets
- C. A tree that orders elements statistically
- D. A heap that tracks insertion order of elements

Answer: A

Q820. What is a succinct data structure?

- A. A brief textual description of a data structure
- B. A structure using near information-theoretic minimum space
- C. A small fixed-capacity array with limited size
- D. A compressed hash table designed for very dense data

Answer: B

Q821. What is the Aho-Corasick algorithm?

- A. A multi-pattern matcher building a trie-based automaton
- B. A compression algorithm for pattern libraries
- C. A single pattern matching algorithm only
- D. A string sorting algorithm for multiple strings

Answer: A

Q822. What is Manacher's algorithm?

- A. A string sorting algorithm for arrays
- B. A pattern matching algorithm for text
- C. A string compression algorithm for storage
- D. An $O(n)$ algorithm finding all palindromic substrings

Answer: D

Q823. What is the time complexity of building a suffix array using the DC3/skew algorithm?

- A. $O(n \log^2 n)$
- B. $O(n \log n)$
- C. $O(n^2)$
- D. $O(n)$

Answer: D

Q824. What is the LCP (Longest Common Prefix) array?

- A. An array storing longest common prefix lengths between consecutive sorted suffixes
- B. An array storing character frequencies found within a given string
- C. An array storing the lengths of all strings within a collection
- D. An array storing prefix sums computed over numeric data elements

Answer: A

Q825. What is Burrows-Wheeler Transform (BWT)?

- A. A string matching algorithm for patterns
- B. A reversible transformation aiding text compression
- C. A string encryption method for security
- D. A hash function designed for string data

Answer: B

Q826. What is a suffix automaton?

- A. An automaton that generates all suffixes
- B. A pattern matching machine for any text
- C. A minimal DFA recognizing all suffixes of a string
- D. A suffix tree variant with compressed nodes

Answer: C

Q827. What is the time complexity of the Boyer-Moore algorithm in the best case?

- A. $O(n/m)$
- B. $O(n \times m)$
- C. $O(n)$
- D. $O(n + m)$

Answer: A

Q828. What is polynomial string hashing?

- A. Hashing using polynomial long division
- B. Computing hash as sum of characters times powers of a base
- C. Hashing that produces polynomial-length hash output
- D. Using polynomial time algorithms for hash computation

Answer: B

Q829. What is the longest repeated substring problem's optimal solution?

- A. $O(n^3)$ brute force approach
- B. $O(n^2)$ using dynamic programming
- C. $O(n \log^2 n)$ using rolling hashing
- D. $O(n)$ using suffix array with LCP array

Answer: D

Q830. What is the number of distinct substrings of a string of length n ?

- A. $n(n+1)/2$ minus $\text{sum}(\text{LCP})$
- B. 2^n (exponential count)
- C. n^2 (quadratic count)
- D. $n!$ (factorial of n)

Answer: A

Q831. What is the complexity class PSPACE?

- A. Problems solvable using polynomial space amount
- B. Problems requiring exponential space to solve
- C. Problems solvable in constant space only
- D. Problems with no space requirement at all

Answer: A

Q832. What is Cook's theorem?

- A. P equals NP has been conclusively proven
- B. Every NP problem is fundamentally unsolvable
- C. Every problem is solvable in polynomial time
- D. The Boolean satisfiability problem (SAT) is NP-complete

Answer: D

Q833. What is the difference between a decision problem and an optimization problem?

- A. Decision has yes/no answer; optimization seeks best solution
- B. They are exactly the same type of problem
- C. Optimization problems are always solvable in P
- D. Decision problems are always computationally harder

Answer: A

Q834. What is pseudo-polynomial time?

- A. An algorithm that pretends to be efficient
- B. Time polynomial in input value rather than input size
- C. False polynomial time that does not exist
- D. The same as polynomial time in all cases

Answer: B

Q835. What is the 3-SAT problem?

- A. A Boolean satisfiability problem with 3 literals per clause
- B. A problem with exactly 3 valid solutions
- C. A sorting problem with 3 distinct elements
- D. Satisfying exactly 3 linear equations

Answer: A

Q836. What is the vertex cover problem?

- A. Coloring all vertices with minimum colors
- B. Finding all connected vertices in a graph
- C. Finding smallest vertex set covering all edges
- D. Covering all vertices with color assignments

Answer: C

Q837. What is the approximation ratio of the 2-approximation algorithm for vertex cover?

- A. The solution is at most twice the optimal
- B. The solution is at most 1.5 times optimal
- C. The solution is at most 3 times the optimal
- D. The solution always equals the optimal exactly

Answer: A

Q838. What is the class co-NP?

- A. The class whose complement problems are in NP
- B. The complement of NP problems
- C. Problems that are harder than NP class
- D. Problems that are not in NP at all

Answer: A

Q839. What is a parameterized algorithm?

- A. An algorithm that dynamically changes its parameters at runtime
- B. An algorithm whose complexity depends on input size and a fixed parameter
- C. A tunable sorting algorithm that uses an adjustable pivot element
- D. An algorithm that has many user-tunable parameter values

Answer: B

Q840. What is the time hierarchy theorem?

- A. All problems can be solved in the same time
- B. Time complexity is always fixed and unchanging
- C. More time allows a Turing machine to solve more problems
- D. There is no hierarchy in time complexity classes

Answer: C

Q841. What is the time complexity of $T(n) = 3T(n/4) + n \log n$ by Master Theorem?

- A. $O(n \log^2 n)$
- B. $O(n^2 \log n)$
- C. $O(n \log n)$
- D. $O(n^{\log_4 3})$

Answer: C

Q842. Which method handles recurrences not solvable by the Master Theorem?

- A. Simple iteration counting method
- B. Akra-Bazzi generalized method
- C. Direct formula application method
- D. Basic substitution guess method

Answer: B

Q843. What is the complexity of $T(n) = 2T(n/2) + n \log n$?

- A. $O(n^2 \log n)$
- B. $O(n \log n)$
- C. $O(n \log^2 n)$
- D. $O(n^2)$

Answer: C

Q844. In amortized analysis, what is the potential method?

- A. Assigning actual cost to each operation directly
- B. Assigning a potential energy function to data structure
- C. Counting total operations divided by total time
- D. Measuring memory usage across all operations run

Answer: B

Q845. What is the significance of the P vs NP problem?

- A. It asks if space complexity equals time complexity always
- B. It asks if all algorithms can run in constant time always
- C. It asks if polynomial verification implies polynomial solving
- D. It asks if every problem has a brute force solution path

Answer: C

Q846. What does the recursion tree method visualize?

- A. The sorted order of elements after each comparison
- B. The memory allocation pattern of the algorithm run
- C. The branching cost at each level of the recurrence
- D. The input-output mapping of the algorithm function

Answer: C

Q847. What is the time complexity of $T(n) = T(n/3) + T(2n/3) + n$?

- A. $O(n)$
- B. $O(n \log^2 n)$
- C. $O(n^2)$
- D. $O(n \log n)$

Answer: D

Q848. When does the Master Theorem not apply directly?

- A. When b is greater than a in the recurrence
- B. When the recurrence has exactly two subproblems
- C. When subproblems have different unequal sizes
- D. When $f(n)$ is a simple polynomial function

Answer: C

Q849. What is the aggregate method in amortized analysis?

- A. Using a potential function to track stored credits
- B. Assigning different costs to different operation types
- C. Counting only the most expensive operation each time
- D. Summing total cost of n operations then dividing by n

Answer: D

Q850. What is the difference between deterministic and nondeterministic algorithms?

- A. Deterministic ones use randomness; nondeterministic do not
- B. Deterministic ones are slower; nondeterministic are faster
- C. Nondeterministic ones always find the optimal solution path
- D. Nondeterministic ones can explore multiple paths at once

Answer: D

Q851. How does the XOR linked list reduce memory usage?

- A. It stores XOR of next and previous addresses per node
- B. It stores only even-indexed nodes and skips odd ones
- C. It removes half the nodes and links remaining ones
- D. It compresses data values using XOR encoding scheme

Answer: A

Q852. What is the time complexity of reversing a singly linked list iteratively?

- A. $O(1)$
- B. $O(n^2)$
- C. $O(n \log n)$
- D. $O(n)$

Answer: D

Q853. How can you detect a cycle in a linked list in $O(1)$ space?

- A. Use Floyd's cycle detection with two pointer speeds
- B. Sort the list and check for duplicate node values
- C. Use a hash set to store all visited node addresses
- D. Store visited flags inside each node data field

Answer: A

Q854. What is the worst-case time for dynamic array resizing?

- A. $O(\log n)$ due to binary copy optimization
- B. $O(1)$ per single resize operation performed
- C. $O(n)$ to copy all elements to new array
- D. $O(n^2)$ due to nested copy and shift steps

Answer: C

Q855. How is a min-stack implemented with $O(1)$ getMin?

- A. Store min value only at the bottom of the stack
- B. Sort the stack after every push operation call
- C. Scan the entire stack each time getMin is called
- D. Maintain an auxiliary stack tracking current minimums

Answer: D

Q856. What is the skip list and how does it improve linked list search?

- A. A multi-level linked list with express lane pointers
- B. A balanced tree structure embedded in a linked list
- C. A sorted array overlay that replaces the linked list
- D. A hash table combined with a singly linked list node

Answer: A

Q857. What is an unrolled linked list?

- A. A linked list where each node stores an array of elements
- B. A linked list with randomized pointer assignments only
- C. A linked list where nodes are stored in reverse order
- D. A linked list that has been converted into a flat array

Answer: A

Q858. How do you find the intersection point of two linked lists efficiently?

- A. Reverse both lists and compare from the tail forward
- B. Hash all elements of both lists and compare buckets
- C. Sort both lists and perform a binary search on each
- D. Compute lengths, align starts, then traverse together

Answer: D

Q859. What is the time complexity of merging k sorted linked lists using a min-heap?

- A. $O(nk \log k)$
- B. $O(nk^2)$
- C. $O(nk)$
- D. $O(n \log k)$

Answer: A

Q860. What is a gap buffer and where is it used?

- A. An array with a movable gap used in text editors
- B. A circular array used for network packet buffering
- C. A stack-based buffer used for function call storage
- D. A linked list variant used in database indexing

Answer: A

Q861. What is a bipartite graph?

- A. A graph whose vertices can be split into two disjoint independent sets
- B. A directed graph with edges going in exactly two directions only
- C. A graph where every vertex has exactly two edges connected to it
- D. A graph that contains exactly two connected components in total

Answer: A

Q862. How many edges does a tree with n vertices have?

- A. $n + 1$
- B. n
- C. $n - 1$
- D. $n / 2$

Answer: C

Q863. What is the worst-case height of an unbalanced BST with n nodes?

- A. $O(\sqrt{n})$
- B. $O(n \log n)$
- C. $O(\log n)$
- D. $O(n)$

Answer: D

Q864. What is a strongly connected component in a directed graph?

- A. A subgraph where every vertex is reachable from every other vertex
- B. A subgraph containing exactly one vertex with no outgoing edges
- C. A subgraph where all vertices have the same in-degree value
- D. A maximal set of vertices with no edges between any of them

Answer: A

Q865. What is the Euler's formula for connected planar graphs?

- A. $V + E - F = 2$
- B. $V + E + F = 2$
- C. $V - E - F = 2$
- D. $V - E + F = 2$

Answer: D

Q866. What is a threaded binary tree?

- A. A binary tree where null pointers are replaced with in-order links
- B. A binary tree where each node stores a thread identifier
- C. A binary tree stored across multiple threads of execution
- D. A binary tree that requires multi-threaded access for traversal

Answer: A

Q867. What is a hypergraph?

- A. A graph where edges can connect any number of vertices
- B. A three-dimensional representation of a standard graph
- C. A graph where vertices are organized in a grid pattern
- D. A graph where each edge connects exactly two vertices

Answer: A

Q868. What is the chromatic number of a graph?

- A. The maximum degree of any single vertex in the graph
- B. The minimum number of colors needed to color all vertices
- C. The number of connected components in the given graph
- D. The total count of edges present in the entire graph

Answer: B

Q869. What is the maximum number of edges in a simple undirected graph with n vertices?

- A. n^2
- B. $n(n-1)$
- C. $n(n+1)/2$
- D. $n(n-1)/2$

Answer: D

Q870. What is a k -ary tree?

- A. A tree with k root nodes and shared leaf structure
- B. A tree where each node has exactly k children always
- C. A tree where each node has at most k children nodes
- D. A tree that has exactly k levels from root to leaves

Answer: C

Q871. What is the worst-case time complexity of interpolation search?

- A. $O(n)$
- B. $O(\sqrt{n})$
- C. $O(\log \log n)$
- D. $O(\log n)$

Answer: A

Q872. How does exponential search combine with binary search?

- A. It uses exponentiation to compute hash for direct access
- B. It finds a range by doubling then binary searches within
- C. It applies binary search on exponentially growing subarrays
- D. It sorts the array exponentially then searches linearly

Answer: B

Q873. What is the time complexity of search in a skip list?

- A. $O(\log n)$ expected with $O(n)$ worst case
- B. $O(n \log n)$ expected and $O(n^2)$ worst case
- C. $O(n)$ in both average and worst cases
- D. $O(1)$ in all cases due to direct access

Answer: A

Q874. What is the lower bound for comparison-based searching in sorted arrays?

- A. $O(\log n)$
- B. $O(1)$
- C. $O(\sqrt{n})$
- D. $O(n)$

Answer: A

Q875. How does search work in a van Emde Boas tree?

- A. It exploits universe size for $O(\log \log u)$ operations
- B. It uses comparison-based binary search on tree nodes
- C. It performs breadth-first search on a balanced binary tree
- D. It hashes each element for constant time access always

Answer: A

Q876. What is a fractional cascading technique?

- A. Dividing search space into fractional geometric sections
- B. Cascading hash functions for multi-level collision handling
- C. Sharing elements between catalogs to speed up searches
- D. Splitting arrays into fractions for parallel search

Answer: C

Q877. What is the time complexity of searching in an optimal binary search tree?

- A. $O(1)$ since the tree is precomputed optimally
- B. $O(n)$ for every possible search query
- C. $O(n^2)$ due to the overhead of optimal structure
- D. $O(\log n)$ or better based on access frequencies

Answer: D

Q878. What is the two-pointer search technique?

- A. Using two separate arrays for parallel search
- B. Using two recursive calls to divide the search space
- C. Using two hash tables to resolve search collisions
- D. Using two indices moving toward each other or same direction

Answer: D

Q879. What is the complexity of searching in a perfectly balanced k-d tree?

- A. $O(k * \log n)$
- B. $O(n)$
- C. $O(\log n)$
- D. $O(n^{(1-1/d)} + k)$

Answer: D

Q880. What is the sentinel linear search optimization?

- A. Adding a sentinel at the end to eliminate bound checking
- B. Using multiple sentinels to divide array into segments
- C. Replacing all elements with sentinels then binary search
- D. Sorting the array first and placing sentinel in middle

Answer: A

Q881. What is the lower bound for comparison-based sorting?

- A. $O(\log n)$
- B. $O(n)$
- C. $O(n \log n)$
- D. $O(n^2)$

Answer: C

Q882. What is the worst-case time complexity of randomized quick sort?

- A. $O(n \log^2 n)$ amortized
- B. $O(n \log n)$ guaranteed always
- C. $O(n^2)$ but extremely unlikely
- D. $O(n)$ with high probability

Answer: C

Q883. What is introsort and why is it used?

- A. A sort using interpolation search within sorting passes
- B. A sort based on introspective analysis of data patterns
- C. A sort combining insertion, heap, and quick sort strategies
- D. A sort that introduces randomness into merge sort steps

Answer: C

Q884. How does Tim sort combine merge sort and insertion sort?

- A. It merge sorts odd indices and insertion sorts even ones
- B. It finds natural runs, extends them with insertion sort, merges
- C. It uses merge sort first then refines with insertion sort
- D. It alternates between merge and insertion on each pass

Answer: B

Q885. What is the time complexity of bucket sort in the average case?

- A. $O(n + k)$
- B. $O(n \log n)$
- C. $O(n^2)$
- D. $O(k \log k)$

Answer: A

Q886. What is the median-of-three strategy in quick sort?

- A. Sorting the first three elements as a preprocessing step
- B. Dividing the array into three parts instead of two parts
- C. Choosing the median of first, middle, and last elements
- D. Running three separate quick sorts and taking the median

Answer: C

Q887. What is the worst-case time of radix sort for n d -digit numbers?

- A. $O(n \log n)$
- B. $O(d * n^2)$
- C. $O(n^2)$
- D. $O(d * n)$

Answer: D

Q888. Why is merge sort preferred for sorting linked lists?

- A. Merge sort automatically handles duplicate values better
- B. Merge sort is simpler to implement than other sorts
- C. Merge sort has better time complexity for linked lists
- D. Merge sort needs no random access and merges in $O(1)$ space

Answer: D

Q889. What is the Dutch National Flag problem in sorting?

- A. A problem of sorting flags in alphabetical country order
- B. Sorting elements from multiple countries by region codes
- C. Arranging colored blocks using merge sort partitioning
- D. Sorting elements by color: red, white, and blue efficiently

Answer: D

Q890. What is the adaptive property in sorting algorithms?

- A. The algorithm automatically selects the best data type
- B. The algorithm takes advantage of existing order in input
- C. The algorithm adapts to different programming languages
- D. The algorithm changes its memory usage based on input

Answer: B

Q891. What is universal hashing?

- A. A single hash function that works for all input types
- B. Using the same hash function universally in all tables
- C. A hash function standardized across all programming uses
- D. Randomly selecting a hash function from a designed family

Answer: D

Q892. What is perfect hashing?

- A. A hash function that produces perfectly sequential values
- B. A hashing method that uses perfectly balanced tree nodes
- C. A hashing scheme with zero collisions for a static set
- D. A hash table with load factor of exactly one point zero

Answer: C

Q893. What is a cuckoo hash table?

- A. A table using bird-like migration patterns for elements
- B. A table that randomly evicts elements when it gets full
- C. A table using cuckoo search optimization for hash values
- D. A table using two hash functions with element displacement

Answer: D

Q894. What is consistent hashing used for?

- A. Maintaining sorted order of keys in the hash table
- B. Ensuring hash values are always consecutive integers
- C. Distributing keys across nodes with minimal reassignment
- D. Guaranteeing the same hash for same input every time

Answer: C

Q895. What is a Bloom filter?

- A. A sorting filter applied before inserting into hash table
- B. A hash table variant that filters out duplicate key entries
- C. A deterministic set membership data structure using hashing
- D. A probabilistic structure that may give false positives only

Answer: D

Q896. What is Robin Hood hashing?

- A. A chaining method where collisions are sorted by hash value
- B. Stealing slots from rich elements to give to poor elements
- C. Using a tree structure within each hash table bucket slot
- D. Probing by comparing displacement distances and swapping

Answer: D

Q897. What is the expected longest probe sequence with linear probing at load factor 0.5?

- A. $O(1)$
- B. $O(\sqrt{n})$
- C. $O(\log n)$
- D. $O(n)$

Answer: C

Q898. What is the space complexity of a Bloom filter with n elements and false positive rate p ?

- A. $O(n * \log(1/p))$
- B. $O(\log n)$
- C. $O(n^2)$
- D. $O(n)$

Answer: A

Q899. What is hopscotch hashing?

- A. Each element must be within a fixed neighborhood of its hash
- B. Elements are sorted by hop count from their hash position
- C. Probing follows a randomized hopping pattern each time
- D. Elements hop between two completely separate tables

Answer: A

Q900. What is the worst-case lookup time in a hash table with chaining?

- A. $O(1)$
- B. $O(\log n)$
- C. $O(n \log n)$
- D. $O(n)$

Answer: D

Q901. What is a splay tree?

- A. A multi-way search tree used for external data storage
- B. A self-adjusting BST that moves accessed nodes to root
- C. A tree that splays leaves outward for balanced structure
- D. A tree that spreads nodes evenly across all levels always

Answer: B

Q902. What is a treap?

- A. A tree with trap-based deletion for security purposes
- B. A combination of a BST and a heap using random priorities
- C. A binary tree that operates as both stack and queue
- D. A tree that traps duplicate elements in separate nodes

Answer: B

Q903. What is the time complexity of range query in a segment tree?

- A. $O(1)$
- B. $O(\log n)$
- C. $O(n)$
- D. $O(n \log n)$

Answer: B

Q904. What is a B+ tree and how does it differ from a B-tree?

- A. B+ trees have fewer keys per node than standard B-trees
- B. B+ trees store all data in leaves with linked leaf nodes
- C. B+ trees do not maintain sorted order within their nodes
- D. B+ trees store data only in internal nodes not in leaves

Answer: B

Q905. What is lazy propagation in a segment tree?

- A. Deferring updates to child nodes until they are accessed
- B. Propagating updates to all nodes immediately upon change
- C. Lazily constructing the tree only when queries are made
- D. Propagating values from leaves to root after every update

Answer: A

Q906. What is an Euler tour of a tree?

- A. A traversal listing each node when entering and leaving
- B. A tour that visits only leaf nodes of the given tree
- C. A traversal that processes nodes in reverse level order
- D. A path visiting every edge exactly once in entire tree

Answer: A

Q907. What is the amortized time of splay tree operations?

- A. $O(\log n)$ per operation in the amortized analysis
- B. $O(n)$ per operation in the amortized analysis
- C. $O(1)$ per operation in the amortized analysis
- D. $O(n \log n)$ per operation in the amortized analysis

Answer: A

Q908. What is a Fenwick tree (Binary Indexed Tree)?

- A. A tree that indexes binary representations of all keys
- B. A full binary tree storing prefix sums at each node
- C. A binary tree indexed by Fenwick encoding of data keys
- D. A tree using binary indexing for prefix sum queries

Answer: D

Q909. What is a persistent data structure in the context of trees?

- A. A tree that persists on disk rather than in memory
- B. A tree that cannot be modified once it is created
- C. A tree stored in a persistent database management system
- D. A tree that preserves all previous versions after updates

Answer: D

Q910. What is the order of a B-tree?

- A. The total number of keys stored in the entire tree
- B. The maximum number of children each node can have
- C. The minimum number of keys in the root node only
- D. The height of the tree from root to deepest leaf

Answer: B

Q911. What is Tarjan's algorithm used for?

- A. Finding shortest paths in weighted directed graph
- B. Finding minimum spanning tree of an undirected graph
- C. Finding strongly connected components in directed graph
- D. Finding maximum matching in a bipartite graph only

Answer: C

Q912. What is an articulation point in a graph?

- A. A vertex that lies on the shortest path between two nodes
- B. A vertex whose removal disconnects the graph components
- C. A vertex with the highest degree in the entire graph
- D. A vertex that connects two spanning trees together only

Answer: B

Q913. What is the Ford-Fulkerson method?

- A. Computing maximum flow in a flow network iteratively
- B. Finding shortest path between two vertices in a graph
- C. Finding minimum cut in an unweighted undirected graph
- D. Computing minimum spanning tree using edge relaxation

Answer: A

Q914. What is the max-flow min-cut theorem?

- A. Maximum flow equals the minimum cut capacity value
- B. Minimum flow equals the maximum cut capacity value
- C. Maximum cut always exceeds the minimum flow amount
- D. Maximum flow equals minimum number of edges in graph

Answer: A

Q915. What is Kosaraju's algorithm?

- A. A two-pass DFS algorithm for strongly connected components
- B. A dynamic programming approach for all-pairs shortest paths
- C. A greedy algorithm for finding minimum spanning trees
- D. A divide and conquer method for planar graph coloring

Answer: A

Q916. What is a bridge in a graph?

- A. A vertex that connects two separate components together
- B. An edge whose removal disconnects the graph into parts
- C. The shortest edge in the entire graph by edge weight
- D. An edge that forms part of every spanning tree found

Answer: B

Q917. What is the A* search algorithm?

- A. A brute-force search examining all paths exhaustively
- B. A random walk search with probabilistic path selection
- C. An uninformed search using only edge weight for priority
- D. A heuristic-guided search combining cost and estimate

Answer: D

Q918. What is the time complexity of finding all bridges in a graph?

- A. $O(V + E)$
- B. $O(V * E)$
- C. $O(E^2)$
- D. $O(V^2)$

Answer: A

Q919. What is a minimum cut in a graph?

- A. The smallest set of edges whose removal disconnects source from sink
- B. The edge with the minimum weight in the entire graph structure
- C. The minimum number of vertices needed to cover all edges
- D. The shortest path between any two vertices in graph

Answer: A

Q920. What is the Johnson's algorithm for all-pairs shortest paths?

- A. Running Bellman-Ford from every vertex in the given graph
- B. Using Floyd-Warshall with optimized matrix multiplication
- C. Using BFS from each vertex for unweighted graph paths
- D. Reweighting edges then running Dijkstra from each vertex

Answer: D

Q921. What is the matrix chain multiplication problem?

- A. Multiplying all matrices in a chain from left to right
- B. Multiplying chains of numbers using matrix representation
- C. Finding optimal parenthesization to minimize scalar multiplications
- D. Finding the chain of matrices with the largest determinant

Answer: C

Q922. What is the traveling salesman problem (TSP)?

- A. Finding the fastest way to deliver mail in one city only
- B. Finding the shortest route visiting all cities exactly once
- C. Finding the minimum spanning tree of a city road network
- D. Finding the cheapest flight between two specific airports

Answer: B

Q923. What is the time complexity of the DP solution for TSP?

- A. $O(n!)$
- B. $O(n^2 * 2^n)$
- C. $O(n^3)$
- D. $O(2^n)$

Answer: B

Q924. What is the principle of optimality in dynamic programming?

- A. The first feasible solution found is always the optimal one
- B. An optimal solution contains optimal solutions to subproblems
- C. Every subproblem must be solved before the main problem
- D. The optimal solution always uses the greedy choice at start

Answer: B

Q925. What is a randomized algorithm?

- A. An algorithm that only works on randomly generated inputs
- B. An algorithm that uses random numbers to guide decisions
- C. An algorithm that produces random incorrect output values
- D. An algorithm that randomly skips steps for faster speed

Answer: B

Q926. What is the edit distance problem?

- A. Finding the distance between two nodes in an edit graph
- B. Counting the number of edits made to a document over time
- C. Finding minimum operations to transform one string to another
- D. Measuring the physical distance between edit buttons on screen

Answer: C

Q927. What is the difference between Las Vegas and Monte Carlo algorithms?

- A. Las Vegas always gives correct answer; Monte Carlo may not
- B. Monte Carlo always gives correct answer; Las Vegas may not
- C. Both may give wrong answers but differ in probability only
- D. Both always give correct answers but differ in time only

Answer: A

Q928. What is the subset sum problem?

- A. Finding the sum of all elements in a given set of numbers
- B. Partitioning a set into subsets with equal element counts
- C. Determining if a subset exists that sums to a target value
- D. Finding the largest subset with maximum element sum total

Answer: C

Q929. What is the time complexity of the edit distance DP solution?

- A. $O(\max(m, n))$
- B. $O(m^2 * n^2)$
- C. $O(m + n)$
- D. $O(m * n)$

Answer: D

Q930. What is the approximation ratio of a 2-approximation algorithm?

- A. It runs in twice the time of the optimal algorithm speed
- B. It finds a solution at most twice the optimal cost value
- C. It finds a solution exactly half of the optimal cost value
- D. It uses twice the memory of the optimal algorithm space

Answer: B

Q931. What is a van Emde Boas tree?

- A. A BST variant with van Emde Boas balancing rotations
- B. A hash tree using van Emde Boas hashing for collisions
- C. A tree supporting $O(\log \log u)$ operations on integers
- D. A tree named after its visual Van-like branching pattern

Answer: C

Q932. What is a splay tree's working set property?

- A. Elements in the working set are stored in a separate list
- B. The working set of the program determines tree structure
- C. Accessing recently accessed elements is faster than others
- D. Recently accessed elements are always at leaf positions

Answer: C

Q933. What is a link-cut tree?

- A. A binary tree where links between nodes can be toggled
- B. A linked list that can be cut and reconnected efficiently
- C. A tree that measures the number of cuts needed to split
- D. A tree that links and cuts edges to maintain forest paths

Answer: D

Q934. What is the inverse Ackermann function's role in Union-Find?

- A. It generates the sequence of ranks for union operations
- B. It determines the initial size of the disjoint set array
- C. It computes the Ackermann value used for hashing keys
- D. It describes the amortized time per operation with optimizations

Answer: D

Q935. What is a persistent segment tree?

- A. A segment tree stored persistently on disk for backup
- B. A segment tree that persists in memory and never freed
- C. A segment tree that preserves all versions after updates
- D. A segment tree with persistent root node that never changes

Answer: C

Q936. What is a suffix tree?

- A. A binary tree where values are string suffix indices
- B. A tree storing all prefixes of a given input string
- C. A tree that sorts strings by their suffix characters
- D. A compressed trie of all suffixes of a given string

Answer: D

Q937. What is the time complexity to build a suffix array?

- A. $O(n \log^2 n)$
- B. $O(n \log n)$
- C. $O(n^2)$
- D. $O(n)$

Answer: D

Q938. What is a range tree?

- A. A tree storing ranges of IP addresses for routing
- B. A multi-level tree for orthogonal range queries on points
- C. A binary tree where nodes represent value range intervals
- D. A tree that stores the range of values at each level

Answer: B

Q939. What is the merge operation complexity in a binomial heap?

- A. $O(n \log n)$
- B. $O(\log n)$
- C. $O(1)$
- D. $O(n)$

Answer: B

Q940. What is a weight-balanced tree?

- A. A tree where each node stores its subtree weight for balancing
- B. A tree that assigns weights to nodes based on depth level
- C. A tree balanced by the total data size in each subtree
- D. A tree where all edges have equal weight values assigned

Answer: A

Q941. What is Manacher's algorithm?

- A. An algorithm for computing edit distance in linear time
- B. An algorithm for sorting strings in optimal $O(n \log n)$ time
- C. An algorithm for building suffix trees in constant space
- D. An algorithm for finding all palindromic substrings in $O(n)$

Answer: D

Q942. What is the Aho-Corasick algorithm?

- A. A text indexing algorithm using Aho-style balanced trees
- B. A single-pattern matching algorithm using dynamic tables
- C. A multi-pattern matching algorithm using automaton and trie
- D. A string compression algorithm using coarse approximation

Answer: C

Q943. What is the suffix automaton?

- A. A pushdown automaton that processes suffix expressions
- B. An automaton that generates all possible suffixes randomly
- C. A finite automaton accepting all suffixes of a given string
- D. A deterministic automaton that sorts suffixes alphabetically

Answer: C

Q944. What is the Burrows-Wheeler Transform used for?

- A. Transforming binary data into readable text string format
- B. Sorting strings using Wheeler's specialized comparison
- C. Encrypting strings for secure data transmission purposes
- D. Transforming strings to improve compressibility of data

Answer: D

Q945. What is the LCP array in suffix array construction?

- A. An array mapping each character position to its prefix identifier
- B. An array of longest common prefixes between consecutive sorted suffixes
- C. An array storing the length of each suffix in the original string
- D. An array of least common patterns found across all string suffixes

Answer: B

Q946. What is the time complexity of building a suffix tree using Ukkonen's algorithm?

- A. $O(n^2)$
- B. $O(n^3)$
- C. $O(n)$
- D. $O(n \log n)$

Answer: C

Q947. What is the longest repeated substring problem?

- A. Finding the longest substring that occurs at least twice
- B. Finding the longest string that can be formed by repeating
- C. Finding any substring that repeats in consecutive positions
- D. Finding the substring that appears most frequently overall

Answer: A

Q948. What is the worst-case time of Boyer-Moore algorithm?

- A. $O(n / m)$
- B. $O(n + m)$
- C. $O(n * m)$
- D. $O(n)$

Answer: C

Q949. What is the difference between Rabin-Karp and KMP algorithms?

- A. Rabin-Karp is recursive; KMP is always purely iterative
- B. Rabin-Karp is for numbers; KMP is only for text strings
- C. Rabin-Karp uses hashing; KMP uses prefix table comparison
- D. Rabin-Karp uses sorting; KMP uses dynamic programming tables

Answer: C

Q950. What is the time complexity of computing the Z-array?

- A. $O(n^2)$
- B. $O(n \log n)$
- C. $O(n)$
- D. $O(n * m)$

Answer: C

Q951. What is Strassen's algorithm and what complexity does it achieve?

- A. A matrix multiplication algorithm achieving $O(n^{2.81})$ time
- B. A string matching algorithm achieving $O(n + m)$ time bound
- C. A graph search algorithm achieving $O(V + E)$ complexity
- D. A sorting algorithm achieving $O(n \log n)$ complexity always

Answer: A

Q952. What is an NP-complete problem?

- A. A problem that can only be solved using nondeterministic steps
- B. A problem that has been completely solved in polynomial time
- C. A problem in NP that is at least as hard as all NP problems
- D. A problem with no possible solution by any known algorithm

Answer: C

Q953. What is the significance of NP-hardness?

- A. NP-hard problems can only be solved using quantum computers
- B. NP-hard problems are easy to solve in polynomial time
- C. NP-hard problems are at least as hard as all problems in NP
- D. NP-hard problems have exactly one unique correct solution

Answer: C

Q954. What is the Cook-Levin theorem?

- A. It proves that SAT is the first NP-complete problem known
- B. It proves that P equals NP for all practical problems
- C. It proves that all NP-hard problems are also in the class P
- D. It proves that no NP problem can be solved in polynomial time

Answer: A

Q955. What is the complexity class PSPACE?

- A. Problems where space equals time complexity for all inputs
- B. Problems requiring exponential space for any solution found
- C. Problems solvable in polynomial time with unlimited space
- D. Problems solvable using polynomial amount of memory space

Answer: D

Q956. What is an FPT (Fixed-Parameter Tractable) algorithm?

- A. An algorithm with fixed running time for all parameter values
- B. An algorithm where exponential cost depends only on parameter k
- C. An algorithm using fixed-point arithmetic for optimization
- D. An algorithm that fixes parameters to reduce space complexity

Answer: B

Q957. What is the gap between P and NP problems practically?

- A. P problems are deterministic; NP problems are always randomized
- B. P problems use less memory; NP problems use more memory only
- C. P problems are solved in seconds; NP takes millennia for large input
- D. There is no practical gap; all NP problems run fast in practice

Answer: C

Q958. What is a polynomial-time reduction?

- A. Reducing the polynomial degree of a time complexity function
- B. Transforming one problem into another in polynomial time steps
- C. Converting exponential algorithms into polynomial ones always
- D. Reducing the running time of any algorithm to polynomial time

Answer: B

Q959. What is the complexity of the best known matrix multiplication algorithm?

- A. $O(n^2)$ matching the lower bound for matrix multiplication
- B. $O(n^3)$ with no improvement possible over standard method
- C. $O(n^{2.373})$ using advanced algebraic technique approaches
- D. $O(n^{2.5})$ using a combination of Strassen and standard methods

Answer: C

Q960. What is the role of reductions in proving NP-completeness?

- A. Reductions simplify NP-complete problems to make them solvable
- B. Reductions convert NP-complete problems into polynomial ones
- C. Reductions eliminate the need to prove problem membership in NP
- D. Reductions show a known NP-complete problem transforms to new one

Answer: D

Q961. What is the solution to the recurrence $T(n) = 2T(n/4) + n$?

- A. $O(n)$
- B. $O(n \log n)$
- C. $O(n \log n)$
- D. $O(n^2)$

Answer: B

Q962. What is the Akra-Bazzi method used for?

- A. Sorting elements in an array
- B. Solving recurrences with unequal subproblem sizes that the Master Theorem cannot handle
- C. Finding shortest paths in graphs
- D. Balancing binary search trees

Answer: B

Q963. In amortized analysis, what does the potential function represent?

- A. The actual cost of each operation
- B. The stored-up energy in the data structure that can pay for future expensive operations
- C. The maximum possible running time
- D. The probability of an expensive operation occurring

Answer: B

Q964. What complexity class does an algorithm with running time $O(n^{\log_2 3})$ belong to when compared to $O(n^2)$?

- A. Faster than $O(n^2)$ since $\log_2 3 \approx 1.585 < 2$
- B. Slower than $O(n^2)$ since $\log_2 3 > 2$
- C. Exactly $O(n^2)$
- D. Incomparable to $O(n^2)$

Answer: A

Q965. What is the smoothed complexity analysis of an algorithm?

- A. Analysis after rounding input sizes to powers of 2
- B. Analysis of the algorithm on inputs with small random perturbations added to worst-case inputs
- C. Analysis that ignores the first n operations
- D. Analysis only on uniformly random inputs

Answer: B

Q966. What is the time complexity of $T(n) = T(n-1) + T(n-2) + O(1)$?

- A. $O(n)$
- B. $O(n^2)$
- C. $O(2^n)$
- D. $O(n \log n)$

Answer: C

Q967. What is the difference between output-sensitive and input-sensitive algorithms?

- A. Output-sensitive algorithms depend on both input and output size, while input-sensitive depend only on input size
- B. They are the same thing
- C. Output-sensitive algorithms are always faster
- D. Input-sensitive algorithms cannot be analyzed

Answer: A

Q968. What is the complexity of $T(n) = T(\lceil n \rceil) + O(1)$ after the substitution $m = \log n$?

- A. $O(\log n)$
- B. $O(\log \log n)$
- C. $O(\lceil n \rceil)$
- D. $O(n)$

Answer: B

Q969. What does the Extended Church-Turing thesis imply about algorithm complexity?

- A. All problems can be solved in polynomial time
- B. Any model of computation can be simulated by a Turing machine with at most a polynomial slowdown
- C. Quantum computers can solve NP-complete problems efficiently
- D. Recursion is always faster than iteration

Answer: B

Q970. What is the worst-case space complexity of a recursive algorithm with branching factor b and depth d ?

- A. $O(b^d)$
- B. $O(b * d)$
- C. $O(d)$
- D. $O(b + d)$

Answer: C

Q971. What is the time complexity of sorting a stack using only one additional stack?

- A. $O(n)$
- B. $O(n \log n)$
- C. $O(n^2)$
- D. $O(2^n)$

Answer: C

Q972. What is the time complexity of the LRU cache eviction using a doubly linked list and a hash map?

- A. $O(n)$ for get and $O(1)$ for put
- B. $O(1)$ for both get and put
- C. $O(\log n)$ for both get and put
- D. $O(n)$ for both get and put

Answer: B

Q973. What is a self-organizing list using the move-to-front heuristic?

- A. A list that sorts elements on each access
- B. A list that moves the most recently accessed element to the front
- C. A list that randomly rearranges elements
- D. A list that deletes rarely accessed elements

Answer: B

Q974. What is the worst-case time complexity of the skip list's insertion operation?

- A. $O(1)$
- B. $O(\log n)$
- C. $O(n)$
- D. $O(n \log n)$

Answer: C

Q975. How does a lock-free concurrent queue typically achieve thread safety?

- A. Using mutexes on every operation
- B. Using Compare-And-Swap (CAS) atomic operations
- C. By making copies of the entire queue for each thread
- D. By allowing only one thread to use the queue

Answer: B

Q976. What is the amortized time complexity of the multipop operation on a stack that pops k elements?

- A. $O(k)$ amortized per multipop
- B. $O(1)$ amortized per element across all operations
- C. $O(n)$ amortized per multipop
- D. $O(k^2)$ amortized per multipop

Answer: B

Q977. What problem does the 'Next Greater Element' algorithm solve using a stack?

- A. Finding the maximum element in an array
- B. For each element, finding the next element in the array that is strictly greater
- C. Sorting the array in descending order
- D. Finding the median of the array

Answer: B

Q978. What is the cache performance advantage of arrays over linked lists?

- A. Arrays support faster deletion
- B. Arrays store elements contiguously in memory, enabling better spatial locality and prefetching
- C. Arrays use less total memory
- D. Arrays have faster pointer operations

Answer: B

Q979. What is the maximum number of elements that can be stored in a circular queue of array size n ?

- A. n
- B. $n - 1$
- C. $n + 1$
- D. $2n$

Answer: B

Q980. What is the time complexity of flattening a multilevel doubly linked list?

- A. $O(n)$ where n is the total number of nodes across all levels
- B. $O(n^2)$
- C. $O(n \log n)$
- D. $O(2^n)$

Answer: A

Q981. What is the minimum number of nodes in a Red-Black tree of height h (counting black height)?

- A. $2^h - 1$
- B. $2^{(h/2)}$
- C. $h + 1$
- D. 2^h

Answer: A

Q982. What is the time complexity of checking if an undirected graph is bipartite?

- A. $O(V^2)$
- B. $O(V + E)$
- C. $O(V * E)$
- D. $O(2^V)$

Answer: B

Q983. How many distinct binary trees can be formed with n nodes?

- A. $n!$
- B. 2^n
- C. The n th Catalan number $C(2n, n)/(n+1)$
- D. n^2

Answer: C

Q984. What is the time complexity of building a min-heap from an array of n elements using the bottom-up approach?

- A. $O(n \log n)$
- B. $O(n^2)$
- C. $O(n)$
- D. $O(\log n)$

Answer: C

Q985. What is a multigraph?

- A. A graph with multiple connected components
- B. A graph that allows multiple edges between the same pair of vertices
- C. A graph where all vertices have the same degree
- D. A graph that is both directed and undirected

Answer: B

Q986. What is the constant factor in the upper bound of AVL tree height relative to $\log_2(n)$?

- A. Exactly 1.0
- B. Approximately 1.44
- C. Approximately 2.0
- D. Approximately 0.5

Answer: B

Q987. What is the minimum number of edges to remove to make a connected graph a tree?

- A. $V - 1$
- B. $E - V + 1$
- C. $E - 1$
- D. V

Answer: B

Q988. What property does a B-tree of order m guarantee about node occupancy?

- A. Every node has exactly m children
- B. Every non-root internal node has at least $\text{ceil}(m/2)$ children
- C. Every node has at most 2 children
- D. Leaf nodes store no keys

Answer: B

Q989. What is the isomorphism problem for trees?

- A. Checking if two trees have the same number of nodes
- B. Determining if two trees are structurally identical by finding a bijection between their nodes
- C. Finding the largest common subtree
- D. Checking if a tree is balanced

Answer: B

Q990. What is the graph density and how is it calculated for an undirected simple graph?

- A. The number of vertices divided by the number of edges
- B. $2E / (V(V-1))$, representing the ratio of actual edges to maximum possible edges
- C. $E * V$
- D. V^2 / E

Answer: B

Q991. What is the worst-case time complexity of the Quickselect algorithm?

- A. $O(n)$
- B. $O(n \log n)$
- C. $O(n^2)$
- D. $O(\log n)$

Answer: C

Q992. What is the Median of Medians algorithm used for?

- A. Sorting arrays efficiently
- B. Guaranteeing $O(n)$ worst-case selection of the k -th smallest element
- C. Finding the mode of an array
- D. Building a balanced binary search tree

Answer: B

Q993. How can you search in a row-wise and column-wise sorted 2D matrix in optimal time?

- A. Binary search every row: $O(n \log n)$
- B. Start from the top-right corner and eliminate a row or column at each step: $O(n + m)$
- C. Linear search: $O(n * m)$
- D. Sort the entire matrix first: $O(n * m * \log(n * m))$

Answer: B

Q994. What is the time complexity of searching for a key in a sorted and rotated array using modified binary search?

- A. $O(n)$
- B. $O(n \log n)$
- C. $O(\log n)$
- D. $O(n)$

Answer: C

Q995. What is a van Emde Boas tree's search time complexity?

- A. $O(\log n)$
- B. $O(\log \log U)$ where U is the universe size
- C. $O(1)$
- D. $O(n)$

Answer: B

Q996. What is the lower bound on the number of comparisons needed for searching in a sorted array?

- A. $O(1)$
- B. $O(\log n)$
- C. $\Omega(\log n)$ — worst case for comparison-based algorithms
- D. $O(n)$

Answer: C

Q997. What is the time complexity of searching in a balanced k-d tree with n points in d dimensions?

- A. $O(n)$
- B. $O(\log n)$ always
- C. $O(n^{1-1/d} + k)$ for range search returning k points
- D. $O(d * n)$

Answer: C

Q998. What is the fractional cascading technique and what speedup does it provide?

- A. It sorts fractional numbers faster
- B. It speeds up searching the same key in k sorted lists from $O(k \log n)$ to $O(k + \log n)$
- C. It divides the search space into fractions
- D. It cascades hash functions for faster lookup

Answer: B

Q999. What is the expected search time in a skip list with n elements?

- A. $O(n)$
- B. $O(n)$
- C. $O(\log n)$ with high probability
- D. $O(1)$

Answer: C

Q1000. What is the time complexity of the A* search algorithm in the worst case?

- A. $O(V + E)$ always
- B. $O(b^d)$ where b is the branching factor and d is the solution depth
- C. $O(V \log V)$
- D. $O(E \log V)$

Answer: B

Q1001. What is the time complexity of TimSort in the worst case?

- A. $O(n^2)$
- B. $O(n \log n)$
- C. $O(n)$
- D. $O(n \cdot n)$

Answer: B

Q1002. What is the worst-case time complexity of introsort?

- A. $O(n^2)$
- B. $O(n \log n)$
- C. $O(n \log^2 n)$
- D. $O(n)$

Answer: B

Q1003. What is the information-theoretic lower bound for comparison-based sorting expressed in terms of comparisons?

- A. n comparisons
- B. $n \log n$ comparisons
- C. $\lceil \log_2(n!) \rceil$ comparisons, which is $\Theta(n \log n)$
- D. n^2 comparisons

Answer: C

Q1004. What is the three-way partitioning (Dutch National Flag) algorithm's advantage for quicksort?

- A. It eliminates the need for recursion
- B. It handles arrays with many duplicate elements efficiently by partitioning into three groups: less, equal, and greater
- C. It always selects the optimal pivot
- D. It reduces space complexity to $O(1)$ for merge sort

Answer: B

Q1005. What is external sorting and when is it used?

- A. Sorting data stored in external variables
- B. Sorting large datasets that do not fit in main memory by using disk-based merge strategies
- C. Sorting data received from external APIs
- D. Sorting arrays using an external library

Answer: B

Q1006. What is the time complexity of patience sorting?

- A. $O(n^2)$
- B. $O(n \log n)$
- C. $O(n)$
- D. $O(n \cdot n)$

Answer: B

Q1007. How does block sort (WikiSort) achieve $O(n \log n)$ sorting with $O(1)$ extra memory and stability?

- A. By using quicksort internally
- B. By using block-based merging strategies that rotate and swap fixed-size blocks in-place
- C. By sacrificing stability for space efficiency
- D. By using hash tables for comparisons

Answer: B

Q1008. What is the optimal number of comparisons needed to sort 5 elements?

- A. 5
- B. 7
- C. 8
- D. 10

Answer: B

Q1009. What is the worst-case time complexity of the best known gap sequence for shell sort?

- A. $O(n^2)$
- B. $O(n \log^2 n)$
- C. $O(n^{4/3})$
- D. $O(n \log n)$

Answer: C

Q1010. What is a sorting network and what is the depth of the optimal sorting network for n inputs?

- A. A neural network for sorting with $O(n)$ depth
- B. A fixed sequence of comparators; optimal depth is $O(\log^2 n)$ by the AKS network
- C. A graph that represents sorting comparisons; depth is $O(n)$
- D. A distributed system for parallel sorting; depth is $O(n \log n)$

Answer: B

Q1011. What is the expected maximum chain length in a hash table with n elements and n slots using universal hashing?

- A. $O(1)$
- B. $O(\log n / \log \log n)$
- C. $O(n)$
- D. $O(n)$

Answer: B

Q1012. How does cuckoo hashing achieve $O(1)$ worst-case lookup?

- A. By using a single hash function with a large table
- B. By using two hash functions and two tables, where each key is in exactly one of two possible locations
- C. By sorting the keys before inserting
- D. By maintaining a sorted linked list at each slot

Answer: B

Q1013. What is the false positive rate of a Bloom filter with m bits, k hash functions, and n inserted elements?

- A. $(1 - e^{-(kn/m)})^k$ approximately
- B. k/m
- C. n/m
- D. $1 - (1 - 1/m)^n$

Answer: A

Q1014. What is consistent hashing and why is it important in distributed systems?

- A. A hash function that always returns the same value
- B. A technique where adding or removing a server only requires remapping $O(K/n)$ keys on average, where K is keys and n is servers
- C. Hashing that guarantees zero collisions
- D. A hash function that works across different programming languages

Answer: B

Q1015. What is the amortized insertion time of cuckoo hashing?

- A. $O(1)$ amortized
- B. $O(\log n)$ amortized
- C. $O(n)$ amortized
- D. $O(n^2)$ amortized

Answer: A

Q1016. What is a minimal perfect hash function?

- A. A hash function with minimal computation
- B. A hash function that maps n keys to n consecutive integers with no collisions
- C. The simplest possible hash function
- D. A hash function that only works for small inputs

Answer: B

Q1017. What is the count-min sketch data structure used for?

- A. Exact counting of elements
- B. Approximate frequency estimation of elements in a data stream using sub-linear space
- C. Sorting elements by frequency
- D. Storing minimum and maximum values

Answer: B

Q1018. What is tabulation hashing and what property does it guarantee?

- A. A hash function stored in a table that guarantees $O(n \log n)$ sorting
- B. A hash function that XORs values from lookup tables for each byte of the key, providing 3-wise independence
- C. A hashing method that creates a truth table
- D. A method of hashing that only works for table-structured data

Answer: B

Q1019. What is the space complexity of a Bloom filter achieving false positive rate ϵ ?

- A. $O(n)$
- B. $O(n \log(1/\epsilon))$
- C. $O(1/\epsilon)$
- D. $O(n^2)$

Answer: B

Q1020. What is Robin Hood hashing's key invariant?

- A. Each slot has exactly one element
- B. Elements with longer probe distances displace elements with shorter probe distances to equalize probe lengths
- C. The hash table is always sorted
- D. Each element is hashed twice

Answer: B

Q1021. What is the time complexity of LCA computation using Euler tour and sparse table after $O(n \log n)$ preprocessing?

- A. $O(\log n)$ per query
- B. $O(1)$ per query
- C. $O(n)$ per query
- D. $O(n^2)$ per query

Answer: B

Q1022. What is Heavy-Light Decomposition (HLD) used for?

- A. Balancing binary trees
- B. Decomposing a tree into chains so that any root-to-leaf path crosses $O(\log n)$ chains, enabling efficient path queries
- C. Finding the heaviest subtree
- D. Removing heavy nodes from the tree

Answer: B

Q1023. Without lazy propagation, what is the time complexity of a range update on a segment tree, and how does lazy propagation improve it?

- A. $O(n)$ without, $O(\log n)$ with lazy propagation
- B. $O(n \log n)$ without, $O(n)$ with lazy propagation
- C. $O(\log n)$ without, $O(1)$ with lazy propagation
- D. Both are $O(n)$

Answer: A

Q1024. What is the amortized time complexity of splay tree operations?

- A. $O(1)$
- B. $O(\log n)$
- C. $O(n)$
- D. $O(n)$

Answer: B

Q1025. What is a Fenwick tree's advantage over a segment tree?

- A. Fenwick trees support more operations
- B. Fenwick trees use less memory (n elements vs $4n$) and have smaller constant factors while supporting prefix operations
- C. Fenwick trees have better worst-case complexity
- D. Fenwick trees support range updates natively

Answer: B

Q1026. How much extra space does each update operation use in a persistent segment tree?

- A. $O(1)$
- B. $O(n)$
- C. $O(\log n)$
- D. $O(n \log n)$

Answer: C

Q1027. What is the Centroid Decomposition of a tree used for?

- A. Finding the center of mass of weighted nodes
- B. Recursively finding and removing centroids to create a decomposition tree enabling efficient path-based queries
- C. Balancing unbalanced trees
- D. Finding the median node

Answer: B

Q1028. What is the time complexity of range update and point query in a Fenwick tree?

- A. $O(n)$ for each
- B. $O(\log n)$ for each
- C. $O(1)$ for update, $O(n)$ for query
- D. $O(n \log n)$ for each

Answer: B

Q1029. What is the time complexity of constructing an optimal BST given n keys with known access probabilities?

- A. $O(n)$
- B. $O(n^2)$
- C. $O(n^3)$
- D. $O(2^n)$

Answer: C

Q1030. How does a treap maintain both BST and heap properties simultaneously?

- A. By using two separate trees
- B. Each node has a key (BST-ordered) and a random priority (heap-ordered); rotations maintain both properties
- C. By sorting elements twice
- D. By storing elements in two arrays

Answer: B

Q1031. What is the time complexity of Tarjan's algorithm for finding strongly connected components?

- A. $O(V^2)$
- B. $O(V + E)$
- C. $O(V * E)$
- D. $O(V^3)$

Answer: B

Q1032. According to the duality principle in network flow, what is equal to the capacity of the minimum s - t cut?

- A. The number of edges in the graph
- B. The maximum flow from source s to sink t
- C. The minimum spanning tree weight
- D. The number of vertices on the shortest path

Answer: B

Q1033. What is the time complexity of the Edmonds-Karp algorithm for maximum flow?

- A. $O(V * E)$
- B. $O(V * E^2)$
- C. $O(V^2 * E)$
- D. $O(V^3)$

Answer: C

Q1034. What is an Euler path in a graph?

- A. A path that visits every vertex exactly once
- B. A path that traverses every edge exactly once
- C. The shortest path between two vertices
- D. A path with minimum total weight

Answer: B

Q1035. What is the time complexity of finding all articulation points in an undirected graph?

- A. $O(V^2)$
- B. $O(V + E)$
- C. $O(V * E)$
- D. $O(E^2)$

Answer: B

Q1036. What is Johnson's algorithm for all-pairs shortest paths?

- A. Running Dijkstra from every vertex without preprocessing
- B. Reweighting edges using Bellman-Ford to eliminate negative weights, then running Dijkstra from every vertex
- C. Using Floyd-Warshall on a reduced graph
- D. A divide-and-conquer approach for shortest paths

Answer: B

Q1037. What is a Hamiltonian path and why is finding one computationally hard?

- A. A shortest path in a weighted graph, solvable in polynomial time
- B. A path visiting every vertex exactly once; the decision problem is NP-complete
- C. A path traversing every edge exactly once; solvable in $O(V + E)$
- D. The longest path in a tree

Answer: B

Q1038. What is the Dinic's algorithm time complexity for maximum flow?

- A. $O(V * E)$
- B. $O(V^2 * E)$
- C. $O(V * E^2)$
- D. $O(V^3)$

Answer: B

Q1039. What is the 2-SAT problem and how is it related to graph algorithms?

- A. Finding two satisfying assignments using brute force
- B. A satisfiability problem solvable in polynomial time by building an implication graph and finding SCCs
- C. A graph coloring problem
- D. An NP-complete problem with no efficient solution

Answer: B

Q1040. What is the time complexity of finding a minimum spanning tree using Borůvka's algorithm?

- A. $O(V^2)$
- B. $O(E \log V)$
- C. $O(V + E)$
- D. $O(E \log E)$

Answer: B

Q1041. What is the time complexity of the matrix chain multiplication DP solution?

- A. $O(n^2)$
- B. $O(n^3)$
- C. $O(2^n)$
- D. $O(n \log n)$

Answer: B

Q1042. How does the meet-in-the-middle technique reduce the Subset Sum time from $O(2^n)$ to approximately $O(2^{(n/2)})$?

- A. By sorting the subsets before checking
- B. By splitting elements into two halves, enumerating all sums in each half, then combining using binary search or hashing
- C. By using dynamic programming instead of brute force
- D. By reducing the problem to a graph problem

Answer: B

Q1043. What is the difference between Las Vegas and Monte Carlo randomized algorithms?

- A. Las Vegas algorithms are deterministic
- B. Las Vegas algorithms always produce correct results (with variable runtime), while Monte Carlo may produce incorrect results (with bounded runtime)
- C. Monte Carlo algorithms are always faster
- D. They are the same type of algorithm

Answer: B

Q1044. What is the time complexity of solving the edit distance problem using DP?

- A. $O(n)$
- B. $O(n * m)$
- C. $O(n^3)$
- D. $O(2^n)$

Answer: B

Q1045. Which of the following problems violates Bellman's principle of optimality and therefore cannot be solved by standard DP?

- A. Shortest path problem
- B. Longest simple path in a general graph
- C. Matrix chain multiplication
- D. Edit distance

Answer: B

Q1046. What is the time complexity of the Held-Karp algorithm for the Traveling Salesman Problem?

- A. $O(n!)$
- B. $O(n^2 * 2^n)$
- C. $O(n^3)$
- D. $O(2^n)$

Answer: B

Q1047. What is the Knuth optimization for DP?

- A. Using hash tables for memoization
- B. An optimization for DP problems with the property that the optimal split point is monotone, reducing $O(n^3)$ to $O(n^2)$
- C. Parallelizing DP computations
- D. Converting recursion to iteration

Answer: B

Q1048. What is the Convex Hull Trick optimization in dynamic programming?

- A. Finding the convex hull of points
- B. An optimization for DP transitions of the form $dp[i] = \min(dp[j] + b[j] * a[i])$ using a set of lines
- C. Computing convex combinations of solutions
- D. A geometric algorithm for nearest neighbors

Answer: B

Q1049. What is the difference between online and offline algorithms?

- A. Online algorithms require internet access
- B. Online algorithms process input piece by piece as it arrives, while offline algorithms have access to the entire input upfront
- C. Offline algorithms are always slower
- D. Online algorithms cannot be analyzed

Answer: B

Q1050. What is the time complexity of solving the Subset Sum problem using DP?

- A. $O(n * S)$ where S is the target sum (pseudo-polynomial)
- B. $O(2^n)$
- C. $O(n^2)$
- D. $O(n \log n)$

Answer: A

Q1051. What is a van Emde Boas tree's space complexity?

- A. $O(n)$
- B. $O(U)$ where U is the universe size
- C. $O(\log U)$
- D. $O(n \log n)$

Answer: B

Q1052. What is a rope data structure and what advantage does it offer over arrays for string operations?

- A. A compressed string format
- B. A balanced binary tree of strings that supports $O(\log n)$ concatenation, split, and insertion instead of $O(n)$ for arrays
- C. A linked list of characters
- D. An encrypted string storage

Answer: B

Q1053. What is the order-statistic tree and what additional information does it maintain?

- A. A tree that only stores statistics
- B. An augmented BST that stores subtree sizes to support $O(\log n)$ rank and select queries
- C. A tree that orders nodes by creation time
- D. A tree that counts the number of queries

Answer: B

Q1054. What is a wavelet tree used for?

- A. Audio signal processing
- B. Answering rank, select, and range frequency queries on sequences in $O(\log \sigma)$ where σ is the alphabet size
- C. Compressing images
- D. Sorting wavelets

Answer: B

Q1055. Which underlying data structure does the link-cut tree use internally to maintain preferred paths?

- A. Red-Black trees
- B. AVL trees
- C. Splay trees
- D. B-trees

Answer: C

Q1056. How much space does a succinct data structure use relative to the information-theoretic minimum Z ?

- A. $O(Z)$ bits
- B. $Z + o(Z)$ bits, meaning within a lower-order additive term of the minimum
- C. Exactly Z bits
- D. $2Z$ bits

Answer: B

Q1057. What is a suffix tree and how does it relate to a suffix array?

- A. They are completely unrelated data structures
- B. A suffix tree is a compressed trie of all suffixes; a suffix array stores the same suffixes sorted, and they can be converted between in $O(n)$ time
- C. A suffix array is always preferred over a suffix tree
- D. A suffix tree stores prefix information while a suffix array stores suffix information

Answer: B

Q1058. What is the time complexity of the decrease-key operation in a Fibonacci heap?

- A. $O(\log n)$ amortized
- B. $O(1)$ amortized
- C. $O(n)$
- D. $O(\log \log n)$

Answer: B

Q1059. What is an x-fast trie and what time complexity does it achieve for predecessor queries?

- A. $O(\log n)$
- B. $O(\log \log U)$ using hash tables at each level of a binary trie over the universe
- C. $O(1)$
- D. $O(n)$

Answer: B

Q1060. What is the weight-balanced tree's balance condition?

- A. Left and right subtrees have equal height
- B. The ratio of sizes of left and right subtrees is bounded by a constant ; $\frac{1}{2} \leq \frac{B}{L} \leq 2$
- C. Each node has the same weight
- D. The total weight equals n

Answer: B

Q1061. What is the time complexity of Manacher's algorithm for finding the longest palindromic substring?

- A. $O(n^2)$
- B. $O(n \log n)$
- C. $O(n)$
- D. $O(n^3)$

Answer: C

Q1062. What is the Aho-Corasick algorithm's time complexity for matching k patterns of total length m against a text of length n ?

- A. $O(n * k)$
- B. $O(n + m + z)$ where z is the number of matches
- C. $O(n * m)$
- D. $O(k * m * n)$

Answer: B

Q1063. What is the suffix automaton of a string?

- A. An automaton that accepts only the suffixes of the string
- B. A minimal DFA that accepts exactly all substrings of the string, built in $O(n)$ time and space
- C. An automaton that generates all permutations
- D. A finite automaton for sorting strings

Answer: B

Q1064. What is the Burrows-Wheeler Transform (BWT) primarily used for?

- A. Encrypting strings for security
- B. Rearranging a string to group similar characters together, enabling better compression
- C. Sorting strings in reverse order
- D. Finding palindromes in strings

Answer: B

Q1065. What key optimization technique does Ukkonen's suffix tree algorithm use to achieve linear time?

- A. Hash tables for each node
- B. Suffix links, implicit extensions, and the active point trick to avoid redundant traversal
- C. Divide and conquer on the string
- D. Binary search on suffix positions

Answer: B

Q1066. How can the number of distinct substrings of a string be computed using a suffix array and LCP array?

- A. Count = $n(n+1)/2$
- B. Count = $n(n+1)/2$ - sum of all LCP values
- C. Count = sum of all LCP values
- D. Count = n^2

Answer: B

Q1067. What is the worst-case time complexity of the Boyer-Moore algorithm?

- A. $O(n/m)$
- B. $O(n + m)$
- C. $O(n * m)$
- D. $O(n)$

Answer: C

Q1068. What is the LCP (Longest Common Prefix) array and how is it used with suffix arrays?

- A. An array of string lengths
- B. $LCP[i]$ stores the length of the longest common prefix between the i -th and $(i-1)$ -th sorted suffixes; it enables efficient pattern matching and string analysis
- C. An array of character frequencies
- D. An array mapping characters to positions

Answer: B

Q1069. What is the time complexity of the Kasai algorithm for constructing the LCP array from a suffix array?

- A. $O(n^2)$
- B. $O(n \log n)$
- C. $O(n)$
- D. $O(n * m)$

Answer: C

Q1070. What is the palindromic tree (Eertree) data structure?

- A. A binary tree of palindrome lengths
- B. A data structure that stores all distinct palindromic substrings of a string in $O(n)$ time and space
- C. A trie of reversed strings
- D. A balanced BST of palindromes

Answer: B

Q1071. What is Cook's theorem and its significance?

- A. It proves $P = NP$
- B. It proves that the Boolean satisfiability problem (SAT) is NP-complete, establishing the first NP-complete problem
- C. It proves that all NP problems are unsolvable
- D. It provides a polynomial algorithm for SAT

Answer: B

Q1072. What is pseudo-polynomial time and why is it misleading?

- A. It is polynomial time disguised as exponential
- B. An algorithm runs in time polynomial in the numeric value of the input but exponential in the input size (number of bits)
- C. It is always faster than polynomial time
- D. It means the algorithm is approximately polynomial

Answer: B

Q1073. What is the complexity class co-NP?

- A. Problems whose complements are in P
- B. Problems whose complements are in NP; equivalently, problems with polynomial-time disproofs
- C. Problems harder than NP
- D. Problems that are not in NP

Answer: B

Q1074. What fundamental consequence does the time hierarchy theorem have for the P vs NP question?

- A. It proves $P = NP$
- B. It proves that strictly more time allows solving strictly more problems, implying $P \neq EXPTIME$
- C. It shows all problems are equally hard
- D. It proves $NP = PSPACE$

Answer: B

Q1075. What is the best known time complexity for matrix multiplication of two $n \times n$ matrices?

- A. $O(n^3)$
- B. $O(n^{2.807})$ (Strassen)
- C. $O(n^{2.371})$ approximately (current best)
- D. $O(n^2)$

Answer: C

Q1076. What is a Fixed-Parameter Tractable (FPT) algorithm?

- A. An algorithm with fixed running time
- B. An algorithm with time $O(f(k) * n^c)$ where k is a parameter, f is any computable function, and c is a constant independent of k
- C. An algorithm that fixes errors in parameters
- D. An algorithm with $O(n)$ time for any parameters

Answer: B

Q1077. What is the role of polynomial-time reductions in proving NP-completeness?

- A. They simplify the problem to make it solvable
- B. They show that if problem A reduces to problem B in polynomial time and A is NP-hard, then B is also NP-hard
- C. They reduce the time complexity of algorithms
- D. They convert NP problems to P problems

Answer: B

Q1078. What is the PSPACE complexity class?

- A. Problems solvable in polynomial time
- B. Problems solvable using polynomial space on a deterministic Turing machine
- C. Problems that use no space at all
- D. Problems solvable on parallel computers

Answer: B

Q1079. What is the approximation ratio and when is it used?

- A. The ratio of input size to output size
- B. For NP-hard optimization problems, the ratio of the algorithm's solution value to the optimal solution value, guaranteed to be within a bound
- C. The ratio of time complexity to space complexity
- D. The ratio of average case to worst case complexity

Answer: B

Q1080. What is the Exponential Time Hypothesis (ETH) and what does it imply?

- A. All problems require exponential time
- B. 3-SAT cannot be solved in $2^{o(n)}$ time, implying many NP-hard problems have no sub-exponential algorithms
- C. $P = NP$ is true
- D. Exponential algorithms are always impractical

Answer: B

SUBJECT

Operating Systems

1,080 Multiple Choice Questions

Operating Systems — Easy

360 questions

Q1. What is an operating system?

- A. System software that manages hardware and resources
- B. A physical component installed inside the computer
- C. A programming language for writing source code
- D. An application designed for end-user productivity

Answer: A

Q2. Which of the following is NOT a function of an operating system?

- A. Process management
- B. File system management
- C. Compiling user programs
- D. Memory management

Answer: C

Q3. Which of the following is an example of a mobile operating system?

- A. Android mobile OS
- B. Ubuntu desktop OS
- C. MS-DOS legacy OS
- D. Windows Server OS

Answer: A

Q4. What does the term 'booting' mean in the context of operating systems?

- A. Shutting down the computer safely
- B. Loading the OS into main memory
- C. Formatting the hard disk drive
- D. Installing new software packages

Answer: B

Q5. In a batch operating system, jobs are processed:

- A. Sequentially in submitted order
- B. Based on assigned priority
- C. Randomly without any order
- D. Simultaneously on all cores

Answer: A

Q6. Which type of operating system allows multiple users to use the computer simultaneously?

- A. Single-user OS
- B. Embedded OS
- C. Batch-job OS
- D. Time-sharing OS

Answer: D

Q7. The interface between the user and the operating system is called:

- A. The BIOS
- B. The API
- C. Shell
- D. Kernel

Answer: C

Q8. Which of the following is a real-time operating system characteristic?

- A. It does not support multitasking at all
- B. It only supports one single user at once
- C. It processes all jobs in large batches
- D. It guarantees response within strict time

Answer: D

Q9. UNIX was originally developed at:

- A. IBM Corporation
- B. Apple Computer
- C. AT&T Bell Labs
- D. Microsoft Research

Answer: C

Q10. The core component of an operating system is called the:

- A. Terminal
- B. Kernel
- C. Desktop
- D. Shell

Answer: B

Q11. Which of the following is a type of system call?

- A. Compiler optimization
- B. Syntax validation
- C. Process control
- D. Code refactoring

Answer: C

Q12. In a layered OS structure, the innermost layer is:

- A. Hardware layer
- B. Applications
- C. File systems
- D. User interface

Answer: A

Q13. Which of the following is an example of a command-line interpreter?

- A. Microsoft Word
- B. Bash shell
- C. Google Chrome
- D. Windows Explorer

Answer: B

Q14. System programs provide:

- A. Direct access to hardware resources
- B. Environment for program development
- C. Only file management service calls
- D. Gaming and multimedia functionality

Answer: B

Q15. The OS component that provides a graphical way to interact with the computer is called:

- A. SDK
- B. CLI
- C. API
- D. GUI

Answer: D

Q16. Which system call is used to create a new process in UNIX?

- A. exec()
- B. open()
- C. fork()
- D. wait()

Answer: C

Q17. A monolithic kernel is one where:

- A. The entire OS runs as one kernel program
- B. The OS is divided into separate modules
- C. Only the scheduler runs in kernel mode
- D. Each function runs as a separate process

Answer: A

Q18. Which of the following manages how data is stored and retrieved on a disk?

- A. Network manager
- B. File system
- C. Memory manager
- D. Process manager

Answer: B

Q19. What is the role of device drivers in an OS?

- A. To interface between the OS and hardware
- B. To manage user accounts and passwords
- C. To compile programs into machine code
- D. To execute user-level application programs

Answer: A

Q20. Linux is an example of which type of kernel?

- A. Monolithic kernel
- B. Nanokernel design
- C. Microkernel design
- D. Hybrid kernel

Answer: A

Q21. A process is:

- A. A single CPU register
- B. A program in execution
- C. A file held in memory
- D. A program stored on disk

Answer: B

Q22. Which of the following is NOT a state of a process?

- A. Compiling
- B. Ready
- C. Waiting
- D. Running

Answer: A

Q23. The Process Control Block (PCB) contains:

- A. Only the executable program code of the process
- B. Process state, program counter, CPU registers, etc
- C. Only the virtual and physical memory addresses
- D. Only the process name and its unique identifier

Answer: B

Q24. What system call creates a new process in UNIX/Linux?

- A. new()
- B. spawn()
- C. create()
- D. fork()

Answer: D

Q25. A context switch occurs when:

- A. A file is saved to secondary storage
- B. The user logs out of the session
- C. The CPU switches between processes
- D. A new program is installed on disk

Answer: C

Q26. An orphan process is:

- A. A process that has no memory allocated
- B. A child process whose parent has terminated
- C. A process that is waiting for I/O
- D. A process that is in the ready queue

Answer: B

Q27. A zombie process is:

- A. A process that is stuck inside an infinite loop
- B. A process running with very high execution priority
- C. A terminated process whose exit status is uncollected
- D. A process consuming excessive CPU cycles continuously

Answer: C

Q28. The program counter in a process stores:

- A. The address of the next instruction to execute
- B. The scheduling priority assigned to the process
- C. The total amount of memory currently in use
- D. The total count of programs currently running

Answer: A

Q29. Inter-process communication (IPC) allows:

- A. A process to access hardware directly
- B. A process to run faster on the CPU
- C. A process to compile source code files
- D. Processes to exchange data and synchronize

Answer: D

Q30. When a fork() system call is executed successfully, what does it return to the child process?

- A. Parent PID value
- B. Zero is returned
- C. Child PID value
- D. Negative one val

Answer: B

Q31. CPU scheduling decides:

- A. How much memory to allocate to each process
- B. Which ready process gets the CPU execution next
- C. Which files to open for the running processes
- D. Which device to use for the I/O operations

Answer: B

Q32. In First-Come, First-Served (FCFS) scheduling, processes are executed:

- A. In the order they arrive in the queue
- B. Based on their assigned priority level
- C. With the shortest job running first
- D. Randomly without any specific order

Answer: A

Q33. Turnaround time of a process is defined as:

- A. Total time from submission to completion of process
- B. Time the process spends waiting in the ready queue
- C. Time spent performing input and output operations
- D. Time the CPU is busy executing process instructions

Answer: A

Q34. Waiting time of a process is:

- A. The total time from arrival to completion
- B. The time spent waiting in the ready queue
- C. The time spent in I/O operations
- D. The time spent executing on the CPU

Answer: B

Q35. Round Robin scheduling uses:

- A. Random process allocation
- B. Longest job first execution
- C. A fixed time quantum per process
- D. Priority-based CPU allocation

Answer: C

Q36. Preemptive scheduling means:

- A. The OS can interrupt and reassign the CPU to another
- B. Processes schedule themselves without any OS control
- C. Only one process can ever run at a time on the CPU
- D. A process always runs until it completes execution

Answer: A

Q37. Non-preemptive scheduling means:

- A. The CPU can be taken away from a running process at any time
- B. Once a process starts, it runs until it finishes or blocks
- C. The scheduler never runs and the CPU sits idle indefinitely
- D. Multiple processes share the CPU at the same exact instant

Answer: B

Q38. Which scheduling algorithm gives the minimum average waiting time for a given set of processes?

- A. First-Come First-Served
- B. Priority Scheduling
- C. Shortest Job First SJF
- D. Round Robin scheduling

Answer: C

Q39. The dispatcher is responsible for:

- A. Selecting the next process to execute on CPU
- B. Managing main memory allocation for process
- C. Giving control of CPU to the selected process
- D. Handling file system operations for process

Answer: C

Q40. CPU utilization refers to:

- A. The number of processes completed per unit time
- B. The number of context switches performed per second
- C. The percentage of time the CPU is busy executing
- D. The total time a process waits in the ready queue

Answer: C

Q41. A thread is:

- A. A type of volatile memory block
- B. A lightweight execution unit in process
- C. A physical hardware component part
- D. A complete standalone process

Answer: B

Q42. Threads within the same process share all of the following EXCEPT:

- A. Data section
- B. Stack space
- C. Open files
- D. Code section

Answer: B

Q43. What is the main advantage of using threads over processes?

- A. Threads can run on different machines
- B. Threads are inherently more secure overall
- C. Thread creation and switching are faster
- D. Threads use significantly more memory

Answer: C

Q44. User-level threads are managed by:

- A. The file system layer
- B. A user-level thread library
- C. The hardware platform
- D. The kernel directly

Answer: B

Q45. Kernel-level threads are managed by:

- A. A web browser
- B. A compiler tool
- C. User applications
- D. The OS kernel

Answer: D

Q46. Multithreading benefits include all EXCEPT:

- A. Resource sharing across threads
- B. Responsiveness to user input
- C. Economy of creation and switching
- D. Guaranteed elimination of bugs

Answer: D

Q47. POSIX Pthreads is:

- A. A standard API for thread management
- B. A file system for multithreading
- C. A hardware component for threads
- D. An operating system distribution

Answer: A

Q48. In a single-threaded process, how many threads of control exist?

- A. One thread
- B. Many threads
- C. Two threads
- D. Zero threads

Answer: A

Q49. Which of the following is a benefit of multithreading in web server applications?

- A. It handles multiple client requests concurrently
- B. It reduces overall system security
- C. It simplifies the server source code logic
- D. It eliminates the need for any networking

Answer: A

Q50. A process can contain:

- A. One or more threads
- B. No threads at all
- C. Only two threads
- D. Only one thread

Answer: A

Q51. A race condition occurs when:

- A. Two processes have the same scheduling priority level
- B. Outcome depends on execution order of shared data access
- C. A process runs too fast for the system to handle
- D. Two processes run on the CPU simultaneously

Answer: B

Q52. The critical section is:

- A. The most important section of a running program
- B. Code segment accessing shared resources exclusively
- C. The first section executed when a program starts
- D. A section within the operating system kernel code

Answer: B

Q53. Mutual exclusion means:

- A. Processes must exclude each other from the system
- B. Two processes can be in critical sections at once
- C. No process is ever allowed to enter any section
- D. Only one process can be in its critical section

Answer: D

Q54. A mutex lock is:

- A. A hardware component on the motherboard
- B. A type of running process in the system
- C. A synchronization tool using acquire and release
- D. A type of physical memory module

Answer: C

Q55. A semaphore is:

- A. A type of hardware interrupt signal
- B. A CPU scheduling algorithm technique
- C. A type of boolean flag variable
- D. A variable accessed via wait and signal

Answer: D

Q56. A binary semaphore can have values:

- A. Only negative value range
- B. Any positive integer value
- C. Only zero and one values
- D. Any integer value at all

Answer: C

Q57. A counting semaphore can have values:

- A. Any non-negative integer value
- B. Only negative value ranges
- C. Only zero and one values
- D. Unrestricted integer domain

Answer: D

Q58. Busy waiting (spinlock) means:

- A. A process is waiting for I/O completion
- B. A process is terminated by the operating system
- C. A process continuously loops testing a condition
- D. A process is sleeping and waiting for a signal

Answer: C

Q59. What are the three requirements for a solution to the critical section problem?

- A. Atomicity, consistency, and isolation
- B. Fairness, priority, and preemption
- C. Mutual exclusion, progress, bounded waiting
- D. Speed, efficiency, and simplicity

Answer: C

Q60. The wait() operation on a semaphore is also known as:

- A. P operation (proberen)
- B. V operation (verhogen)
- C. Lock operation call
- D. Signal operation call

Answer: A

Q61. A deadlock occurs when:

- A. Processes wait indefinitely for resources held by each other
- B. A process enters an infinite loop consuming all CPU time
- C. The CPU is overloaded with too many running processes
- D. A process terminates unexpectedly due to an error

Answer: A

Q62. How many necessary conditions must hold simultaneously for a deadlock to occur?

- A. Three
- B. Four
- C. Two
- D. Five

Answer: B

Q63. Which of the following is NOT a necessary condition for deadlock?

- A. Mutual exclusion
- B. Hold and wait
- C. Circular wait
- D. Preemption

Answer: D

Q64. The 'hold and wait' condition means:

- A. A process waits without holding any resource
- B. A process holds no resources at all currently
- C. A process releases all resources before waiting
- D. A process holds resources while waiting for more

Answer: D

Q65. Circular wait means:

- A. Processes are scheduled in a circular fashion
- B. Each process waits for the next, forming a cycle
- C. Processes simply wait in a standard queue
- D. Only one single process is actively waiting

Answer: B

Q66. Deadlock prevention works by:

- A. Ignoring deadlocks entirely and hoping they pass
- B. Ensuring at least one necessary condition fails
- C. Recovering from deadlocks by terminating tasks
- D. Detecting deadlocks after they already occurred

Answer: B

Q67. A resource allocation graph is used to:

- A. Display the memory layout of processes
- B. Represent resource allocation and requests
- C. Show the CPU scheduling execution order
- D. Show the file system directory structure

Answer: B

Q68. If a resource allocation graph has no cycle, then:

- A. Deadlock definitely exists in system
- B. Deadlock may or may not exist here
- C. Deadlock definitely does not exist
- D. The resource allocation graph invalid

Answer: C

Q69. Deadlock avoidance requires:

- A. Advance info about resource requests
- B. Using only one resource at a time
- C. Killing processes periodically needed
- D. No additional information at all

Answer: A

Q70. The Banker's Algorithm is used for:

- A. Checking if granting request is safe
- B. Processing banking transactions safely
- C. Memory allocation for new processes
- D. CPU scheduling of running processes

Answer: A

Q71. Main memory (RAM) is:

- A. Volatile storage the CPU accesses directly
- B. A backup device for storing archives
- C. Non-volatile persistent storage on disk
- D. A type of secondary storage for backup

Answer: A

Q72. Logical address is:

- A. The network address assigned to a device
- B. An address generated by CPU for translation
- C. The physical location in hardware memory
- D. The address of data on the hard disk drive

Answer: B

Q73. Physical address is:

- A. An address generated by the CPU logic
- B. A virtual address used by programs
- C. An IP address on the local network
- D. The actual address in memory hardware

Answer: D

Q74. The Memory Management Unit (MMU) is responsible for:

- A. Managing I/O device controllers
- B. Scheduling processes on the CPU
- C. Translating logical to physical addresses
- D. Managing the file system on disk

Answer: C

Q75. Paging divides:

- A. The disk into sectors and track clusters
- B. Processes into variable-size logical segments
- C. Physical memory into frames, logical into pages
- D. The CPU into separate processing partitions

Answer: C

Q76. A page table stores:

- A. CPU register contents
- B. File names and paths
- C. Process priority values
- D. Page to frame mappings

Answer: D

Q77. External fragmentation occurs when:

- A. The disk is completely full and cannot store data
- B. The allocated pages are much too large for use
- C. Internal memory within blocks is wasted by padding
- D. Free memory is scattered, cannot satisfy requests

Answer: D

Q78. Internal fragmentation occurs when:

- A. No memory is available to allocate to processes
- B. The file system is corrupted and cannot allocate blocks
- C. Allocated block is larger than requested, wasting space
- D. Free memory is scattered in small pieces across RAM

Answer: C

Q79. Virtual memory allows:

- A. Only one process to run at any given time
- B. Programs larger than physical memory to execute
- C. Faster CPU execution by clock overclocking
- D. Direct access to underlying hardware devices

Answer: B

Q80. A page fault occurs when:

- A. A referenced page is not in physical memory
- B. A page has become corrupted in memory
- C. The page table is completely full of entries
- D. A page is too large to fit in any frame

Answer: A

Q81. A file is:

- A. A physical storage device like a hard drive
- B. A hardware component on the computer motherboard
- C. A type of process running in the operating system
- D. A named collection of data on secondary storage

Answer: D

Q82. A directory in a file system:

- A. Is the same as a disk partition or volume
- B. Contains file info like names and locations
- C. Is a file type that stores only text content
- D. Stores raw data directly on the disk surface

Answer: B

Q83. Which file access method reads records in order, one after another?

- A. Random access
- B. Indexed access
- C. Sequential access
- D. Direct access

Answer: C

Q84. The absolute path of a file starts from:

- A. The root directory
- B. The home directory
- C. The current directory
- D. The parent directory

Answer: A

Q85. File attributes typically include:

- A. Name, type, size, location, timestamps
- B. Only the file creation date stored
- C. Only the file size in bytes stored
- D. Only the file name string value

Answer: A

Q86. An inode in UNIX/Linux file systems stores:

- A. Metadata: permissions, size, pointers
- B. The actual file content data itself
- C. The file name string value
- D. The directory name for this entry

Answer: A

Q87. A hard link in UNIX:

- A. Creates a complete copy of the original file
- B. Links to a file located on another disk drive
- C. Creates a shortcut file with a different path
- D. Creates another entry pointing to same inode

Answer: D

Q88. A symbolic (soft) link is:

- A. A compressed version of original file
- B. A file containing path to another file
- C. A full copy of the original source file
- D. The same as a hard link to a file

Answer: B

Q89. File system mounting is the process of:

- A. Deleting a file system from a partition
- B. Compressing a file system for archiving
- C. Creating a new file system on disk
- D. Making a file system accessible at a point

Answer: D

Q90. FAT (File Allocation Table) file system uses:

- A. Inodes for each file entry
- B. A table of linked block lists
- C. Hash tables for fast lookup
- D. B-trees for indexing entries

Answer: B

Q91. Secondary storage refers to:

- A. Processor register files
- B. Non-volatile storage like disks
- C. Random access memory chips
- D. CPU cache memory modules

Answer: B

Q92. A hard disk drive (HDD) stores data on:

- A. Flash memory components
- B. Semiconductor memory chips
- C. Rotating magnetic platters
- D. Optical disc technology

Answer: C

Q93. Seek time in a disk drive is:

- A. Time for the disk to spin to correct sector
- B. Time to transfer requested data from disk
- C. Time for the head to move to correct track
- D. Time to process and queue the I/O request

Answer: C

Q94. Rotational latency is:

- A. Time to start the disk motor from stopped
- B. Time for desired sector to rotate under head
- C. Time for the head to move to correct track
- D. Time to format a disk sector with structure

Answer: B

Q95. A Solid State Drive (SSD) stores data using:

- A. Magnetic tape reels
- B. Magnetic platters
- C. Flash memory chips
- D. Optical technology

Answer: C

Q96. FCFS disk scheduling processes disk I/O requests:

- A. Based on closest track first
- B. In the order they arrive
- C. In reverse arrival order
- D. Based on request priority

Answer: B

Q97. The SSTF (Shortest Seek Time First) disk scheduling algorithm:

- A. Services the closest request to head
- B. Always moves the head toward outer
- C. Services the largest request first
- D. Services all requests in arrival order

Answer: A

Q98. Disk formatting prepares a disk for use by:

- A. Deleting all viruses from the disk
- B. Creating file system structure on disk
- C. Increasing the disk rotational speed
- D. Adding more storage space to a disk

Answer: B

Q99. A disk partition is:

- A. A physical disk drive unit in chassis
- B. A type of data file stored on the disk
- C. A backup copy of the entire disk drive
- D. An independent section with own file system

Answer: D

Q100. Disk access time is the sum of:

- A. Only the seek time component alone
- B. Only the data transfer time component
- C. Only the rotational latency component
- D. Seek time, rotational latency, transfer time

Answer: D

Q101. I/O devices communicate with the CPU through:

- A. The file system interface
- B. Direct wiring to CPU
- C. Shared memory regions only
- D. Device controllers and buses

Answer: D

Q102. Polling in I/O means:

- A. CPU repeatedly checks device status register
- B. Voting on which device to use for the task
- C. Automatically detecting new devices on the bus
- D. Prioritizing I/O requests based on their urgency

Answer: A

Q103. Interrupt-driven I/O means:

- A. I/O operations are interrupted frequently
- B. The CPU is constantly being interrupted
- C. The CPU stops all I/O until processing done
- D. The device signals CPU when I/O is complete

Answer: D

Q104. DMA (Direct Memory Access) allows:

- A. Direct access to the motherboard components
- B. Devices to transfer data to memory without CPU
- C. The CPU to access main memory much faster
- D. Memory to access the CPU directly for data

Answer: B

Q105. A device driver is:

- A. Software providing OS-to-hardware interface
- B. A type of cable connecting to the bus
- C. A person who physically operates devices
- D. A diagnostic tool for testing hardware

Answer: A

Q106. Block devices transfer data in:

- A. Variable packets
- B. Single characters
- C. Continuous streams
- D. Fixed-size blocks

Answer: D

Q107. Character devices transfer data:

- A. One byte at a time in a stream
- B. In fixed-size data blocks
- C. In variable-size network packets
- D. In fixed-size memory page frames

Answer: A

Q108. A bus in computer architecture is:

- A. A type of vehicle for transport use
- B. A type of volatile memory used in computers
- C. A processing unit that executes instructions
- D. A shared pathway transferring data between parts

Answer: D

Q109. Buffering in I/O systems means:

- A. Encrypting I/O data for security purposes
- B. Compressing I/O data for disk storage
- C. Temporarily storing data during transfer
- D. Slowing down I/O operations intentionally

Answer: C

Q110. Spooling is commonly used for:

- A. CPU scheduling for processes
- B. Network routing of data packets
- C. Memory management of processes
- D. Managing print jobs via spooling

Answer: D

Q111. Protection in operating systems refers to:

- A. Controlling process and user access to resources
- B. Antivirus software scanning for malware threats
- C. Firewall configuration for network protection
- D. Physical security of the computer hardware

Answer: A

Q112. Security in operating systems deals with:

- A. Only network security from remote attacks
- B. Defending against internal and external attacks
- C. Only internal threats from local users
- D. Only physical security of the server room

Answer: B

Q113. Authentication is the process of:

- A. Verifying identity of a user
- B. Encrypting data for security
- C. Authorizing resource access
- D. Backing up critical data

Answer: A

Q114. Authorization determines:

- A. When to back up data for disaster recovery
- B. How to encrypt data for secure communication
- C. Who the user is based on their credentials
- D. What resources a verified user may access

Answer: D

Q115. The principle of least privilege states that:

- A. All users should have full administrator access rights
- B. The least privileged user should have no access at all
- C. Users get only the minimum access needed for tasks
- D. All privileges should be equal for all system users

Answer: C

Q116. A virus in computer security is:

- A. Self-replicating malicious program on host
- B. A network communication protocol standard
- C. A hardware malfunction in the CPU unit
- D. A type of data encryption algorithm used

Answer: A

Q117. A worm differs from a virus in that:

- A. A worm only affects the physical hardware parts
- B. A worm replicates across networks without host
- C. A worm is less harmful than a virus overall
- D. A worm cannot spread to other network systems

Answer: B

Q118. A Trojan horse in computer security is:

- A. A system backup utility tool for data
- B. A type of network firewall tool used
- C. A type of hardware peripheral device
- D. Malicious program disguised as legitimate

Answer: D

Q119. Encryption is the process of:

- A. Converting plaintext to ciphertext with key
- B. Backing up data to a remote server
- C. Deleting data permanently from disk
- D. Compressing data for efficient storage

Answer: A

Q120. A firewall is:

- A. A system monitoring and controlling network traffic
- B. A physical wall in the server room for safety
- C. A type of computer virus that spreads via network
- D. A backup system for disaster recovery procedures

Answer: A

Q121. Which component loads the operating system into memory during startup?

- A. The application manager
- B. The device controller
- C. The bootstrap loader
- D. The system compiler

Answer: C

Q122. What does a distributed operating system manage?

- A. Multiple networked computers together
- B. A single standalone desktop machine
- C. Only mobile phone applications here
- D. Exclusively embedded device controllers

Answer: A

Q123. Which of the following best describes firmware?

- A. Programs permanently stored in hardware
- B. Software stored in volatile RAM chips
- C. Temporary data written to disk drives
- D. User applications installed externally

Answer: A

Q124. What is the primary purpose of an interrupt in a computer system?

- A. To delete all currently running tasks
- B. To signal the CPU about an event
- C. To format all storage device media
- D. To permanently halt all processors

Answer: B

Q125. Which type of OS is designed for devices with very limited resources?

- A. Embedded system OS
- B. Desktop general OS
- C. Mainframe batch OS
- D. Enterprise server OS

Answer: A

Q126. What does GUI stand for in the context of operating systems?

- A. Graphical User Interface
- B. General Unified Interface
- C. Global Usage Indicator
- D. Generated Utility Index

Answer: A

Q127. Which OS function manages the allocation and deallocation of main memory?

- A. Network management
- B. User authentication
- C. Disk formatting
- D. Memory management

Answer: D

Q128. What is spooling in an operating system context?

- A. Buffering data for slow output devices
- B. Compressing files to save disk storage
- C. Defragmenting the hard disk sectors
- D. Encrypting data before transmitting out

Answer: A

Q129. Which of the following is NOT a type of operating system?

- A. Real-time operating system
- B. Network operating system
- C. Distributed operating system
- D. Compiled operating system

Answer: D

Q130. What is the role of device drivers in an operating system?

- A. They enable OS-hardware communication
- B. They translate high-level languages
- C. They manage user account passwords
- D. They compile user application code

Answer: A

Q131. What is a monolithic kernel operating system?

- A. An OS where all services run in kernel space
- B. An OS that only runs a single application task
- C. An OS designed for mobile devices exclusively
- D. An OS with no kernel component present at all

Answer: A

Q132. Which of the following is an example of a system call category?

- A. Process control
- B. Web browsing
- C. Graphics rendering
- D. Text formatting

Answer: A

Q133. What is the purpose of an API in an operating system?

- A. It manages physical hardware components alone
- B. It replaces the kernel entirely in modern OS
- C. It provides functions for programmers to use
- D. It handles only network communication protocol

Answer: C

Q134. What does the command interpreter do in an operating system?

- A. It controls hardware interrupt processing now
- B. It compiles all programs into machine code
- C. It manages virtual memory page tables only
- D. It reads and executes user commands directly

Answer: D

Q135. Which layer is at the bottom in a layered operating system design?

- A. File system layer
- B. Hardware layer
- C. Application layer
- D. User interface layer

Answer: B

Q136. What is a system program in the context of an operating system?

- A. A utility program bundled with the OS itself
- B. A virus that attacks the operating system now
- C. A program that only runs during system boot
- D. A diagnostic tool built into the CPU hardware

Answer: A

Q137. What type of system call is used to create a new process?

- A. Communication call
- B. File management call
- C. Device management call
- D. Process control call

Answer: D

Q138. What is the function of a linker in program execution?

- A. It translates source code to assembly code
- B. It handles input and output device operation
- C. It manages memory allocation during runtime
- D. It combines object modules into executable

Answer: D

Q139. Which of the following is a communication system call?

- A. open() call
- B. exit() call
- C. pipe() call
- D. fork() call

Answer: C

Q140. What is the role of a loader in program execution?

- A. It manages the network connections system
- B. It designs the user interface for programs
- C. It writes source code for the programmer
- D. It loads executable files into main memory

Answer: D

Q141. What is a process in an operating system?

- A. A program in execution in main memory
- B. A type of network protocol standard
- C. A hardware component inside the CPU
- D. A file stored on the hard disk drive

Answer: A

Q142. Which of the following is a valid process state?

- A. Formatting
- B. Installing
- C. Waiting
- D. Compiling

Answer: C

Q143. What does the Process Control Block contain?

- A. All user passwords and credentials
- B. Only the program source code text
- C. Information about a specific process
- D. The entire operating system kernel

Answer: C

Q144. What system call creates a new process in Unix?

- A. exec()
- B. fork()
- C. kill()
- D. wait()

Answer: B

Q145. What is a parent process in operating system terminology?

- A. The first process created during boot only
- B. A process with the highest priority in queue
- C. A process that created one or more children
- D. A process managing all hardware device drivers

Answer: C

Q146. What happens when a process moves from running to waiting state?

- A. The process begins on a different CPU core
- B. The process is permanently terminated now
- C. The process waits for an event like I/O
- D. The process moves back to the ready queue

Answer: C

Q147. What is the ready queue in process management?

- A. Processes that are ready and waiting for CPU
- B. A list of processes waiting to load from disk
- C. Blocked processes waiting for user keyboard input
- D. A queue of completed processes to be deleted

Answer: A

Q148. What is the purpose of the process scheduler?

- A. To handle network packet routing operations
- B. To manage file system directory structures now
- C. To select a process from ready queue for CPU
- D. To compile programs into executable files

Answer: C

Q149. What is process termination in operating systems?

- A. Pausing a process temporarily for later
- B. Ending a process and releasing its resources
- C. Changing the priority of a running process
- D. Moving a process to another computer node

Answer: B

Q150. What is an orphan process in Unix/Linux systems?

- A. A process consuming no memory or CPU at all
- B. A process whose parent has terminated first
- C. A process running only during system shutdown
- D. A process that has no executable code to run

Answer: B

Q151. What is CPU scheduling in an operating system?

- A. Overclocking the CPU for better processing performance
- B. Installing new processors into the computer motherboard
- C. Deciding which process gets the CPU next for execution
- D. Measuring the temperature of the CPU chip continuously

Answer: C

Q152. What does FCFS stand for in CPU scheduling?

- A. First Cache First Store
- B. Fast CPU Flexible Scheduling
- C. Fully Controlled File System
- D. First Come First Served

Answer: D

Q153. What is turnaround time in CPU scheduling?

- A. Time to switch between two processes
- B. Total time from submission to completion
- C. The speed at which the CPU clock runs
- D. The time to restart the computer

Answer: B

Q154. What is waiting time in CPU scheduling?

- A. Time needed to compile a program from source
- B. Time required to boot the operating system
- C. Time a process spends in the ready queue
- D. Time taken to save a file to the disk drive

Answer: C

Q155. What is the time quantum in Round Robin scheduling?

- A. The time required for a context switch now
- B. The clock speed of the processor chip now
- C. The total time to complete all processes
- D. A fixed time slice given to each process

Answer: D

Q156. Which scheduling algorithm selects the shortest job first?

- A. First Come First Served
- B. Round Robin scheduling
- C. Shortest Job First
- D. Priority scheduling only

Answer: C

Q157. What is throughput in CPU scheduling metrics?

- A. The speed of data transfer to the disk drive
- B. Number of processes completed per time unit
- C. The temperature of CPU during operation
- D. The size of the ready queue at a given moment

Answer: B

Q158. What is a preemptive scheduling algorithm?

- A. One where OS can take CPU from a process
- B. One that only works with a single process
- C. One that requires manual user intervention
- D. One that never interrupts a running process

Answer: A

Q159. What is a CPU burst in process execution?

- A. A period of continuous CPU computation time
- B. A malfunction causing the CPU to overheat now
- C. A type of interrupt from external hardware now
- D. A period when the CPU is powered off safely

Answer: A

Q160. In priority scheduling, which process gets the CPU first?

- A. The process with lowest priority value
- B. The process with highest priority value
- C. The process arriving most recently now
- D. The process with longest burst time

Answer: B

Q161. What is a thread in operating system terminology?

- A. A network connection between two machines
- B. A type of file stored on the hard disk
- C. A lightweight execution unit in a process
- D. A separate physical CPU core in system

Answer: C

Q162. What do threads within the same process share?

- A. Nothing at all between them
- B. Individual stack memory only
- C. Code, data, and file resources
- D. Only the process identifier

Answer: C

Q163. What is multithreading in operating systems?

- A. Running multiple threads within a single process
- B. Running multiple operating systems simultaneously
- C. Using multiple monitors for display output only
- D. Connecting multiple keyboards to one computer

Answer: A

Q164. Which of the following is a benefit of using threads?

- A. Threads prevent all forms of parallel execution
- B. Thread creation is faster than process creation
- C. Threads consume more resources than processes
- D. Threads restrict programs to a single CPU core

Answer: B

Q165. What is a user-level thread managed by?

- A. The operating system kernel directly
- B. A user-level thread library only
- C. The network interface card alone
- D. The hardware CPU scheduler chip

Answer: B

Q166. What is a kernel-level thread?

- A. A thread that manages hardware device interrupts
- B. A thread that only runs during system boot time
- C. A thread managed and scheduled by the kernel
- D. A thread that exists only in embedded OS systems

Answer: C

Q167. What does each thread have that is unique to itself?

- A. Its own copy of all open file handles now
- B. Its own operating system kernel instance
- C. Its own stack and program counter value
- D. Its own copy of entire process memory

Answer: C

Q168. What is the main advantage of multithreading in a web server?

- A. It prevents any client from accessing the server
- B. It reduces the server network bandwidth usage
- C. It eliminates the need for any network protocol
- D. It handles multiple client requests concurrently

Answer: D

Q169. What is the Pthreads API?

- A. A POSIX standard for thread management
- B. A network protocol for file transferring
- C. A graphics library for rendering images
- D. A database query language for data access

Answer: A

Q170. What happens when one thread in a process calls exit()?

- A. Other threads continue running without any impact
- B. Only that specific thread is terminated safely
- C. The entire process and all its threads terminate
- D. The thread is moved to a different process now

Answer: C

Q171. What is a race condition in concurrent programming?

- A. A technique for sorting data in parallel
- B. A way to synchronize clocks between PCs
- C. When output depends on timing of threads
- D. A method to speed up program execution

Answer: C

Q172. What is mutual exclusion in process synchronization?

- A. Allowing multiple processes to access resources
- B. Preventing any process from executing on CPU
- C. Deleting shared resources to prevent any access
- D. Ensuring only one process in critical section now

Answer: D

Q173. What is a critical section in concurrent programming?

- A. The first section of code executed by program
- B. The most important function in program code
- C. A section of code needing no synchronization
- D. A code segment accessing shared resources here

Answer: D

Q174. What is a semaphore in operating systems?

- A. A synchronization variable with atomic ops
- B. A memory allocation technique for processes
- C. A scheduling algorithm for real-time tasks
- D. A type of hardware interrupt signal line

Answer: A

Q175. What is the purpose of a mutex lock?

- A. To handle network packet routing faster
- B. To speed up single-thread execution
- C. To ensure mutual exclusion in critical section
- D. To manage memory allocation for processes

Answer: C

Q176. What does the wait() operation do on a semaphore?

- A. It creates a new semaphore with given value
- B. It decrements value and blocks if negative
- C. It increments the semaphore value by one
- D. It destroys the semaphore variable for good

Answer: B

Q177. What does signal() do on a semaphore?

- A. It increments value and may wake a process
- B. It deletes the semaphore from the system
- C. It resets semaphore to its initial value now
- D. It decrements semaphore value by one unit

Answer: A

Q178. What is a binary semaphore?

- A. A semaphore requiring two processes only
- B. A semaphore restricted to zero or one
- C. A semaphore holding any integer value
- D. A semaphore for binary search algorithms

Answer: B

Q179. What is busy waiting in the context of synchronization?

- A. A thread sleeps until explicitly woken up
- B. A thread exits when a resource is locked
- C. A thread continuously checks a condition
- D. A thread creates more threads to help out

Answer: C

Q180. What is a counting semaphore used for?

- A. Measuring CPU clock speed in gigahertz
- B. Counting total processes in the OS system
- C. Controlling access to finite resource instances
- D. Counting the number of files on the disk

Answer: C

Q181. What is a deadlock in an operating system?

- A. When processes wait for each other forever
- B. When memory is fully available to all tasks
- C. When the CPU overheats and shuts down
- D. When a program runs faster than expected

Answer: A

Q182. How many conditions must hold simultaneously for deadlock?

- A. One
- B. Four
- C. Two
- D. Three

Answer: B

Q183. What is the mutual exclusion condition for deadlock?

- A. Processes can preempt resources from other ones
- B. Resources can be shared by multiple processes
- C. All resources are always available to every process
- D. At least one resource must be held non-sharably

Answer: D

Q184. What is the hold and wait condition for deadlock?

- A. Processes never hold more than one resource at a time
- B. Processes wait without holding any resources at all
- C. Processes hold resources while waiting for others
- D. Processes release resources before requesting new ones

Answer: C

Q185. What is a resource allocation graph used for?

- A. Designing the layout of network infrastructure
- B. Measuring CPU utilization over time periods
- C. Drawing flowcharts for program logic design
- D. Visualizing allocation and detecting deadlock

Answer: D

Q186. What is the circular wait condition for deadlock?

- A. Processes wait for resources in random order
- B. A circular chain where each waits for the next
- C. Processes wait in a linear queue for resources
- D. All processes wait for the same single resource

Answer: B

Q187. What is the no preemption condition for deadlock?

- A. Only the OS kernel can hold resources permanently
- B. Resources cannot be preempted from holders now
- C. Resources can be forcibly taken from processes
- D. All resources are preemptible at any time freely

Answer: B

Q188. What is deadlock prevention as a strategy?

- A. Recovering by terminating all running processes
- B. Ignoring deadlocks and hoping they never occur
- C. Detecting deadlocks after they already occurred
- D. Ensuring at least one deadlock condition fails

Answer: D

Q189. What is starvation in resource allocation?

- A. A process runs out of memory and crashes fast
- B. A process is denied needed resources indefinitely
- C. A process completes too quickly and wastes resources
- D. A process gets all resources it needs quickly

Answer: B

Q190. What is a safe state in deadlock avoidance?

- A. A state where resources can be allocated avoiding deadlock
- B. A state where all resources are fully allocated already
- C. A state where the system has no processes running
- D. A state where CPU temperature is within safe limits

Answer: A

Q191. What is the main purpose of memory management?

- A. To permanently store data like a hard drive
- B. To increase the physical size of RAM chips
- C. To manage network bandwidth between hosts
- D. To allocate and track memory for processes

Answer: D

Q192. What is virtual memory in an operating system?

- A. A type of cache memory inside the CPU chip
- B. A technique giving illusion of larger memory
- C. Physical RAM installed in the computer
- D. A backup copy of files stored on the disk

Answer: B

Q193. What is a page in the context of virtual memory?

- A. A type of hardware interrupt signal to CPU
- B. A section of the operating system kernel
- C. A unit of data stored on the hard disk now
- D. A fixed-size block of virtual memory space

Answer: D

Q194. What is a page table used for?

- A. Displaying a table of contents for docs
- B. Storing user passwords and credentials now
- C. Mapping virtual addresses to physical ones
- D. Scheduling CPU time for running processes

Answer: C

Q195. What is a page fault in virtual memory systems?

- A. A syntax error in the running program code
- B. An access to a page not in main memory now
- C. A hardware malfunction in memory module
- D. A failure of CPU to execute an instruction

Answer: B

Q196. What is a frame in physical memory?

- A. A fixed-size block of physical RAM memory
- B. A unit of data transmitted over a network
- C. A type of CPU register storing addresses
- D. A decorative border around monitor screen

Answer: A

Q197. What is swapping in memory management?

- A. Exchanging CPU cores between processes
- B. Replacing old hardware with new hardware
- C. Switching between different user accounts
- D. Moving process between memory and disk

Answer: D

Q198. What is internal fragmentation in memory?

- A. Memory corrupted due to hardware failure
- B. File fragments scattered across the disk
- C. Free memory between allocated blocks here
- D. Wasted space inside an allocated partition

Answer: D

Q199. What is external fragmentation in memory?

- A. A process using more than available memory
- B. Memory corruption from external hardware
- C. Free memory scattered in small pieces now
- D. Memory allocated to already exited processes

Answer: C

Q200. What is the logical address space of a process?

- A. Address of the process control block here
- B. Set of all addresses generated by process
- C. Total amount of RAM installed in computer
- D. Physical location of process on the disk

Answer: B

Q201. What is a file in operating system terminology?

- A. A network protocol for data transfer
- B. A physical storage device like a disk
- C. A type of CPU instruction for processing
- D. A named collection of related data items

Answer: D

Q202. What is a directory in a file system?

- A. A hardware device for reading media
- B. A physical partition on the hard disk
- C. A structure containing info about files
- D. A file type that cannot be modified now

Answer: C

Q203. What is a file extension used for?

- A. To indicate the type or format of a file
- B. To increase the file size beyond limits
- C. To encrypt the file contents securely now
- D. To compress files for saving disk space

Answer: A

Q204. What is the root directory in a file system?

- A. Where all temporary files are stored
- B. A directory only for administrators here
- C. A directory storing the OS kernel files
- D. The topmost directory in the hierarchy

Answer: D

Q205. What is a file's access permission used for?

- A. Determining the file size on the disk
- B. Setting the creation date of the entry
- C. Controlling who can read, write, execute
- D. Choosing which app opens the file type

Answer: C

Q206. What is sequential file access?

- A. Writing files in reverse order from the end
- B. Accessing files from multiple disks at once
- C. Accessing any record directly by position
- D. Reading records one after another in order

Answer: D

Q207. What is the purpose of file system mounting?

- A. Creating a backup copy of the file system
- B. Making a file system accessible at a location
- C. Physically installing new hard drive hardware
- D. Removing a file system from the OS entirely

Answer: B

Q208. What does the open() system call do for files?

- A. It permanently deletes a file from system
- B. It creates a brand new empty file on disk
- C. It prepares a file for reading or writing
- D. It changes the permissions of a file now

Answer: C

Q209. What is an absolute path in a file system?

- A. A path containing only the file name itself
- B. A path relative to current working directory
- C. A complete path starting from root directory
- D. A path that points to a deleted file now

Answer: C

Q210. What is a symbolic link in a file system?

- A. A file containing a path to another file
- B. A physical copy of a file in another place
- C. A special file storing encrypted disk data
- D. A temporary file created during system boot

Answer: A

Q211. What is secondary storage in a computer system?

- A. Non-volatile storage for permanent data
- B. The main memory used by running procs
- C. The CPU cache memory for fast access
- D. Registers inside the CPU for computation

Answer: A

Q212. What is a hard disk drive composed of?

- A. Optical discs using laser beams to read the data
- B. Solid-state flash memory chips sealed in case
- C. Spinning magnetic platters with read-write heads
- D. Volatile memory modules losing data when off now

Answer: C

Q213. What is disk seek time?

- A. Time to transfer data from disk to memory
- B. Time to format the entire disk from scratch
- C. Time to move head to the correct track now
- D. Time for platter to complete one rotation

Answer: C

Q214. What is a solid-state drive (SSD)?

- A. A mechanical drive with spinning platters
- B. A volatile device requiring constant power now
- C. Storage using flash memory, no moving parts
- D. A type of optical disc using laser for data

Answer: C

Q215. What is disk rotational latency?

- A. Total time to read an entire disk platter now
- B. Time for head to move to the correct track
- C. Time for desired sector to rotate under head
- D. Time required to start spinning disk platters

Answer: C

Q216. What is a disk partition?

- A. A hardware device connecting disks to CPU
- B. A logical division of disk into sections
- C. A type of disk error corrupting stored data
- D. A physical break in disk platter surface

Answer: B

Q217. What is the transfer rate of a disk?

- A. Speed at which disk platters rotate per minute
- B. Time taken to partition the disk into sections
- C. Rate of data flow between disk and computer now
- D. Rate at which disk head moves across all tracks

Answer: C

Q218. What is a disk track on a magnetic hard drive?

- A. A type of error correction code on the disk
- B. A line connecting disk to the motherboard now
- C. A physical groove cut into disk for head path
- D. A concentric circle on disk platter surface

Answer: D

Q219. What is a sector on a hard disk drive?

- A. Total storage capacity of one disk platter now
- B. A complete circular track on the platter now
- C. Smallest addressable unit of storage on disk
- D. Hardware component controlling the read head

Answer: C

Q220. What is RAID in secondary storage?

- A. A method to encrypt data on single disk
- B. A type of RAM for caching disk operations
- C. A protocol for transferring data on network
- D. Combining multiple disks for redundancy

Answer: D

Q221. What is an I/O device in a computer system?

- A. A device that only stores the OS kernel
- B. A device generating electrical power supply
- C. A device used for input, output, or both
- D. A device that only processes internal data

Answer: C

Q222. What is a device driver in an operating system?

- A. Software controlling a specific hardware device
- B. A person who operates hardware devices
- C. A physical cable connecting devices to the PC
- D. A tool for testing hardware functionality now

Answer: A

Q223. What is programmed I/O for device communication?

- A. Data transfers directly between device and memory
- B. OS ignores device until a user requests the data
- C. CPU continuously checks device status for data
- D. Device notifies CPU when it completes operation

Answer: C

Q224. What is interrupt-driven I/O?

- A. CPU continuously polls device for data status
- B. OS disables all interrupts during I/O operations
- C. Data transfers without any CPU involvement here
- D. Device signals CPU via interrupt when it is ready

Answer: D

Q225. What is Direct Memory Access used for?

- A. Transferring data between I/O devices and memory
- B. Directly accessing CPU registers from a program
- C. Providing direct access to the OS kernel code now
- D. Managing memory allocation for running processes

Answer: A

Q226. What is a block device in I/O classification?

- A. A device transferring data one character at a time
- B. A device that blocks all other devices from running
- C. A device transferring data in fixed-size blocks now
- D. A device working only when the system is blocked

Answer: C

Q227. What is a character device in I/O classification?

- A. A device storing character encoding tables on disk
- B. A device displaying only text characters on screen
- C. A device transferring data one character at a time
- D. A device used for character recognition in documents

Answer: C

Q228. What is the purpose of an I/O port?

- A. An address for CPU to communicate with devices
- B. A physical hole for ventilation of the case
- C. A port used only for internet connections now
- D. A type of software app for managing printers

Answer: A

Q229. What is buffering in I/O systems?

- A. Encrypting data before sending to output devices
- B. Temporarily storing data in transit between devices
- C. Permanently storing data on the hard disk drive now
- D. Deleting data no longer needed by any process here

Answer: B

Q230. What is the I/O request lifecycle?

- A. Time from device manufacture to its disposal
- B. Lifespan of cable connecting device to computer
- C. Sequence from application request to I/O completion
- D. Period when I/O device is powered on and running

Answer: C

Q231. What is the difference between protection and security?

- A. Protection controls internal; security guards external
- B. Neither is relevant to modern operating systems
- C. Security controls internal; protection guards external
- D. Protection and security mean the same thing

Answer: A

Q232. What is authentication in OS security?

- A. Installing antivirus software on the system
- B. Compressing data to save storage on disk
- C. Verifying the identity of a user or process
- D. Encrypting all files stored on the disk

Answer: C

Q233. What is an access control list in operating systems?

- A. A list specifying user permissions for objects
- B. A list of all network connections currently
- C. A list of all hardware devices attached now
- D. A list of all software installed on PC

Answer: A

Q234. What is a password used for in OS security?

- A. Encrypting the entire hard disk contents
- B. Compressing files to save storage on disk
- C. Increasing the processing speed of the CPU
- D. Authenticating user identity before access

Answer: D

Q235. What is malware in computer security?

- A. A type of hardware failure in the computer
- B. Malicious software designed to harm systems
- C. Software improving computer performance
- D. A legitimate tool for administration tasks

Answer: B

Q236. What is encryption used for in OS security?

- A. Converting data to unreadable format for privacy
- B. Deleting sensitive files permanently from the disk
- C. Compressing files to reduce their storage size now
- D. Increasing the speed of data transfer on network

Answer: A

Q237. What is a firewall in computer security?

- A. A tool for formatting hard drives securely now
- B. A device supplying emergency power to the PC
- C. A physical wall preventing computer overheating
- D. A system filtering network traffic based on rules

Answer: D

Q238. What is a computer virus?

- A. A tool for legitimate system maintenance tasks
- B. A type of operating system for mobile phones
- C. A self-replicating program attaching to files
- D. A hardware malfunction from overheating chips

Answer: C

Q239. What is the principle of least privilege?

- A. Granting users only minimum permissions needed
- B. Assigning same permissions to every system user
- C. Giving every user full administrator access
- D. Preventing all users from accessing the system

Answer: A

Q240. What is a Trojan horse in computer security?

- A. A security tool protecting against all viruses
- B. A hardware device used for network security
- C. A backup system storing copies of vital files
- D. Malware disguised as legitimate software here

Answer: D

Q241. Which of the following is a primary goal of an operating system?

- A. Convenience for the user
- B. Designing hardware circuits
- C. Writing application software
- D. Manufacturing peripheral devices

Answer: A

Q242. What is a batch operating system?

- A. An OS that processes jobs in groups without user interaction
- B. An OS that allows real-time user input during execution
- C. An OS designed only for mobile devices
- D. An OS that runs a single program at a time with user control

Answer: A

Q243. Which of the following best describes a time-sharing operating system?

- A. An OS that allocates dedicated hardware to each user
- B. An OS that allows multiple users to share CPU time interactively
- C. An OS that processes only one task at a time
- D. An OS used exclusively for scientific calculations

Answer: B

Q244. What does the term 'kernel mode' mean in an operating system?

- A. A privileged execution mode with full hardware access
- B. A mode where only user applications can run
- C. A power-saving mode that reduces CPU speed
- D. A mode used exclusively during system shutdown

Answer: A

Q245. Which component of a computer system directly executes instructions?

- A. Central Processing Unit (CPU)
- B. Hard disk drive
- C. Network interface card
- D. Power supply unit

Answer: A

Q246. What is multiprogramming in operating systems?

- A. Running multiple programs in memory to maximize CPU utilization
- B. Running the same program on multiple computers
- C. Writing programs in multiple programming languages
- D. Using multiple monitors for display output

Answer: A

Q247. What is the function of an operating system's resource manager?

- A. Allocating and managing hardware and software resources
- B. Designing new hardware components
- C. Writing user application programs
- D. Connecting to the internet

Answer: A

Q248. Which of the following is a network operating system?

- A. Windows Server
- B. MS-DOS
- C. Android
- D. Embedded Linux

Answer: A

Q249. What is the purpose of a command-line interface in an operating system?

- A. To allow users to interact with the OS by typing text commands
- B. To display graphical windows and icons
- C. To manage the file allocation table
- D. To compile source code into machine language

Answer: A

Q250. What is an embedded operating system?

- A. An OS designed for specialized devices with limited resources
- B. An OS used only on desktop computers
- C. An OS that requires a graphical user interface
- D. An OS that runs only in cloud environments

Answer: A

Q251. What is a system call used for in an operating system?

- A. Requesting services from the OS kernel
- B. Directly accessing hardware without the OS
- C. Compiling programs into machine code
- D. Formatting a hard disk drive

Answer: A

Q252. What is a shell in an operating system?

- A. A command interpreter that accepts and executes user commands
- B. The innermost part of the operating system kernel
- C. A hardware component on the motherboard
- D. A type of system memory

Answer: A

Q253. What type of system call is 'read()' in Unix?

- A. File management system call
- B. Process control system call
- C. Device management system call
- D. Communication system call

Answer: A

Q254. In a layered OS design, which layer interacts directly with the hardware?

- A. The lowest layer (Layer 0)
- B. The highest layer (user interface)
- C. The middle layer
- D. All layers equally

Answer: A

Q255. What is a device management system call used for?

- A. Requesting and releasing hardware devices
- B. Creating new files on disk
- C. Terminating running processes
- D. Sending messages between computers

Answer: A

Q256. What is the purpose of a boot loader?

- A. To load the operating system kernel into memory during startup
- B. To install new application software
- C. To manage user login credentials
- D. To defragment the hard disk

Answer: A

Q257. Which of the following is a process control system call?

- A. fork()
- B. open()
- C. read()
- D. chmod()

Answer: A

Q258. What is an operating system kernel?

- A. The core component of the OS that is always resident in memory
- B. A user-level application for managing files
- C. An external hardware device controller
- D. A network communication protocol

Answer: A

Q259. What is a system utility program?

- A. A program that provides common system functions like file management
- B. The core kernel of the operating system
- C. A hardware diagnostic chip on the motherboard
- D. An application for browsing the internet

Answer: A

Q260. What does the 'write()' system call do?

- A. Writes data to a file or output device
- B. Creates a new directory
- C. Deletes a file from disk
- D. Changes file permissions

Answer: A

Q261. What is a process state in operating systems?

- A. The current activity or condition of a process during its lifecycle
- B. The physical location of the process on disk
- C. The programming language used to write the process
- D. The amount of memory installed in the system

Answer: A

Q262. What happens when a running process requests an I/O operation?

- A. It moves from the running state to the waiting state
- B. It moves from the running state to the terminated state
- C. It continues executing without any state change
- D. It is permanently removed from memory

Answer: A

Q263. What information does the process ID (PID) provide?

- A. A unique numerical identifier assigned to each process by the OS
- B. The memory address where the process is stored
- C. The priority level of the process
- D. The name of the user who started the process

Answer: A

Q264. What is a child process?

- A. A process created by another process using a system call like fork()
- B. A process that has the highest priority in the system
- C. A process that manages hardware devices
- D. A process that is waiting for user input

Answer: A

Q265. What is the 'new' state of a process?

- A. The process is being created and has not yet been admitted to the ready queue
- B. The process is currently executing on the CPU
- C. The process is waiting for an I/O event to complete
- D. The process has finished execution

Answer: A

Q266. What is the 'terminated' state of a process?

- A. The process has finished execution and is being cleaned up by the OS
- B. The process is actively running on the CPU
- C. The process is waiting in the ready queue
- D. The process is blocked on an I/O request

Answer: A

Q267. What does the wait() system call do in Unix?

- A. Causes the parent process to wait until a child process terminates
- B. Pauses the CPU clock for a specified duration
- C. Waits for user keyboard input
- D. Delays the boot sequence of the operating system

Answer: A

Q268. What is a process table in an operating system?

- A. A data structure that stores information about all active processes
- B. A physical table of hardware components in the computer
- C. A list of all installed software applications
- D. A directory of files stored on the hard disk

Answer: A

Q269. What is the ready state of a process?

- A. The process is loaded in memory and waiting for CPU allocation
- B. The process is currently executing instructions
- C. The process is waiting for an I/O device to respond
- D. The process has not yet been created

Answer: A

Q270. What does the exit() system call do?

- A. Terminates the calling process and returns a status to the parent
- B. Restarts the operating system
- C. Closes all open files on the system
- D. Logs the user out of the system

Answer: A

Q271. What is the main goal of CPU scheduling?

- A. To maximize CPU utilization and ensure fair process execution
- B. To increase the physical clock speed of the CPU
- C. To add more cores to the processor
- D. To reduce the amount of installed RAM

Answer: A

Q272. What is a scheduling algorithm in operating systems?

- A. A method that determines which process gets the CPU next
- B. A formula for calculating memory addresses
- C. A technique for compressing files on disk
- D. A procedure for installing device drivers

Answer: A

Q273. In Round Robin scheduling, what happens when a process's time quantum expires?

- A. The process is preempted and moved to the end of the ready queue
- B. The process is terminated immediately
- C. The process continues running until it finishes
- D. The process is moved to a higher priority queue

Answer: A

Q274. What is a non-preemptive scheduling algorithm?

- A. An algorithm where a running process keeps the CPU until it voluntarily releases it
- B. An algorithm that can interrupt any process at any time
- C. An algorithm that runs all processes simultaneously
- D. An algorithm used only for real-time systems

Answer: A

Q275. What does response time measure in CPU scheduling?

- A. The time from submission of a request until the first response is produced
- B. The total time a process takes to complete execution
- C. The time a process spends waiting in the ready queue
- D. The time the CPU is idle between processes

Answer: A

Q276. Which scheduling algorithm is the simplest to implement?

- A. First-Come, First-Served (FCFS)
- B. Multilevel Feedback Queue
- C. Completely Fair Scheduler
- D. Rate Monotonic Scheduling

Answer: A

Q277. What is CPU burst in the context of process execution?

- A. A period during which a process is actively using the CPU for computation
- B. The maximum speed at which the CPU can operate
- C. A hardware failure that causes the CPU to overheat
- D. The time it takes to boot the CPU

Answer: A

Q278. In priority scheduling, what happens when two processes have the same priority?

- A. They are typically scheduled using FCFS among equal-priority processes
- B. The process with the larger PID always runs first
- C. Both processes run simultaneously on the same core
- D. One process is terminated to resolve the conflict

Answer: A

Q279. What is an I/O burst in process execution?

- A. A period during which a process waits for an I/O operation to complete
- B. A spike in CPU processing speed
- C. A burst of data transmitted over the network
- D. A sudden increase in available memory

Answer: A

Q280. What does the dispatcher module do in CPU scheduling?

- A. Gives control of the CPU to the process selected by the scheduler
- B. Selects which process should run next from the ready queue
- C. Calculates the priority of each process
- D. Manages memory allocation for processes

Answer: A

Q281. Which component is NOT shared between threads of the same process?

- A. Global variables
- B. Stack
- C. Code section
- D. Open file descriptors

Answer: B

Q282. What is the primary reason thread creation is faster than process creation?

- A. Threads use faster CPUs
- B. Threads share the address space and resources of their parent process
- C. Threads do not need an operating system
- D. Threads are stored on the hard disk

Answer: B

Q283. Which of the following best describes a single-threaded process?

- A. A process that can run on multiple CPUs simultaneously
- B. A process with exactly one thread of execution
- C. A process that cannot perform I/O operations
- D. A process that runs only in kernel mode

Answer: B

Q284. What is the role of a thread library?

- A. To manage hardware interrupts
- B. To provide an API for creating and managing threads
- C. To compile multithreaded programs
- D. To allocate disk space for threads

Answer: B

Q285. Context switching between threads of the same process is generally faster than between processes because:

- A. Threads have smaller program counters
- B. No need to switch the memory address space
- C. Threads do not use the CPU
- D. The operating system does not track threads

Answer: B

Q286. Which threading model maps many user-level threads to many kernel threads?

- A. One-to-one model
- B. Many-to-one model
- C. Many-to-many model
- D. One-to-many model

Answer: C

Q287. In Java, which interface must a class implement to define a task that can be executed by a thread?

- A. Serializable
- B. Runnable
- C. Callable only
- D. Throwable

Answer: B

Q288. What does the term 'multicore programming' refer to?

- A. Programming with multiple monitors
- B. Writing programs that can utilize multiple processing cores
- C. Running multiple operating systems simultaneously
- D. Using multiple keyboards for input

Answer: B

Q289. What is the difference between concurrency and parallelism in threading?

- A. They are identical concepts
- B. Concurrency means tasks make progress over time; parallelism means tasks run simultaneously
- C. Parallelism requires only one CPU core
- D. Concurrency is only possible with GPUs

Answer: B

Q290. Which register is unique to each thread and tracks the next instruction to execute?

- A. Base register
- B. Program counter
- C. Memory address register
- D. I/O register

Answer: B

Q291. What is the purpose of process synchronization?

- A. To make all processes run at the same speed
- B. To coordinate the execution of processes that share resources to prevent data inconsistency
- C. To synchronize the system clock with an external server
- D. To ensure all processes start at the same time

Answer: B

Q292. The 'progress' requirement for the critical section problem states that:

- A. All processes must execute at the same speed
- B. Only processes not in their remainder section can participate in deciding which process enters the critical section next
- C. A process must never leave the critical section
- D. All processes must enter the critical section simultaneously

Answer: B

Q293. The 'bounded waiting' requirement for the critical section problem ensures that:

- A. The critical section has a fixed size
- B. There is a limit on the number of times other processes can enter the critical section after a process has made a request
- C. Processes wait for exactly 10 seconds
- D. Only bounded variables are used

Answer: B

Q294. What does a mutex lock provide?

- A. Unlimited concurrent access to shared data
- B. Mutual exclusion so that only one thread can hold the lock at a time
- C. Automatic parallelization of code
- D. Network synchronization between computers

Answer: B

Q295. What operation decreases the value of a semaphore?

- A. signal()
- B. post()
- C. wait()
- D. release()

Answer: C

Q296. What is a spinlock?

- A. A lock that puts the waiting thread to sleep
- B. A lock where the waiting thread continuously checks if the lock is available in a loop
- C. A lock that only works on single-core systems
- D. A lock that spins the hard disk for faster access

Answer: B

Q297. In the producer-consumer problem, what shared resource do the producer and consumer access?

- A. The CPU registers
- B. A bounded buffer
- C. The page table
- D. The file system journal

Answer: B

Q298. What is an atomic operation?

- A. An operation related to nuclear physics
- B. An operation that completes entirely without interruption
- C. An operation that can be partially executed
- D. An operation that requires multiple CPU cycles

Answer: B

Q299. What does the signal() operation do on a semaphore?

- A. Decreases its value by one
- B. Increases its value by one and may wake a waiting process
- C. Destroys the semaphore
- D. Resets the semaphore to zero

Answer: B

Q300. What is the key difference between a binary semaphore and a counting semaphore?

- A. Binary semaphores are faster
- B. A binary semaphore has only values 0 and 1, while a counting semaphore can range over an unrestricted domain
- C. Counting semaphores cannot be used for mutual exclusion
- D. Binary semaphores work only in user mode

Answer: B

Q301. Which of the following is a necessary condition for deadlock?

- A. Preemption of all resources
- B. Mutual exclusion
- C. Unlimited resources
- D. Single-threaded execution

Answer: B

Q302. In a resource allocation graph, what does an edge from a process to a resource indicate?

- A. The process has finished using the resource
- B. The process is requesting the resource
- C. The resource is allocated to the process
- D. The process has released the resource

Answer: B

Q303. What are the four necessary conditions for deadlock?

- A. Speed, memory, I/O, CPU
- B. Mutual exclusion, hold and wait, no preemption, circular wait
- C. Scheduling, paging, caching, buffering
- D. Authentication, authorization, encryption, logging

Answer: B

Q304. What is deadlock avoidance?

- A. Ignoring the possibility of deadlock
- B. Using additional information about future resource requests to ensure the system never enters an unsafe state
- C. Terminating all processes periodically
- D. Removing all shared resources from the system

Answer: B

Q305. What does deadlock recovery involve?

- A. Preventing deadlock from ever occurring
- B. Detecting that a deadlock has occurred and taking action to break it
- C. Ignoring the deadlock permanently
- D. Adding more resources to the system preemptively

Answer: B

Q306. An unsafe state means:

- A. Deadlock has already occurred
- B. The system may potentially lead to deadlock; there is no guarantee of a safe sequence
- C. All processes have terminated
- D. The system has run out of memory

Answer: B

Q307. In deadlock handling, what does the 'no preemption' condition mean?

- A. Resources can be forcibly taken from processes at any time
- B. Resources cannot be forcibly taken away from a process; they must be released voluntarily
- C. Processes cannot be scheduled on the CPU
- D. No process has higher priority than another

Answer: B

Q308. Which of the following is an example of a reusable resource?

- A. A CPU interrupt signal
- B. A printer or a memory block
- C. A network packet
- D. A one-time password

Answer: B

Q309. What is the simplest method to deal with deadlocks?

- A. Use the Banker's Algorithm
- B. Ignore the problem entirely (ostrich algorithm)
- C. Preempt all resources
- D. Allow only one process to run at a time

Answer: B

Q310. What type of graph is used specifically for deadlock detection with single-instance resources?

- A. A binary tree
- B. A wait-for graph
- C. A B+ tree
- D. A hash table

Answer: B

Q311. What is the primary function of a Memory Management Unit (MMU)?

- A. To increase the size of physical RAM
- B. To translate logical addresses to physical addresses at runtime
- C. To manage the file system
- D. To schedule CPU processes

Answer: B

Q312. What is a frame in the context of paging?

- A. A division of logical memory
- B. A fixed-size block of physical memory
- C. A type of CPU register
- D. A section of the hard disk

Answer: B

Q313. What is the difference between a logical address and a physical address?

- A. They are always the same value
- B. A logical address is generated by the CPU and a physical address is the actual location in memory hardware
- C. A physical address is virtual and a logical address is real
- D. Logical addresses are used only for disk storage

Answer: B

Q314. What is the purpose of a base register in memory management?

- A. To store the process's name
- B. To define the starting physical address of a process's memory space
- C. To count the number of pages
- D. To store the program counter

Answer: B

Q315. What is the role of the limit register in memory protection?

- A. It stores the maximum speed of memory access
- B. It specifies the size of the range of legal addresses for a process
- C. It limits the number of processes in the system
- D. It controls the clock speed of the CPU

Answer: B

Q316. What is dynamic loading?

- A. Loading all routines into memory at program start
- B. Loading a routine into memory only when it is called, reducing memory usage
- C. Loading programs onto the hard disk
- D. Loading the operating system into the GPU

Answer: B

Q317. What is the difference between first-fit and best-fit memory allocation?

- A. They always produce the same result
- B. First-fit allocates the first hole large enough; best-fit allocates the smallest hole large enough
- C. Best-fit is always slower than first-fit
- D. First-fit only works with paging

Answer: B

Q318. What problem does compaction solve?

- A. Internal fragmentation
- B. External fragmentation by shuffling memory contents to place all free memory together
- C. Page faults
- D. CPU cache misses

Answer: B

Q319. What does demand paging mean?

- A. All pages are loaded into memory at startup
- B. Pages are loaded into memory only when they are accessed and needed
- C. The CPU demands faster memory
- D. Pages are stored only on the hard disk

Answer: B

Q320. What is the valid-invalid bit in a page table entry used for?

- A. To indicate whether the data in the page has errors
- B. To indicate whether the page is currently in physical memory (valid) or not (invalid)
- C. To store the page number
- D. To count the number of accesses to the page

Answer: B

Q321. What is the purpose of a file system?

- A. To manage CPU scheduling
- B. To organize, store, and retrieve files on storage media
- C. To control network connections
- D. To manage power distribution to hardware

Answer: B

Q322. What is a relative path in a file system?

- A. A path that always starts from the root directory
- B. A path specified relative to the current working directory
- C. A path that points to a network resource
- D. A path that can only reference directories, not files

Answer: B

Q323. What file attribute indicates the number of bytes in a file?

- A. File type
- B. File size
- C. File name
- D. File permissions

Answer: B

Q324. What does the close() system call do in file operations?

- A. Deletes the file permanently
- B. Releases the file's resources and removes the entry from the open-file table
- C. Creates a backup of the file
- D. Changes the file's permissions

Answer: B

Q325. What is a directory entry?

- A. A CPU register that stores directory information
- B. A record in a directory that maps a file name to its location on storage
- C. A type of file that cannot be opened
- D. A network address for remote files

Answer: B

Q326. What is a single-level directory structure?

- A. A hierarchy of nested directories
- B. A single directory containing all files in the system
- C. A directory structure with exactly two levels
- D. A directory that stores only one file

Answer: B

Q327. What does 'mounting' a file system mean?

- A. Physically installing a hard drive
- B. Making a file system accessible by attaching it to a point in the existing directory tree
- C. Formatting a disk partition
- D. Encrypting the file system

Answer: B

Q328. What type of file access reads data sequentially from beginning to end?

- A. Random access
- B. Sequential access
- C. Indexed access
- D. Direct access

Answer: B

Q329. What is a two-level directory structure?

- A. A directory with two files
- B. Each user has a separate directory under a master file directory
- C. A directory stored on two different disks
- D. A directory accessible by only two users

Answer: B

Q330. What are file permissions used for?

- A. To compress files to save storage space
- B. To control which users can read, write, or execute a file
- C. To speed up file access times
- D. To encrypt file contents automatically

Answer: B

Q331. What is the main advantage of secondary storage over primary memory (RAM)?

- A. It is faster than RAM
- B. It is non-volatile and retains data when power is off
- C. It is directly accessible by the CPU
- D. It has lower latency than RAM

Answer: B

Q332. What is a cylinder in the context of hard disk drives?

- A. The motor that spins the disk
- B. The set of tracks at the same position across all platters
- C. A single sector on the disk
- D. The disk controller chip

Answer: B

Q333. What is the FCFS (First-Come, First-Served) disk scheduling algorithm?

- A. Requests are served based on their distance from the current head position
- B. Requests are served in the order they arrive in the queue
- C. Requests are served from the innermost track to the outermost track
- D. Requests are served based on their priority

Answer: B

Q334. Why do SSDs have no seek time or rotational latency?

- A. They use faster motors
- B. They have no moving parts; data is accessed electronically from flash memory cells
- C. They use optical technology instead of magnetic
- D. They spin at infinite speed

Answer: B

Q335. What does RAID stand for?

- A. Random Array of Independent Data
- B. Redundant Array of Independent Disks
- C. Rapid Access to Internal Drives
- D. Reliable Access to Indexed Data

Answer: B

Q336. What is the purpose of disk formatting?

- A. To make the disk spin faster
- B. To prepare a disk for use by creating the necessary data structures for a file system
- C. To encrypt all data on the disk
- D. To physically clean the disk platters

Answer: B

Q337. What is the function of the disk controller?

- A. To physically manufacture the disk
- B. To interface between the disk drive and the computer, managing data transfer and disk operations
- C. To supply power to the disk
- D. To cool the disk platters

Answer: B

Q338. What is swap space used for?

- A. Swapping data between two hard drives
- B. Extending virtual memory by storing pages that are not currently in physical RAM
- C. Exchanging files between users
- D. Temporarily storing CPU registers

Answer: B

Q339. What is a platter in a hard disk drive?

- A. The cable connecting the disk to the motherboard
- B. A circular, flat disk coated with magnetic material that stores data
- C. The electronic circuit board of the drive
- D. The protective case of the drive

Answer: B

Q340. What is an I/O port?

- A. A physical port for connecting displays
- B. A set of registers through which the CPU communicates with an I/O device controller
- C. A network port used for internet connections
- D. A type of RAM used for I/O caching

Answer: B

Q341. What is the main disadvantage of polling (programmed I/O)?

- A. It does not work with any devices
- B. The CPU wastes cycles repeatedly checking the device status instead of doing useful work
- C. It is too fast for most devices
- D. It requires special hardware

Answer: B

Q342. How does DMA (Direct Memory Access) benefit the system?

- A. It eliminates the need for memory entirely
- B. It allows devices to transfer data directly to or from memory without CPU involvement for each byte
- C. It doubles the CPU's clock speed during I/O
- D. It replaces the operating system's I/O management

Answer: B

Q343. What is spooling in the context of I/O?

- A. Rewinding magnetic tape
- B. Storing output data in a buffer (usually on disk) for a device that cannot accept interleaved data from multiple sources
- C. Speeding up I/O operations
- D. Connecting multiple printers together

Answer: B

Q344. What is the difference between a block device and a character device?

- A. Block devices are slower than character devices
- B. Block devices transfer data in fixed-size blocks and support random access; character devices transfer data one character at a time in a stream
- C. Character devices can only handle text files
- D. Block devices do not use device drivers

Answer: B

Q345. What is an interrupt in the context of I/O?

- A. A power outage affecting the system
- B. A signal from a device to the CPU indicating that an I/O event needs attention
- C. A type of memory error
- D. A command from the CPU to stop a device

Answer: B

Q346. What role does a device driver play in the I/O system?

- A. It physically connects devices to the computer
- B. It provides a uniform interface for the OS to interact with a specific hardware device
- C. It generates electricity for the device
- D. It manufactures new devices

Answer: B

Q347. What is I/O buffering?

- A. Removing data from I/O devices
- B. Temporarily storing data in memory during transfer between a device and an application to handle speed differences
- C. Encrypting I/O data for security
- D. Compressing data before sending it to a device

Answer: B

Q348. What is a bus in a computer system?

- A. A vehicle for transporting computer equipment
- B. A shared communication pathway that transfers data between components of a computer
- C. A type of secondary storage
- D. A software program for managing I/O

Answer: B

Q349. What happens when an interrupt occurs during program execution?

- A. The program is permanently deleted
- B. The CPU saves the current state and transfers control to the appropriate interrupt handler routine
- C. The computer shuts down
- D. The interrupt is ignored until the program finishes

Answer: B

Q350. What is authorization in operating system security?

- A. Verifying a user's identity
- B. Determining what resources a verified user is permitted to access
- C. Encrypting user data
- D. Logging user activity

Answer: B

Q351. What is a worm in computer security?

- A. A program that modifies other programs by inserting its code
- B. A self-replicating program that spreads across networks without needing a host program
- C. A program that disguises itself as legitimate software
- D. A hardware device that protects against malware

Answer: B

Q352. What is a denial-of-service (DoS) attack?

- A. Stealing data from a computer
- B. An attack that overwhelms a system's resources to make it unavailable to legitimate users
- C. Installing malware on a computer
- D. Breaking encryption to access protected data

Answer: B

Q353. What is the purpose of a user account password?

- A. To encrypt files automatically
- B. To authenticate the user's identity before granting system access
- C. To speed up the login process
- D. To store the user's preferences

Answer: B

Q354. What is the difference between a virus and a Trojan horse?

- A. Viruses are hardware; Trojans are software
- B. A virus self-replicates by attaching to other programs; a Trojan disguises itself as legitimate software but does not self-replicate
- C. A Trojan self-replicates; a virus does not
- D. There is no difference between them

Answer: B

Q355. What does the principle of least privilege mean in practice?

- A. Users should have the maximum possible privileges
- B. Users and programs should only have the minimum privileges necessary to complete their tasks
- C. All users should have equal privileges
- D. Privileges should be assigned randomly

Answer: B

Q356. What is the purpose of a security audit log?

- A. To speed up system performance
- B. To record security-relevant events for later review and analysis
- C. To encrypt system files
- D. To authenticate users automatically

Answer: B

Q357. What is biometric authentication?

- A. Authentication using passwords only
- B. Authentication using unique physical characteristics such as fingerprints, facial recognition, or iris patterns
- C. Authentication using security tokens only
- D. Authentication using encryption keys

Answer: B

Q358. What is a protection domain in operating system security?

- A. A geographic region where the computer is located
- B. A set of access rights that defines what resources a process can access and what operations it can perform
- C. A physical barrier around the server room
- D. A type of antivirus software

Answer: B

Q359. What is the function of a firewall?

- A. To cool down the CPU to prevent overheating
- B. To filter network traffic based on security rules, blocking unauthorized access while permitting legitimate communication
- C. To encrypt all files on the hard drive
- D. To scan files for viruses

Answer: B

Q360. What is the relationship between disk access time and its components?

- A. Access time equals seek time only
- B. Access time is the sum of seek time, rotational latency, and data transfer time
- C. Access time equals rotational latency only
- D. Access time equals transfer time minus seek time

Answer: B

Operating Systems — Medium

360 questions

Q361. What is the difference between multiprogramming and multitasking?

- A. Multiprogramming uses multiple CPUs; multitasking one
- B. Multiprogramming maximizes CPU; multitasking adds sharing
- C. There is no meaningful difference between them
- D. Multitasking is an older technique than multiprogramming

Answer: B

Q362. What is a system call?

- A. A call made from one user to another user
- B. A function call within user-level programs
- C. A programmatic interface to request OS services
- D. A type of hardware interrupt signal to CPU

Answer: C

Q363. Which mode of CPU operation has unrestricted access to all hardware resources?

- A. User mode
- B. Safe mode
- C. Debug mode
- D. Kernel mode

Answer: D

Q364. What is the purpose of the mode bit in the CPU?

- A. To store the current instruction address
- B. To manage the CPU cache memory size
- C. To count the total CPU clock cycles
- D. To distinguish user mode from kernel mode

Answer: D

Q365. Spooling in operating systems refers to:

- A. A type of interrupt handling method
- B. Overlapping I/O with CPU computation
- C. A type of CPU scheduling algorithm
- D. A memory management technique used

Answer: B

Q366. An interrupt is:

- A. A type of software error in code
- B. A signal needing immediate CPU attention
- C. A specific type of system call request
- D. A process scheduling technique in OS

Answer: B

Q367. What is the difference between a trap and an interrupt?

- A. A trap is software-generated, interrupt is hardware
- B. There is no real difference between them
- C. A trap executes faster than an interrupt does
- D. A trap is hardware-generated, interrupt is software

Answer: A

Q368. DMA (Direct Memory Access) is used to:

- A. Speed up internal CPU processing tasks
- B. Transfer data to memory without the CPU
- C. Increase the processor clock speed rate
- D. Manage virtual memory paging on disk

Answer: B

Q369. A distributed operating system:

- A. Makes independent computers appear as one system
- B. Runs only on a single standalone machine
- C. Cannot share resources between multiple machines
- D. Is the same as a standard network operating system

Answer: A

Q370. What is the bootstrap program?

- A. A program that manages user accounts and groups
- B. A program that handles active network connections
- C. A program that compiles other user-level programs
- D. The first program run at power-on to load the OS

Answer: D

Q371. In a microkernel OS, which of the following typically runs in user space?

- A. Interrupt handling routines
- B. Core CPU scheduling
- C. Hardware-level memory mgmt
- D. File system services

Answer: D

Q372. What is the primary mechanism microkernels use for communication between services?

- A. Message passing IPC
- B. Shared memory only
- C. Continuous polling
- D. Direct function calls

Answer: A

Q373. System calls are typically invoked by user programs through:

- A. Direct hardware manipulation
- B. Application Programming APIs
- C. Bootloader entry points
- D. Device driver interfaces

Answer: B

Q374. What is the main disadvantage of the layered approach to OS design?

- A. It does not provide any file system support
- B. Difficulty defining layers and traversal overhead
- C. It is too simple and basic to be practical
- D. It cannot support any form of multitasking

Answer: B

Q375. What is a loadable kernel module (LKM)?

- A. Code loaded into the kernel at runtime dynamically
- B. A user application loaded on demand by the OS
- C. A type of system call for interprocess messaging
- D. A bootloader component for starting the OS kernel

Answer: A

Q376. Which of the following is a category of system calls?

- A. Compilation control
- B. Code optimization
- C. Information maintenance
- D. Syntax analysis

Answer: C

Q377. Parameters to system calls can be passed by placing them in:

- A. Only through CPU registers
- B. Only via the system stack
- C. Registers, stack, or memory block
- D. Only through shared memory

Answer: C

Q378. The POSIX API is:

- A. A Windows-specific system API
- B. A file format specification
- C. A portable UNIX interface standard
- D. A networking protocol standard

Answer: C

Q379. What is the purpose of the system call interface?

- A. To link API calls to kernel implementations
- B. To format disks and storage partitions
- C. To manage user passwords and accounts
- D. To execute hardware instructions directly

Answer: A

Q380. Windows NT uses which type of kernel architecture?

- A. Hybrid kernel
- B. Pure microkernel
- C. Pure monolithic
- D. Exokernel type

Answer: A

Q381. In the producer-consumer problem using shared memory, what role does the buffer play?

- A. It handles hardware interrupt processing
- B. It is a shared region for producer-consumer
- C. It stores the executable program code
- D. It manages the CPU scheduling algorithm

Answer: B

Q382. What is the difference between shared memory and message passing IPC?

- A. Shared memory faster for large data; message passing for remote
- B. There is no difference between these two approaches
- C. Shared memory uses common region; message passing uses kernel
- D. Both B and C are correct descriptions of their differences

Answer: D

Q383. What happens during a context switch?

- A. Current process state is saved, next is loaded
- B. All of main memory is cleared completely
- C. All running processes are terminated at once
- D. The CPU is shut down and then restarted

Answer: A

Q384. In UNIX, what does the exec() family of system calls do?

- A. Terminates the currently running process
- B. Replaces process image with new program
- C. Suspends the process until a signal
- D. Creates a brand new child process

Answer: B

Q385. What is cascading termination?

- A. Terminating processes one by one sequentially
- B. OS terminates all children when parent exits
- C. Restarting processes in a specific sequence
- D. Suspending processes in a particular order

Answer: B

Q386. A pipe in UNIX provides:

- A. Bidirectional communication between any processes
- B. Network communication across remote machines
- C. Direct communication with the hardware devices
- D. Unidirectional communication between parent-child

Answer: D

Q387. What is a named pipe (FIFO) and how does it differ from an ordinary pipe?

- A. It only works on Microsoft Windows platforms
- B. It provides faster data communication speed
- C. It has a name and works between unrelated processes
- D. It is identical to an ordinary pipe

Answer: C

Q388. Which scheduling queue holds processes that are in main memory and ready to execute?

- A. Wait queue
- B. Device queue
- C. Ready queue
- D. Job queue

Answer: C

Q389. The long-term scheduler (job scheduler) controls:

- A. Which process gets I/O device access next
- B. Which process gets swapped out to the disk
- C. Which process gets the CPU execution next
- D. Which processes are admitted to the system

Answer: D

Q390. What is the role of the short-term scheduler (CPU scheduler)?

- A. To swap processes between memory and disk
- B. To select which ready process gets the CPU
- C. To admit new jobs into the system queue
- D. To terminate processes that have finished

Answer: B

Q391. What is the convoy effect in FCFS scheduling?

- A. Processes are grouped into convoys for batch runs
- B. Short processes wait for long ones, raising wait time
- C. CPU utilization increases dramatically and suddenly
- D. All processes finish at exactly the same time

Answer: B

Q392. In Shortest Remaining Time First (SRTF) scheduling, what happens when a new process arrives?

- A. The current process is always preempted regardless of timing
- B. The new process always waits until the current one completes
- C. The scheduler compares remaining times and preempts if shorter
- D. The new process automatically gets the highest priority level

Answer: C

Q393. What is starvation in the context of CPU scheduling?

- A. A process running completely out of allocated memory
- B. A process consuming too much CPU time from the system
- C. A deadlock situation between two or more blocked tasks
- D. A low-priority process waiting indefinitely for the CPU

Answer: D

Q394. Aging is a technique used to:

- A. Track how long each process has been running on the CPU
- B. Increase process priority over time to prevent starvation
- C. Decrease the priority of a process as it runs much longer
- D. Remove old processes from the system to free up resources

Answer: B

Q395. If the time quantum in Round Robin is very large, it behaves like:

- A. Multilevel queues
- B. Priority scheduling
- C. SJF scheduling
- D. FCFS scheduling

Answer: D

Q396. If the time quantum in Round Robin is very small, what problem occurs?

- A. Excessive switching
- B. System deadlocking
- C. Memory overflowing
- D. Process starvation

Answer: A

Q397. In Multilevel Queue Scheduling, processes are:

- A. Permanently assigned to queues by type
- B. All placed in one single queue
- C. Moved between queues based on priority
- D. Randomly assigned to available queues

Answer: A

Q398. What distinguishes Multilevel Feedback Queue from Multilevel Queue scheduling?

- A. It only uses the Round Robin algorithm
- B. Processes can move between queues by behavior
- C. It does not support priority scheduling
- D. It uses fewer scheduling queues overall

Answer: B

Q399. Dispatch latency is:

- A. The time a process waits in the ready queue
- B. The time between two consecutive CPU interrupts
- C. The time to stop one process and start another
- D. The total execution time of a single process

Answer: C

Q400. Response time in scheduling is defined as:

- A. The time between two context switches
- B. The time to complete all system processes
- C. The total CPU execution time of process
- D. Time from submission until first response

Answer: D

Q401. What is the many-to-one threading model?

- A. Each user thread maps to one kernel thread
- B. Many kernel threads map to one user thread
- C. Many user-level threads map to one kernel thread
- D. One user thread maps to many kernel threads

Answer: C

Q402. What is the one-to-one threading model?

- A. One process only ever has one thread
- B. One kernel thread maps to many user threads
- C. One user thread maps to one kernel thread
- D. All threads share one single CPU together

Answer: C

Q403. What is the disadvantage of the many-to-one threading model?

- A. If one thread blocks, all threads in process block
- B. It cannot create more than two threads total
- C. It uses too many kernel resources overall
- D. It is too slow for practical usage scenarios

Answer: A

Q404. What is a thread pool?

- A. A memory allocation area used for threads
- B. Pre-created threads that wait for work tasks
- C. A collection of currently unused threads
- D. A debugging tool for analyzing thread bugs

Answer: B

Q405. What happens to threads when fork() is called in a multithreaded program?

- A. The child process is created with no threads
- B. All threads are always duplicated in the child
- C. Fork cannot be used in multithreaded programs
- D. It depends on the system; some dup all, some one

Answer: D

Q406. What is thread cancellation?

- A. Terminating a thread before completion
- B. Creating a brand new thread from scratch
- C. Moving a thread to a different process
- D. Suspending a thread temporarily in place

Answer: A

Q407. What is the difference between asynchronous and deferred thread cancellation?

- A. Asynchronous is safer because it checks all resources before acting
- B. Asynchronous terminates immediately; deferred checks at safe points
- C. Deferred cancellation is always much faster than asynchronous type
- D. There is no difference between the two cancellation types

Answer: B

Q408. Thread-Local Storage (TLS) allows:

- A. Threads to communicate with each other via messages
- B. Each thread to have its own copy of certain data
- C. All threads to share the exact same data copies
- D. Data to be stored persistently on disk for threads

Answer: B

Q409. What is implicit threading?

- A. Threads that remain invisible to the programmer
- B. Threads that run implicitly without any OS scheduling
- C. Threads that do not use any system resources at all
- D. Transferring thread management to compilers and runtimes

Answer: D

Q410. OpenMP is used for:

- A. File compression utilities
- B. Parallel shared-memory programming
- C. Database management systems
- D. Network programming via APIs

Answer: B

Q411. What is Peterson's solution for the critical section problem?

- A. A solution implemented using monitors
- B. A hardware-based solution only
- C. A solution implemented using semaphores
- D. A software solution using turn and flag

Answer: D

Q412. The test-and-set instruction is:

- A. A specific type of semaphore implementation
- B. An atomic hardware instruction to test and set
- C. A scheduling algorithm for critical sections
- D. A software-only algorithm for locking

Answer: B

Q413. The compare-and-swap (CAS) instruction:

- A. Compares two files and swaps their contents
- B. Compares the speeds of two different CPUs
- C. Atomically compares and swaps a memory value
- D. Swaps two processes in the scheduling queue

Answer: C

Q414. A monitor is:

- A. A specific type of counting semaphore lock mechanism
- B. A construct encapsulating shared data with mutual exclusion
- C. A hardware lock on the system memory bus interface
- D. A display device or screen monitor unit

Answer: B

Q415. Condition variables in monitors are used with which operations?

- A. push() and pop()
- B. acquire() and release()
- C. wait() and signal()
- D. lock() and unlock()

Answer: C

Q416. In the bounded buffer problem, what are the semaphores typically used?

- A. One counting semaphore for the buffer
- B. Only a single mutex lock for the buffer
- C. Mutex plus empty and full semaphores
- D. Two binary semaphores for the buffer

Answer: C

Q417. In the readers-writers problem, what is the key synchronization challenge?

- A. Only one reader can read the data at a time
- B. Only one process can ever access the shared database
- C. Writers can write while readers are actively reading
- D. Multiple readers can read; writers need exclusive access

Answer: D

Q418. What is a spinlock most useful for?

- A. Short critical sections on multiprocessor systems
- B. Any critical section on single processor only
- C. I/O-bound processes waiting for disk access
- D. Long critical sections on single processors

Answer: A

Q419. What is priority inversion?

- A. Changing process priorities dynamically at runtime
- B. Assigning priorities in reverse numerical order
- C. A scheduling bug in all operating system kernels
- D. High-priority blocked by low-priority holding resource

Answer: D

Q420. Priority inheritance protocol solves priority inversion by:

- A. Raising low-priority to match highest waiting
- B. Blocking all medium-priority from running
- C. Terminating the low-priority resource holder
- D. Removing all priorities from every process

Answer: A

Q421. A safe state is one where:

- A. A safe sequence exists for all processes
- B. No process is waiting for any resource
- C. All system resources are fully available
- D. No process is currently running at all

Answer: A

Q422. What is the relationship between safe state and deadlock?

- A. An unsafe state always guarantees that deadlock exists
- B. Safe state means deadlock currently exists in system
- C. Safe and unsafe states are unrelated to deadlock at all
- D. Safe means no deadlock; unsafe means deadlock may occur

Answer: D

Q423. To prevent the 'hold and wait' condition, a process must:

- A. Use only shared resources that do not need locking
- B. Never request any resources from the operating system
- C. Request all resources at once, or release before requesting
- D. Hold all needed resources indefinitely without releasing

Answer: C

Q424. To prevent circular wait, one approach is:

- A. Remove all shared resources from system
- B. Allow processes to request in any order
- C. Impose ordering and require increasing order
- D. Use only a single resource type allowed

Answer: C

Q425. In the Banker's Algorithm, the 'Need' matrix represents:

- A. Currently allocated resources to each process
- B. The total number of resources in whole system
- C. Remaining resources each process may still need
- D. Total resources currently available to allocate

Answer: C

Q426. What is deadlock detection?

- A. Allowing deadlocks then detecting and recovering
- B. Ignoring deadlocks completely in all situations
- C. Avoiding deadlocks using the Banker Algorithm
- D. Preventing deadlocks from occurring in system

Answer: A

Q427. For deadlock detection with single-instance resources, which algorithm is used?

- A. Wait-for graph cycling
- B. Banker's Algorithm test
- C. Round Robin scheduling
- D. Best Fit allocating

Answer: A

Q428. Which deadlock recovery method terminates processes?

- A. Resource preemption from deadlocked processes
- B. Restarting the entire system from the beginning
- C. Ignoring the deadlock and hoping it resolves
- D. Process termination - abort all or one at a time

Answer: D

Q429. When recovering from deadlock by process termination, which factor is NOT typically considered?

- A. Time it has been running
- B. Process priority level
- C. Resources it now holds
- D. The color of terminal

Answer: D

Q430. The ostrich algorithm for handling deadlocks refers to:

- A. Ignoring deadlocks when rare and costly to handle
- B. A recovery mechanism for resolving deadlocks
- C. A sophisticated deadlock detection algorithm
- D. A prevention technique eliminating conditions

Answer: A

Q431. The Translation Lookaside Buffer (TLB) is:

- A. A disk cache for swap file contents
- B. A fast cache for page table entries
- C. A type of random access memory chip
- D. A backup copy of the main page table

Answer: B

Q432. In demand paging, pages are loaded into memory:

- A. All at once when the process first starts
- B. Based on a predetermined schedule
- C. Only when they are actually referenced
- D. Randomly without any specific order

Answer: C

Q433. The FIFO page replacement algorithm replaces:

- A. The page that has been in memory longest
- B. The least recently used page in memory
- C. The most frequently used page in memory
- D. A random page selected from the frames

Answer: A

Q434. The LRU (Least Recently Used) page replacement algorithm replaces:

- A. The page that was loaded into memory first
- B. The page not used for the longest time period
- C. The largest page currently loaded in a frame
- D. The page with the lowest assigned page number

Answer: B

Q435. Belady's anomaly states that:

- A. More frames can increase faults for some algorithms
- B. More memory always reduces total page faults
- C. LRU always outperforms FIFO in every situation
- D. Optimal page replacement causes most page faults

Answer: A

Q436. The optimal page replacement algorithm (OPT) replaces:

- A. The page unused for longest future time
- B. The smallest page currently in a frame
- C. The least recently used page in memory
- D. The first page that was loaded in memory

Answer: A

Q437. Thrashing occurs when:

- A. The CPU is idle and not running tasks
- B. Memory is completely free and available
- C. A process pages more than it executes
- D. Too few processes are currently running

Answer: C

Q438. Segmentation divides a program into:

- A. Equal-sized partitions of memory
- B. Fixed-size pages of equal length
- C. Variable-size segments like code and data
- D. Random blocks of varying sizes

Answer: C

Q439. The first-fit memory allocation algorithm:

- A. Allocates memory blocks at random
- B. Searches for the smallest suitable hole
- C. Allocates the first hole large enough
- D. Allocates the largest available hole

Answer: C

Q440. The best-fit memory allocation algorithm:

- A. Is the fastest allocation algorithm
- B. Allocates the first hole it finds
- C. Allocates the largest available hole
- D. Searches for the smallest sufficient hole

Answer: D

Q441. What is the difference between contiguous, linked, and indexed file allocation?

- A. They only differ in the speed of read and write operations
- B. Contiguous uses adjacent; linked chains; indexed uses pointers
- C. They only differ in the security settings for file access
- D. They are all identical in behavior and performance

Answer: B

Q442. Free space on disk can be managed using:

- A. Only the FAT table for tracking
- B. Only bit vectors for tracking
- C. Only linked lists for tracking
- D. Bit vectors, lists, or counting

Answer: D

Q443. A tree-structured directory organizes files in:

- A. A randomly organized structure
- B. A flat single-level structure
- C. A hierarchical parent-child structure
- D. A circular connected structure

Answer: C

Q444. An acyclic-graph directory structure allows:

- A. Only one single link per file in the system
- B. Circular file references between directories
- C. No file sharing between any user accounts
- D. Sharing files via links while preventing cycles

Answer: D

Q445. What is the Virtual File System (VFS) layer?

- A. An abstraction for uniform file system interface
- B. A file system stored in virtual memory
- C. A compressed file system saving disk space
- D. A file system designed for virtual machines

Answer: A

Q446. NTFS (New Technology File System) is primarily used by:

- A. Linux
- B. macOS
- C. Android
- D. Windows

Answer: D

Q447. ext4 is:

- A. A database system format
- B. The default Linux file system
- C. A Windows file system format
- D. A macOS file system format

Answer: B

Q448. File system consistency checking is needed because:

- A. Users make frequent file handling errors
- B. Files grow too large for disk capacity
- C. Hardware never fails in normal practice
- D. A crash during ops leaves inconsistent state

Answer: D

Q449. Access control lists (ACLs) in file systems provide:

- A. File versioning for tracking changes over time
- B. File data encryption for security and privacy
- C. File data compression for saving disk space
- D. Fine-grained permissions for users and groups

Answer: D

Q450. Direct access (random access) to a file means:

- A. Reading or writing any block directly without prior reads
- B. Reading the file sequentially from start to end
- C. Accessing the file at random without any specific pattern
- D. Accessing the file exclusively through the OS kernel layer

Answer: A

Q451. The SCAN (elevator) disk scheduling algorithm:

- A. Moves head one direction servicing, then reverses
- B. Only scans tracks that have pending I/O requests
- C. Scans files for viruses and malware threats
- D. Scans all tracks in a random selection order

Answer: A

Q452. The C-SCAN (Circular SCAN) algorithm differs from SCAN in that:

- A. It only works on solid state drives not on HDDs
- B. It is slower than the standard SCAN algorithm is
- C. Returns to start without servicing for uniform wait
- D. It never reverses the head direction at all

Answer: C

Q453. The LOOK disk scheduling algorithm differs from SCAN in that:

- A. It is significantly slower than the SCAN algorithm
- B. Reverses when no more requests ahead, not at end
- C. It only handles one single request at a time
- D. It looks at all pending requests before starting

Answer: B

Q454. Why is disk scheduling less important for SSDs compared to HDDs?

- A. SSDs use a completely different data interface
- B. SSDs are actually slower than HDDs overall
- C. SSDs have no moving parts, so seek is negligible
- D. SSDs cannot handle multiple concurrent requests

Answer: C

Q455. RAID 0 provides:

- A. Data redundancy only for safety
- B. Performance via striping, no redundancy
- C. Mirroring of data across two disks
- D. Parity-based protection from failure

Answer: B

Q456. RAID 1 provides:

- A. No data protection of any kind
- B. Striping data across disks only
- C. Parity-based protection on disks
- D. Mirroring via duplicate on disks

Answer: D

Q457. RAID 5 uses:

- A. Double parity across all the disks
- B. Block striping with distributed parity
- C. No redundancy of any kind at all
- D. Only mirroring for redundancy

Answer: B

Q458. Low-level (physical) formatting creates:

- A. Logical partitions and volume table structures
- B. A high-level file system and all its metadata
- C. Sector structure with headers and error correction
- D. Directories and user files on the disk surface

Answer: C

Q459. A swap space on a disk is used for:

- A. Disk defragmentation temporary workspace
- B. Exchanging files between different users
- C. Temporary file storage for applications
- D. Storing pages swapped out from memory

Answer: D

Q460. The boot block on a disk contains:

- A. Bootstrap code to load OS
- B. User data and documents
- C. Error correction for sectors
- D. File system metadata tables

Answer: A

Q461. What is the difference between programmed I/O, interrupt-driven I/O, and DMA?

- A. Programmed polls; interrupt frees CPU; DMA bypasses CPU
- B. They only work with specific hardware device types
- C. They only differ in speed of the data transfer
- D. They are all identical in behavior and speed

Answer: A

Q462. What is memory-mapped I/O?

- A. Storing I/O data into a temporary buffer
- B. Mapping device registers to memory addresses
- C. A type of virtual memory for page tables
- D. Mapping files from disk into virtual memory

Answer: B

Q463. What is the purpose of the I/O subsystem in the OS kernel?

- A. To manage user accounts and authentication
- B. To compile programs into executable binaries
- C. To provide games and multimedia for users
- D. Uniform interfaces for buffering and scheduling

Answer: D

Q464. Double buffering uses:

- A. No buffers needed at all ever
- B. Three buffers in a rotation cycle
- C. One buffer for all data transfers
- D. Two buffers: fill one, process other

Answer: D

Q465. Caching in the I/O system differs from buffering in that:

- A. They are identical in purpose and behavior
- B. Buffering stores during transfer; cache keeps copies
- C. Caching is always slower than buffering in practice
- D. Buffering is used only for output data operations

Answer: B

Q466. What is a device controller's role?

- A. To manage hardware and translate OS commands
- B. To manage virtual memory paging and swapping
- C. To compile device driver source code files
- D. To schedule processes for CPU execution time

Answer: A

Q467. What is the difference between blocking and non-blocking I/O?

- A. Blocking is always faster than non-blocking I/O
- B. They are identical in every way and behavior
- C. Non-blocking prevents all I/O from starting
- D. Blocking suspends until done; non-blocking returns

Answer: D

Q468. What is the purpose of an interrupt controller?

- A. To block all incoming interrupt signals
- B. To generate interrupt signals to CPU
- C. To manage, prioritize, route interrupts
- D. To count total number of interrupts

Answer: C

Q469. What is the unified buffer cache?

- A. A cache located inside the CPU for all instructions
- B. Single cache for files and mapped memory, no duplication
- C. A cache shared by all system users equally
- D. A universal cache for every device in the system

Answer: B

Q470. Error handling in I/O systems typically involves:

- A. Shutting down the entire system on any error
- B. Deleting the faulty device from the system table
- C. Returning codes, retrying, logging, informing users
- D. Ignoring all errors that occur during I/O

Answer: C

Q471. An access control matrix is:

- A. A type of encryption algorithm for data security
- B. A type of visual display device for output
- C. A hardware protection mechanism in the CPU unit
- D. A model specifying subject access rights to objects

Answer: D

Q472. Access Control Lists (ACLs) are a column-based implementation of the access control matrix where:

- A. Each object stores subjects and their access rights
- B. A central authority stores all access control rights
- C. Access rights are not stored anywhere in the system
- D. Each subject stores its own access rights list

Answer: A

Q473. Capability lists are a row-based implementation of the access control matrix where:

- A. Capabilities are stored in a central directory
- B. Each object stores its own capability tokens
- C. Each subject holds objects and allowed operations
- D. Capabilities are not used in any modern system

Answer: C

Q474. A denial-of-service (DoS) attack:

- A. Modifies system files to gain higher access levels
- B. Steals confidential data from the system quietly
- C. Overwhelms a system to deny service to users
- D. Encrypts user files and demands ransom payment

Answer: C

Q475. A buffer overflow attack exploits:

- A. Too much available memory in the system overall
- B. Overwriting memory beyond a buffer boundary
- C. Network buffer limits causing packet loss
- D. Disk buffer overflow causing data corruption

Answer: B

Q476. In UNIX file permissions, the permission mode 755 means:

- A. Owner has rwx, group and others have r-x access
- B. Everyone has full unrestricted access to file
- C. The file is hidden from all users in system
- D. Only the owner can access and modify the file

Answer: A

Q477. What is the difference between symmetric and asymmetric encryption?

- A. Asymmetric is faster than symmetric in practice
- B. Symmetric uses one key; asymmetric uses key pair
- C. They are the same algorithm for encryption use
- D. Symmetric is always more secure than asymmetric

Answer: B

Q478. What is a protection domain?

- A. A set of access rights for process resources
- B. A physical area in the server room building
- C. A network domain managed by a DNS server
- D. A type of firewall zone for network security

Answer: A

Q479. Multifactor authentication combines:

- A. Using multiple different usernames
- B. Combining different auth factor types
- C. Allowing multiple login attempts
- D. Using multiple passwords for login

Answer: B

Q480. What is the role of a digital certificate in security?

- A. To encrypt files for security
- B. To manage user passwords securely
- C. To bind public key to identity via CA
- D. To compress data for efficient storage

Answer: C

Q481. What distinguishes a microkernel architecture from a monolithic kernel?

- A. Microkernels run all services in kernel space for speed
- B. Microkernels move most services to user space processes
- C. Monolithic kernels keep only minimal services in kernel
- D. Both architectures are identical in design and approach

Answer: B

Q482. What is the purpose of dual-mode operation in an operating system?

- A. To double the processing speed of the CPU cores
- B. To protect the system from errant user programs
- C. To run two operating systems simultaneously here
- D. To allow two separate users to log in together

Answer: B

Q483. How does a hypervisor differ from a traditional operating system?

- A. A hypervisor is just another name for an OS kernel
- B. A hypervisor only runs on mobile phone platforms
- C. A hypervisor manages virtual machines on hardware
- D. A hypervisor cannot manage hardware resources ever

Answer: C

Q484. What is the key advantage of time-sharing over batch processing?

- A. Batch processing provides faster interactive response
- B. Time-sharing eliminates the need for any memory
- C. Time-sharing provides interactive response to users
- D. Time-sharing consumes less CPU than batch does here

Answer: C

Q485. What role does the DMA controller play in an operating system?

- A. It compiles source code into machine language
- B. It schedules processes for execution on processors
- C. It manages user authentication and access control
- D. It transfers data between memory and I/O directly

Answer: D

Q486. Which statement about system programs is correct?

- A. They replace the need for an operating system here
- B. They provide convenient environment for development
- C. System programs only include text editor programs
- D. System programs are the same as the OS kernel

Answer: B

Q487. What is the trap instruction used for in operating systems?

- A. To physically trap faulty hardware in the system
- B. To prevent users from logging into the system
- C. To switch from user mode to kernel mode safely
- D. To catch and eliminate computer virus threats

Answer: C

Q488. What is the difference between policy and mechanism in OS design?

- A. They are interchangeable terms with the same meaning
- B. Policy refers to hardware; mechanism to software here
- C. Mechanism decides what; policy decides how it is done
- D. Policy decides what to do; mechanism decides how

Answer: D

Q489. What happens during a context switch in an operating system?

- A. The computer restarts and reloads the entire system
- B. All running processes are terminated and then removed
- C. The OS saves one process state and loads another
- D. The kernel is replaced with a newer version of code

Answer: C

Q490. What is the purpose of the OS system call interface?

- A. It controls voltage supply to all peripheral devices
- B. It connects the monitor cable directly to CPU chip
- C. It manages the physical layout of the motherboard
- D. It provides a layer between user programs and kernel

Answer: D

Q491. What is the main advantage of a layered OS structure?

- A. It eliminates the need for any system call interfaces
- B. It provides the fastest possible system performance now
- C. It requires less memory than all other architectures
- D. Each layer can be debugged and verified independently

Answer: D

Q492. How do loadable kernel modules benefit the operating system?

- A. They allow extending kernel functionality at runtime easily
- B. They eliminate all system calls from the OS call interface
- C. They require the entire kernel to be recompiled each time
- D. They replace the need for any device drivers completely now

Answer: A

Q493. What is the role of the system call interface in OS architecture?

- A. It only manages file system operations and nothing else
- B. It replaces the kernel in all microkernel architectures
- C. It serves as boundary between user programs and the OS
- D. It provides direct hardware access to user-level programs

Answer: C

Q494. What is a key disadvantage of the simple monolithic structure?

- A. It cannot access any hardware devices at all currently
- B. It is very slow due to many excessive mode switches
- C. It does not support any form of multitasking operation
- D. It is hard to maintain because parts are tightly coupled

Answer: D

Q495. How does message passing work in a microkernel OS design?

- A. Services exchange messages through microkernel's IPC
- B. Services call each other like regular function calls
- C. All communication happens directly in shared memory
- D. Services communicate through shared global variables

Answer: A

Q496. What is the parameter table approach for passing system call arguments?

- A. It limits system calls to only three parameters total
- B. It stores system call results on the hard disk drive
- C. Parameters are stored in a table; address passed in register
- D. It eliminates the need for any CPU registers entirely

Answer: C

Q497. What advantage does a microkernel have over a monolithic kernel?

- A. Microkernels have more code running inside kernel space
- B. Microkernels always have better raw system performance
- C. Microkernels are easier to extend and port to new hardware
- D. Microkernels eliminate the need for device driver support

Answer: C

Q498. What is the difference between static and dynamic linking?

- A. Dynamic libraries must be recompiled with each application
- B. Static libraries are linked at compile time into binary
- C. Static libraries are always smaller than dynamic libraries
- D. Dynamic libraries are slower but use more total disk space

Answer: B

Q499. Why is debugging a layered OS more straightforward than others?

- A. Because layered systems do not use system calls at all
- B. Because errors in one layer are isolated from other ones
- C. Because all layers share the exact same address space
- D. Because layered systems have no kernel component in them

Answer: B

Q500. What role does the POSIX standard play in operating system design?

- A. It describes physical network topologies for computers
- B. It specifies a standard OS interface for portability use
- C. It defines a mandatory hardware specification for PCs
- D. It is a language used specifically to write OS kernels

Answer: B

Q501. What differs between the short-term and long-term scheduler?

- A. Short-term selects from ready queue; long-term from disk
- B. Long-term runs more frequently than short-term scheduler
- C. Short-term controls the degree of multiprogramming level
- D. They are the same scheduler with two different OS names

Answer: A

Q502. What is a zombie process and why does it occur?

- A. A process created but never actually started running
- B. A terminated process whose parent has not called wait
- C. A process paused by the operating system permanently
- D. A process that runs forever consuming all resources

Answer: B

Q503. How does exec() differ from fork() in Unix systems?

- A. Neither affects the process structure in any way here
- B. exec() creates a new process; fork() replaces image
- C. Both create new processes in exactly identical fashion
- D. fork() creates new process; exec() replaces the image

Answer: D

Q504. What is inter-process communication (IPC) used for?

- A. Connecting a computer to external network devices
- B. Compiling multiple source files at the same time here
- C. Managing CPU scheduling algorithms exclusively now
- D. Allowing processes to exchange data and synchronize

Answer: D

Q505. How do shared memory and message passing IPC differ?

- A. Neither mechanism allows processes to exchange any data
- B. Both mechanisms work identically in all systems always
- C. Message passing uses shared region; shared uses kernel
- D. Shared memory uses common region; messages use kernel

Answer: D

Q506. What is the purpose of the medium-term scheduler?

- A. It schedules I/O operations for device driver code
- B. It manages the file system directory tree structure
- C. It assigns priorities to all newly created processes
- D. It swaps processes between main memory and the disk

Answer: D

Q507. How does a pipe work as an IPC mechanism in Unix?

- A. It provides unidirectional byte stream between procs
- B. It only works between processes on different machines
- C. It replaces the need for any process synchronization
- D. It provides bidirectional communication always here

Answer: A

Q508. What is a process address space composed of?

- A. Exclusively the operating system kernel code only
- B. Only the stack and no other memory segment data
- C. Text, data, heap, and stack segments combined
- D. Only the program code and nothing else at all

Answer: C

Q509. What is cascading termination in process management?

- A. Restarting a process multiple times before stopping
- B. Terminating all children when parent process exits
- C. Gradually slowing down a process before termination
- D. Saving the process state to disk before termination

Answer: B

Q510. What is the role of the dispatcher in process management?

- A. It gives CPU control to the process the scheduler chose
- B. It manages file system allocation tables on the disk
- C. It handles all network communication between hosts
- D. It creates new processes from executable binaries

Answer: A

Q511. What is the convoy effect in FCFS scheduling?

- A. All processes complete at exactly the same time
- B. Short processes wait behind a long CPU-bound one
- C. The CPU is idle while waiting for process arrival
- D. Processes are sorted by priority before execution

Answer: B

Q512. How does the Shortest Remaining Time First algorithm work?

- A. It preempts if a new process has shorter remaining time
- B. It always runs the longest process first for throughput
- C. It is identical to non-preemptive SJF in every way
- D. It uses Round Robin with variable time quantum sizes

Answer: A

Q513. What problem does aging solve in priority scheduling?

- A. It prevents context switching between processes now
- B. It increases the time quantum in Round Robin schemes
- C. It solves starvation for low-priority waiting processes
- D. It eliminates the need for any scheduling algorithm

Answer: C

Q514. What happens with a very small time quantum in Round Robin?

- A. It reduces context switches significantly overall
- B. It increases context switch overhead and degrades speed
- C. It eliminates the need for a ready queue in the system
- D. It converts Round Robin into FCFS scheduling behavior

Answer: B

Q515. How does multilevel queue scheduling organize processes?

- A. Each process gets its own dedicated scheduling queue
- B. Processes are divided into classes with separate queues
- C. All processes share a single common ready queue only
- D. Processes are randomly assigned to available CPU cores

Answer: B

Q516. What is dispatch latency in CPU scheduling?

- A. Time to transfer data from main memory to the CPU
- B. Time to stop one process and start running another
- C. Time for a process to complete its entire execution
- D. Time between submitting a job and its first output

Answer: B

Q517. What is response time and why is it important?

- A. It counts total number of context switches performed
- B. It is time from submission to first response produced
- C. It is the total time to complete the process entirely
- D. It measures amount of memory a process consumes now

Answer: B

Q518. How does SJF achieve optimal average waiting time?

- A. By using a large time quantum to minimize switching
- B. By assigning equal CPU time to every queued process
- C. By running the longest jobs first to clear the queue
- D. By placing shortest jobs first reducing total wait time

Answer: D

Q519. What distinguishes multilevel feedback queue from multilevel queue?

- A. Multilevel queues allow movement but feedback does not
- B. Feedback queues allow processes to move between queues
- C. Both scheduling algorithms are identical in all behavior
- D. Feedback queues do not allow preemption of processes

Answer: B

Q520. What is the main disadvantage of Shortest Job First?

- A. It needs CPU burst time which is hard to predict ahead
- B. It cannot work with more than two processes at a time
- C. It always has the worst average waiting time possible
- D. It causes excessive context switching between processes

Answer: A

Q521. What is the many-to-one threading model?

- A. Many user threads map to one kernel thread
- B. Each user thread maps to one kernel thread
- C. Many kernel threads map to one user thread
- D. Multiple kernel threads share one CPU core

Answer: A

Q522. What is the one-to-one threading model's main advantage?

- A. It eliminates the need for any synchronization at all
- B. It uses less memory than all other threading models
- C. It allows true concurrency on multiprocessor systems
- D. It does not need any operating system kernel support

Answer: C

Q523. What is the purpose of thread pools in multithreaded systems?

- A. To pre-create threads and reuse them for new requests
- B. To prevent any thread from being created in a system
- C. To merge all threads into a single execution context
- D. To permanently store thread objects on the hard disk

Answer: A

Q524. What is thread-local storage (TLS) used for?

- A. Storing thread data on a remote network file server
- B. Sharing data between all threads in a process
- C. Encrypting data that threads access in shared memory
- D. Giving each thread its own copy of specific data

Answer: D

Q525. How does the many-to-many threading model work?

- A. It eliminates both user and kernel threads entirely
- B. It maps one user thread to many kernel threads here
- C. It maps many kernel threads to one user thread now
- D. It maps many user threads to many kernel threads

Answer: D

Q526. What is the signal handling challenge in multithreaded programs?

- A. Signals are only relevant to single-threaded programs
- B. Deciding which thread should receive and handle signal
- C. Signals cannot be generated in multithreaded programs
- D. All signals are handled automatically by the OS kernel

Answer: B

Q527. What is the difference between implicit and explicit threading?

- A. Both approaches are identical in their implementation
- B. Implicit threading needs manual thread creation code
- C. Explicit threading is handled automatically by system
- D. Implicit threading is done by compilers and libraries

Answer: D

Q528. What happens during deferred thread cancellation?

- A. Thread cancellation is impossible in deferred mode here
- B. The target thread is terminated immediately on request
- C. Thread checks for cancellation at cancellation points
- D. The entire process is terminated along with all threads

Answer: C

Q529. Why is fork() complex in multithreaded programs?

- A. fork() terminates the calling thread before the child
- B. It must decide whether to duplicate all or one thread
- C. fork() cannot be called from multithreaded programs
- D. fork() always duplicates all threads without exception

Answer: B

Q530. What is OpenMP and how does it support threading?

- A. OpenMP is a hardware spec for multicore CPUs only
- B. OpenMP is a network protocol for distributed systems
- C. OpenMP provides compiler directives for parallel code
- D. OpenMP is an operating system for embedded controllers

Answer: C

Q531. What are the three requirements for a correct critical section solution?

- A. Speed, memory efficiency, and low power consumption
- B. Simplicity, readability, and backward compatibility now
- C. Atomicity, consistency, isolation, and full durability
- D. Mutual exclusion, progress, and bounded waiting time

Answer: D

Q532. What is the producer-consumer problem in synchronization?

- A. A technique for optimizing database query performance
- B. Coordinating producer adding and consumer removing
- C. A problem about consumer electronics device drivers
- D. A problem about manufacturing physical hardware

Answer: B

Q533. How does Peterson's solution achieve mutual exclusion?

- A. It uses a centralized server to manage all locks
- B. It uses two shared variables: turn and flag array
- C. It uses hardware test-and-set atomic instructions
- D. It relies on disabling all interrupts on processor

Answer: B

Q534. How do spinlocks differ from blocking locks?

- A. Both types always put waiting threads to sleep now
- B. Spinlocks sleep; blocking locks use busy waiting
- C. Both types always use busy waiting to check locks
- D. Spinlocks busy-wait; blocking locks put threads asleep

Answer: D

Q535. What is a monitor in concurrent programming?

- A. A display screen for viewing program output
- B. A network device monitoring all traffic on LAN
- C. A high-level synchronization construct with mutex
- D. A tool for monitoring CPU temperature and speed

Answer: C

Q536. What is the purpose of condition variables in monitors?

- A. They store values of conditional expressions only
- B. They replace the need for any mutual exclusion now
- C. They count threads inside the monitor at any time
- D. They let threads wait for specific conditions safely

Answer: D

Q537. What is the readers-writers problem in synchronization?

- A. A problem about reading user input from a keyboard
- B. Allowing concurrent reads but exclusive writes to data
- C. A technique for writing data to multiple disk drives
- D. A problem about reading and writing files in order

Answer: B

Q538. What is the test-and-set instruction used for?

- A. Testing network and setting bandwidth limits on it
- B. Setting CPU clock speed after testing the hardware
- C. Testing programs and setting breakpoints in them
- D. Atomically testing and modifying a word for locks

Answer: D

Q539. What is the compare-and-swap (CAS) instruction?

- A. A sorting algorithm comparing and swapping values
- B. A technique to compare files and swap if different
- C. A method to swap two processes in scheduling queue
- D. Atomic instruction conditionally updating memory word

Answer: D

Q540. What is a barrier synchronization mechanism?

- A. A way to permanently block a thread from running
- B. A point where all threads must arrive before any go
- C. A physical barrier between two system processors
- D. A mechanism preventing threads from being created

Answer: B

Q541. How does the Banker's algorithm determine if a state is safe?

- A. It finds a safe sequence where all processes complete
- B. It only checks if CPU utilization is above fifty percent
- C. It checks if total resources exceed process count
- D. It verifies no process has any resource allocated now

Answer: A

Q542. How can the hold-and-wait condition be prevented?

- A. By removing mutual exclusion from all shared resources
- B. By giving every process all existing system resources
- C. By requiring processes to request all resources at once
- D. By allowing processes to request at any time freely

Answer: C

Q543. How do deadlock prevention and avoidance differ?

- A. Prevention ignores deadlocks; avoidance terminates all tasks
- B. Prevention detects deadlocks; avoidance recovers from them
- C. Prevention constrains requests; avoidance checks each allocation
- D. They are exactly the same strategy with different names now

Answer: C

Q544. How does total ordering on resources prevent circular wait?

- A. It allows processes to request resources in any order
- B. It makes all resources shareable between all processes
- C. It eliminates the need for any resource allocation logic
- D. Processes must request resources in increasing order now

Answer: D

Q545. What is deadlock detection using a wait-for graph?

- A. A graph showing CPU utilization over time periods
- B. A diagram of physical network connections between PCs
- C. A chart showing memory usage of each running process
- D. A variant with only processes; a cycle means deadlock

Answer: D

Q546. What are the recovery options after detecting a deadlock?

- A. Only restarting the entire computer system works now
- B. Process termination or resource preemption can be used
- C. Deadlocks always resolve themselves automatically here
- D. Adding more physical memory always resolves all deadlock

Answer: B

Q547. What is the ostrich algorithm for handling deadlocks?

- A. A sophisticated algorithm detecting all deadlocks fast
- B. An algorithm preventing deadlock with resource ordering
- C. It ignores the problem assuming deadlocks rarely occur
- D. A technique recovering from deadlocks automatically now

Answer: C

Q548. What is a livelock and how does it differ from deadlock?

- A. Livelock means all processes are blocked and inactive
- B. In livelock processes are active but make no progress
- C. Livelock only occurs in single-threaded applications
- D. Livelock is identical to deadlock in every way here

Answer: B

Q549. What does the Banker's algorithm need from processes?

- A. The total disk space used by each process on system
- B. Only the current CPU usage of each active process
- C. Maximum resource needs declared by each process first
- D. Only the priority level assigned to each process now

Answer: C

Q550. Why might deadlock prevention lead to low resource utilization?

- A. Because conservative restrictions leave resources idle now
- B. Because it allocates all resources to one process always
- C. Because it detects deadlocks only after they have formed
- D. Because it allows processes to hold resources indefinitely

Answer: A

Q551. How does the LRU page replacement algorithm work?

- A. It replaces the page not used for the longest time
- B. It randomly selects any page for replacement now
- C. It replaces the page used furthest in the future
- D. It replaces pages in the order they were loaded in

Answer: A

Q552. What is the Translation Lookaside Buffer (TLB)?

- A. A register storing the base address of process
- B. A buffer for storing data being written to disk
- C. A queue for processes waiting for memory alloc
- D. A fast cache for recent page table entries only

Answer: D

Q553. What is Belady's anomaly in page replacement?

- A. Page faults are constant regardless of the frame count
- B. More frames always reduce the number of page faults
- C. Belady's anomaly only occurs with the optimal algorithm
- D. More frames can increase page faults with FIFO scheme

Answer: D

Q554. How does segmentation differ from paging?

- A. Neither technique uses any address translation now
- B. Both use exactly the same fixed-size memory blocks
- C. Segmentation uses fixed-size; paging uses variable
- D. Paging uses fixed-size; segmentation uses variable

Answer: D

Q555. What is demand paging in virtual memory?

- A. Loading all pages into memory at startup
- B. Requiring users to manually load needed pages
- C. Pre-loading all pages from disk before start
- D. Loading pages only when they are accessed now

Answer: D

Q556. What is the working set model in virtual memory?

- A. Pages a process actively uses in time window
- B. Set of all pages stored on the hard disk now
- C. Total physical memory available in computer
- D. Total number of processes in ready queue

Answer: A

Q557. What is thrashing in virtual memory systems?

- A. A process running very fast and finishing quickly
- B. The process of defragmenting virtual memory space
- C. A technique to improve memory utilization efficiency
- D. Excessive page faults causing most time spent swapping

Answer: D

Q558. How does the FIFO page replacement algorithm work?

- A. It replaces the page with fewest access counts
- B. It replaces pages by their physical address now
- C. It replaces the most recently loaded page here
- D. It replaces the oldest page that was loaded in

Answer: D

Q559. What is the dirty bit used for in paging?

- A. It marks pages containing errors or corrupted data
- B. It indicates whether a page was modified in RAM
- C. It shows whether a page is currently locked in RAM
- D. It counts the number of times a page was accessed

Answer: B

Q560. What is a multi-level page table and why is it used?

- A. A page table stored on multiple different disk drives
- B. A hierarchical page table reducing memory overhead
- C. A page table storing entries for multiple processes
- D. A table mapping pages to multiple physical frames

Answer: B

Q561. How does indexed allocation manage file disk blocks?

- A. Blocks are linked using pointers in each block
- B. An index block contains pointers to data blocks
- C. Blocks are allocated randomly without tracking
- D. All blocks are stored consecutively on disk

Answer: B

Q562. What differs between a hard link and a symbolic link?

- A. Hard links work across file systems but soft do not
- B. Hard link points to inode; soft link points to path
- C. Both hard and symbolic links are exactly identical
- D. Symbolic links point to inodes; hard links to paths

Answer: B

Q563. What is the inode structure in Unix file systems?

- A. A structure storing file metadata and block pointers
- B. A network node handling file transfer protocols
- C. A hardware component managing disk read operations
- D. A process monitoring file system integrity at boot

Answer: A

Q564. How does the FAT file system organize file allocation?

- A. Using a file allocation table chaining disk blocks
- B. Using B-tree indexes for each file on the disk now
- C. Using extent-based allocation with balanced trees
- D. Using inodes and indirect blocks like Unix systems

Answer: A

Q565. What is journaling in modern file systems?

- A. Recording user activity logs for security audit
- B. Writing metadata changes to a log before commit
- C. Storing journal entries for file content searching
- D. Maintaining daily backup schedule of all files now

Answer: B

Q566. What is the free-space management bitmap method?

- A. A method to compress free space for more storage
- B. A technique to visualize the directory tree layout
- C. Each bit represents a block allocation status bit
- D. A graphical representation of files on screen

Answer: C

Q567. What is direct access file organization?

- A. Any record can be accessed directly by position
- B. Files can only be read sequentially from start
- C. Files are accessed by searching all records first
- D. Files can only be accessed through symbolic links

Answer: A

Q568. What is the purpose of the Virtual File System layer?

- A. It provides uniform interface to file system types
- B. It creates virtual copies of files for backup
- C. It manages virtual memory not file system at all
- D. It encrypts all files stored on virtual disk drives

Answer: A

Q569. What is file system consistency checking used for?

- A. Ensuring all files have correct extensions set here
- B. Checking if files contain valid program source
- C. Verifying integrity and repairing corruption now
- D. Checking if file system has enough free space left

Answer: C

Q570. How does contiguous allocation organize file storage?

- A. Files are indexed using a multi-level tree structure
- B. Each file occupies contiguous disk blocks together
- C. Files are in random blocks scattered across the disk
- D. Files are stored using linked list of distant blocks

Answer: B

Q571. How does the SCAN disk scheduling algorithm work?

- A. Head moves one direction servicing then reverses
- B. It randomly selects the next request from queue
- C. It services requests in arrival order always
- D. It always jumps to the closest request first now

Answer: A

Q572. How does SSTF scheduling differ from FCFS?

- A. SSTF services requests in the exact arrival order
- B. SSTF always moves the head to outermost track first
- C. SSTF and FCFS produce identical scheduling results
- D. SSTF services the request nearest to current head

Answer: D

Q573. What is the difference between RAID 0 and RAID 1?

- A. RAID 0 provides redundancy; RAID 1 stripes only
- B. RAID 0 stripes for speed; RAID 1 mirrors for safety
- C. Both provide identical performance and safety here
- D. RAID 0 uses parity; RAID 1 uses error correction

Answer: B

Q574. What is the C-SCAN disk scheduling algorithm?

- A. It services the closest request to minimize seeking
- B. It processes requests in the order submitted always
- C. It scans in both directions alternately each time
- D. Head services one direction then jumps back to start

Answer: D

Q575. What is the purpose of disk defragmentation for HDDs?

- A. Encrypting all data stored on the hard disk drive
- B. Removing all deleted files permanently from disk
- C. Reorganizing fragmented files into contiguous blocks
- D. Increasing physical speed of disk platter rotation

Answer: C

Q576. How does RAID 5 provide fault tolerance?

- A. It keeps three copies of every data block on disk
- B. It distributes parity information across all drives
- C. It mirrors all data on two separate disk drives
- D. It uses error-correcting codes on a separate disk

Answer: B

Q577. Why is SSTF disk scheduling prone to starvation?

- A. Far requests may wait indefinitely for any service
- B. Because it services all requests equally in order
- C. It services farthest request first causing slow times
- D. SSTF scheduling never causes starvation of requests

Answer: A

Q578. What is the disk controller's role?

- A. It manages CPU scheduling for all running processes
- B. It handles network routing for file sharing on LAN
- C. It controls user interface for file management
- D. It translates logical addresses to physical locations

Answer: D

Q579. What is wear leveling in solid-state drives?

- A. Technique to physically level SSD circuit boards
- B. Distributing writes evenly to extend cell lifespan
- C. A process to cool down overheated memory chips now
- D. A method to reduce power consumption of SSD drive

Answer: B

Q580. What is the boot block and where is it stored?

- A. It stores the file system journal for crash recovery
- B. It stores user data in the first track of disk
- C. It contains bootstrap code in a fixed disk location
- D. It contains disk encryption keys for secure access

Answer: C

Q581. How do memory-mapped and port-mapped I/O differ?

- A. Memory-mapped requires special CPU instructions for access
- B. Port-mapped I/O is faster than memory-mapped in every case
- C. Both methods use exactly the same addressing mechanism now
- D. Memory-mapped uses address space; port-mapped uses separate

Answer: D

Q582. How does double buffering improve I/O performance?

- A. One buffer fills while other is being processed by CPU
- B. It eliminates the need for any buffering in the system
- C. It slows down I/O by requiring twice the memory space
- D. It doubles the speed of the physical I/O device hardware

Answer: A

Q583. What is device-independent I/O software?

- A. It handles only network device communication protocol
- B. It makes all devices behave like character devices
- C. It eliminates the need for any device drivers at all
- D. It provides uniform interfacing for diverse devices

Answer: D

Q584. What is an interrupt vector table used for?

- A. Storing data to be written to I/O devices later
- B. Tracking which I/O devices are powered on currently
- C. Mapping interrupt numbers to handler routines fast
- D. Counting total number of interrupts since boot time

Answer: C

Q585. What is the role of the I/O scheduler?

- A. It schedules manufacturing of I/O device hardware
- B. It orders and merges I/O requests for optimal access
- C. It schedules network packets for wire transmission
- D. It determines which user can log into the system now

Answer: B

Q586. How does spooling differ from buffering?

- A. Spooling uses disk as buffer for slow device queues
- B. Buffering uses disk while spooling uses only main memory
- C. Spooling is only used for input devices like keyboards
- D. Spooling and buffering are the same I/O technique here

Answer: A

Q587. What is the device status table?

- A. Storing serial numbers of all connected devices here
- B. Maintaining firmware update history for every device
- C. Displaying device information to end user on screen
- D. Tracking status and request queue for each device now

Answer: D

Q588. How do blocking and non-blocking I/O differ?

- A. Blocking suspends process; non-blocking returns immediately
- B. Both blocking and non-blocking I/O work the same way
- C. Blocking I/O is always faster than non-blocking I/O now
- D. Non-blocking suspends process until I/O completes fully

Answer: A

Q589. What is the DMA controller's role during transfer?

- A. It performs all CPU computations during the transfer
- B. It compiles programs while data is being transferred
- C. It encrypts all data before transferring it to memory
- D. It manages data transfer between device and memory now

Answer: D

Q590. How does asynchronous I/O differ from non-blocking?

- A. Asynchronous I/O is identical to non-blocking I/O now
- B. Non-blocking sends completion notification to caller
- C. Asynchronous notifies process upon I/O completion later
- D. Asynchronous blocks process until data is fully ready

Answer: C

Q591. What is a buffer overflow attack?

- A. It fills printer buffer with too many print jobs
- B. It fills the disk buffer causing file system crash
- C. It overwrites memory beyond buffer to inject code
- D. It overflows network buffer causing packet loss now

Answer: C

Q592. What is role-based access control (RBAC)?

- A. Assigning permissions to roles then roles to users
- B. Assigning permissions directly to individual users
- C. Assigning all permissions to every user in system
- D. Removing all access control from the OS entirely

Answer: A

Q593. How does the access matrix represent protection?

- A. Rows are resources and columns are users only now
- B. It is a physical matrix of hardware protection now
- C. Rows are domains, columns are objects with rights
- D. It only represents network access permissions here

Answer: C

Q594. What is a denial-of-service attack?

- A. Deleting all files from the target system forever
- B. Denying users permission to install software
- C. Encrypting data and demanding ransom for the key
- D. Overwhelming system with requests to make unavailable

Answer: D

Q595. What is a rootkit and why is it dangerous?

- A. Malware hiding itself providing persistent root access
- B. A security scanner checking for root vulnerabilities
- C. A toolkit for administrators to manage root access
- D. A legitimate tool for recovering forgotten passwords

Answer: A

Q596. How does sandboxing improve application security?

- A. It isolates apps in restricted execution environments
- B. It prevents applications from being installed at all
- C. It buries security threats deep in system storage
- D. It encrypts all data processed by sandboxed apps now

Answer: A

Q597. What is two-factor authentication (2FA)?

- A. Using same password on two different computers
- B. Logging into two different accounts at same time
- C. Using two different passwords for one account
- D. Requiring two independent verification for access

Answer: D

Q598. What is an intrusion detection system (IDS)?

- A. Installing security patches on system software now
- B. Managing user passwords and auth credentials only
- C. Monitoring system activity to detect malicious acts
- D. Detecting physical break-ins at office buildings

Answer: C

Q599. How do symmetric and asymmetric encryption differ?

- A. Symmetric uses two keys; asymmetric uses one shared
- B. Both encryption types use the same key mechanism now
- C. Symmetric encryption is always weaker than asymmetric
- D. Symmetric uses one shared key; asymmetric a key pair

Answer: D

Q600. What is privilege escalation in system security?

- A. Exploiting flaws to gain higher privileges than allowed
- B. Reducing user privileges for security enforcement now
- C. Escalating a security ticket to management team here
- D. Legitimately granting more permissions by an admin

Answer: A

Q601. What is the difference between a hard real-time system and a soft real-time system?

- A. Hard real-time guarantees deadlines while soft real-time allows occasional misses
- B. Hard real-time uses hardware timers while soft real-time uses software timers
- C. Hard real-time supports more users than soft real-time
- D. Hard real-time is faster than soft real-time in all scenarios

Answer: A

Q602. How does a multiprocessor system differ from a multicomputer system?

- A. Multiprocessor shares memory while multicomputer has distributed memory
- B. Multiprocessor uses one OS while multicomputer uses multiple keyboards
- C. Multiprocessor is slower than multicomputer in all workloads
- D. Multiprocessor only runs batch jobs while multicomputer handles interactive tasks

Answer: A

Q603. What is the role of a timer interrupt in operating system design?

- A. To prevent user processes from monopolizing the CPU indefinitely
- B. To speed up arithmetic calculations
- C. To increase the amount of available memory
- D. To manage network packet routing

Answer: A

Q604. What is the purpose of privileged instructions in a CPU?

- A. To restrict certain dangerous operations to kernel mode only
- B. To make programs execute faster in user mode
- C. To allow users to directly access I/O devices
- D. To enable parallel processing on multiple cores

Answer: A

Q605. How does symmetric multiprocessing (SMP) differ from asymmetric multiprocessing?

- A. SMP treats all processors equally while asymmetric assigns a master processor
- B. SMP uses one processor while asymmetric uses many
- C. SMP is only for real-time systems while asymmetric is for batch systems
- D. SMP requires identical software on each processor while asymmetric does not

Answer: A

Q606. What is the purpose of a system call in the operating system?

- A. To provide an interface for user programs to request OS kernel services
- B. To compile high-level code into machine language
- C. To physically connect hardware devices to the motherboard
- D. To partition the hard disk into volumes

Answer: A

Q607. What does Direct Memory Access (DMA) accomplish in a computer system?

- A. Allows devices to transfer data to memory without CPU intervention for each byte
- B. Directly increases the size of physical memory
- C. Enables memory to execute instructions without the CPU
- D. Provides encryption for data stored in memory

Answer: A

Q608. What is the memory hierarchy in a computer system?

- A. Arrangement of storage from fast/small registers to slow/large secondary storage
- B. A method for encrypting data in memory
- C. The process of allocating memory to user processes
- D. A technique for compressing files on disk

Answer: A

Q609. What is the purpose of an interrupt vector in an operating system?

- A. To store addresses of interrupt service routines for each interrupt type
- B. To count the number of interrupts that have occurred
- C. To prioritize user processes over system processes
- D. To buffer data during I/O transfers

Answer: A

Q610. What does it mean for an operating system to provide abstraction?

- A. It hides complex hardware details behind simpler interfaces
- B. It makes the hardware run at faster clock speeds
- C. It reduces the physical size of hardware components
- D. It eliminates the need for device drivers entirely

Answer: A

Q611. What is the main trade-off between monolithic and microkernel architectures?

- A. Monolithic has better performance but less modularity than microkernel
- B. Monolithic uses less memory than microkernel in all cases
- C. Microkernel is always faster because it has a smaller kernel
- D. Monolithic supports more hardware devices than microkernel

Answer: A

Q612. How are parameters passed to system calls via registers?

- A. Parameter values are placed directly in CPU registers before the system call
- B. Parameters are written to a file that the kernel reads
- C. Parameters are encoded in the system call number itself
- D. Parameters are passed through network sockets

Answer: A

Q613. What is the key benefit of loadable kernel modules over a static monolithic kernel?

- A. New functionality can be added at runtime without recompiling the kernel
- B. They eliminate the need for device drivers entirely
- C. They make the kernel run in user space
- D. They remove the need for system calls

Answer: A

Q614. What is the run-time environment (RTE) in the context of system calls?

- A. A software layer that intercepts API calls and invokes the correct system calls
- B. A special hardware unit that executes system calls
- C. A debugging tool for testing operating system code
- D. A network protocol for remote procedure calls

Answer: A

Q615. Why is inter-process communication slower in microkernel systems?

- A. Services communicate via message passing through the kernel rather than direct calls
- B. The microkernel uses slower processors
- C. User space runs at reduced clock speed
- D. Microkernel systems cannot use DMA for I/O

Answer: A

Q616. What is the purpose of the /proc file system in Linux?

- A. It provides a virtual file interface to kernel data structures and process info
- B. It stores user documents and personal files
- C. It contains compiled application programs
- D. It holds backup copies of the boot loader

Answer: A

Q617. What is the difference between a policy and a mechanism in OS design?

- A. Mechanism defines how to do something, policy decides what should be done
- B. Mechanism is software while policy is hardware
- C. Policy runs in kernel mode while mechanism runs in user mode
- D. There is no practical difference between the two

Answer: A

Q618. How does the stack-based method of passing system call parameters work?

- A. Parameters are pushed onto the program stack and the kernel pops them off
- B. Parameters are stored in a special CPU cache
- C. Parameters are passed as part of the instruction opcode
- D. Parameters are sent via interrupt signals

Answer: A

Q619. What are information maintenance system calls used for?

- A. Transferring information between the user program and the OS
- B. Creating new user accounts
- C. Formatting storage devices
- D. Establishing network connections

Answer: A

Q620. What advantage does a hybrid kernel offer over a pure microkernel?

- A. It keeps performance-critical services in kernel space while maintaining some modularity
- B. It completely eliminates context switches
- C. It runs all services in user space for maximum reliability
- D. It requires no device drivers

Answer: A

Q621. What is the difference between a process and a program?

- A. A process is an active instance of a program in execution with its own state
- B. A program runs on the CPU while a process is stored on disk
- C. A process is always larger in size than a program
- D. There is no difference between a process and a program

Answer: A

Q622. What is the purpose of the process address space?

- A. It defines the memory layout including text, data, heap, and stack sections
- B. It specifies the physical location of the process on the hard disk
- C. It determines the network address used by the process
- D. It sets the display resolution for the process window

Answer: A

Q623. How does message passing IPC differ from shared memory IPC in terms of synchronization?

- A. Message passing provides implicit synchronization while shared memory requires explicit synchronization
- B. Shared memory is always faster for large data transfers
- C. Message passing cannot be used between processes on the same machine
- D. Shared memory does not require any OS support

Answer: A

Q624. What is a message queue as an IPC mechanism?

- A. A linked list of messages stored in the kernel that processes can read and write
- B. A hardware buffer on the network card
- C. A queue of CPU instructions waiting to execute
- D. A list of print jobs waiting for the printer

Answer: A

Q625. What is the difference between direct and indirect inter-process communication?

- A. Direct IPC names the recipient process explicitly while indirect uses shared mailboxes
- B. Direct IPC is faster because it skips the kernel
- C. Indirect IPC can only be used in distributed systems
- D. Direct IPC requires shared memory while indirect does not

Answer: A

Q626. What information is saved during a context switch?

- A. The CPU registers, program counter, and process state of the outgoing process
- B. Only the process name and user ID
- C. The contents of the hard disk
- D. The screen display buffer

Answer: A

Q627. What is the role of the job queue in process scheduling?

- A. It contains all processes in the system regardless of their state
- B. It only holds processes currently executing on the CPU
- C. It stores terminated processes waiting for cleanup
- D. It manages disk I/O requests

Answer: A

Q628. How does the Unix kill() system call work?

- A. It sends a specified signal to a process identified by its PID
- B. It immediately destroys the process without any cleanup
- C. It only works on child processes of the calling process
- D. It deletes the executable file of the target process

Answer: A

Q629. What is the difference between I/O-bound and CPU-bound processes?

- A. I/O-bound processes spend more time on I/O operations while CPU-bound spend more time computing
- B. I/O-bound processes are faster than CPU-bound processes
- C. CPU-bound processes use more memory than I/O-bound processes
- D. There is no meaningful difference between them

Answer: A

Q630. What is swapping in the context of process management?

- A. Moving a process between main memory and a backing store to free memory
- B. Exchanging data between two running processes
- C. Switching the CPU between user mode and kernel mode
- D. Replacing one operating system with another

Answer: A

Q631. Why does Shortest Job First (SJF) scheduling produce the minimum average waiting time?

- A. Shorter jobs finish quickly, reducing the waiting time for all subsequent jobs
- B. It uses the fastest CPU clock speed for short jobs
- C. It skips I/O operations to save time
- D. It runs all processes simultaneously

Answer: A

Q632. What is the problem of starvation in priority scheduling?

- A. Low-priority processes may never execute because higher-priority processes keep arriving
- B. All processes starve for memory simultaneously
- C. The CPU runs out of power during execution
- D. Processes consume too much disk space

Answer: A

Q633. How does the aging technique prevent starvation?

- A. It gradually increases the priority of waiting processes over time
- B. It reduces the priority of all running processes
- C. It terminates processes that have waited too long
- D. It allocates more memory to starving processes

Answer: A

Q634. What is the effect of time quantum size on Round Robin performance?

- A. Too large behaves like FCFS; too small causes excessive context switching overhead
- B. Larger quanta always improve performance
- C. Smaller quanta always improve response time without any downside
- D. Time quantum size has no effect on system performance

Answer: A

Q635. In multilevel queue scheduling, how are processes assigned to queues?

- A. Based on process characteristics such as type, priority, or memory requirements
- B. Randomly distributed among available queues
- C. All placed in a single queue sorted by arrival time
- D. Assigned based on the day of the week they were created

Answer: A

Q636. What advantage does multilevel feedback queue scheduling have over multilevel queue?

- A. Processes can move between queues based on their behavior allowing dynamic adaptation
- B. It uses fewer CPU resources for scheduling decisions
- C. It eliminates the need for context switching
- D. It guarantees that all processes finish in the same amount of time

Answer: A

Q637. What is the difference between preemptive SJF and non-preemptive SJF?

- A. Preemptive SJF interrupts a running process if a shorter one arrives; non-preemptive does not
- B. Non-preemptive SJF is always faster than preemptive SJF
- C. Preemptive SJF does not require knowing burst times
- D. Non-preemptive SJF causes more context switches

Answer: A

Q638. What is the role of the medium-term scheduler in CPU scheduling?

- A. It swaps processes in and out of memory to control the degree of multiprogramming
- B. It selects which process runs on the CPU next
- C. It decides which jobs enter the system from the job pool
- D. It manages disk I/O scheduling

Answer: A

Q639. Why is predicting the next CPU burst length important for SJF scheduling?

- A. SJF requires knowing burst lengths to select the shortest job but future bursts are unknown
- B. Burst length determines how much memory a process needs
- C. Longer bursts indicate higher priority processes
- D. Burst prediction is only needed for real-time systems

Answer: A

Q640. What is the difference between throughput and turnaround time as scheduling metrics?

- A. Throughput measures processes completed per time unit while turnaround time measures total time for one process
- B. Throughput is always larger than turnaround time
- C. They measure the same thing using different units
- D. Turnaround time applies only to I/O-bound processes

Answer: A

Q641. What challenge does Amdahl's Law highlight for multicore systems?

- A. Adding more memory always improves performance linearly
- B. The serial portion of a program limits the maximum speedup achievable through parallelism
- C. All programs can be fully parallelized
- D. Multicore systems are slower than single-core systems

Answer: B

Q642. What is the key disadvantage of the one-to-one threading model?

- A. It does not support kernel-level threading
- B. Creating a user thread requires creating a corresponding kernel thread, adding overhead
- C. It cannot run on multiprocessor systems
- D. It merges all threads into a single execution context

Answer: B

Q643. What is the purpose of the pthread_join() function?

- A. To create a new thread
- B. To wait for a specific thread to terminate before continuing
- C. To send a signal to another thread
- D. To detach a thread from the process

Answer: B

Q644. What happens when a user-level thread makes a blocking system call in a many-to-one model?

- A. Only that thread blocks while others continue
- B. The entire process blocks because the single kernel thread is blocked
- C. The thread is automatically converted to a kernel thread
- D. The system call is ignored

Answer: B

Q645. What is the difference between task parallelism and data parallelism?

- A. Task parallelism distributes data across cores; data parallelism distributes tasks
- B. Data parallelism distributes subsets of data across cores performing the same operation; task parallelism distributes different tasks across cores
- C. They are different names for the same concept
- D. Data parallelism only works with databases

Answer: B

Q646. What is the purpose of pthread_cancel() in the Pthreads API?

- A. To create a cancelable thread
- B. To request termination of a target thread
- C. To pause a thread temporarily
- D. To increase a thread's priority

Answer: B

Q647. Why is asynchronous thread cancellation considered more dangerous than deferred cancellation?

- A. It uses more memory
- B. The target thread is terminated immediately, potentially leaving shared data in an inconsistent state
- C. It requires kernel mode exclusively
- D. It can only cancel one thread at a time

Answer: B

Q648. What advantage does a thread pool provide over creating threads on demand?

- A. Thread pools use more memory
- B. Pre-created threads reduce the overhead of thread creation and limit the total number of threads
- C. Thread pools eliminate the need for synchronization
- D. Thread pools allow only one thread at a time

Answer: B

Q649. In the context of threading, what is the fork-join model?

- A. A model where threads can only fork new processes
- B. A model where a parent thread creates child threads that execute tasks and then join back with the parent
- C. A model specific to database transactions
- D. A model that prevents any thread from terminating

Answer: B

Q650. What is the relationship between threads and signals in UNIX?

- A. Signals can only be sent to processes, never to individual threads
- B. A signal may be delivered to the specific thread to which it applies, to every thread, or to a designated thread
- C. Threads automatically ignore all signals
- D. Each thread has its own independent signal handler table

Answer: B

Q651. How does the test-and-set instruction help solve the critical section problem?

- A. It tests whether the CPU is available
- B. It atomically reads a variable and sets it to true, enabling mutual exclusion without interruption
- C. It sets all variables in memory to zero
- D. It tests the speed of the system bus

Answer: B

Q652. In the bounded buffer problem using semaphores, what are the three semaphores typically used?

- A. read, write, execute
- B. mutex (for mutual exclusion), full (counts filled slots), empty (counts empty slots)
- C. start, stop, pause
- D. lock, unlock, wait

Answer: B

Q653. What is the key challenge in the readers-writers problem?

- A. Ensuring that files can only be read once
- B. Allowing multiple readers simultaneous access while ensuring writers get exclusive access
- C. Preventing any process from reading data
- D. Synchronizing disk rotational speed

Answer: B

Q654. How does priority inversion occur in synchronization?

- A. A high-priority process runs before a low-priority process
- B. A high-priority process is blocked waiting for a lock held by a low-priority process, which is itself preempted by medium-priority processes
- C. All processes run with the same priority
- D. The scheduler inverts the order of the ready queue

Answer: B

Q655. What is the purpose of a condition variable in a monitor?

- A. To store the return value of a function
- B. To allow a thread to wait until a particular condition is met, and to be notified when it changes
- C. To count the number of active monitors
- D. To monitor CPU temperature conditions

Answer: B

Q656. What is the difference between wait() and signal() on a condition variable versus on a semaphore?

- A. They are identical in behavior
- B. Condition variable signal() has no effect if no thread is waiting, while semaphore signal() always increments the value
- C. Semaphore wait() never blocks
- D. Condition variables can only be used in Java

Answer: B

Q657. What problem can occur when semaphores are used incorrectly, such as swapping the order of wait() operations?

- A. Faster execution
- B. Deadlock, where two or more processes are each waiting on semaphores held by the other
- C. Automatic error correction
- D. Improved throughput

Answer: B

Q658. What is a reentrant lock and how does it differ from a standard mutex?

- A. A reentrant lock can only be acquired once
- B. A reentrant lock allows the same thread to acquire it multiple times without deadlocking
- C. A reentrant lock is slower than a standard mutex in all cases
- D. A reentrant lock provides no mutual exclusion

Answer: B

Q659. Why are spinlocks more efficient than blocking locks on multiprocessor systems for short critical sections?

- A. Spinlocks use less memory
- B. Spinning avoids the overhead of context switching, which would take longer than the short wait time
- C. Spinlocks automatically parallelize the critical section
- D. Blocking locks are not supported on multiprocessors

Answer: B

Q660. How does imposing a total ordering on resource types prevent circular wait?

- A. It requires processes to release all resources before requesting new ones
- B. Processes must request resources in increasing order of their assigned numbers, making circular dependencies impossible
- C. It limits the number of resources available
- D. It assigns equal priority to all resources

Answer: B

Q661. In the Banker's Algorithm, what does the 'Available' vector represent?

- A. The total number of resources in the system
- B. The number of instances of each resource type currently available for allocation
- C. The maximum resources any process can request
- D. The resources that have been permanently allocated

Answer: B

Q662. What factors should be considered when selecting a victim process for deadlock recovery?

- A. Only the process name
- B. Priority, computation time completed, resources held, remaining time, and number of rollbacks
- C. Only the process ID number
- D. The color of the user's terminal

Answer: B

Q663. What is the difference between deadlock prevention and deadlock avoidance?

- A. They are identical approaches
- B. Prevention ensures one of the four necessary conditions can never hold; avoidance dynamically checks each request to stay in a safe state
- C. Avoidance eliminates resources; prevention adds resources
- D. Prevention is used only in distributed systems

Answer: B

Q664. What is the safe sequence in the context of the Banker's Algorithm?

- A. A random ordering of all processes
- B. An ordering of processes such that each process can obtain its maximum required resources from currently available resources plus resources held by all preceding processes
- C. The order in which processes were created
- D. A sequence where no process ever requests resources

Answer: B

Q665. How can the hold-and-wait condition be prevented by requiring processes to request all resources at once?

- A. Processes request resources one at a time
- B. A process must request all needed resources before execution begins; it cannot hold some and wait for others
- C. Processes share all resources simultaneously
- D. The OS holds all resources and never allocates them

Answer: B

Q666. What is the difference between process termination and resource preemption as deadlock recovery methods?

- A. They are the same approach
- B. Process termination kills processes to break deadlock; resource preemption forcibly takes resources from processes and reassigns them
- C. Resource preemption always terminates all processes
- D. Process termination only pauses processes temporarily

Answer: B

Q667. Why is the Banker's Algorithm considered impractical for many real systems?

- A. It is too simple to implement
- B. It requires processes to declare their maximum resource needs in advance, which is often unknown
- C. It only works with one resource type
- D. It causes deadlock instead of preventing it

Answer: B

Q668. In deadlock detection with multiple resource instances, what algorithm is used?

- A. Dijkstra's shortest path algorithm
- B. A variant of the Banker's Algorithm that checks if the current state can complete all processes
- C. Binary search on resource allocation
- D. Depth-first search on a binary tree

Answer: B

Q669. What is starvation in the context of deadlock recovery, and how can it be prevented?

- A. Starvation is when the CPU runs out of power
- B. The same process is repeatedly selected as the victim for rollback; prevented by including rollback count in victim selection criteria
- C. Starvation means processes do not have enough memory
- D. It is a type of deadlock that cannot be prevented

Answer: B

Q670. How does a Translation Lookaside Buffer (TLB) improve memory access performance?

- A. By increasing the physical memory size
- B. By caching recently used page table entries to avoid the overhead of accessing the page table in memory for every address translation
- C. By compressing data in memory
- D. By reducing the CPU clock speed

Answer: B

Q671. What is the effective memory access time if the TLB hit ratio is 90%, TLB access time is 10ns, and memory access time is 100ns?

- A. 100ns
- B. 120ns
- C. 110ns
- D. 200ns

Answer: B

Q672. What is the advantage of a multi-level page table over a single-level page table?

- A. It is faster to look up
- B. It reduces the memory required for the page table by only allocating entries for used portions of the address space
- C. It eliminates the need for a TLB
- D. It increases the page size

Answer: B

Q673. How does the optimal (OPT) page replacement algorithm work?

- A. It replaces the most recently used page
- B. It replaces the page that will not be used for the longest period of time in the future
- C. It replaces pages in FIFO order
- D. It replaces the page with the lowest page number

Answer: B

Q674. Why does the FIFO page replacement algorithm suffer from Belady's anomaly?

- A. Because it uses too much memory
- B. Increasing the number of frames can lead to more page faults because FIFO does not consider page usage frequency
- C. Because it always replaces the most recently used page
- D. Because it requires future knowledge of page references

Answer: B

Q675. What is the working set of a process?

- A. The total number of pages in the process
- B. The set of pages that the process has referenced within a recent time window
- C. The set of all physical frames in the system
- D. The set of processes currently in the ready queue

Answer: B

Q676. How does segmentation differ from paging in memory management?

- A. Segmentation uses fixed-size blocks; paging uses variable-size blocks
- B. Segmentation divides memory into variable-size logical units (code, data, stack); paging divides into fixed-size physical blocks
- C. Paging is visible to the programmer; segmentation is not
- D. They are identical approaches with different names

Answer: B

Q677. What is the purpose of the dirty (modified) bit in a page table entry?

- A. To indicate the page contains errors
- B. To indicate the page has been modified since it was loaded, so it must be written back to disk before replacement
- C. To mark the page as deleted
- D. To count how many times the page was accessed

Answer: B

Q678. What causes thrashing and what is its primary symptom?

- A. Too many CPU cores causing overheating
- B. A process spends more time paging (swapping pages in and out) than executing, causing severe performance degradation
- C. The hard disk runs out of space
- D. Too many users are logged in simultaneously

Answer: B

Q679. What is the inverted page table and how does it save memory?

- A. A page table stored upside down in memory
- B. It has one entry per physical frame instead of one per virtual page, reducing table size for large address spaces
- C. It inverts the logical-to-physical mapping to physical-to-logical
- D. It is a page table that grows dynamically

Answer: B

Q680. How does linked allocation manage disk blocks for a file?

- A. All blocks are stored contiguously on disk
- B. Each block contains a pointer to the next block, forming a linked list of disk blocks
- C. A central index block stores pointers to all data blocks
- D. Blocks are stored in a hash table

Answer: B

Q681. What is the advantage of indexed allocation over linked allocation?

- A. It uses less disk space per file
- B. It supports direct (random) access by storing all block pointers in a single index block
- C. It never wastes any disk space
- D. It is simpler to implement

Answer: B

Q682. What is journaling in file systems and what problem does it solve?

- A. It keeps a human-readable log of user activity
- B. It records file system changes in a log before committing them, enabling recovery from crashes without full consistency checks
- C. It creates daily backups of all files
- D. It journals CPU usage for each file

Answer: B

Q683. How does the bitmap (bit vector) method manage free disk space?

- A. It stores free blocks in a linked list
- B. Each bit represents a disk block — 1 for free, 0 for allocated (or vice versa), allowing quick identification of free blocks
- C. It counts the total number of free blocks without tracking which ones
- D. It maps file names to free blocks

Answer: B

Q684. What is the difference between a hard link and a symbolic link in terms of inode behavior?

- A. Both create a new inode for the linked file
- B. A hard link creates another directory entry pointing to the same inode; a symbolic link creates a new inode containing the path to the target
- C. A symbolic link shares the same inode as the original file
- D. Hard links work across different file systems

Answer: B

Q685. What does the Virtual File System (VFS) layer provide?

- A. A virtual disk for storage
- B. A common interface that allows the OS to support multiple file system types transparently
- C. Virtual memory for file operations
- D. Encryption for all file systems

Answer: B

Q686. What is an Access Control List (ACL) in the context of file systems?

- A. A list of all files in a directory
- B. A per-file list specifying which users or groups have which permissions on that file
- C. A list of CPU access times for files
- D. A list of all mounted file systems

Answer: B

Q687. What is the purpose of the superblock in a Unix file system?

- A. To store user passwords
- B. To store metadata about the file system itself, including its size, block size, free blocks count, and inode information
- C. To store the largest file in the system
- D. To manage network connections

Answer: B

Q688. What is contiguous allocation's main advantage and disadvantage?

- A. Advantage: no fragmentation; Disadvantage: slow access
- B. Advantage: fast sequential and direct access; Disadvantage: external fragmentation and difficulty growing files
- C. Advantage: no wasted space; Disadvantage: complex implementation
- D. Advantage: supports journaling; Disadvantage: requires large inodes

Answer: B

Q689. What is the difference between ordered mode and writeback mode journaling?

- A. They are identical journaling modes
- B. Ordered mode writes data to disk before journaling metadata; writeback mode journals metadata without guaranteeing data order, risking stale data after crash
- C. Writeback mode is always slower than ordered mode
- D. Ordered mode only journals file names

Answer: B

Q690. How does the SCAN (elevator) algorithm work for disk scheduling?

- A. It randomly selects the next request to service
- B. The disk arm moves in one direction servicing requests until it reaches the end, then reverses direction and services requests going back
- C. It services only requests on the outermost tracks
- D. It services requests in alphabetical order of file names

Answer: B

Q691. What advantage does C-SCAN have over SCAN?

- A. C-SCAN is simpler to implement
- B. C-SCAN provides more uniform wait times by only servicing requests in one direction and returning to the beginning without servicing
- C. C-SCAN uses less disk space
- D. C-SCAN eliminates all seek time

Answer: B

Q692. How does RAID 5 distribute parity information?

- A. Parity is stored on a dedicated disk
- B. Parity blocks are distributed across all disks in the array, avoiding the bottleneck of a single parity disk
- C. Parity is not used in RAID 5
- D. Each disk stores its own parity only

Answer: B

Q693. What is wear leveling in SSDs?

- A. Evenly distributing weight on the drive housing
- B. Distributing write operations evenly across all flash memory cells to prevent some cells from wearing out faster than others
- C. Reducing the speed of write operations to extend drive life
- D. Leveling the surface of the flash memory chips

Answer: B

Q694. What is the LOOK variant of disk scheduling?

- A. The disk arm looks ahead one request at a time
- B. The arm moves in one direction only as far as the last pending request, then reverses, instead of going to the disk's end
- C. It looks at all requests simultaneously before choosing
- D. It skips every other request

Answer: B

Q695. What is the boot block on a disk?

- A. A block that prevents booting
- B. The disk block(s) containing the bootstrap loader program that initializes the system and loads the operating system
- C. A block reserved for backup data
- D. The largest block on the disk

Answer: B

Q696. What is disk defragmentation and when is it beneficial?

- A. It encrypts disk data for security
- B. It reorganizes scattered file fragments into contiguous blocks, beneficial for HDDs to reduce seek times but unnecessary for SSDs
- C. It deletes duplicate files from the disk
- D. It increases the disk's storage capacity

Answer: B

Q697. What is the purpose of a hot spare disk in a RAID configuration?

- A. A disk that stores temporary data
- B. An idle disk that automatically replaces a failed disk in the array, reducing the window of vulnerability
- C. A disk that runs at a higher temperature
- D. A disk used only for read operations

Answer: B

Q698. What is memory-mapped I/O and how does it differ from port-mapped I/O?

- A. Memory-mapped I/O is slower than port-mapped I/O in all cases
- B. Memory-mapped I/O maps device registers into the regular memory address space; port-mapped I/O uses separate I/O address space with special instructions
- C. Port-mapped I/O uses regular memory addresses
- D. Memory-mapped I/O requires dedicated I/O instructions

Answer: B

Q699. What is double buffering and why is it used?

- A. Using two CPUs for I/O operations
- B. Using two buffers so the device can fill one while the process reads from the other, overlapping I/O and computation
- C. Buffering data twice for error checking
- D. Using buffers of double the normal size

Answer: B

Q700. How does an interrupt vector table work?

- A. It stores all device drivers in a single table
- B. It is an array of pointers to interrupt handler routines, indexed by interrupt number, allowing quick dispatch to the correct handler
- C. It converts analog signals to digital interrupts
- D. It counts the number of interrupts per second

Answer: B

Q701. What is the role of the DMA controller during a data transfer?

- A. It replaces the CPU entirely during the transfer
- B. It manages the transfer of data between I/O devices and memory, handling addressing, counting, and bus arbitration while the CPU runs other tasks
- C. It only monitors the transfer without participating
- D. It encrypts data during transfer

Answer: B

Q702. How does asynchronous I/O differ from non-blocking I/O?

- A. They are identical concepts
- B. Non-blocking I/O returns immediately with available data; asynchronous I/O initiates the operation and notifies the process when it completes, fully decoupling the I/O from the process
- C. Asynchronous I/O blocks the process until completion
- D. Non-blocking I/O uses more memory than asynchronous I/O

Answer: B

Q703. What is the device status table maintained by the OS?

- A. A table listing all devices that need repair
- B. A table that tracks the type, address, status, and pending I/O requests for each device in the system
- C. A table of device prices for inventory management
- D. A table mapping device names to driver versions

Answer: B

Q704. What is device-independent I/O software responsible for?

- A. Manufacturing hardware devices
- B. Providing uniform interfaces, buffering, error reporting, and device allocation that work regardless of the specific device
- C. Installing device drivers automatically
- D. Converting device signals to audio

Answer: B

Q705. What is an interrupt controller and what does it do?

- A. A software program that generates interrupts
- B. Hardware that manages multiple interrupt lines, prioritizes interrupts, and routes them to the CPU
- C. A controller that disables all interrupts permanently
- D. A device that converts interrupts to I/O requests

Answer: B

Q706. What is the access matrix model and how is it used for protection?

- A. A matrix that stores encryption keys
- B. A conceptual model with rows for domains (subjects) and columns for objects, where each cell contains the set of operations the domain can perform on the object
- C. A physical matrix of hardware protection circuits
- D. A matrix that maps memory addresses

Answer: B

Q707. How does Role-Based Access Control (RBAC) simplify security management?

- A. By eliminating all access control
- B. By assigning permissions to roles instead of individual users, then assigning users to roles, reducing the complexity of managing per-user permissions
- C. By giving all users the same role
- D. By using biometrics exclusively

Answer: B

Q708. What is a buffer overflow attack and why is it dangerous?

- A. Filling a printer buffer with too much data
- B. Writing data beyond a buffer's allocated memory to overwrite adjacent memory, potentially hijacking program control flow to execute malicious code
- C. Downloading too many files at once
- D. Overflowing the CPU's instruction cache

Answer: B

Q709. What is sandboxing in operating system security?

- A. Running programs on beach-themed desktops
- B. Running a program in an isolated environment with restricted access to system resources, limiting the damage if the program is compromised
- C. Encrypting program data with sand-based algorithms
- D. A technique to speed up program execution

Answer: B

Q710. What is privilege escalation?

- A. A user requesting a password reset
- B. An attacker exploiting a vulnerability to gain higher access privileges than initially granted, such as obtaining root access from a regular user account
- C. The process of granting administrator access to new employees
- D. Upgrading the operating system to a newer version

Answer: B

Q711. How does the capability list approach implement the access matrix?

- A. It stores one list per object with authorized subjects
- B. Each process holds a list of capabilities (tokens) specifying objects it can access and the permitted operations
- C. It stores a single global list of all permissions
- D. It uses a hash table indexed by user names

Answer: B

Q712. What is a rootkit and why is it particularly dangerous?

- A. A toolkit for system administrators to manage root access
- B. Malware that hides its presence by modifying the OS or firmware, making it extremely difficult to detect while maintaining persistent unauthorized access
- C. A set of root vegetables used as a screen saver
- D. A backup tool for the root partition

Answer: B

Q713. What are protection rings and how do they enforce security?

- A. Physical rings placed around the CPU for protection
- B. Hierarchical privilege levels (Ring 0 being most privileged) enforced by hardware, where inner rings have more privileges and outer rings are restricted
- C. Software-only protection mechanisms in the firewall
- D. Network security zones for different departments

Answer: B

Q714. What is the purpose of a turnstile synchronization mechanism used in some kernels?

- A. To limit the number of threads in the system
- B. A queue structure where blocked threads waiting for a lock are placed, allowing priority inheritance and ordered wakeup
- C. To rotate CPU access among threads
- D. To count the number of active synchronization primitives

Answer: B

Q715. What is the C-LOOK disk scheduling algorithm?

- A. A variant of C-SCAN that reverses at the end of the disk
- B. A variant of C-SCAN where the arm only goes as far as the last pending request in each direction before wrapping around, avoiding unnecessary travel to disk edges
- C. A scheduling algorithm that looks at every other request
- D. An algorithm identical to FCFS but circular

Answer: B

Q716. What is the difference between low-level formatting and high-level formatting of a disk?

- A. They are the same process with different names
- B. Low-level formatting creates the physical sector structure on the disk surface; high-level formatting creates the file system data structures within a partition
- C. High-level formatting erases the physical platters
- D. Low-level formatting only applies to SSDs

Answer: B

Q717. What is the purpose of the I/O request queue in the OS?

- A. To store the output of completed I/O operations
- B. To hold pending I/O requests for a device, allowing the scheduler to reorder them for efficiency
- C. To count the total number of I/O operations performed
- D. To buffer all keyboard inputs

Answer: B

Q718. How does caching differ from buffering in the context of I/O?

- A. They are identical concepts with different names
- B. A buffer holds the only copy of data during transfer; a cache holds a faster duplicate copy of data that also exists elsewhere
- C. Caching is always slower than buffering
- D. Buffering uses CPU cache while caching uses RAM

Answer: B

Q719. What is the purpose of a digital certificate in security?

- A. To compress encrypted data
- B. To bind a public key to an entity's identity, verified by a trusted Certificate Authority, enabling secure key exchange
- C. To digitally sign physical documents
- D. To store all of a user's passwords

Answer: B

Q720. What is a man-in-the-middle (MITM) attack?

- A. A physical attack on a server's hardware
- B. An attacker secretly intercepts and potentially alters communication between two parties who believe they are communicating directly with each other
- C. A type of virus that hides in the middle of a file
- D. An attack that only targets the middle layer of the network stack

Answer: B

Operating Systems — Hard

360 questions

Q721. In the context of OS design, what is the key advantage of a microkernel architecture over a monolithic kernel?

- A. Better performance from fewer context switches in kernel
- B. Direct hardware access available for all system services
- C. Higher reliability since most services run in user space
- D. Simpler implementation compared to other kernel designs

Answer: C

Q722. What is the purpose of a hypervisor in virtualization?

- A. To serve as a real-time antivirus scanning tool
- B. To create and manage virtual machines on hardware
- C. To increase the total RAM capacity of a system
- D. To compile user programs at a much faster speed

Answer: B

Q723. What distinguishes a Type 1 hypervisor from a Type 2 hypervisor?

- A. Type 1 runs on top of host OS, Type 2 on hardware
- B. There is no real or meaningful difference between them
- C. Type 1 is inherently slower in performance than Type 2
- D. Type 1 runs directly on hardware, Type 2 on host OS

Answer: D

Q724. In a hard real-time system, what happens when a deadline is missed?

- A. The task priority gets increased
- B. The system gradually slows down
- C. The task is automatically retried
- D. The result is a total system failure

Answer: D

Q725. What is the main purpose of the interrupt vector table?

- A. To hold addresses of interrupt service routines
- B. To schedule and dispatch CPU tasks to run
- C. To manage the overall file system structure
- D. To permanently store user programs on disk

Answer: A

Q726. What is the concept of 'dual-mode operation' in operating systems and why is it essential?

- A. It separates user and kernel modes for protection
- B. It doubles the overall CPU processing speed rate
- C. It enables running two operating systems at once
- D. It allows two users to work at the same time

Answer: A

Q727. In clustered systems, what is the difference between asymmetric clustering and symmetric clustering?

- A. Asymmetric has standby node; symmetric all nodes share
- B. Asymmetric uses different OS, symmetric uses same OS
- C. There is no meaningful difference between these approaches
- D. Asymmetric clustering is inherently faster than symmetric

Answer: A

Q728. What mechanism does the OS use to prevent a user program from running indefinitely and never returning control?

- A. Timer interrupt
- B. Watchdog timer
- C. Busy waiting
- D. Device polling

Answer: A

Q729. What is the purpose of caching in the memory hierarchy of an operating system?

- A. To increase the size of main memory
- B. To store frequent data in faster memory
- C. To permanently store data on disk
- D. To compress data for efficient storage

Answer: B

Q730. What problem does cache coherency address in multiprocessor systems?

- A. Preventing the processor cache from overflowing
- B. Reducing the overall processor cache miss rate
- C. Ensuring all processor caches are the same size
- D. Ensuring shared data updates reflect in all caches

Answer: D

Q731. In the Mach microkernel, what are the primary abstractions provided?

- A. Files, directories, and named pipes
- B. Processes, semaphores, and mutexes
- C. Blocks, inodes, and superblocks
- D. Tasks, threads, ports, and messages

Answer: D

Q732. What is an exokernel?

- A. A kernel that runs outside the main CPU
- B. A minimal kernel giving apps direct hardware control
- C. A kernel designed only for external USB devices
- D. A kernel that runs entirely in user space only

Answer: B

Q733. What is the purpose of a policy/mechanism separation in OS design?

- A. To separate what is done from how it is done
- B. To divide the OS into two equal-sized halves
- C. To separate user programs from system programs
- D. To directly improve hardware processing speed

Answer: A

Q734. What is the main concept behind a unikernel?

- A. A kernel that runs on one processor core only
- B. A kernel that supports only one single user
- C. A kernel composed of one single code module
- D. A single-image OS compiled with the application

Answer: D

Q735. In OS debugging, what is a core dump?

- A. A memory snapshot of a process at crash time
- B. A network diagnostic and monitoring utility
- C. A listing of all currently running processes
- D. A disk formatting and partitioning operation

Answer: A

Q736. What is the strace command used for in Linux?

- A. To trace network packets on the wire
- B. To trace system calls made by a process
- C. To trace file system changes on the disk
- D. To trace memory allocations in the heap

Answer: B

Q737. What is the advantage of the modular kernel approach over a pure layered approach?

- A. It uses significantly less memory than layered design
- B. It is much simpler to implement and maintain
- C. Modules call any other module directly with flexibility
- D. It does not require loading any device drivers at all

Answer: C

Q738. In system generation (SYSGEN), the OS is configured for a specific hardware platform. Which of the following is NOT typically determined during SYSGEN?

- A. CPU type and count
- B. User access passwords
- C. Connected peripherals
- D. Available memory size

Answer: B

Q739. What is the role of the Hardware Abstraction Layer (HAL) in Windows?

- A. To manage and track user login sessions
- B. To provide a graphical user interface
- C. To hide hardware differences for portability
- D. To compile and link device driver programs

Answer: C

Q740. What distinguishes a library OS from a traditional OS design?

- A. Apps link directly with OS library code
- B. It manages library catalog structures
- C. It is a collection of shared libraries
- D. It only manages shared library files

Answer: A

Q741. In Chrome's multi-process architecture, what types of processes does it use?

- A. Only renderer display processes
- B. Background and foreground processes
- C. Only a single monolithic process
- D. Browser, renderer, and plug-in processes

Answer: D

Q742. What is the difference between direct and indirect communication in message passing?

- A. Direct names sender/receiver; indirect uses mailboxes
- B. Direct is synchronous only; indirect is asynchronous
- C. Direct uses mailboxes; indirect names each process
- D. There is no difference between the two

Answer: A

Q743. What is the rendezvous mechanism in synchronous message passing?

- A. A method to forcefully terminate processes
- B. Both sender and receiver block until ready
- C. A strategy for memory allocation
- D. A technique for CPU load balancing

Answer: B

Q744. In UNIX, how does the init process (PID 1) handle orphan processes?

- A. It adopts them and collects their exit status
- B. It moves them to a special waiting queue
- C. It terminates orphan processes immediately
- D. It ignores them and frees their resources

Answer: A

Q745. What is a socket in the context of IPC?

- A. A type of distributed file system
- B. A physical hardware port connector
- C. An endpoint defined by IP and port
- D. A special-purpose CPU register set

Answer: C

Q746. What is Remote Procedure Call (RPC) in the context of distributed process communication?

- A. A file transfer protocol for distributed systems
- B. Calling a remote procedure as if it were local
- C. A method to restart processes on remote machines
- D. A type of CPU scheduling algorithm for clusters

Answer: B

Q747. What is the purpose of the medium-term scheduler?

- A. To select the next process for CPU execution
- B. To create new processes from program images
- C. To manage file systems on storage devices
- D. To swap processes in and out of main memory

Answer: D

Q748. In Android, what is a Binder and what role does it play?

- A. A kernel memory allocation system
- B. An IPC framework using shared memory
- C. A CPU scheduler for Android apps
- D. A file management utility tool

Answer: B

Q749. What is the copy-on-write (COW) optimization used with fork()?

- A. Creating a complete full copy of the parent immediately
- B. Parent and child share pages; copy only on modification
- C. Writing all process data to disk before forking begins
- D. Preventing the child process from writing to any memory

Answer: B

Q750. What is the purpose of the vfork() system call?

- A. Child shares parent space; parent suspends until exec
- B. It creates a virtual process with no memory mapping
- C. It is identical to the standard fork() call
- D. It creates a process with higher priority than parent

Answer: A

Q751. In the Completely Fair Scheduler (CFS) used by Linux, how is CPU time allocated?

- A. Based only on arrival time of each process
- B. Based on vruntime, picking the smallest vruntime
- C. Using fixed time quanta for all processes equally
- D. Using strict static priority scheduling for all

Answer: B

Q752. What is the problem with using SJF in practice?

- A. Next burst length is unknown and must be estimated
- B. It always causes deadlock among the running tasks
- C. It is far too slow for practical use
- D. It does not support multitasking at all in practice

Answer: A

Q753. In exponential averaging for SJF burst prediction, the formula is $B_{t+1} = \alpha t(n) + (1-\alpha) \hat{A}(n)$. What happens when $\alpha = 0$?

- A. Only the most recent actual burst is considered here
- B. All previous burst lengths are weighted equally here
- C. Recent history has no effect; uses initial estimate only
- D. The prediction always becomes zero in all situations

Answer: C

Q754. In multiprocessor scheduling, what is processor affinity?

- A. Keeping a process on the same processor for cache
- B. The clock speed of a specific processor core
- C. The power consumption of a specific processor
- D. The total number of processors in the system

Answer: A

Q755. What is load balancing in multiprocessor scheduling?

- A. Dividing a single process into equal-sized parts
- B. Running all processes on a single CPU processor
- C. Reducing the CPU clock speed to save energy
- D. Distributing workload evenly across all processors

Answer: D

Q756. What is the difference between asymmetric and symmetric multiprocessing scheduling?

- A. Asymmetric uses more processors than symmetric
- B. They are identical in behavior and design
- C. Symmetric is always faster than asymmetric mode
- D. Asymmetric has one master; symmetric self-schedules

Answer: D

Q757. In real-time CPU scheduling, what is the difference between rate-monotonic and earliest-deadline-first (EDF) scheduling?

- A. Rate-monotonic is non-preemptive; EDF is preemptive
- B. Rate-monotonic uses fixed priority; EDF uses deadlines
- C. EDF uses fixed priority; rate-monotonic uses dynamic
- D. They are identical algorithms for real-time systems

Answer: B

Q758. What is the maximum CPU utilization bound for rate-monotonic scheduling with n processes?

- A. Always exactly 75% utilization
- B. Always 100% utilization
- C. Always exactly 50% utilization
- D. $n(2^{1/n}-1)$, approaches ~69.3%

Answer: D

Q759. What is NUMA-aware scheduling?

- A. Scheduling that ignores all memory placement
- B. Scheduling based on network usage patterns
- C. A scheduling algorithm for non-uniform CPUs
- D. Placing processes near their memory for speed

Answer: D

Q760. What is proportional share scheduling?

- A. Strict priority-based scheduling approach
- B. Scheduling based on memory usage amounts
- C. CPU time proportional to allocated shares
- D. Equal CPU time for all running processes

Answer: C

Q761. What is the many-to-many threading model and what advantage does it offer?

- A. It creates many processes for running threads
- B. It is the simplest threading model to implement
- C. It multiplexes user threads onto kernel threads
- D. It maps many threads to many individual CPUs

Answer: C

Q762. What is the two-level threading model?

- A. A model with exactly two priority levels only
- B. A model with exactly two scheduling queues
- C. Many-to-many plus binding user thread to kernel
- D. A model that uses exactly two processor CPUs

Answer: C

Q763. What is a Lightweight Process (LWP) in the context of user and kernel thread mapping?

- A. A process that runs and completes very quickly
- B. A type of user-level thread implementation
- C. An intermediate structure between user and kernel
- D. A process with very low memory usage footprint

Answer: C

Q764. How does signal handling work in a multithreaded program?

- A. Signals are ignored in multithreaded programs
- B. Only the main thread ever receives any signal
- C. Signals are always delivered to every thread
- D. Signals go to applicable, all, or designated thread

Answer: D

Q765. What is the scheduler activation mechanism?

- A. A method to start or initialize the operating system
- B. A timer-based scheduling technique for real-time use
- C. A way to activate the CPU scheduling component
- D. Kernel upcalls maintaining correct kernel thread count

Answer: D

Q766. What is Grand Central Dispatch (GCD)?

- A. A database system for concurrency
- B. Apple technology managing dispatch queues
- C. A network routing protocol for devices
- D. A Linux kernel module for scheduling

Answer: B

Q767. What is the fork-join parallelism pattern?

- A. A pattern for optimizing disk I/O performance
- B. A pattern for managing file system operations
- C. Parent forks child threads, then joins on completion
- D. A pattern for managing network communication tasks

Answer: C

Q768. What problem does the clone() system call solve in Linux that fork() doesn't?

- A. It creates child processes much faster than fork
- B. It allows more than 256 concurrent running processes
- C. It prevents zombie processes from accumulating in table
- D. It controls exactly which resources parent and child share

Answer: D

Q769. In the context of Intel Threading Building Blocks (TBB), what is a task?

- A. A fixed-size dedicated thread object
- B. A system call for thread management
- C. A work unit mapped to threads by runtime
- D. A hardware timer for task scheduling

Answer: C

Q770. Why have most modern operating systems moved away from the many-to-many threading model to the one-to-one model?

- A. Many-to-many has known security vulnerabilities
- B. Systems now efficiently support many kernel threads
- C. Many-to-many is more expensive to run overall
- D. One-to-one uses significantly less memory overall

Answer: B

Q771. What is the dining philosophers problem and why is it significant?

- A. A problem about memory allocation fragmentation
- B. A problem about restaurant management logistics
- C. A concurrency problem illustrating deadlock and starvation
- D. A problem about network packet routing efficiency

Answer: C

Q772. How can the dining philosophers problem be solved to prevent deadlock?

- A. Add more chopsticks than the total number of diners
- B. Give each philosopher two dedicated chopsticks to use
- C. Limit sitting philosophers, use asymmetric or monitors
- D. Remove one philosopher from the dining table entirely

Answer: C

Q773. What is a read-write lock and when is it preferred over a mutex?

- A. A lock used only for writing to shared data
- B. A lock only used for file read operations
- C. A lock that prevents all access to a resource
- D. Allows concurrent reads but requires exclusive writes

Answer: D

Q774. What is a memory barrier (memory fence)?

- A. A physical barrier inside memory chip hardware
- B. A limit on the total memory allocation size
- C. An instruction enforcing memory operation ordering
- D. A type of virtual memory address translation

Answer: C

Q775. What are lock-free algorithms and why are they important?

- A. Algorithms that use no synchronization at all
- B. Algorithms that run without any operating system
- C. Algorithms using atomic ops like CAS, no locks
- D. Algorithms that are never blocked by anything

Answer: C

Q776. In the readers-writers problem, what is the difference between the first and second variants?

- A. There is no difference between the two at all
- B. First is for single CPU; second for multiprocessor
- C. First uses semaphores; the second uses only monitors
- D. First favors readers no wait; second favors writers

Answer: D

Q777. What is the difference between Hoare-style and Mesa-style monitors?

- A. Hoare signal transfers control; Mesa is just a hint
- B. They use different programming languages entirely
- C. Mesa-style monitors are older than Hoare-style
- D. Hoare-style monitors are faster than Mesa-style

Answer: A

Q778. What is transactional memory?

- A. Memory transferred between running processes
- B. Atomic memory operations similar to databases
- C. A type of CPU cache for fast computations
- D. Memory used for financial transactions only

Answer: B

Q779. What is the ABA problem in lock-free programming?

- A. A deadlock scenario involving two blocked threads
- B. A naming convention issue in the source code
- C. A problem with kernel memory allocation routines
- D. CAS succeeds incorrectly due to A-to-B-to-A change

Answer: D

Q780. What is a futex (fast userspace mutex) in Linux?

- A. User-space mutex with kernel support on contention
- B. A fast file transfer utility for networks
- C. A type of file lock for exclusive access
- D. A network protocol for data synchronization

Answer: A

Q781. In the Banker's Algorithm safety check, what is the time complexity for n processes and m resource types?

- A. $O(n)$ linear
- B. $O(n!)$ factorial
- C. $O(n * m)$ time
- D. $O(n^2 * m)$ time

Answer: D

Q782. What is the problem with resource preemption as a deadlock recovery method?

- A. It uses too much memory for tracking state
- B. Selecting victim, rollback, preventing starvation
- C. It only works on single-processor systems
- D. It runs too fast and is hard to control

Answer: B

Q783. In a resource allocation graph with multiple instances per resource type, a cycle:

- A. Necessary but not sufficient
- B. Never indicates any deadlock
- C. Always indicates deadlock
- D. Neither necessary nor sufficient

Answer: A

Q784. What is a livelock and how does it differ from deadlock?

- A. There is no difference between livelock and deadlock
- B. Livelock only occurs in single-threaded programs only
- C. In livelock processes change state without any progress
- D. Livelock is simply a faster form of deadlock occurring

Answer: C

Q785. How does the 'no preemption' condition get addressed in deadlock prevention for resources like CPU registers and memory?

- A. All held resources are preempted if request fails
- B. The process is immediately terminated by the OS
- C. These resources can never be preempted at all
- D. The resources are duplicated for both processes

Answer: A

Q786. What is the main limitation of the Banker's Algorithm in practice?

- A. It runs too fast for real world use cases
- B. It only works with exactly two processes
- C. Processes rarely know max needs in advance
- D. It does not detect any resource cycles found

Answer: C

Q787. How frequently should a deadlock detection algorithm be invoked?

- A. Only when the system is shutting down
- B. Depends on frequency and tradeoff of cost
- C. Every single millisecond without exception
- D. Once at initial system startup only ever

Answer: B

Q788. In distributed systems, what additional challenges does deadlock detection face?

- A. Deadlocks cannot occur in distributed systems
- B. Only overall performance is affected here
- C. No additional challenges at all exist
- D. Delays, no global state, phantom deadlocks

Answer: D

Q789. What is a two-phase locking protocol and how does it relate to deadlocks?

- A. A protocol using only two shared resources
- B. A protocol with exactly two locks only
- C. Growing acquires, shrinking releases; can deadlock
- D. A protocol that is completely deadlock-free

Answer: C

Q790. What is the wound-wait scheme for deadlock prevention in database systems?

- A. Older wounds younger; younger waits for older
- B. A recovery technique for crashed transactions
- C. A method to wound and terminate processes
- D. A scheduling algorithm for query processing

Answer: A

Q791. What is the working set model and how does it relate to thrashing?

- A. It is a disk scheduling model for I/O requests
- B. It is a CPU scheduling model for processes
- C. It models CPU utilization across all processors
- D. Defines active pages; fewer frames causes thrashing

Answer: D

Q792. What is a multi-level page table and why is it used?

- A. A page table with multiple entries per page
- B. Hierarchical structure saving memory for sparse spaces
- C. A page table stored entirely on secondary disk
- D. A table shared by multiple running processes

Answer: B

Q793. What is an inverted page table?

- A. A page table stored in reverse order
- B. One entry per physical frame, reducing overhead
- C. A page table stored entirely on the disk drive
- D. A table that maps physical to logical addresses

Answer: B

Q794. What is the second-chance (clock) page replacement algorithm?

- A. Modified FIFO checking reference bit, gives second chance
- B. An algorithm that always replaces the second loaded page
- C. A two-pass algorithm scanning memory blocks twice each
- D. A modified version of the FIFO algorithm only

Answer: A

Q795. What is copy-on-write in the context of virtual memory?

- A. Copying all pages to disk during every write
- B. Writing data twice to provide data redundancy
- C. Sharing pages until a write, then making a copy
- D. Copying all pages before any read operation

Answer: C

Q796. What are huge pages (large pages) and when are they beneficial?

- A. Pages that span across multiple physical disk drives
- B. Pages assigned the highest priority access level
- C. Pages used to store very large files on disk
- D. Larger pages reducing TLB misses for big workloads

Answer: D

Q797. What is memory-mapped I/O?

- A. Mapping device registers into address space
- B. A technique for allocating memory to tasks
- C. An I/O scheduling algorithm for requests
- D. Storing I/O data in regular memory buffers

Answer: A

Q798. What is the page-fault frequency (PFF) approach to preventing thrashing?

- A. Setting a maximum number of page faults allowed
- B. Monitoring fault rate and adjusting frame allocation
- C. Counting total page faults across all processes
- D. Preventing all page faults from occurring ever

Answer: B

Q799. What is the buddy system for memory allocation?

- A. A system with two separate pools for allocation
- B. A system where two processes share memory pool
- C. A system for backing up memory to the disk
- D. Power-of-2 blocks, splitting and coalescing buddies

Answer: D

Q800. What is the slab allocator used in Linux kernel memory management?

- A. A page replacement algorithm for virtual memory
- B. A disk allocation method for file storage
- C. A virtual memory technique for address mapping
- D. Pre-allocating object caches to reduce overhead

Answer: D

Q801. What is a journaling file system and how does it improve reliability?

- A. Logs changes to journal before applying for crash recovery
- B. A read-only file system that never allows modifications
- C. A file system that keeps a diary of user activities
- D. A file system that journals all network activity on disk

Answer: A

Q802. What is a log-structured file system (LFS)?

- A. Writes all data sequentially to a log for write speed
- B. A file system that only stores system log files
- C. A file system with built-in logging capability
- D. A file system used only for system log storage

Answer: A

Q803. In UNIX, what is the purpose of the superblock?

- A. To store the largest file in the system
- B. To store file system metadata like size and counts
- C. To speed up file access through data caching
- D. To store user passwords and authentication info

Answer: B

Q804. How do UNIX file systems handle files larger than what direct block pointers in an inode can address?

- A. By splitting the file into multiple smaller separate files
- B. Using indirect, double indirect, and triple indirect pointers
- C. They cannot handle files larger than direct pointers allow
- D. By compressing the file data to reduce total block count

Answer: B

Q805. What is a copy-on-write file system like ZFS or Btrfs?

- A. Never overwrites data; writes to new locations always
- B. A file system that prevents all write operations
- C. A file system that duplicates all data for safety
- D. A file system that copies files when writing to disk

Answer: A

Q806. What is the Network File System (NFS)?

- A. A database for networked storage
- B. A firewall for file server protection
- C. A local file system for one machine
- D. A protocol for remote files as local

Answer: D

Q807. What are extents in modern file systems like ext4?

- A. Extra storage space added to disk
- B. External file links to other systems
- C. Contiguous block range by start and length
- D. Extended file attributes for metadata

Answer: C

Q808. What is file system defragmentation and why is it needed?

- A. Checking file system integrity and repairing errors
- B. Breaking files into smaller fragments for storage
- C. Deleting all fragmented files from the disk drive
- D. Reorganizing blocks to be contiguous for performance

Answer: D

Q809. What is the buffer cache in file system implementation?

- A. A buffer used for printer output spooling
- B. Memory caching disk blocks to reduce I/O ops
- C. A buffer used for storing network packet data
- D. A cache for graphical user interface buffers

Answer: B

Q810. What is a RAID system and what is its purpose?

- A. A debugging tool for disk diagnostics
- B. A memory management technique for allocation
- C. Multiple disks combined for performance or reliability
- D. A type of file system for special purposes

Answer: C

Q811. What is wear leveling in SSDs and why is it necessary?

- A. Equalizing read speeds across memory blocks
- B. Distributing physical wear on the drive casing
- C. Distributing write/erase cycles evenly on cells
- D. Leveling the drive in its mounting bracket

Answer: C

Q812. What is the TRIM command in the context of SSDs?

- A. A command to reduce overall file size used
- B. A command to trim long file names shorter
- C. Tells SSD which blocks are free for erasure
- D. A command to resize existing disk partitions

Answer: C

Q813. What is the write amplification problem in SSDs?

- A. Data being written too quickly to the drive
- B. Amplified noise during write operations on the drive
- C. Physical writes exceed logical due to garbage collection
- D. Writing same data multiple times for data redundancy

Answer: C

Q814. What is the difference between RAID 5 and RAID 6?

- A. RAID 6 uses mirroring instead of distributed parity
- B. There is no significant difference between them at all
- C. RAID 6 is faster in overall performance than RAID 5
- D. RAID 6 uses double parity, surviving two disk failures

Answer: D

Q815. What is the I/O scheduler's role in Linux and how do different schedulers (CFQ, deadline, noop) differ?

- A. They all work in an identical way overall
- B. CFQ ensures fairness; deadline prevents starvation; noop for SSDs
- C. They only differ in their names not in actual behavior
- D. They are all deprecated and no longer used in modern Linux

Answer: B

Q816. What is disk striping and how does it improve performance?

- A. Painting visual stripes on the disk surface
- B. Removing unused sectors from the disk drive
- C. Compressing data stored on the disk drive
- D. Distributing data across disks for parallel I/O

Answer: D

Q817. What is the purpose of a disk controller's command queue (NCQ/TCQ)?

- A. To queue commands issued by the user at terminal
- B. To encrypt all disk commands for data security
- C. To reorder pending I/O for reduced head movement
- D. To limit the number of concurrent disk requests

Answer: C

Q818. What is a hot spare disk in RAID configurations?

- A. A disk that runs at higher operating temperature
- B. An unused disk that auto-replaces a failed disk
- C. A backup disk stored off-site for recovery use
- D. A disk storing the most frequently accessed files

Answer: B

Q819. What is the difference between NVMe and SATA interfaces for SSDs?

- A. SATA is faster than NVMe in all standard benchmarks
- B. NVMe uses PCIe for higher bandwidth and lower latency
- C. NVMe only works with traditional hard disk drive units
- D. They are the exact same interface standard overall

Answer: B

Q820. What is the significance of the 4K sector size (Advanced Format) in modern disk drives?

- A. It slows down the disk read and write speed
- B. It reduces overall disk capacity significantly
- C. It is designed only for SSD drives and not HDDs
- D. Larger sectors improve error correction and efficiency

Answer: D

Q821. What is vectored I/O (scatter-gather I/O)?

- A. I/O operations on vector graphics rendering
- B. Broadcasting I/O to multiple devices at once
- C. Single I/O on multiple non-contiguous buffers
- D. Parallel I/O on vector processors for tasks

Answer: C

Q822. What is the purpose of I/O scheduling at the OS level?

- A. To encrypt I/O data for security and privacy
- B. To slow down all I/O operations for safety
- C. To prevent all I/O operations from executing
- D. To order and optimize I/O for throughput and fairness

Answer: D

Q823. What is the kernel I/O subsystem's role in device reservation?

- A. Updating device firmware and driver software
- B. Reserving devices for future hardware purchase
- C. Backing up all device data to secondary storage
- D. Granting exclusive device access to prevent conflicts

Answer: D

Q824. What is the life cycle of an I/O request from a user application to hardware?

- A. Application goes through network stack then to hardware
- B. Application directly accesses hardware with no intermediary
- C. App calls kernel, driver, controller, then hardware device
- D. Application goes through file system then to the hardware

Answer: C

Q825. What is asynchronous I/O and how does it differ from non-blocking I/O?

- A. Non-blocking uses interrupts while async does not
- B. They are the same thing in every single way
- C. Asynchronous I/O is actually a blocking I/O form
- D. Non-blocking returns available; async notifies on completion

Answer: D

Q826. What is IOCTL (I/O Control) used for?

- A. A syscall for device-specific operations beyond read/write
- B. Input and output data formatting utility
- C. I/O error correction and data recovery mechanism
- D. I/O traffic control and bandwidth management tool

Answer: A

Q827. What is zero-copy I/O and why is it important for performance?

- A. I/O that produces no output at all ever
- B. Eliminating unnecessary copies between kernel and user
- C. I/O that uses no disk space for data storage
- D. I/O with zero errors during the transfer process

Answer: B

Q828. What is the role of the select()/poll()/epoll() system calls?

- A. To monitor multiple file descriptors for I/O readiness
- B. To select which CPU scheduling algorithm to apply
- C. To select files to delete from the disk drive
- D. To poll system users for their preference settings

Answer: A

Q829. What is the difference between port-mapped I/O and memory-mapped I/O?

- A. Port-mapped I/O is always faster than memory-mapped I/O
- B. Port-mapped uses separate space; memory-mapped uses memory
- C. Memory-mapped requires special hardware not in modern CPUs
- D. They are identical in every respect and behavior

Answer: B

Q830. What is the transformer model in I/O subsystem layering?

- A. A machine learning model for I/O prediction
- B. A power supply model for device power management
- C. A database model for storing I/O request metadata
- D. Layered approach where each layer transforms requests

Answer: D

Q831. What is Address Space Layout Randomization (ASLR)?

- A. Randomizing data area positions to prevent exploits
- B. A network addressing scheme for distributed nodes
- C. A disk layout optimization for faster read access
- D. A memory allocation algorithm used by OS kernel

Answer: A

Q832. What is a rootkit?

- A. A set of root user access tools for admin
- B. A root user configuration file set used
- C. A system administration utility toolkit
- D. Malware providing hidden privileged access

Answer: D

Q833. What is the concept of security rings (protection rings) in OS architecture?

- A. Hierarchical privilege levels with inner having more
- B. Physical rings around the processor die chip
- C. A type of ring network topology for communication
- D. A form of data encryption using circular key bits

Answer: A

Q834. What is a side-channel attack?

- A. Exploiting physical info like timing or cache behavior
- B. An attack through a secondary network connection
- C. An attack on backup systems stored at off-site
- D. An attack on auxiliary devices connected to system

Answer: A

Q835. What are Spectre and Meltdown vulnerabilities?

- A. Software bugs found in major web browser apps
- B. Hardware flaws exploiting speculative execution to leak data
- C. Network protocol flaws in the TCP/IP stack implementation
- D. File system vulnerabilities found in modern operating systems

Answer: B

Q836. What is mandatory access control (MAC) and how does it differ from discretionary access control (DAC)?

- A. DAC is mandatory and MAC is discretionary
- B. They are the same access control model used
- C. MAC is system-enforced; DAC is owner-controlled
- D. MAC is weaker than DAC in all situations

Answer: C

Q837. What is Role-Based Access Control (RBAC)?

- A. Access control based on user physical location
- B. A type of encryption for access control data
- C. Controlling access based on time of day only
- D. Permissions assigned to roles; users gain via roles

Answer: D

Q838. What is a sandbox in operating system security?

- A. A debugging tool for code analysis testing
- B. A test database for development and staging
- C. Isolated environment restricting program access
- D. A type of malware infection on the system

Answer: C

Q839. What is the Bell-LaPadula security model?

- A. A network security protocol for data transfer
- B. A database security model for SQL prevention
- C. Confidentiality model: no read up, no write down
- D. A physical security model for server rooms

Answer: C

Q840. What is the Biba integrity model and how does it differ from Bell-LaPadula?

- A. Biba focuses on integrity: no read down, no write up
- B. They are the exact same model under different names
- C. Biba is a networking model for packet integrity checks
- D. Biba replaces Bell-LaPadula in all modern OS systems

Answer: A

Q841. In a hybrid kernel architecture, which services typically remain in kernel space?

- A. Only user applications remain inside the kernel space
- B. Networking stacks exclusively run in the kernel space
- C. All device drivers and file systems run in user space
- D. Core IPC and some drivers stay in kernel for speed

Answer: D

Q842. What security vulnerability does the Meltdown attack exploit in CPUs?

- A. It exploits buffer overflow in user application code
- B. It exploits speculative execution to read kernel data
- C. It exploits faulty device drivers in the OS subsystem
- D. It exploits weak encryption in network protocol sets

Answer: B

Q843. What is the key design trade-off in exokernel operating systems?

- A. Exokernels sacrifice all security for maximum speed
- B. They trade app portability for resource efficiency now
- C. Exokernels trade memory capacity for CPU performance
- D. They sacrifice networking for faster disk I/O throughput

Answer: B

Q844. How does a Type-1 hypervisor differ from a Type-2 hypervisor?

- A. Type-2 runs on bare metal while Type-1 needs a host
- B. Type-1 runs directly on bare metal without host OS
- C. Type-1 only supports Linux; Type-2 supports all OSes
- D. Both types require a host operating system to function

Answer: B

Q845. What problem does the UEFI boot process solve over legacy BIOS?

- A. UEFI only supports smaller disk sizes than BIOS can
- B. UEFI is slower but more compatible with older hardware
- C. UEFI eliminates the need for any operating system now
- D. UEFI removes the 2TB disk limit and adds secure boot

Answer: D

Q846. In capability-based OS design, what does a capability represent?

- A. A type of scheduling algorithm for real-time tasks
- B. A measure of CPU processing speed in megahertz
- C. A token granting specific access rights to objects
- D. A hardware register for storing interrupt addresses

Answer: C

Q847. What challenge does the OS face in managing NUMA architectures?

- A. Memory access times vary by processor-memory distance
- B. All memory accesses take the same time in NUMA setups
- C. NUMA systems cannot support more than one processor
- D. NUMA eliminates the need for cache memory in all CPUs

Answer: A

Q848. What is the fundamental principle behind separating mechanism and policy?

- A. It ensures the OS can never be modified or updated
- B. It prevents any user from customizing the OS at all
- C. It allows changing decisions without changing the code
- D. It requires all policies to be hard-coded in hardware

Answer: C

Q849. What is a unikernel and how does it differ from a traditional OS?

- A. A unikernel bundles app and OS in one address space
- B. A unikernel is a multi-user general-purpose OS now
- C. A unikernel requires a full OS installation underneath
- D. A unikernel runs many applications simultaneously here

Answer: A

Q850. Why is the halting problem relevant to operating system design?

- A. It guarantees all programs terminate in finite time span
- B. It proves the OS can always detect all infinite loops
- C. It shows no algorithm can decide if programs will halt
- D. It applies only to database management systems today

Answer: C

Q851. In the L4 microkernel family, what design decision improves IPC performance?

- A. L4 uses asynchronous message passing exclusively always
- B. L4 batches all IPC messages and processes them in bulk
- C. L4 implements synchronous IPC with register-based transfer
- D. L4 eliminates IPC entirely and uses shared memory only

Answer: C

Q852. What is the key challenge in defining layers for a layered OS?

- A. Finding hardware that supports the layered architecture
- B. Determining correct layer ordering with dependencies here
- C. Making sure each layer has the exact same code size now
- D. Ensuring every layer can communicate with all other ones

Answer: B

Q853. How does the Mach microkernel handle memory management?

- A. It requires only hardware-specific memory management
- B. It manages all memory entirely within the microkernel
- C. It does not support virtual memory in any form at all
- D. It uses external pagers in user space for page handling

Answer: D

Q854. What is a system generation (SYSGEN) process in OS deployment?

- A. It automatically writes new operating system source code
- B. It generates documentation for the operating system APIs
- C. It configures the OS for specific hardware at installation
- D. It creates backup copies of the OS kernel on other disk

Answer: C

Q855. Why do modern OS designs favor hybrid kernel over pure microkernel?

- A. Pure microkernels provide better performance in all cases
- B. Hybrid kernels balance modularity with better performance
- C. Pure microkernels cannot support any device drivers now
- D. Hybrid kernels are simpler to implement than microkernels

Answer: B

Q856. What is the role of the Hardware Abstraction Layer in Windows NT?

- A. HAL manages user interface themes and visual appearance
- B. HAL compiles C++ code into native machine instructions
- C. HAL hides hardware differences from upper OS layers now
- D. HAL provides a game rendering engine for DirectX apps

Answer: C

Q857. How does the Android architecture differ from standard Linux?

- A. Android adds middleware like ART runtime over Linux kernel
- B. Android does not use any Linux components in its design
- C. Android replaces the Linux kernel with a custom one here
- D. Android uses Linux only for networking and nothing else

Answer: A

Q858. What problem does eBPF solve in modern Linux kernels?

- A. It replaces the TCP/IP stack with a new protocol stack
- B. It allows safe programmable extensions in the kernel now
- C. It provides a graphical desktop environment for Linux
- D. It manages Bluetooth device connections in Linux only

Answer: B

Q859. What is the significance of MINIX 3 in operating system research?

- A. MINIX 3 eliminated the need for file systems completely
- B. MINIX 3 proved monolithic kernels are always superior
- C. MINIX 3 demonstrated a reliable self-healing microkernel
- D. MINIX 3 was the first OS to support graphical interfaces

Answer: C

Q860. How do C library wrappers simplify system call invocation?

- A. They bypass the kernel entirely for all OS operations
- B. They provide function interface hiding assembly details
- C. They convert system calls into network protocol packets
- D. They eliminate the need for privilege mode transitions

Answer: B

Q861. How does copy-on-write optimize the fork() system call?

- A. It eliminates the need for virtual memory addressing
- B. It copies the entire address space at fork instantly
- C. It prevents child processes from writing any memory
- D. It shares pages until a write then copies that page

Answer: D

Q862. What is vfork() and how does it differ from fork()?

- A. vfork() creates multiple child processes at same time
- B. vfork() bypasses the kernel for process creation now
- C. vfork() is identical to fork() in all design aspects
- D. vfork() shares address space and suspends the parent

Answer: D

Q863. How do process groups and sessions work in Unix systems?

- A. Sessions and process groups are the same exact concept
- B. Every process belongs to exactly one group and session
- C. Process groups exist only for root user processes now
- D. Process groups are only used in Windows and not Unix

Answer: B

Q864. How do preemptive and cooperative multitasking differ?

- A. Preemptive lets OS forcibly take CPU from a process
- B. Cooperative lets the OS forcibly interrupt any process
- C. Both approaches give unlimited CPU time to processes
- D. Preemptive requires processes to yield voluntarily now

Answer: A

Q865. What is the purpose of the clone() system call in Linux?

- A. It converts a user process into a kernel-level process
- B. It merges two existing processes into a single process
- C. It creates a thread or process with fine-grained sharing
- D. It only duplicates files on disk, not any processes

Answer: C

Q866. How does Chrome use process architecture for security?

- A. Chrome uses separate processes per tab for isolation
- B. Chrome runs all tabs in one single process for speed
- C. Chrome does not use processes; it uses threads only
- D. Chrome relies only on the OS kernel for tab isolation

Answer: A

Q867. What is the role of the init process (PID 1) in Unix?

- A. It only manages network interfaces during boot time
- B. It handles only file system mounting and nothing else
- C. It is temporary and exits after boot process completes
- D. It is ancestor of all processes and adopts all orphans

Answer: D

Q868. What challenges arise with process migration in distributed systems?

- A. Only the process ID needs to move between machines now
- B. Process migration is trivial with no challenges at all
- C. Open file handles and network state must be transferred
- D. Process migration eliminates the need for scheduling here

Answer: C

Q869. How does process creation differ between Windows and Unix?

- A. Windows maintains a strict parent-child tree like Unix
- B. Neither system tracks relationships between processes now
- C. Unix has parent-child trees; Windows uses handle tokens
- D. Both systems have identical process creation mechanisms

Answer: C

Q870. What is process checkpointing in fault-tolerant systems?

- A. It saves process state periodically for failure recovery
- B. It permanently deletes process state to free memory
- C. It checks if a process has any source code syntax errors
- D. It counts the number of processes running on the system

Answer: A

Q871. How does the Completely Fair Scheduler in Linux work?

- A. CFS maintains virtual runtime for fair CPU allocation
- B. CFS uses lottery scheduling for random CPU assignment
- C. CFS uses strict priority queues with fixed time slices
- D. CFS implements simple Round Robin with a fixed quantum

Answer: A

Q872. What is the exponential averaging formula for burst prediction?

- A. Next burst is always set to a fixed constant value here
- B. Next burst equals the exact length of the previous burst
- C. $\tau(n+1) = \alpha \cdot t(n) + (1-\alpha) \cdot \tau(n)$ using history
- D. Prediction uses only the average of all prior burst times

Answer: C

Q873. What is priority inversion and how can it be solved?

- A. Low-priority blocks high-priority; use priority inheritance
- B. High-priority tasks always run first without any exception
- C. Priority inversion occurs only in single-core CPU systems
- D. It is solved by removing all priorities from the system

Answer: A

Q874. How does the O(1) scheduler differ from CFS in Linux?

- A. O(1) is fairer than CFS for interactive desktop workloads
- B. O(1) uses active/expired arrays; CFS uses red-black tree
- C. Both schedulers are identical in design and implementation
- D. CFS uses fixed time slices while O(1) uses virtual runtime

Answer: B

Q875. What is gang scheduling in multiprocessor systems?

- A. It only works with single-threaded processes on one core
- B. It co-schedules related threads simultaneously on CPUs
- C. It randomly assigns processes to processors without order
- D. It schedules one process at a time across all CPUs now

Answer: B

Q876. How does processor affinity affect multiprocessor scheduling?

- A. It has no effect on scheduling in multiprocessor systems
- B. It forces processes to migrate between CPUs constantly
- C. It keeps a process on the same CPU to exploit warm cache
- D. It assigns all processes to the same single CPU core here

Answer: C

Q877. What is global vs. partitioned scheduling in multicore systems?

- A. Global and partitioned scheduling are the exact same thing
- B. Global scheduling assigns each process to a fixed core
- C. Partitioned scheduling uses a single queue for all cores
- D. Global uses one shared queue; partitioned uses per-core

Answer: D

Q878. What is rate-monotonic scheduling for real-time systems?

- A. It uses Round Robin for all periodic real-time task sets
- B. It assigns the same fixed priority to every task equally
- C. It assigns priorities randomly to all real-time task sets
- D. It assigns priorities inversely proportional to periods

Answer: D

Q879. How does Earliest Deadline First schedule real-time tasks?

- A. It only works with non-periodic tasks in batch systems
- B. It dynamically prioritizes tasks with nearest deadlines
- C. It schedules the task with the latest deadline first
- D. It ignores deadlines and uses burst time for scheduling

Answer: B

Q880. What is the schedulability test for rate-monotonic scheduling?

- A. All tasks must have priority higher than the OS kernel
- B. Each task must have the same exact period and burst time
- C. CPU utilization must be below 100 percent for all tasks
- D. Sum of C_i/T_i must be at most $n \cdot (2^{1/n} - 1)$ for guarantee

Answer: D

Q881. What is a lightweight process (LWP) in many-to-many threading?

- A. An LWP is a file storing thread state on the hard disk
- B. An LWP is intermediate between user and kernel threads
- C. An LWP is a network packet for inter-thread messaging
- D. An LWP is a physical CPU core in the processor chip

Answer: B

Q882. What is the scheduler activation mechanism in threading?

- A. Kernel notifies user library via upcalls about events
- B. It is a simple timer interrupt for thread scheduling
- C. It replaces the CPU scheduler with a user-space one
- D. It prevents communication between kernel and threads

Answer: A

Q883. How does Windows thread scheduling differ from Linux?

- A. Both systems only schedule processes, never any threads
- B. Windows schedules processes; Linux schedules threads
- C. Windows uses Round Robin; Linux uses SJF exclusively
- D. Windows schedules threads; Linux schedules tasks via CFS

Answer: D

Q884. What is the thundering herd problem in threaded servers?

- A. A single thread monopolizes the CPU preventing others
- B. Threads run too slowly due to excessive synchronization
- C. Too many threads cause total memory exhaustion now
- D. Multiple threads wake for one event but one handles it

Answer: D

Q885. What is the purpose of the futex mechanism in Linux?

- A. It provides a graphics rendering pipeline for threads
- B. It controls CPU frequency scaling for threaded workload
- C. It manages file transfers between threads over network
- D. It implements fast user-space locking with kernel fallback

Answer: D

Q886. How does hyperthreading affect OS thread scheduling?

- A. It eliminates the need for any software threading at all
- B. It reduces the number of CPUs visible to the OS kernel
- C. It has no effect on how the OS schedules threads at all
- D. OS sees additional logical CPUs and schedules on them now

Answer: D

Q887. What are the challenges of implementing user-level threads?

- A. User-level threads are always slower than kernel threads
- B. User-level threads cannot be created by any application
- C. User-level threads require special hardware support only
- D. A blocking syscall blocks all threads in the process now

Answer: D

Q888. What is the two-level model in thread implementation?

- A. It combines many-to-many with one-to-one mapping
- B. It requires two separate kernels for thread management
- C. It uses exactly two threads per process, no more here
- D. It maps all threads to exactly two kernel-level threads

Answer: A

Q889. How does Grand Central Dispatch handle threading on macOS?

- A. GCD only supports single-threaded application execution
- B. GCD requires programmers to manually create threads
- C. GCD uses blocks and dispatch queues managed by runtime
- D. GCD is a hardware threading mechanism in Apple chips

Answer: C

Q890. What is a fiber and how does it differ from a thread?

- A. Fibers are identical to threads in every way possible
- B. Fibers are hardware threads built into processor cores
- C. Fibers are cooperatively scheduled user-level threads
- D. Fibers are kernel-level threads with preemptive scheduling

Answer: C

Q891. What is the dining philosophers problem and its significance?

- A. It models network routing algorithms in large data centers
- B. It is about optimizing database query execution plans now
- C. It models resource allocation leading to potential deadlock
- D. It describes scheduling of philosophers teaching classes

Answer: C

Q892. How does Read-Copy-Update (RCU) achieve lock-free reads?

- A. RCU prevents any writes from occurring while reading
- B. RCU uses spinlocks internally for all read operations
- C. RCU locks all readers while writer updates the data
- D. Readers access old data while writers prepare new copy

Answer: D

Q893. What is the ABA problem in lock-free data structures?

- A. ABA is an algorithm for balanced binary tree building
- B. It describes automatic backup and archival of data
- C. It refers to alphabetical ordering of lock acquisition
- D. Value changes from A to B to A, hiding the updates

Answer: D

Q894. What is a memory barrier instruction used for?

- A. It ensures memory operations complete in proper order
- B. It prevents the CPU from accessing some memory regions
- C. It doubles available memory by compressing all data now
- D. It creates a physical barrier in RAM to separate data

Answer: A

Q895. What differs between Mesa and Hoare monitor semantics?

- A. Hoare uses signal-and-continue; Mesa signal-and-wait
- B. Mesa requires busy waiting; Hoare uses blocking only
- C. Mesa and Hoare semantics are completely identical here
- D. Mesa uses signal-and-continue; Hoare signal-and-wait

Answer: D

Q896. How does transactional memory simplify concurrent programming?

- A. It only works with single-threaded programs in practice
- B. It requires more locks than traditional synchronization
- C. It eliminates all shared memory between threads fully
- D. It lets code blocks execute atomically without locks now

Answer: D

Q897. What is the priority ceiling protocol and how does it work?

- A. Each mutex gets ceiling equal to highest priority user now
- B. It only lets highest priority thread acquire locks at all
- C. It lowers all thread priorities to zero to stop conflicts
- D. It eliminates all priorities making every thread equal here

Answer: A

Q898. What is lock-free vs. wait-free in concurrent algorithms?

- A. Wait-free ensures system progress but not per-thread
- B. Lock-free ensures system progress; wait-free per-thread
- C. Lock-free and wait-free algorithms are exactly the same
- D. Lock-free uses hidden locks but hides them from users

Answer: B

Q899. What does the sleeping barber problem model?

- A. It models network packet queuing in router memory now
- B. It describes file system journaling for crash recovery
- C. A scheduling algorithm for barber shop appointments
- D. It models bounded-buffer producer-consumer with sleep

Answer: D

Q900. How does Linux handle preemption within critical sections?

- A. Linux disables preemption using preemption counters now
- B. Linux never allows any kernel-level preemption at all
- C. Linux always allows preemption inside critical sections
- D. Linux uses hardware locks to prevent all preemption

Answer: A

Q901. What is the time complexity of the Banker's algorithm?

- A. $O(n \times n \times m)$ for determining safe sequence each time here
- B. $O(n)$ linear time complexity for each resource request
- C. $O(\log n)$ logarithmic time for each allocation decision
- D. $O(1)$ constant time regardless of process or resources

Answer: A

Q902. How does the Linux kernel handle deadlock in practice?

- A. Linux uses lockdep to detect potential deadlock patterns
- B. Linux ignores deadlocks and restarts the system if stuck
- C. Linux prevents all four deadlock conditions at all times
- D. Linux implements full Banker's algorithm at runtime

Answer: A

Q903. What is two-phase locking's relationship to deadlock?

- A. Two-phase locking completely eliminates all deadlocks
- B. Two-phase locking causes deadlock in every single case
- C. Two-phase locking ensures serializability allows deadlock
- D. Two-phase locking has no relationship to deadlock at all

Answer: C

Q904. What is a knot in a resource allocation graph?

- A. A knot is simply another name for any edge in graph
- B. A knot means the graph has no cycles present at all
- C. A cycle where no outside process can break the cycle
- D. A knot indicates that the system is in a safe state

Answer: C

Q905. How does wound-wait differ from wait-die for databases?

- A. Wait-die: older waits or dies; wound-wait wounds younger
- B. Both schemes are identical and behave the same always
- C. Wound-wait: older wounds younger; wait-die: older waits
- D. Wound-wait preempts younger; wait-die makes older wait

Answer: C

Q906. How are victim selection strategies used in deadlock recovery?

- A. Cost-based selection minimizes rollback overhead impact
- B. Victim selection only applies to file system operations
- C. All victims are always the newest process in the system
- D. Victim selection is random and cannot be optimized now

Answer: A

Q907. How can deadlock occur in distributed systems?

- A. Processes on different nodes can form distributed cycles
- B. Deadlock cannot occur across multiple distributed nodes
- C. Only single-node deadlocks are possible in any system
- D. Distributed systems automatically prevent all deadlocks

Answer: A

Q908. What is the phantom deadlock problem in distributed detection?

- A. It describes deadlocks that only occur during system startup
- B. It is a deadlock that is invisible to all detection methods
- C. It occurs when a detected deadlock does not actually exist
- D. It refers to deadlocks in virtual machine environments only

Answer: C

Q909. What is hierarchical deadlock prevention using lock ordering?

- A. It assigns random order to locks at runtime dynamically
- B. It groups locks into levels and enforces strict ordering
- C. It eliminates all locks from the system to prevent issues
- D. It uses only one global lock for the entire system here

Answer: B

Q910. How does the edge-chasing algorithm detect distributed deadlocks?

- A. It uses a centralized server to collect all wait-for info
- B. It prevents deadlock rather than detecting it after fact
- C. It broadcasts all lock info to every node all the time
- D. It propagates probe messages along wait-for edges across

Answer: D

Q911. How does the inverted page table reduce memory usage?

- A. It eliminates the need for any page table entirely
- B. It uses one entry per virtual page for each process
- C. It compresses page table entries to half their size
- D. It uses one entry per physical frame not per page

Answer: D

Q912. What is the clock (second-chance) page replacement algorithm?

- A. It gives each page a second chance via reference bit
- B. It always replaces the page with highest frame number
- C. It schedules page replacements at fixed clock intervals
- D. It replaces pages based on the actual wall clock time

Answer: A

Q913. What is copy-on-write in virtual memory management?

- A. Copying entire memory space when a process is created
- B. Sharing pages until write occurs then making private copy
- C. Writing all memory contents to disk after every instruction
- D. Preventing any process from writing to memory at all

Answer: B

Q914. How does ASLR improve security?

- A. It disables virtual memory to improve system performance
- B. It encrypts all data stored in process address space
- C. It randomizes code and data locations to prevent exploits
- D. It prevents all processes from accessing shared memory

Answer: C

Q915. What is the optimal page replacement algorithm's limitation?

- A. It does not support virtual memory in any implementation
- B. It causes more page faults than any other algorithm
- C. It can only work with a single page frame in system
- D. It requires future knowledge of page references to work

Answer: D

Q916. How does huge page support improve TLB performance?

- A. Huge pages reduce TLB size to improve lookup speed
- B. Huge pages reduce TLB miss frequency with fewer entries
- C. Huge pages eliminate the need for any TLB hardware now
- D. Huge pages make TLB slower but more accurate in lookup

Answer: B

Q917. What is the page-buffering algorithm used for?

- A. Buffering network data in page-sized blocks only
- B. Compressing pages to fit more data into memory now
- C. Preventing page faults by locking all pages in RAM
- D. Keeping a pool of free frames to speed replacement

Answer: D

Q918. What is memory-mapped I/O and how does it work?

- A. It stores all I/O data exclusively in the CPU registers
- B. It maps physical memory addresses to network addresses
- C. It maps I/O device registers into memory address space
- D. It eliminates the need for any I/O devices in system

Answer: C

Q919. What is kernel same-page merging (KSM) in Linux?

- A. It merges all kernel modules into a single binary
- B. It identifies and merges identical pages across procs
- C. It prevents any page sharing between virtual machines
- D. It splits each page into two smaller pages for speed

Answer: B

Q920. How does the buddy system manage memory allocation?

- A. It assigns memory randomly without any structure here
- B. It pairs every process with a buddy for backup now
- C. It splits memory into power-of-two blocks recursively
- D. It allocates memory in fixed blocks without splitting

Answer: C

Q921. How does ext4 use extents instead of block pointers?

- A. Extents are only used for files smaller than 4 KB
- B. An extent describes a contiguous run with one entry
- C. Extents eliminate the need for any metadata on disk
- D. Extents store each block address individually here

Answer: B

Q922. What is a copy-on-write file system and its advantage?

- A. It copies all data twice for redundancy on every write
- B. It writes data in place and copies only on deletion
- C. It never overwrites existing data ensuring consistency
- D. It only allows reads and prevents all write operations

Answer: C

Q923. How does the B+ tree structure benefit NTFS?

- A. It provides efficient $O(\log n)$ lookups and range scans
- B. It only works with files smaller than one megabyte now
- C. It stores file data directly inside B-tree leaf nodes
- D. It eliminates the need for any disk space management

Answer: A

Q924. What is the log-structured file system design philosophy?

- A. LFS stores data in traditional inode-based tree layout
- B. LFS writes all data and metadata sequentially to log
- C. LFS requires all data to be stored in main memory now
- D. LFS uses random writes to maximize disk head movement

Answer: B

Q925. How does ZFS ensure data integrity using checksums?

- A. ZFS checksums are stored alongside data in same block
- B. ZFS stores checksums in parent blocks enabling healing
- C. ZFS does not use any checksums for integrity checking
- D. ZFS only checksums metadata, not actual file data now

Answer: B

Q926. What is the TRIM command for SSDs in file systems?

- A. It trims the physical size of the SSD device now
- B. It informs SSD which blocks are no longer in use
- C. It compresses data on SSD to save storage space
- D. It increases read speed by defragmenting the SSD

Answer: B

Q927. How do deduplication techniques work in enterprise storage?

- A. They duplicate important files for safety purposes
- B. They compress each file individually to reduce size
- C. They identify and eliminate duplicate blocks to save
- D. They remove all files not accessed in recent months

Answer: C

Q928. What differs between ordered and writeback journaling?

- A. Ordered mode does not use a journal; writeback does
- B. Writeback writes data first; ordered metadata first now
- C. Both modes write data and metadata in the same order
- D. Ordered writes data before metadata; writeback either

Answer: D

Q929. What is the purpose of the superblock in Unix file systems?

- A. It stores the boot loader program for the OS kernel
- B. It contains critical file system metadata like size now
- C. It stores the contents of the largest file in system
- D. It is extra storage reserved for the root user only

Answer: B

Q930. How does FUSE work in Linux?

- A. FUSE moves all file system code into kernel space
- B. FUSE allows implementing file systems in user space
- C. FUSE replaces the VFS layer with a simpler version
- D. FUSE is a hardware device for connecting ext drives

Answer: B

Q931. How does LOOK improve upon the SCAN algorithm?

- A. LOOK eliminates the need for any head movement
- B. LOOK moves head to the physical end of disk
- C. LOOK reverses at the last request, not disk end
- D. LOOK processes requests in strict FCFS order now

Answer: C

Q932. What is the write amplification problem in SSDs?

- A. Writes are amplified to improve SSD performance
- B. SSDs amplify signals to write data faster on chip
- C. Writes are duplicated to multiple chips for safety
- D. More physical data is written than logical requested

Answer: D

Q933. How does RAID 6 improve on RAID 5 fault tolerance?

- A. RAID 6 uses triple mirroring of all data blocks
- B. RAID 6 uses dual parity allowing two disk failures
- C. RAID 6 stores four copies of parity on one drive
- D. RAID 6 and RAID 5 offer exactly the same safety

Answer: B

Q934. What is the FTL in solid-state drives?

- A. A network layer for transferring data between SSDs
- B. A physical layer connecting SSD to motherboard
- C. Firmware mapping logical addresses to flash pages
- D. A software tool for formatting flash drives on PC

Answer: C

Q935. What is shingled magnetic recording for HDDs?

- A. It uses circular tracks with no overlap at all
- B. It uses multiple read heads for faster read speed
- C. Tracks overlap like shingles for higher density now
- D. It eliminates platters using solid-state memory

Answer: C

Q936. What is NCQ in SATA drives?

- A. It prevents commands from executing out of order
- B. It queues user interface commands for the OS now
- C. It lets drives reorder commands for optimal execution
- D. It queues network commands for remote disk access

Answer: C

Q937. How does NVMe improve upon AHCI for SSDs?

- A. NVMe uses the same command set as legacy IDE chips
- B. NVMe and AHCI provide identical performance for SSDs
- C. NVMe uses deeper queues and lower latency than AHCI
- D. NVMe is designed for spinning disks limiting SSDs

Answer: C

Q938. What is the NOOP I/O scheduler and when is it used?

- A. NOOP completely eliminates all I/O scheduling here
- B. NOOP is the most advanced scheduler for disk drives
- C. NOOP uses simple FIFO queue with minimal ordering
- D. NOOP performs extensive sorting of I/O for HDDs

Answer: C

Q939. What is the RAID write hole problem?

- A. Power failure during write leaves parity inconsistent
- B. It prevents any write from completing on RAID now
- C. It is a physical hole in disk platter causing damage
- D. It only occurs when reading data from RAID array now

Answer: A

Q940. How does host-managed zoned storage work with ZNS SSDs?

- A. SSD firmware manages zones without any host input
- B. ZNS SSDs work identically to traditional block SSDs
- C. The host has no control over data placement here
- D. Host manages zones and writes sequentially in them

Answer: D

Q941. How does scatter-gather DMA improve I/O performance?

- A. It transfers data to non-contiguous memory regions now
- B. It eliminates DMA and uses programmed I/O instead
- C. It gathers all data into single buffer before transfer
- D. It scatters data randomly across memory for security

Answer: A

Q942. What is the IOCP model in Windows?

- A. A port number used for network socket communication
- B. A scalable model for handling many async I/O events
- C. A model that blocks all threads until I/O finishes
- D. A physical port for connecting I/O devices to PC

Answer: B

Q943. How does Linux epoll improve over select()?

- A. epoll uses O(1) notification instead of O(n) scanning
- B. epoll only works with network sockets not file I/O
- C. epoll and select have identical performance overall
- D. epoll monitors fewer file descriptors than select

Answer: A

Q944. What is IOMMU and how does it improve I/O security?

- A. It is a type of MMU used exclusively for memory mgmt
- B. It encrypts all data transferred between devices and CPU
- C. It translates device addresses providing memory isolation
- D. It is a memory unit storing I/O driver code only

Answer: C

Q945. How do edge-triggered and level-triggered I/O notifications differ?

- A. Edge-triggered notifications are always slower than level
- B. Both notification types trigger in the same situation here
- C. Level-triggered only works with sockets not with disk I/O
- D. Edge notifies on state transition; level while state holds

Answer: D

Q946. What is the io_uring interface in modern Linux?

- A. It provides high-performance async I/O via ring buffers
- B. It creates circular network connections between nodes
- C. It manages circular disk platters in HDD controllers
- D. It is a graphical interface for configuring I/O drivers

Answer: A

Q947. How does RDMA work?

- A. It requires CPU to copy all data for remote transfers
- B. It allows direct memory access between remote machines
- C. It only works within a single machine for local DMA
- D. It uses traditional TCP/IP for remote memory operations

Answer: B

Q948. What is interrupt coalescing and why is it used?

- A. It eliminates all interrupts using only polling for I/O
- B. It combines multiple events into fewer interrupt signals
- C. Each I/O event generates a separate interrupt always
- D. It splits one interrupt into multiple signals for detail

Answer: B

Q949. What is the SPDK approach to storage?

- A. It uses user-space drivers and polling for low latency
- B. It is a standard kernel-based storage driver framework
- C. It is a hardware kit for assembling storage arrays now
- D. It provides a graphical toolkit for storage management

Answer: A

Q950. How does virtio enable efficient I/O virtualization?

- A. It uses hardware passthrough without any abstraction
- B. It eliminates the need for any I/O drivers in the VM
- C. It provides standard paravirtual I/O interface for VMs
- D. It emulates physical hardware perfectly for each VM

Answer: C

Q951. How does return-oriented programming bypass NX protection?

- A. ROP disables the NX bit in CPU page table entries
- B. ROP chains existing code gadgets to execute payload
- C. ROP injects new executable code into stack directly
- D. ROP only works on systems without memory protection

Answer: B

Q952. What is the Trusted Platform Module and its security role?

- A. A hardware chip providing cryptographic security ops
- B. A kernel module managing user authentication only
- C. A software module scanning for viruses on system
- D. A network module filtering traffic like a firewall

Answer: A

Q953. How does mandatory access control differ from discretionary?

- A. MAC uses system-enforced labels; DAC lets owners control
- B. MAC and DAC are different names for the same mechanism
- C. Both give resource owners full control over all access
- D. MAC lets owners set permissions; DAC uses system policy

Answer: A

Q954. What is a side-channel attack and how does Spectre work?

- A. Spectre exploits speculative execution via cache timing
- B. Spectre exploits buffer overflows in browser apps only
- C. It attacks through unused network ports on system
- D. Side-channel attacks directly modify target source code

Answer: A

Q955. How does SELinux implement mandatory access control?

- A. SELinux assigns security labels to subjects and objects
- B. SELinux only encrypts file system data for privacy
- C. SELinux manages firewalls but not file access at all
- D. SELinux replaces Linux kernel with a secure version

Answer: A

Q956. What is a zero-day vulnerability?

- A. A flaw affecting zero users due to limited deployment
- B. A flaw existing only on the day software is released
- C. A flaw known for zero days and still unpatched now
- D. A flaw taking zero effort to exploit using any tool

Answer: C

Q957. How does control flow integrity protect against code-reuse?

- A. CFI encrypts all control flow instructions in binary
- B. CFI ensures execution follows predetermined control graph
- C. CFI only protects against buffer overflow not ROP now
- D. CFI prevents any branching instructions from executing

Answer: B

Q958. What is the confused deputy problem in security?

- A. A user confused about which password to use for login
- B. A privileged program tricked into misusing its authority
- C. An admin who misconfigures the security settings now
- D. A process confusing virtual and physical memory addrs

Answer: B

Q959. How does memory tagging (MTE) prevent memory safety bugs?

- A. MTE tags memory with metadata detecting out-of-bounds
- B. MTE increases the amount of physical RAM available here
- C. MTE encrypts all memory contents using hardware now
- D. MTE only works with interpreted languages not compiled

Answer: A

Q960. What is attestation in trusted computing?

- A. A technique for testing software under heavy load stress
- B. A legal document certifying software license compliance
- C. Cryptographic proof that system is in a known trusted state
- D. A user declaring identity through a sworn statement now

Answer: C

Q961. What is the key difference between a Type-0 and a Type-2 hypervisor in virtualization?

- A. Type-0 is firmware-based while Type-2 runs as an application on a host OS
- B. Type-0 is software while Type-2 is hardware
- C. Type-0 supports only Linux while Type-2 supports only Windows
- D. Type-0 is slower than Type-2 in all scenarios

Answer: A

Q962. What is paravirtualization and how does it improve VM performance?

- A. Guests are modified to make hypercalls instead of privileged instructions, reducing overhead
- B. Guests run unmodified but the hypervisor interprets all instructions
- C. The host OS is removed and guests run directly on hardware
- D. Multiple VMs share the same guest operating system image

Answer: A

Q963. What role does the IOMMU play in modern computer architectures?

- A. It translates device I/O addresses to physical memory addresses and provides isolation
- B. It manages CPU cache coherence across multiple cores
- C. It schedules interrupt handling among processors
- D. It encrypts all data transferred between CPU and memory

Answer: A

Q964. Why is the separation of user mode and kernel mode considered insufficient for modern security?

- A. Side-channel attacks like Spectre can leak kernel data even from user mode
- B. User mode programs cannot make system calls
- C. Kernel mode does not have access to hardware
- D. There are no privileged instructions in modern CPUs

Answer: A

Q965. What is the multikernel approach to OS design for many-core processors?

- A. It treats the machine as a network of independent cores communicating via messages
- B. It runs a single monolithic kernel replicated on each core
- C. It assigns one core to the kernel and all others to user processes
- D. It uses virtual machines to isolate each core

Answer: A

Q966. What is a container and how does it differ from a virtual machine?

- A. A container shares the host OS kernel while a VM runs its own full OS
- B. A container requires a hypervisor while a VM does not
- C. A container provides stronger isolation than a VM
- D. A container can only run Linux while a VM supports any OS

Answer: A

Q967. What is binary translation in the context of full virtualization?

- A. The hypervisor dynamically rewrites guest privileged instructions at runtime
- B. The compiler converts source code to a different CPU architecture
- C. The OS translates binary files between file formats
- D. The CPU converts binary to decimal during arithmetic operations

Answer: A

Q968. What design challenge does heterogeneous multiprocessing (big.LITTLE) introduce for operating systems?

- A. The scheduler must match task requirements to cores with different performance and power profiles
- B. All cores must run at the same clock speed
- C. The OS must use a separate kernel for each core type
- D. Applications must be compiled separately for each core type

Answer: A

Q969. What is the principle behind a library OS in the context of cloud computing?

- A. Application and OS are compiled together into a single-purpose image for minimal overhead
- B. The OS loads libraries dynamically from a central cloud repository
- C. All OS functions are implemented as shared libraries on a network
- D. Applications use library calls instead of system calls for I/O

Answer: A

Q970. What is the primary purpose of hardware-assisted virtualization extensions like Intel VT-x?

- A. To provide CPU hardware support for efficiently running virtual machines
- B. To increase the clock speed of the CPU
- C. To add more physical cores to the processor
- D. To enable GPUs to execute operating system code

Answer: A

Q971. How does the Darwin kernel used in macOS combine different architectural approaches?

- A. It uses a Mach microkernel core with BSD monolithic services in the same address space
- B. It runs entirely as a microkernel with all services in user space
- C. It uses a pure monolithic design identical to Linux
- D. It relies on a hypervisor to separate kernel components

Answer: A

Q972. What is the difference between a system call and a library call in terms of execution?

- A. A system call transitions to kernel mode while a library call stays in user mode
- B. A library call is always slower than a system call
- C. System calls cannot return values to the caller
- D. Library calls bypass the operating system entirely in all cases

Answer: A

Q973. What is the vDSO mechanism in Linux and why is it used?

- A. It maps kernel code into user space to avoid system call overhead for certain functions
- B. It provides virtual device drivers for emulated hardware
- C. It creates virtual disk storage from RAM
- D. It is a debugging tool for tracing system calls

Answer: A

Q974. What is the challenge of maintaining ABI compatibility in operating system development?

- A. Changes to system call interfaces or data structures can break existing compiled applications
- B. ABI compatibility only affects interpreted languages
- C. Maintaining ABI compatibility slows down hardware performance
- D. ABI compatibility is automatically handled by the CPU

Answer: A

Q975. How do system calls work on x86-64 Linux using the SYSCALL instruction?

- A. The syscall number is placed in RAX and SYSCALL transfers control to the kernel entry point
- B. The CPU automatically determines the system call from the program context
- C. System calls are made by writing to a special memory address
- D. The SYSCALL instruction is only available in kernel mode

Answer: A

Q976. What is the seccomp mechanism in Linux kernel architecture?

- A. A security facility that restricts the system calls a process can make
- B. A tool for compressing the kernel image during boot
- C. A scheduler for secondary storage devices
- D. A method for secure communication between microkernels

Answer: A

Q977. What problem does the Windows Subsystem for Linux (WSL2) solve architecturally?

- A. It runs a real Linux kernel in a lightweight VM to provide native Linux binary compatibility
- B. It translates Linux source code to Windows executables at compile time
- C. It emulates x86 Linux on ARM Windows devices
- D. It replaces the Windows kernel with the Linux kernel

Answer: A

Q978. What is the role of the init system (e.g., systemd) in modern Linux architecture?

- A. It is the first user-space process that initializes services and manages the system lifecycle
- B. It is part of the kernel that handles hardware interrupts
- C. It compiles the kernel from source code during boot
- D. It manages the BIOS firmware settings

Answer: A

Q979. What is the difference between a stable and unstable ABI in kernel module development?

- A. A stable ABI guarantees compatibility across versions while unstable ABI may change at any time
- B. Stable ABI is faster while unstable ABI is more secure
- C. Stable ABI runs in user mode while unstable ABI runs in kernel mode
- D. There is no practical difference in modern operating systems

Answer: A

Q980. How does the Fuchsia OS Zircon kernel differ from traditional Unix-like kernels?

- A. It uses a capability-based object system instead of Unix file descriptors for resource access
- B. It is a monolithic kernel identical to Linux in design
- C. It does not support user-space applications
- D. It runs exclusively on x86 architecture

Answer: A

Q981. What is the double fork technique in Unix and what problem does it solve?

- A. Forking twice makes the grandchild an orphan adopted by init to avoid zombies
- B. It creates two identical copies of a process for load balancing
- C. It forks a process on two different CPUs simultaneously
- D. It creates a backup copy of the process for fault tolerance

Answer: A

Q982. How does the Linux cgroup mechanism control process resource usage?

- A. It groups processes and enforces limits on CPU, memory, and I/O usage per group
- B. It encrypts process data for security
- C. It compresses process memory to save space
- D. It pins processes to specific CPU cores only

Answer: A

Q983. What is the difference between posix_spawn() and fork()+exec() for process creation?

- A. posix_spawn creates a process in a single call avoiding the overhead of copying the parent's address space
- B. posix_spawn only works on Windows systems
- C. fork()+exec() is always faster than posix_spawn
- D. posix_spawn cannot set up file descriptors for the child process

Answer: A

Q984. What are Linux namespaces and how do they relate to process isolation?

- A. They isolate system resources so processes see their own independent view of PIDs, network, and mounts
- B. They provide a naming convention for files on disk
- C. They organize kernel source code into modules
- D. They assign human-readable names to IP addresses

Answer: A

Q985. What is the thundering herd problem in process management?

- A. Multiple waiting processes are all woken up for an event but only one can handle it, wasting resources
- B. Too many processes run simultaneously causing CPU overheating
- C. Processes consume too much memory during startup
- D. Network packets arrive faster than processes can handle them

Answer: A

Q986. How does the Windows process creation model differ from the Unix fork/exec model?

- A. Windows uses CreateProcess() which creates a new process from scratch with a specified executable
- B. Windows uses fork() identically to Unix
- C. Windows cannot create child processes
- D. Windows processes do not have process IDs

Answer: A

Q987. What is process affinity in the context of NUMA architectures?

- A. Binding processes to specific NUMA nodes to keep memory accesses local and reduce latency
- B. Measuring how much a process likes its assigned CPU
- C. Preventing processes from accessing network resources
- D. Grouping processes by the user who created them

Answer: A

Q988. What is the OOM (Out-of-Memory) killer in Linux?

- A. A kernel mechanism that selects and terminates processes when the system runs out of memory
- B. A tool that prevents users from allocating too much disk space
- C. A hardware circuit that shuts down when RAM overheats
- D. A security feature that kills processes attempting buffer overflows

Answer: A

Q989. What is a process descriptor in the Linux kernel and what does it contain?

- A. A task_struct that holds all information the kernel needs about a process
- B. A text description of what the process does
- C. A file containing the process source code
- D. A log of all commands the process has executed

Answer: A

Q990. What is the seccomp-bpf mechanism used for in process sandboxing?

- A. It filters system calls using BPF programs to restrict what a process can do
- B. It encrypts process memory using Berkeley Packet Filters
- C. It monitors network packets sent by a process
- D. It compresses process data to reduce memory usage

Answer: A

Q991. How does the Linux CFS use a red-black tree for scheduling decisions?

- A. Processes are ordered by virtual runtime in the tree, and the leftmost node with smallest vruntime runs next
- B. The tree stores processes sorted by PID and the root node runs next
- C. Processes are randomly inserted and the tree auto-balances priorities
- D. The tree is used only for I/O scheduling, not CPU scheduling

Answer: A

Q992. What is the utilization bound test for rate-monotonic scheduling?

- A. The set of tasks is schedulable if total utilization does not exceed $n(2^{1/n} - 1)$
- B. Tasks are schedulable if each task's period is greater than its execution time
- C. The test checks if total memory usage is below the physical memory limit
- D. Tasks are schedulable only if they all have the same period

Answer: A

Q993. What is the priority inversion problem and how did it affect the Mars Pathfinder?

- A. A high-priority task was blocked by a low-priority task holding a shared resource, causing system resets
- B. The rover's CPU overheated due to excessive computation
- C. A software bug caused the scheduler to run processes in reverse priority order
- D. Network communication delays caused missed commands from Earth

Answer: A

Q994. What is the difference between global scheduling and partitioned scheduling on multicore systems?

- A. Global scheduling uses a single queue for all cores while partitioned assigns each task to a fixed core
- B. Global scheduling is always worse than partitioned scheduling
- C. Partitioned scheduling does not support priority-based algorithms
- D. Global scheduling can only be used with two cores

Answer: A

Q995. What is the BFS (Brain F* Scheduler) design philosophy compared to CFS?**

- A. BFS uses a single global run queue with an earliest eligible virtual deadline algorithm for simplicity
- B. BFS uses per-CPU run queues with work stealing for scalability
- C. BFS is designed only for real-time tasks on embedded systems
- D. BFS uses a multilevel feedback queue identical to traditional Unix schedulers

Answer: A

Q996. How does Earliest Deadline First (EDF) differ from rate-monotonic scheduling for real-time tasks?

- A. EDF assigns dynamic priorities based on deadlines achieving 100% utilization while rate-monotonic uses fixed priorities
- B. EDF only works with aperiodic tasks while rate-monotonic handles periodic tasks
- C. Rate-monotonic achieves higher utilization bounds than EDF
- D. EDF requires hardware support that rate-monotonic does not

Answer: A

Q997. What is work stealing in multiprocessor scheduling?

- A. Idle processors take tasks from the run queues of busy processors to balance load
- B. Processes steal CPU cycles from higher-priority processes
- C. The scheduler removes tasks from slow processors and discards them
- D. Memory is reallocated from one core's cache to another

Answer: A

Q998. What is the SCHED_DEADLINE scheduling class in Linux?

- A. An EDF-based scheduler that guarantees CPU bandwidth to tasks using runtime/deadline/period parameters
- B. A scheduler that assigns deadlines based on process creation time
- C. A file system scheduling class for managing disk deadlines
- D. A network scheduler for prioritizing packet delivery

Answer: A

Q999. What is the problem with using lottery scheduling in practice?

- A. Unpredictable short-term behavior due to randomness even though long-term fairness is achieved
- B. It requires specialized hardware random number generators
- C. It cannot handle more than ten processes simultaneously
- D. It always gives equal time to all processes regardless of priority

Answer: A

Q1000. How does the Energy Aware Scheduler (EAS) in Linux optimize for mobile devices?

- A. It uses CPU energy models to place tasks on the most energy-efficient core that meets performance needs
- B. It shuts down the CPU when battery is below 50%
- C. It reduces screen brightness to save power during computation
- D. It limits the number of processes to reduce energy consumption

Answer: A

Q1001. How does the Linux kernel implement threads differently from traditional UNIX thread implementations?

- A. Linux does not support threads at all
- B. Linux uses the clone() system call to create threads as tasks that share the address space, treating threads and processes uniformly
- C. Linux requires all threads to run in user mode only
- D. Linux threads cannot share file descriptors

Answer: B

Q1002. What is the N:M threading model's primary implementation challenge that led most OSes to abandon it?

- A. It uses too much disk space
- B. Coordinating user-level and kernel-level schedulers is complex and the performance benefits rarely justify the complexity
- C. It cannot support more than two threads
- D. It is incompatible with virtual memory

Answer: B

Q1003. What is thread affinity and why does it matter for performance?

- A. Thread affinity refers to how friendly threads are to each other
- B. It binds a thread to a specific CPU core to exploit warm caches and reduce cache migration overhead
- C. Thread affinity means threads must always run on different cores
- D. It is a security feature that isolates threads

Answer: B

Q1004. What problem does the NUMA-aware thread scheduling try to solve?

- A. Reducing the number of threads in the system
- B. Minimizing memory access latency by scheduling threads on cores close to the memory they access
- C. Preventing threads from accessing shared memory
- D. Converting kernel threads to user threads

Answer: B

Q1005. What is cooperative multithreading and what is its main drawback?

- A. Threads are scheduled by the kernel preemptively; the drawback is high overhead
- B. Threads voluntarily yield the CPU; a misbehaving thread can monopolize the CPU indefinitely
- C. Threads share all variables including the stack; the drawback is data corruption
- D. Only one thread can exist per process; the drawback is no parallelism

Answer: B

Q1006. How do green threads differ from native threads?

- A. Green threads use less electricity
- B. Green threads are scheduled by the runtime or virtual machine in user space, not by the OS kernel
- C. Green threads are always faster than native threads
- D. Green threads can only run on green computing hardware

Answer: B

Q1007. What is the purpose of thread-specific data (TSD) and how does it differ from regular global variables?

- A. TSD is stored on disk for persistence
- B. TSD provides each thread with its own copy of a variable, unlike global variables which are shared among all threads
- C. TSD makes variables immutable across all threads
- D. TSD encrypts thread data for security

Answer: B

Q1008. What are coroutines and how do they relate to threads?

- A. Coroutines are a type of hardware thread
- B. Coroutines are cooperatively scheduled lightweight execution units that can suspend and resume at specific points, often multiplexed over fewer OS threads
- C. Coroutines and threads are identical concepts
- D. Coroutines can only execute sequentially with no concurrency

Answer: B

Q1009. What is the thundering herd problem in the context of thread pools?

- A. Too many threads are created, exhausting memory
- B. Multiple threads are woken up to handle a single event, but only one can process it, wasting resources
- C. Threads run too slowly due to excessive synchronization
- D. All threads in the pool terminate simultaneously

Answer: B

Q1010. What is work stealing in the context of thread pool task scheduling?

- A. A malicious technique where threads steal CPU time
- B. An idle thread takes tasks from the queue of a busy thread to balance the workload dynamically
- C. A method to prevent threads from executing tasks
- D. A technique that forces all threads to execute the same task

Answer: B

Q1011. How does the priority ceiling protocol prevent both deadlock and priority inversion?

- A. By removing all locks from the system
- B. By assigning each lock a priority ceiling equal to the highest priority of any thread that may lock it, and only allowing a thread to acquire a lock if its priority is higher than the ceiling of all currently locked mutexes
- C. By running all threads at the same priority
- D. By disabling interrupts permanently

Answer: B

Q1012. What is sequence consistency (sequential consistency) in the context of concurrent memory operations?

- A. All memory operations execute in parallel without ordering
- B. The result of any execution is the same as if the operations of all processors were executed in some sequential order, preserving each processor's program order
- C. Memory is accessed only sequentially by one processor at a time
- D. It means data is stored in sorted sequence

Answer: B

Q1013. What is the difference between lock-free and wait-free algorithms?

- A. They are identical concepts
- B. Lock-free guarantees system-wide progress (at least one thread progresses); wait-free guarantees per-thread progress (every thread completes in bounded steps)
- C. Wait-free algorithms use more locks than lock-free algorithms
- D. Lock-free means no synchronization is needed

Answer: B

Q1014. How does Read-Copy-Update (RCU) allow readers to access shared data without acquiring locks?

- A. By encrypting the data so readers can safely access it
- B. Readers access the current version without locks; writers create a new copy, update it, and atomically switch the pointer, deferring reclamation of old data until all readers finish
- C. Readers must always copy the data before reading
- D. RCU prevents any concurrent writes

Answer: B

Q1015. What is the ABA problem and how is it typically solved?

- A. It is a problem with alphabetical sorting in concurrent programs
- B. A thread reads value A, another thread changes it to B then back to A; the first thread's compare-and-swap succeeds incorrectly; solved using version counters or hazard pointers
- C. It occurs when three threads simultaneously access the same variable
- D. It is a deadlock pattern involving three resources

Answer: B

Q1016. What is a memory barrier (fence) and why is it necessary on modern processors?

- A. A physical barrier in RAM chips that prevents electrical interference
- B. An instruction that enforces ordering constraints on memory operations, preventing the CPU and compiler from reordering reads and writes across the barrier
- C. A software firewall for memory protection
- D. A method to increase memory bandwidth

Answer: B

Q1017. How does transactional memory simplify concurrent programming compared to lock-based approaches?

- A. It eliminates the need for any concurrency control
- B. Programmers designate code regions as transactions that execute atomically; conflicts are detected automatically, eliminating manual lock management and reducing deadlock risk
- C. It makes all memory accesses sequential
- D. It converts concurrent programs into single-threaded programs

Answer: B

Q1018. In the context of Linux kernel synchronization, what are per-CPU variables and why are they efficient?

- A. Variables that count the number of CPUs in the system
- B. Each CPU has its own copy of the variable, eliminating the need for locking since no sharing occurs between CPUs
- C. Variables that are shared equally among all CPUs
- D. Variables that store CPU temperature readings

Answer: B

Q1019. What is the difference between optimistic and pessimistic concurrency control?

- A. Optimistic control is always faster
- B. Pessimistic control acquires locks before accessing data; optimistic control proceeds without locks and validates at commit time, aborting on conflict
- C. They produce different results for the same operations
- D. Optimistic control does not support concurrent access

Answer: B

Q1020. What is the sleeping barber problem and what synchronization challenge does it illustrate?

- A. A problem about scheduling CPU sleep states
- B. A barber sleeps when no customers are present and customers must wake the barber or wait; it illustrates coordination between producers and consumers with a bounded waiting area
- C. A deadlock scenario involving three barber threads
- D. A memory management problem involving page sleeping

Answer: B

Q1021. What is the wait-die scheme for deadlock prevention in database systems?

- A. Processes wait indefinitely until resources become available
- B. An older process waits for a younger one to release a resource, but a younger process requesting a resource held by an older one is aborted (dies)
- C. All processes die when a deadlock is detected
- D. Processes are assigned random wait times

Answer: B

Q1022. How does the edge-chasing algorithm detect deadlocks in distributed systems?

- A. By checking all resource allocation graphs simultaneously
- B. Probe messages are forwarded along edges of the distributed wait-for graph; if a probe returns to its initiator, a cycle (deadlock) exists
- C. By centralizing all resource information on one node
- D. By terminating all processes on all nodes periodically

Answer: B

Q1023. What is a phantom deadlock in distributed deadlock detection?

- A. A deadlock that occurs in virtual machines only
- B. A falsely detected deadlock caused by delays in propagating status information across nodes in a distributed system
- C. A deadlock involving phantom processes that do not exist
- D. A deadlock that occurs during system startup

Answer: B

Q1024. What is the relationship between deadlock and livelock in resource contention?

- A. Livelock is a more severe form of deadlock
- B. In deadlock, processes are blocked and make no progress; in livelock, processes continuously change state in response to each other but still make no progress
- C. Livelock is identical to deadlock
- D. Livelock occurs only with non-preemptive scheduling

Answer: B

Q1025. How does lock ordering (hierarchical locking) prevent deadlock in practice?

- A. By requiring all locks to be held for the same duration
- B. By defining a global order on all locks and requiring every thread to acquire locks only in that order, preventing cycles in the wait-for graph
- C. By limiting the number of locks in the system to one
- D. By making all locks recursive

Answer: B

Q1026. What is the knot in a resource allocation graph?

- A. A cycle involving exactly two processes
- B. A set of nodes such that every node in the set has a path to every other node in the set with no outgoing edges leaving the set
- C. A single node with no edges
- D. An edge that connects a resource to itself

Answer: B

Q1027. What is two-phase locking and how does it relate to deadlock in database systems?

- A. It prevents deadlock by ensuring locks are acquired in two distinct phases
- B. In the growing phase all needed locks are acquired and in the shrinking phase locks are released; it ensures serializability but does not prevent deadlock
- C. It acquires and releases each lock twice for safety
- D. It uses two separate lock managers to prevent deadlock

Answer: B

Q1028. How does the Linux kernel detect potential deadlocks using lockdep?

- A. By running the Banker's Algorithm at each lock acquisition
- B. By tracking lock acquisition order at runtime and checking for potential cycles in the lock dependency graph, even if the actual deadlock hasn't occurred
- C. By limiting the number of locks to 10
- D. By preventing any two threads from running simultaneously

Answer: B

Q1029. What is the convoy effect in the context of resource scheduling, and how does it relate to deadlock-like behavior?

- A. A high-speed scheduling algorithm
- B. Many processes queue behind a slow process holding a resource, causing a chain of delays similar to but distinct from deadlock
- C. All processes execute in parallel without any delays
- D. Processes are permanently blocked and never execute

Answer: B

Q1030. What is the difference between centralized and distributed deadlock detection, and what are the trade-offs?

- A. They are identical approaches applied to different operating systems
- B. Centralized detection maintains a global wait-for graph at one coordinator (simpler but single point of failure); distributed detection uses local algorithms with message passing (more resilient but may detect phantom deadlocks)
- C. Centralized detection is always faster and more accurate
- D. Distributed detection eliminates all false positives

Answer: B

Q1031. How does the clock (second-chance) algorithm approximate LRU replacement?

- A. It uses a hardware timer to track page usage
- B. Pages are arranged in a circular list with a reference bit; the algorithm scans for a page with reference bit 0, clearing bits as it goes, giving referenced pages a second chance
- C. It randomly selects a page for replacement
- D. It replaces the page closest to the clock's 12 o'clock position

Answer: B

Q1032. What is copy-on-write (COW) optimization in fork() and how does it work?

- A. It copies all pages immediately when fork() is called
- B. Parent and child initially share the same physical pages marked read-only; a page is copied only when either process writes to it, reducing unnecessary copying
- C. It prevents child processes from writing to any memory
- D. It copies memory from disk to RAM on write operations

Answer: B

Q1033. What is Address Space Layout Randomization (ASLR) and how does it improve security?

- A. It sorts addresses in alphabetical order
- B. It randomizes the memory layout of a process (stack, heap, libraries) at each execution, making buffer overflow exploits harder by unpredictable addresses
- C. It adds encryption to all memory addresses
- D. It assigns the same layout to all processes for consistency

Answer: B

Q1034. How does the buddy system allocator handle memory fragmentation?

- A. It eliminates fragmentation completely
- B. It splits memory blocks into power-of-two sized buddies and merges adjacent free buddies back together, reducing external fragmentation but causing internal fragmentation
- C. It uses a linked list of arbitrary block sizes
- D. It allocates memory in a single contiguous block

Answer: B

Q1035. What is the slab allocator and why is it efficient for kernel memory allocation?

- A. A disk allocation algorithm for solid-state drives
- B. It pre-allocates caches of objects of specific sizes from contiguous memory slabs, eliminating fragmentation and initialization overhead for frequently allocated kernel objects
- C. A memory compaction technique for user programs
- D. A page replacement algorithm for virtual memory

Answer: B

Q1036. What is page coloring and how does it improve cache performance?

- A. It assigns colors to pages for visual debugging
- B. It assigns physical frames to virtual pages such that pages mapping to the same cache set are distributed across frames, reducing cache conflicts
- C. It marks pages as read-only using color codes
- D. It groups pages by their access frequency

Answer: B

Q1037. How do huge pages (large pages) improve TLB performance?

- A. They reduce the size of each page
- B. Fewer TLB entries cover the same amount of memory, reducing TLB misses for applications with large memory footprints
- C. They increase the number of TLB entries available
- D. They bypass the TLB entirely

Answer: B

Q1038. What is memory-mapped file I/O and what advantage does it offer?

- A. A technique that stores files in CPU registers
- B. It maps a file's contents into a process's virtual address space, allowing file access through memory operations instead of read/write system calls, reducing system call overhead
- C. It encrypts files during I/O operations
- D. It stores all file metadata in physical memory

Answer: B

Q1039. What is the page-fault frequency (PFF) approach to thrashing prevention?

- A. It sets a minimum page fault rate for each process
- B. It monitors each process's page fault rate; if too high, more frames are allocated; if too low, frames are reclaimed, dynamically adjusting allocation
- C. It prevents all page faults by loading every page into memory
- D. It increases the swap space size when faults are frequent

Answer: B

Q1040. What is Kernel Same-page Merging (KSM) and where is it most beneficial?

- A. A technique to merge kernel code pages with user code pages
- B. It scans memory for identical pages across processes and merges them into a single shared copy (with COW), most beneficial in virtualization where VMs run similar OSes
- C. A method to compress kernel memory pages
- D. A page replacement algorithm for kernel pages

Answer: B

Q1041. How does the log-structured file system (LFS) approach differ from traditional file systems?

- A. LFS stores all files in a single log file
- B. LFS writes all changes (data and metadata) sequentially to a log, converting random writes to sequential writes for better write performance, using a cleaner for garbage collection
- C. LFS reads files from the end to the beginning
- D. LFS does not support directories

Answer: B

Q1042. How does ZFS ensure end-to-end data integrity?

- A. By using parity bits only
- B. By storing checksums in parent blocks (not with the data block itself), detecting silent data corruption, bit rot, and phantom writes
- C. By encrypting all data with AES
- D. By writing data twice to the same location

Answer: B

Q1043. What is FUSE and why is it significant for file system development?

- A. A hardware component for fast USB storage
- B. A framework that allows implementing file systems in user space without modifying the kernel, simplifying development and enabling novel file systems
- C. A file compression algorithm
- D. A network protocol for distributed file systems

Answer: B

Q1044. How do extents improve upon traditional block pointers in file systems like ext4?

- A. Extents store data in smaller units for precision
- B. An extent describes a contiguous range of blocks with a single start-block and length, dramatically reducing metadata overhead for large files
- C. Extents eliminate the need for inodes
- D. Extents only work with SSDs

Answer: B

Q1045. What is a copy-on-write file system and what advantages does it offer?

- A. A file system that copies files on every read operation
- B. Instead of overwriting data in place, new data is written to a new location and metadata is updated atomically, enabling snapshots, preventing partial writes, and maintaining consistency
- C. A file system that prevents all write operations
- D. A file system that duplicates the entire disk on each write

Answer: B

Q1046. How does the B+ tree structure benefit NTFS and other modern file systems?

- A. B+ trees are slower but use less memory
- B. B+ trees provide $O(\log n)$ lookup, insertion, and deletion for directory entries and file extents, with all data in leaf nodes for efficient range queries
- C. B+ trees only store file names, not their locations
- D. B+ trees are only used for binary files

Answer: B

Q1047. What is the TRIM command and why is it essential for SSD-backed file systems?

- A. It compresses files to save SSD space
- B. It informs the SSD which blocks are no longer in use, allowing the SSD's FTL to erase and reclaim those blocks proactively for better write performance
- C. It trims file names to save directory space
- D. It limits the number of files per directory

Answer: B

Q1048. What is deduplication in enterprise file systems and what are its trade-offs?

- A. It duplicates data for redundancy
- B. It identifies and eliminates duplicate data blocks, storing only unique blocks with references, saving storage but adding CPU and memory overhead for hash computation
- C. It removes all files older than 30 days
- D. It deletes duplicate file names only

Answer: B

Q1049. How does the Network File System (NFS) handle stateless server design and what are the implications?

- A. NFS requires the server to maintain all client state
- B. The NFS server does not track which clients have opened files; each request is self-contained, simplifying crash recovery but requiring idempotent operations
- C. NFS does not support any crash recovery
- D. NFS stores all state on the client side only

Answer: B

Q1050. What is the purpose of file system snapshots and how are they efficiently implemented in COW file systems?

- A. Snapshots are photographs of the server hardware
- B. Snapshots capture the file system state at a point in time; in COW file systems, they are nearly instant because only the root metadata block is copied, sharing unchanged blocks with the live file system
- C. Snapshots require copying the entire disk
- D. Snapshots delete all files created after a specific date

Answer: B

Q1051. What is write amplification in SSDs and why is it a problem?

- A. Writing louder signals to flash cells for reliability
- B. The actual amount of data written to flash exceeds the logical data written by the host due to garbage collection and wear leveling, reducing SSD lifespan and performance
- C. Writing data to multiple SSDs simultaneously
- D. Increasing the write speed beyond the SSD's rated capacity

Answer: B

Q1052. What is the Flash Translation Layer (FTL) in SSDs?

- A. A physical layer of insulation in the SSD
- B. Firmware that maps logical block addresses to physical flash locations, managing wear leveling, garbage collection, and bad block handling
- C. A network protocol for flash storage
- D. A file system designed specifically for flash memory

Answer: B

Q1053. How does NVMe improve upon the AHCI interface for SSDs?

- A. NVMe uses serial communication while AHCI uses parallel
- B. NVMe supports up to 65,535 I/O queues with 65,536 entries each and uses direct PCIe connection, dramatically reducing latency and increasing throughput compared to AHCI's single queue of 32 entries
- C. NVMe is designed for hard disk drives only
- D. AHCI is faster than NVMe for random access

Answer: B

Q1054. What is shingled magnetic recording (SMR) in HDDs?

- A. A technique using shingle-shaped disk platters
- B. Overlapping write tracks like roof shingles to increase data density, but requiring sequential writes within zones since writing a track partially overwrites adjacent tracks
- C. A method of reading data from multiple tracks simultaneously
- D. A technique to reduce disk noise

Answer: B

Q1055. What is NCQ (Native Command Queuing) in SATA drives?

- A. A method to queue network commands for remote drives
- B. A feature allowing the drive to internally reorder queued I/O commands to minimize seek and rotational delays, improving throughput
- C. A technique that queues commands in the OS driver only
- D. A command that queries the native format of the disk

Answer: B

Q1056. How does RAID 6 improve fault tolerance over RAID 5?

- A. RAID 6 uses triple mirroring
- B. RAID 6 uses two independent parity calculations (P and Q), tolerating two simultaneous disk failures instead of one
- C. RAID 6 has no parity overhead
- D. RAID 6 stripes data without any redundancy

Answer: B

Q1057. What is the NOOP (No Operation) I/O scheduler and when is it appropriate?

- A. It performs no I/O operations at all
- B. It implements a simple FIFO queue with request merging only, appropriate for SSDs and virtualized environments where the device or hypervisor handles scheduling
- C. It is a disk scheduling algorithm for maximum throughput on HDDs
- D. It disables the disk entirely

Answer: B

Q1058. What are Zoned Namespaces (ZNS) in NVMe SSDs and how do they differ from conventional SSDs?

- A. ZNS SSDs have different geographic zones for data storage
- B. ZNS exposes the SSD's internal zone structure to the host, requiring sequential writes within zones but eliminating the need for the FTL's garbage collection, reducing write amplification
- C. ZNS is a security zoning feature for data isolation
- D. ZNS SSDs have larger physical size than conventional SSDs

Answer: B

Q1059. What is the Advanced Format (4K sector) standard and why was it adopted?

- A. A format that stores data in 4KB RAM pages
- B. It increased the physical sector size from 512 bytes to 4096 bytes, improving error correction efficiency, format efficiency, and storage density on modern drives
- C. A new formatting standard that only applies to SSDs
- D. A technique that formats disks four times faster

Answer: B

Q1060. What is scatter-gather DMA and how does it improve I/O performance?

- A. It scatters data across multiple disks randomly
- B. It allows DMA to transfer data between non-contiguous memory regions and a device in a single operation using a list of memory descriptors, avoiding extra copy operations
- C. It gathers interrupts from multiple devices into one
- D. It distributes I/O requests across multiple CPUs

Answer: B

Q1061. How does Linux epoll improve over select() and poll() for I/O multiplexing?

- A. epoll uses more memory than select()
- B. epoll uses an event-driven model that only reports changes, avoids scanning all file descriptors, and scales O(1) with the number of monitored descriptors
- C. epoll works only with network sockets
- D. select() is always faster than epoll

Answer: B

Q1062. What is the IOMMU and how does it enhance I/O security and functionality?

- A. A memory unit for storing I/O device firmware
- B. Hardware that translates I/O device virtual addresses to physical addresses, preventing devices from accessing unauthorized memory regions and enabling DMA for virtualized guests
- C. A specialized I/O processor that replaces the CPU
- D. A memory management unit within the I/O device itself

Answer: B

Q1063. What is the difference between edge-triggered and level-triggered interrupt notifications?

- A. Edge-triggered uses higher voltage levels
- B. Edge-triggered fires on the transition (change) of the signal; level-triggered fires as long as the signal is asserted, requiring different handling strategies to avoid missed or spurious interrupts
- C. Level-triggered is always more efficient
- D. Edge-triggered is only used for software interrupts

Answer: B

Q1064. What is io_uring in modern Linux and why is it significant?

- A. A circular buffer for storing I/O errors
- B. An asynchronous I/O interface using shared ring buffers between user space and kernel, minimizing system call overhead and enabling high-performance I/O without context switches
- C. A debugging tool for I/O operations
- D. A type of network ring topology

Answer: B

Q1065. What is zero-copy I/O and how does sendfile() implement it?

- A. I/O where no data is actually transferred
- B. Data is transferred from disk to network socket directly through kernel buffers without being copied to user space, reducing CPU usage and memory bandwidth consumption
- C. A technique that compresses data to zero bytes during transfer
- D. I/O that creates zero backup copies

Answer: B

Q1066. What is interrupt coalescing and why is it used in high-speed network adapters?

- A. Combining multiple interrupt controllers into one
- B. Delaying and batching multiple interrupts into fewer interrupt notifications to reduce per-packet interrupt overhead at high data rates
- C. Converting interrupts into polling operations
- D. Encrypting interrupt signals for security

Answer: B

Q1067. What is RDMA (Remote Direct Memory Access) and how does it bypass the traditional I/O stack?

- A. A technique to access RAM on a remote machine through SSH
- B. It allows one computer to directly read from or write to another computer's memory without involving either CPU or OS, bypassing the kernel network stack entirely
- C. A remote debugging protocol for memory analysis
- D. A technique to access disk drives over the network

Answer: B

Q1068. What is the SPDK (Storage Performance Development Kit) approach to storage I/O?

- A. A kernel module for legacy storage devices
- B. It moves storage drivers into user space, uses polled mode instead of interrupts, and provides lockless designs to eliminate kernel overhead and achieve near-hardware-level NVMe performance
- C. A standard development kit for building storage hardware
- D. A graphical tool for monitoring storage performance

Answer: B

Q1069. How does virtio enable efficient I/O for virtual machines?

- A. By providing dedicated physical hardware to each VM
- B. It defines a standardized interface between VMs and the hypervisor using shared memory ring buffers (virtqueues), avoiding the overhead of emulating real hardware devices
- C. By running all VMs on the same physical I/O device
- D. By disabling I/O for virtual machines entirely

Answer: B

Q1070. How does return-oriented programming (ROP) bypass Data Execution Prevention (DEP/NX)?

- A. By disabling the CPU's NX bit
- B. By chaining together existing code sequences (gadgets) ending in return instructions from legitimate libraries, executing arbitrary computations without injecting new code
- C. By encrypting the malicious code so DEP cannot detect it
- D. By running code in kernel mode where DEP does not apply

Answer: B

Q1071. What is the Bell-LaPadula security model and what property does it enforce?

- A. A model that allows all users to read and write all data
- B. A mandatory access control model enforcing 'no read up, no write down' to prevent information flow from higher to lower security levels
- C. A model based on role-based access control
- D. A model that enforces data integrity but not confidentiality

Answer: B

Q1072. How does SELinux implement mandatory access control differently from traditional Unix permissions?

- A. SELinux replaces all file permissions with encryption
- B. SELinux assigns security labels to all subjects and objects and enforces policy rules for every access decision, regardless of traditional DAC permissions
- C. SELinux only adds network security controls
- D. SELinux uses the same permission model as standard Unix

Answer: B

Q1073. What is the confused deputy problem in computer security?

- A. When a system administrator is confused about access policies
- B. When a privileged program is tricked by a less-privileged entity into misusing its authority to perform unauthorized actions on the attacker's behalf
- C. When two processes have the same security level and confuse each other
- D. When encryption keys are mixed up between users

Answer: B

Q1074. What is control flow integrity (CFI) and how does it protect against code-reuse attacks?

- A. It controls the flow of network traffic
- B. It restricts the program's control flow transfers to only those targets determined at compile time, preventing ROP and JOP attacks from diverting execution
- C. It encrypts the program's control flow graph
- D. It monitors data flow but not control flow

Answer: B

Q1075. What is the Trusted Platform Module (TPM) and what security functions does it provide?

- A. A software-only security module for web browsers
- B. A dedicated hardware chip that provides secure key generation, storage, platform integrity measurement, and remote attestation capabilities
- C. A removable storage device for encrypted backups
- D. A network firewall appliance

Answer: B

Q1076. How does Memory Tagging Extension (MTE) help prevent memory safety vulnerabilities?

- A. It adds labels to memory for human readability
- B. It associates a 4-bit tag with each memory granule and each pointer; access is only permitted when the pointer's tag matches the memory's tag, detecting use-after-free and buffer overflows
- C. It encrypts all memory with unique tags
- D. It compresses memory to prevent overflow

Answer: B

Q1077. What is a zero-day vulnerability and why is it particularly dangerous?

- A. A vulnerability that takes zero days to fix
- B. A previously unknown vulnerability being exploited before the vendor has released a patch, leaving all systems defenseless
- C. A vulnerability that only exists for one day
- D. A vulnerability in programs that are zero days old

Answer: B

Q1078. What is attestation in the context of trusted computing?

- A. A user testifying in court about a security breach
- B. A process by which a system cryptographically proves its hardware and software configuration to a remote verifier, establishing trust in its integrity
- C. A method of authenticating users via attestation letters
- D. A technique for attesting to the speed of encryption

Answer: B

Q1079. How does the Biba integrity model differ from the Bell-LaPadula model?

- A. The Biba model is a more recent version of Bell-LaPadula
- B. Bell-LaPadula focuses on confidentiality (no read up, no write down); Biba focuses on integrity (no read down, no write up), preventing corruption of higher-integrity data
- C. Biba and Bell-LaPadula are identical models
- D. Biba allows unrestricted reading and writing

Answer: B

Q1080. What is over-provisioning in SSDs and why is it necessary?

- A. Providing more SSDs than needed as backups
- B. Reserving a percentage of flash capacity as spare blocks for garbage collection, wear leveling, and bad block replacement, improving sustained write performance and endurance
- C. Writing more data than the SSD can hold
- D. Providing extra power supply to the SSD

Answer: B

SUBJECT

Problem Solving & Analytical Skills

1,440 Multiple Choice Questions

Problem Solving & Analytical Skills — Easy

480 questions

Q1. What is the first step in problem solving?

- A. Testing the solution for correctness
- B. Understanding the problem thoroughly
- C. Deploying the application to production
- D. Writing code to implement a solution

Answer: B

Q2. Which best describes problem solving in CS?

- A. Memorizing algorithms from a textbook
- B. A systematic approach to finding solutions
- C. Copying solutions from existing projects
- D. Writing programs without any planning

Answer: B

Q3. What does decomposition mean in problem solving?

- A. Making a problem increasingly more complex
- B. Ignoring certain parts of the problem
- C. Combining multiple problems into one task
- D. Breaking a problem into smaller sub-problems

Answer: D

Q4. Which skill is most essential for a problem solver?

- A. Graphic design expertise
- B. Data entry proficiency
- C. Analytical thinking skills
- D. Speed typing on a keyboard

Answer: C

Q5. What is an algorithm?

- A. A database management system design
- B. A step-by-step procedure to solve a problem
- C. A specific programming language syntax
- D. A type of computer hardware component

Answer: B

Q6. Which is NOT a phase of problem solving?

- A. Devising a plan for the solution
- B. Procrastinating and delaying action
- C. Looking back and evaluating results
- D. Understanding the problem clearly

Answer: B

Q7. What is the purpose of brainstorming?

- A. To eliminate all ideas immediately
- B. To finalize the code right away
- C. To test the program for errors
- D. To generate multiple possible solutions

Answer: D

Q8. Who introduced the four-step problem solving process?

- A. George Polya
- B. Alan Turing
- C. Donald Knuth
- D. Ada Lovelace

Answer: A

Q9. What is a heuristic?

- A. A specific programming language tool
- B. A guaranteed optimal solution method
- C. A type of programming syntax error
- D. A practical approach that is sufficient

Answer: D

Q10. Which approach solves a simpler version first?

- A. Randomization approach
- B. Elimination technique
- C. Brute force method
- D. Simplification strategy

Answer: D

Q11. What should you first identify when analyzing a problem?

- A. The user interface color scheme
- B. The input and expected output
- C. The programming language to use
- D. The database engine to deploy

Answer: B

Q12. What is a problem statement?

- A. A set of test cases for validation
- B. The code implementation of the fix
- C. A clear description of what to solve
- D. The final solution to the problem

Answer: C

Q13. Which helps understand a problem better?

- A. Identifying examples and edge cases
- B. Skipping requirements gathering
- C. Ignoring constraints on the system
- D. Immediately coding a solution

Answer: A

Q14. What are requirements?

- A. Conditions the solution must satisfy
- B. Bugs found during the testing phase
- C. Test results from the final review
- D. Optional features to consider later

Answer: A

Q15. What is an edge case?

- A. Unusual input testing boundary limits
- B. The most common input scenario
- C. A syntax error in the source code
- D. A design pattern used in systems

Answer: A

Q16. Why identify constraints?

- A. They serve only as documentation notes
- B. Constraints are completely irrelevant
- C. They define limits the solution must meet
- D. They make the problem harder on purpose

Answer: C

Q17. What does scope mean?

- A. The boundaries of what the problem covers
- B. The programming language being used
- C. The speed of the algorithm execution
- D. The database size for the project

Answer: A

Q18. Most useful analysis question?

- A. What color theme should we apply
- B. What font should the UI display
- C. What are inputs and expected output
- D. How many lines of code are needed

Answer: C

Q19. What is a test case?

- A. A programming paradigm to follow
- B. A particular variable type definition
- C. A debugging tool in the compiler
- D. Scenario with specific input and output

Answer: D

Q20. What are functional requirements?

- A. Appearance and visual requirements
- B. Hardware specifications for the project
- C. Network requirements for deployment
- D. What the system should actually do

Answer: D

Q21. A=true, B=true. A AND B?

- A. True
- B. False
- C. Null
- D. Undefined

Answer: A

Q22. NOT True?

- A. Error
- B. True
- C. False
- D. Null

Answer: C

Q23. A=true, B=false. A OR B?

- A. True
- B. Neither
- C. Undefined
- D. False

Answer: A

Q24. General to specific reasoning?

- A. Abductive
- B. Inductive
- C. Analogical
- D. Deductive

Answer: D

Q25. Specific to general reasoning?

- A. Deductive
- B. Abductive
- C. Inductive
- D. Circular

Answer: C

Q26. Truth table with 2 variables has how many rows?

- A. 16
- B. 4
- C. 2
- D. 8

Answer: B

Q27. What is a logical fallacy?

- A. A correct logical argument
- B. An error making argument invalid
- C. A type of search algorithm
- D. A specific data structure

Answer: B

Q28. All cats are animals. Tom is a cat. Therefore?

- A. Tom is actually a dog
- B. Nothing can be concluded
- C. Tom is not an animal
- D. Tom is an animal too

Answer: D

Q29. What does XOR return?

- A. Always returns true as output
- B. True when exactly one is true
- C. Always returns false as output
- D. True when both inputs are true

Answer: B

Q30. What is a premise?

- A. A conclusion drawn from evidence
- B. A statement assumed true for reasoning
- C. A question posed to the audience
- D. An error found in the argument

Answer: B

Q31. What is a flowchart for?

- A. Creating user interface layouts
- B. Designing database table schemas
- C. Visually representing algorithm steps
- D. Writing executable source code

Answer: C

Q32. Which symbol is a decision?

- A. Diamond
- B. Parallelogram
- C. Rectangle
- D. Oval

Answer: A

Q33. Which symbol is a process?

- A. Arrow
- B. Oval
- C. Diamond
- D. Rectangle

Answer: D

Q34. What does an oval represent?

- A. Process
- B. Decision
- C. Start or End
- D. Input/Output

Answer: C

Q35. What is pseudocode?

- A. A specific programming language
- B. Informal high-level algorithm description
- C. Machine-readable binary code
- D. Directly executable source code

Answer: B

Q36. Three basic control structures?

- A. Start, Process, End
- B. Create, Read, Delete
- C. Sequence, Selection, Iteration
- D. Input, Output, Storage

Answer: C

Q37. What is sequence?

- A. Skipping steps entirely
- B. Repeating a set of steps
- C. Steps in random order
- D. Steps one after another in order

Answer: D

Q38. What does selection do?

- A. Skips all code entirely
- B. Executes all paths at once
- C. Chooses paths based on condition
- D. Repeats code multiple times

Answer: C

Q39. What does iteration do?

- A. Stops the program entirely
- B. Deletes all variable values
- C. Executes code only once
- D. Repeats while condition is true

Answer: D

Q40. Which symbol is input/output?

- A. Circle
- B. Parallelogram
- C. Rectangle
- D. Diamond

Answer: B

Q41. Binary of decimal 5?

- A. 110
- B. 111
- C. 101
- D. 100

Answer: C

Q42. What is abstraction in data?

- A. Making things increasingly complex
- B. Showing all implementation details
- C. Deleting all the stored data
- D. Hiding unnecessary, showing essential

Answer: D

Q43. Bits in a byte?

- A. 4
- B. 16
- C. 32
- D. 8

Answer: D

Q44. Data type for name?

- A. String
- B. Boolean
- C. Float
- D. Integer

Answer: A

Q45. What is an array?

- A. A specific type of loop construct
- B. A callable function definition
- C. A single variable only
- D. Same-type elements in contiguous memory

Answer: D

Q46. What does ASCII stand for?

- A. Automated Standard Code for Input Interfacing
- B. American Standard Code for Information Interchange
- C. American System for Computer Information Indexing
- D. Advanced System Code for Internet Integration

Answer: B

Q47. Decimal of binary 1010?

- A. 12
- B. 10
- C. 8
- D. 14

Answer: B

Q48. Which uses key-value pairs?

- A. Dictionary/Map
- B. Array structure
- C. Queue structure
- D. Stack structure

Answer: A

Q49. What is Boolean?

- A. String text type
- B. Decimal number type
- C. Only true or false
- D. Float number type

Answer: C

Q50. What is hexadecimal?

- A. Base-8 system
- B. Base-2 system
- C. Base-10 system
- D. Base-16 system

Answer: D

Q51. What is pattern recognition?

- A. Drawing visual pattern designs
- B. Copying solutions from others
- C. Identifying similarities and trends
- D. Guessing the answer randomly

Answer: C

Q52. Next in 2,4,6,8,?

- A. 9
- B. 10
- C. 12
- D. 14

Answer: B

Q53. Next in 1,2,4,8,?

- A. 16
- B. 32
- C. 12
- D. 10

Answer: A

Q54. What is generalization?

- A. Ignoring all recognized patterns
- B. Applying pattern to broader problems
- C. Making things more specific only
- D. Focusing on one instance alone

Answer: B

Q55. All sorting uses comparison identifies?

- A. A bug in the code
- B. A syntax error in logic
- C. A specific data type issue
- D. A common shared pattern

Answer: D

Q56. What is an analogy?

- A. A programming concept for loops
- B. Similarities between problems found
- C. A random guess at the answer
- D. An error type in the program

Answer: B

Q57. Next letter: A,C,E,G,?

- A. K
- B. J
- C. H
- D. I

Answer: D

Q58. What is Fibonacci?

- A. A random sequence with no pattern
- B. The sequence of even numbers 2,4,6,8
- C. Each equals sum of two preceding values
- D. The sequence of odd numbers 1,3,5,7

Answer: C

Q59. Which demonstrates pattern recognition?

- A. Ignoring existing working solutions
- B. Never reusing any solution pattern
- C. Noticing all searches scan collections
- D. Coding randomly without a plan

Answer: C

Q60. Why pattern recognition useful?

- A. It always produces longer code
- B. Identifies reusable solutions quickly
- C. It is not useful at all in CS
- D. It always introduces more bugs

Answer: B

Q61. 2^{10} ?

- A. 1024
- B. 2048
- C. 256
- D. 512

Answer: A

Q62. $5!$?

- A. 25
- B. 720
- C. 60
- D. 120

Answer: D

Q63. $17 \% 5$?

- A. 3
- B. 5
- C. 4
- D. 2

Answer: D

Q64. $\log_2(8)$?

- A. 4
- B. 8
- C. 3
- D. 2

Answer: C

Q65. Sum of first 10 naturals?

- A. 50
- B. 55
- C. 45
- D. 100

Answer: B

Q66. $C(n,2)$ represents?

- A. The product n times 2
- B. Ways to choose 2 from n
- C. The quotient n divided by 2
- D. The sum of n plus 2

Answer: B

Q67. $GCD(12,8)$?

- A. 6
- B. 2
- C. 4
- D. 8

Answer: C

Q68. $O(n)$, n from 10 to 100 scales?

- A. Stays the same
- B. About 100x slower
- C. About 10x slower
- D. About 2x slower

Answer: C

Q69. What is a prime?

- A. Any odd number in the range
- B. Divisible by any number at all
- C. Only 1 and itself as divisors
- D. Any even number in the range

Answer: C

Q70. $LCM(4,6)$?

- A. 2
- B. 10
- C. 24
- D. 12

Answer: D

Q71. What is algorithmic thinking?

- A. Memorizing algorithms from textbooks
- B. Avoiding any structured approach
- C. Defining clear steps to solve problems
- D. Thinking about hardware components

Answer: C

Q72. Simplest search?

- A. Hash lookup
- B. Jump search
- C. Linear search
- D. Binary search

Answer: C

Q73. Binary search requires?

- A. Only integer type values
- B. An odd length data array
- C. A sorted list of elements
- D. An empty list as input

Answer: C

Q74. Bubble sort idea?

- A. Insert each into a sorted tree
- B. Select the minimum each time
- C. Swap adjacent wrong-order elements
- D. Divide the array in half first

Answer: C

Q75. What is sorting?

- A. Copying data to a new place
- B. Deleting unwanted elements
- C. Searching for an element
- D. Arranging in a specific order

Answer: D

Q76. Linear search worst case?

- A. $O(n)$
- B. $O(1)$
- C. $O(\log n)$
- D. $O(n^2)$

Answer: A

Q77. What is greedy?

- A. Best choice at each step locally
- B. Always using recursive functions
- C. Using lots of memory resources
- D. Running very slowly on purpose

Answer: A

Q78. What is recursive algorithm?

- A. Calls itself with smaller instances
- B. An algorithm without any loops
- C. An algorithm that runs only once
- D. An algorithm that never repeats

Answer: A

Q79. BFS uses?

- A. Array structure
- B. Stack structure
- C. Hash table
- D. Queue structure

Answer: D

Q80. DFS uses?

- A. Heap structure
- B. Hash table lookup
- C. Queue structure
- D. Stack or recursion

Answer: D

Q81. What is critical thinking?

- A. Criticizing everything you see
- B. Objective analysis to form judgments
- C. Accepting all ideas without thought
- D. Ignoring evidence and data given

Answer: B

Q82. Why important in development?

- A. It is not important at all
- B. Only managers need this skill
- C. It slows everything down overall
- D. Evaluates solutions, identifies flaws

Answer: D

Q83. What is bias?

- A. A type of search algorithm
- B. A programming language name
- C. A systematic thinking error
- D. A specific data structure

Answer: C

Q84. Evidence-based decisions?

- A. Decisions made at random
- B. Blindly following authority
- C. Decisions from verified facts
- D. Decisions based on feelings

Answer: C

Q85. What is confirmation bias?

- A. Considering all viewpoints equally
- B. Seeking confirming evidence only
- C. Being fully objective about data
- D. Testing all hypotheses fairly

Answer: B

Q86. Role of questioning?

- A. It is entirely unnecessary
- B. It slows down the process
- C. It shows ignorance in the team
- D. Challenges assumptions, deepens understanding

Answer: D

Q87. What is a trade-off?

- A. Avoiding all decisions entirely
- B. Having no options to choose from
- C. Getting everything for free
- D. Sacrificing one quality for another

Answer: D

Q88. What is objectivity?

- A. Following the popular opinion only
- B. Evaluating without personal bias
- C. Relying on personal feelings only
- D. Ignoring all facts and evidence

Answer: B

Q89. What is an assumption?

- A. A proven and verified fact
- B. A specific test case to run
- C. An algorithm to implement now
- D. Taken for granted without proof

Answer: D

Q90. Evaluating alternatives?

- A. Choose one completely at random
- B. Always pick the first option
- C. Compare against criteria to select
- D. Ignore all the available options

Answer: C

Q91. What is debugging?

- A. Finding and fixing code errors
- B. Deleting all program files now
- C. Writing new source code files
- D. Compiling the program to binary

Answer: A

Q92. What is syntax error?

- A. A runtime exception in program
- B. A logic mistake in the code
- C. A violation of grammar rules
- D. A design flaw in architecture

Answer: C

Q93. What is runtime error?

- A. Error during code writing phase
- B. Error during code compilation
- C. Error during program execution
- D. Error in the syntax of code

Answer: C

Q94. What is logical error?

- A. A syntax mistake in the code
- B. A missing semicolon in the code
- C. An error during code compilation
- D. Runs fine but gives wrong results

Answer: D

Q95. What is a breakpoint?

- A. A syntax rule for statements
- B. A type of error in the code
- C. Execution pauses for inspection
- D. A point where code is broken

Answer: C

Q96. First debugging step?

- A. Reproduce the error consistently
- B. Rewrite all the code from scratch
- C. Ignore the error and move on
- D. Delete the problematic file now

Answer: A

Q97. What is print debugging?

- A. Printing the error log to a file
- B. Print statements to trace values
- C. Printing the source code on paper
- D. Printing the algorithm description

Answer: B

Q98. What is a stack trace?

- A. Function call sequence at error
- B. A variable type declaration
- C. A stack data structure diagram
- D. A sorting algorithm for stacks

Answer: A

Q99. What is an exception?

- A. A variable declaration statement
- B. A normal expected program flow
- C. Event disrupting flow, indicating error
- D. A design pattern for structure

Answer: C

Q100. Compiler errors identify?

- A. Syntax and type error locations
- B. Runtime errors only in code
- C. Performance bottlenecks found
- D. Logic errors in the algorithm

Answer: A

Q101. What does Big-O describe?

- A. The exact execution time needed
- B. The minimum time for the task
- C. The average case performance
- D. Upper bound of the growth rate

Answer: D

Q102. What is $O(1)$?

- A. Linear time growth
- B. Constant time always
- C. Quadratic time growth
- D. Logarithmic time growth

Answer: B

Q103. What is $O(n)$?

- A. Linear, proportional to input
- B. Quadratic time growth
- C. Constant time always
- D. Exponential time growth

Answer: A

Q104. $O(n)$ vs $O(n^2)$ for large inputs?

- A. They perform exactly equal
- B. $O(n^2)$ is faster overall
- C. $O(n)$ is faster overall
- D. Cannot tell the difference

Answer: C

Q105. Space complexity measures?

- A. Memory relative to input size
- B. Disk storage used by files
- C. Lines of code in the program
- D. Screen pixels for the display

Answer: A

Q106. Array access by index?

- A. $O(1)$
- B. $O(n^2)$
- C. $O(n)$
- D. $O(\log n)$

Answer: A

Q107. Single for loop through n ?

- A. $O(1)$
- B. $O(n)$
- C. $O(\log n)$
- D. $O(n^2)$

Answer: B

Q108. $O(\log n)$ typical of?

- A. Binary search method
- B. Bubble sort algorithm
- C. Linear search scan
- D. Array index access

Answer: A

Q109. Algorithm efficiency?

- A. How readable the code looks
- B. How fast you can type code
- C. Time and space usage vs input
- D. The number of lines of code

Answer: C

Q110. $O(n)$ vs $O(2^n)$?

- A. $O(2^n)$ grows faster than $O(n)$
- B. They grow at the same rate
- C. It depends on the input data
- D. $O(n)$ grows faster than $O(2^n)$

Answer: A

Q111. First step in programming problem?

- A. Choose the programming language
- B. Start coding right away now
- C. Understand requirements and constraints
- D. Open the IDE and create files

Answer: C

Q112. What is a variable?

- A. A constant that never changes
- B. A loop construct in the program
- C. A function definition in code
- D. Named storage with changeable value

Answer: D

Q113. What is a function?

- A. A data structure for storage
- B. Reusable block for a specific task
- C. A variable type declaration
- D. A file format for saving data

Answer: B

Q114. Purpose of return?

- A. Declare a new variable in scope
- B. Create a new loop construct
- C. Send value back, end the function
- D. Print output to the console

Answer: C

Q115. What is a parameter?

- A. An error in the program logic
- B. Variable receiving value when called
- C. A global variable in the program
- D. The return type of a function

Answer: B

Q116. == vs =?

- A. Both are comparison operators
- B. = compares, == assigns values
- C. = assigns, == compares values
- D. They are the same operator

Answer: C

Q117. What is a string?

- A. A numeric data type value
- B. A callable function reference
- C. A boolean true/false value
- D. A sequence of characters stored

Answer: D

Q118. What is if-else for?

- A. Looping through elements repeatedly
- B. Making decisions based on conditions
- C. Importing external code modules
- D. Declaring new variable bindings

Answer: B

Q119. What is list/array for?

- A. Performing math calculations only
- B. Displaying output on the screen
- C. Ordered collection of multiple values
- D. Storing a single value only

Answer: C

Q120. What does DRY mean?

- A. Don't Repeat Yourself ever
- B. Do Repeat Yourself always
- C. Dev Run Yield automation
- D. Do Run Yesterday's builds

Answer: A

Q121. What is data-driven problem solving?

- A. Guessing the answer without facts
- B. Solving problems without any data
- C. Using data to guide all decisions
- D. Relying only on gut intuition

Answer: C

Q122. Purpose of data collection?

- A. Filling databases with random entries
- B. Wasting time gathering useless info
- C. Creating charts for presentations only
- D. Gather info to understand the problem

Answer: D

Q123. What is a dataset?

- A. A programming language to learn
- B. Structured collection of related data
- C. A single data point value only
- D. A server hosting the application

Answer: B

Q124. Mean of 2,4,6,8,10?

- A. 6
- B. 5
- C. 8
- D. 4

Answer: A

Q125. Median of 1,3,5,7,9?

- A. 3
- B. 5
- C. 7
- D. 9

Answer: B

Q126. Bar chart for?

- A. Showing trends over time periods
- B. Comparing quantities across categories
- C. Displaying geographic location data
- D. Showing parts of a whole circle

Answer: B

Q127. What is filtering?

- A. Selecting data meeting criteria
- B. Changing existing data values
- C. Adding more data to the set
- D. Deleting all data from storage

Answer: A

Q128. Pie chart for?

- A. Showing trends over time periods
- B. Displaying individual data points
- C. Proportions of parts to a whole
- D. Comparing many different categories

Answer: C

Q129. Sorting data for?

- A. Organizing for analysis and search
- B. Deleting duplicate data entries
- C. Encrypting data for security
- D. Making data set size smaller

Answer: A

Q130. What is data visualization?

- A. Graphical representation of patterns
- B. An algorithm for data sorting
- C. Raw numbers in a spreadsheet
- D. A database query for records

Answer: A

Q131. What is creative thinking?

- A. Copying solutions from other work
- B. Avoiding anything new or different
- C. Generating novel innovative solutions
- D. Only using standard approaches

Answer: C

Q132. What is brainstorming?

- A. An algorithm for sorting data
- B. A debugging technique for errors
- C. Thinking during a storm outside
- D. Generating many ideas without judgment

Answer: D

Q133. Thinking outside the box?

- A. Following standard procedures strictly
- B. Unconventional angles, challenge assumptions
- C. Thinking about a literal box shape
- D. Only using proven existing methods

Answer: B

Q134. What is lateral thinking?

- A. Thinking very fast without analysis
- B. Indirect creative approaches to solve
- C. Avoiding all creative problem solving
- D. Thinking in a straight line only

Answer: B

Q135. Why creativity in development?

- A. It introduces more bugs in code
- B. Drives innovative solutions and ideas
- C. It makes code harder to maintain
- D. It is not important at all here

Answer: B

Q136. What is a mind map?

- A. A network topology layout diagram
- B. Visual diagram from central concept
- C. A GPS navigation system map
- D. A floor plan of an office building

Answer: B

Q137. What is prototyping?

- A. Building the final product fully
- B. Final testing before the release
- C. Writing documentation for the project
- D. Quick rough version to test ideas

Answer: D

Q138. What is divergent thinking?

- A. Only considering one single idea
- B. Narrowing down to one solution
- C. Making the final decision now
- D. Generating many different ideas

Answer: D

Q139. What is convergent thinking?

- A. Generating many diverse ideas
- B. Narrowing down to best solution
- C. Choosing a solution at random
- D. Avoiding all possible solutions

Answer: B

Q140. What is MVP?

- A. Just enough features to test core
- B. A complete fully finished product
- C. A testing framework for software
- D. Most Valuable Player in sports

Answer: A

Q141. What is a real-world CS problem?

- A. A mathematical proof to write down
- B. A textbook exercise for practice
- C. A purely theoretical proof exercise
- D. Practical problem addressable with tech

Answer: D

Q142. Example of CS real-world solution?

- A. Solving an algebra equation on paper
- B. Philosophical debate on abstract ideas
- C. Writing mathematical proofs by hand
- D. GPS navigation and routing systems

Answer: D

Q143. Real-world vs academic?

- A. Real-world always has exact solutions
- B. Messy data, unclear requirements exist
- C. They are exactly the same thing
- D. Academic problems are always harder

Answer: B

Q144. What is user-centered design?

- A. Ignoring all user feedback given
- B. Design focused on machines only
- C. End user needs as primary focus
- D. Design focused on developers only

Answer: C

Q145. Why requirements gathering?

- A. It is entirely unnecessary to do
- B. It is only for documentation needs
- C. Ensures understanding of what to solve
- D. It slows the project down greatly

Answer: C

Q146. What is scalability?

- A. Handling increasing work or users
- B. The size of the codebase files
- C. The database table size limit
- D. The number of developers hired

Answer: A

Q147. What is ethical consideration?

- A. A type of search algorithm used
- B. A coding standard to follow always
- C. A testing methodology for quality
- D. Moral implications and societal impact

Answer: D

Q148. Why user feedback?

- A. Checks if the solution meets needs
- B. It provides nothing of value here
- C. It is only for documentation uses
- D. It slows down the development work

Answer: A

Q149. What is maintenance?

- A. A one-time task done at launch
- B. Cleaning hardware equipment daily
- C. Ongoing updates, fixes, improvements
- D. Only fixing bugs that are found

Answer: C

Q150. What is accessibility?

- A. Having access to the internet
- B. Design for diverse abilities and needs
- C. Making the application run faster
- D. Securing the application from attacks

Answer: B

Q151. Why documentation important?

- A. It is completely unnecessary to do
- B. Only managers need to read the docs
- C. It wastes time that could be coding
- D. Records solutions for understanding/reuse

Answer: D

Q152. What is a code comment?

- A. A variable declaration in the code
- B. Social feedback on the project
- C. An error in the source code
- D. Text explaining code, compiler ignores

Answer: D

Q153. Good docs include?

- A. Only the final answer or result
- B. Problem, approach, usage, examples
- C. Only the raw source code itself
- D. Only variable names and types

Answer: B

Q154. What is a README?

- A. A personal diary for the developer
- B. A system log of all the events
- C. A configuration file for the app
- D. Introduces project with purpose/setup

Answer: D

Q155. Why clear communication?

- A. Alignment on problems and solutions
- B. It is only needed for presentations
- C. It is not important for the team
- D. It slows down the project progress

Answer: A

Q156. Why clear names?

- A. To confuse other developers reading
- B. They create self-documenting code
- C. To make the code look fancy
- D. Variable names do not matter at all

Answer: B

Q157. What is user manual?

- A. A tutorial for learning to code
- B. End-user software usage instructions
- C. Source code documentation for devs
- D. A developer guide for the APIs

Answer: B

Q158. What is changelog?

- A. A debug log of runtime errors seen
- B. Chronological notable changes per version
- C. An error log of system failures
- D. A log of all user passwords used

Answer: B

Q159. Why document approach?

- A. It is a waste of team time
- B. No one reads documentation anyway
- C. Review thinking, learn, and share
- D. It makes the project much harder

Answer: C

Q160. What is a bug report?

- A. Defect with repro steps and details
- B. A proposal for new feature ideas
- C. A financial report for the project
- D. A report about actual insects found

Answer: A

Q161. What does abstraction help us do in problem solving?

- A. Include every minor detail possible
- B. Focus on relevant details only
- C. Avoid thinking about the problem
- D. Copy solutions from other sources

Answer: B

Q162. Which is an example of a well-defined problem?

- A. Finding the meaning of life itself
- B. Creating the best art masterpiece
- C. Predicting future weather in detail
- D. Sorting a list of numbers in order

Answer: D

Q163. What is the main goal of problem solving?

- A. Memorizing every formula in math
- B. Finding a solution to a challenge
- C. Writing the longest code possible
- D. Avoiding all difficult situations

Answer: B

Q164. Which term means testing a solution works?

- A. Abstraction method
- B. Brainstorming phase
- C. Verification process
- D. Decomposition step

Answer: C

Q165. What do constraints in a problem define?

- A. The limitations and boundaries given
- B. The color scheme of the interface
- C. The number of team members needed
- D. The programming language to be used

Answer: A

Q166. Which of these is a problem-solving tool?

- A. Video streaming app
- B. Online gaming site
- C. Flowchart diagram
- D. Social media post

Answer: C

Q167. What does iteration mean in problem solving?

- A. Solving the problem only one time
- B. Deleting the solution permanently
- C. Ignoring all previous attempt results
- D. Repeating steps to improve results

Answer: D

Q168. Which best describes a systematic approach?

- A. Following an organized step-by-step plan
- B. Random guessing without any planning
- C. Solving only the easiest parts first
- D. Skipping analysis and writing code now

Answer: A

Q169. What is the output of a problem-solving process?

- A. A list of new unrelated problems
- B. A solution to the stated problem
- C. An increase in overall complexity
- D. A decrease in understanding gained

Answer: B

Q170. Why is planning important before solving?

- A. It makes the problem more confusing
- B. It wastes valuable time and effort
- C. It helps organize thoughts and steps
- D. It eliminates all possible solutions

Answer: C

Q171. What is the first step in analyzing a problem?

- A. Writing the solution code immediately now
- B. Identifying what the problem actually is
- C. Presenting findings to the whole team
- D. Choosing a programming language to start

Answer: B

Q172. What are problem requirements?

- A. Random features added without any reason
- B. Personal preferences of the developer only
- C. Conditions the solution must fully satisfy
- D. Optional suggestions that can be ignored

Answer: C

Q173. Which question helps clarify a problem best?

- A. What is the latest technology trend now?
- B. What is my favorite programming language?
- C. What did I eat for breakfast this morning?
- D. What is the expected output or result?

Answer: D

Q174. What does identifying inputs mean in analysis?

- A. Determining what data is given or available
- B. Selecting the server hosting provider to use
- C. Deciding on the color scheme for the design
- D. Choosing the best font for the user interface

Answer: A

Q175. Why is understanding context important in analysis?

- A. Context never affects the problem at all
- B. Context makes every problem harder to solve
- C. Context helps understand the bigger picture
- D. Context only matters in language translation

Answer: C

Q176. What is a problem statement?

- A. A summary of team member qualifications
- B. A budget report for the project finances
- C. A clear description of the issue to solve
- D. A list of all programming languages known

Answer: C

Q177. Which helps separate relevant from irrelevant info?

- A. Critical analysis of the available data
- B. Randomly guessing which facts to use here
- C. Ignoring all the data that was provided
- D. Including every piece of information given

Answer: A

Q178. What are boundary conditions in a problem?

- A. The average expected input values given
- B. The extreme or edge cases to consider
- C. The most common scenarios encountered
- D. The normal typical cases only to test

Answer: B

Q179. What does scope of a problem refer to?

- A. The physical size of the computer screen
- B. The extent and boundaries of the problem
- C. The speed of the internet connection used
- D. The number of developers on the team now

Answer: B

Q180. Why should assumptions be stated explicitly?

- A. To confuse team members with extra content
- B. To prevent misunderstandings about the problem
- C. To delay the start of the actual solution work
- D. To make the documentation longer than needed

Answer: B

Q181. What is deductive reasoning?

- A. Ignoring all available evidence and data given
- B. Drawing general conclusions from specific cases
- C. Making random guesses without any evidence here
- D. Reaching specific conclusions from general rules

Answer: D

Q182. What is inductive reasoning?

- A. Contradicting all evidence with random assertions
- B. Eliminating logical thinking from all processes
- C. Forming general rules from specific observations
- D. Applying general rules to get specific outcomes

Answer: C

Q183. What does a logical AND operation require?

- A. Both conditions must be true to return true
- B. Exactly one condition must always be false now
- C. Neither condition needs to be true at all here
- D. Only one condition needs to be true for this

Answer: A

Q184. What does a logical OR operation require?

- A. Neither condition can ever be true for result
- B. At least one condition must be true for true
- C. Both conditions must be false for result true
- D. Both conditions must always be true for true

Answer: B

Q185. What does the NOT operator do in logic?

- A. It combines two statements together as one
- B. It performs arithmetic addition on numbers
- C. It reverses the truth value of a statement
- D. It checks if two values are exactly equal

Answer: C

Q186. Which conclusion follows from All cats are animals?

- A. All animals are definitely cats themselves
- B. No cats can possibly be animals at all
- C. Animals and cats are completely unrelated
- D. My pet cat is therefore an animal here

Answer: D

Q187. What is a logical fallacy?

- A. A valid and sound logical argument made
- B. A mathematical formula for calculations
- C. A type of programming loop structure
- D. An error in reasoning that weakens logic

Answer: D

Q188. What is a truth table used for?

- A. Displaying data in a formatted spreadsheet view
- B. Organizing project tasks by their due dates
- C. Showing all possible outcomes of logic operations
- D. Storing user information in a database system

Answer: C

Q189. What is a premise in a logical argument?

- A. The final conclusion of the argument made
- B. A question asked at the end of reasoning
- C. An irrelevant fact added for more length
- D. A statement assumed true to support conclusion

Answer: D

Q190. What does if-then represent in logic?

- A. A conditional or implication statement here
- B. A mathematical multiplication operation only
- C. A statement that is always false by nature
- D. An unconditional always true statement only

Answer: A

Q191. What shape represents a decision in a flowchart?

- A. Diamond shape
- B. Rectangle shape
- C. Parallelogram
- D. Oval shape

Answer: A

Q192. What does a rectangle represent in a flowchart?

- A. An input or output operation needed
- B. The start or end point of the flow
- C. A decision that requires a yes or no
- D. A process or action step to perform

Answer: D

Q193. What does sequential flow mean in algorithms?

- A. Steps that skip over each other unpredictably
- B. Steps that execute in completely random order
- C. Steps that all execute at the same exact time
- D. Steps that execute one after another in order

Answer: D

Q194. What is a loop in algorithm flow control?

- A. A repeated execution of a set of steps here
- B. A section of code that never executes at all
- C. An error that crashes the running program now
- D. A single instruction that runs only one time

Answer: A

Q195. What does the oval shape represent in a flowchart?

- A. The start or end of the flow process
- B. A decision requiring a yes or no here
- C. A process or calculation step taken
- D. An input or output data operation now

Answer: A

Q196. What is a conditional statement in algorithms?

- A. A statement that executes based on a condition
- B. A statement that always executes regardless
- C. A statement that randomly decides to execute
- D. A statement that never executes at all here

Answer: A

Q197. What does the arrow in a flowchart indicate?

- A. The size of the data being processed
- B. The color of the interface elements
- C. The direction of the process flow
- D. The speed of program execution rate

Answer: C

Q198. What is a while loop used for?

- A. Executing steps exactly one time only here
- B. Skipping all steps in the given program
- C. Terminating the program immediately right now
- D. Repeating steps while a condition is true

Answer: D

Q199. What is pseudocode?

- A. A specific proprietary programming language
- B. An executable script ready for deployment
- C. An informal description of algorithm steps
- D. Compiled machine-level binary code output

Answer: C

Q200. What is a for loop typically used for?

- A. Iterating a known number of times through code
- B. Declaring variables at the start of program
- C. Defining a function for later use in program
- D. Running code only when an error occurs here

Answer: A

Q201. What is data representation in computer science?

- A. The color of the user interface elements
- B. The physical size of the computer hardware
- C. The speed of the internet connection used
- D. How data is formatted and stored in systems

Answer: D

Q202. What is the binary number system based on?

- A. Sixteen digits including some letters
- B. Ten digits from zero through nine
- C. Two digits which are zero and one
- D. Eight digits from zero through seven

Answer: C

Q203. What does abstraction mean in data representation?

- A. Including every detail about the data given
- B. Making data more complex than it needs to be
- C. Deleting all data from the storage systems
- D. Hiding unnecessary details to show essentials

Answer: D

Q204. What is an array as a data structure?

- A. A type of output display on a screen
- B. A collection of elements stored in order
- C. A random arrangement of unrelated data
- D. A single value stored in one location

Answer: B

Q205. What does ASCII stand for in data encoding?

- A. Advanced System Code for Internet Integration
- B. Automated Script Compiler for Integrated Interfaces
- C. American Standard Code for Information Interchange
- D. Applied Software Configuration for Input Interchange

Answer: C

Q206. What is a variable in programming?

- A. A hardware component inside the computer
- B. An output device like a monitor or screen
- C. A fixed constant that can never change
- D. A named storage location for data values

Answer: D

Q207. What is a string data type used to store?

- A. Only true or false Boolean values given
- B. A sequence of text characters in order
- C. Only numerical integer values and nothing else
- D. Only decimal floating point number values

Answer: B

Q208. What is a Boolean data type?

- A. A type that holds only true or false values
- B. A type that contains image data information
- C. A type that stores very large number values
- D. A type that stores audio or sound file data

Answer: A

Q209. What does a table represent in data organization?

- A. A random collection of unrelated numbers
- B. An algorithm for sorting data elements
- C. Data organized in rows and columns format
- D. A type of computer hardware components

Answer: C

Q210. What is the purpose of data types in programming?

- A. To eliminate variables from program source
- B. To make programs run slower than necessary
- C. To define what kind of data a variable holds
- D. To prevent any data from being stored here

Answer: C

Q211. What is pattern recognition in problem solving?

- A. Ignoring all similarities between problems
- B. Creating random solutions without analysis
- C. Identifying similarities and recurring trends
- D. Memorizing every detail without understanding

Answer: C

Q212. Which is an example of a pattern in numbers?

- A. Letters mixed randomly with the numbers
- B. Each number is double the previous one
- C. Numbers arranged in no particular order
- D. Random numbers with no relation at all

Answer: B

Q213. How does recognizing patterns help solve new problems?

- A. It makes every problem harder to understand
- B. It eliminates the need to think about them
- C. It provides no useful information at all
- D. It allows applying known solutions to cases

Answer: D

Q214. What is generalization in the context of patterns?

- A. Finding exceptions to every rule discovered
- B. Making a rule that only applies to one case
- C. Ignoring all examples and making no rules
- D. Creating a broad rule from specific examples

Answer: D

Q215. What pattern does the sequence 2 4 6 8 10 follow?

- A. Each number decreases by one every time
- B. Each number increases by two every time
- C. Each number is multiplied by two each time
- D. Each number increases by three every time

Answer: B

Q216. What is a template in pattern-based solving?

- A. A random approach with no defined structure
- B. A reusable solution structure for similar problems
- C. A one-time solution that cannot be reused again
- D. A problem that has never been solved before now

Answer: B

Q217. Which skill is essential for pattern recognition?

- A. Careful observation of details and relationships
- B. Avoiding comparison between different examples
- C. Ignoring all details in the data provided
- D. Refusing to look for any similarities at all

Answer: A

Q218. What is a repeating pattern?

- A. A pattern that changes its rule every single time
- B. A sequence that cycles through the same elements
- C. A random arrangement with no repetition at all
- D. A sequence that occurs once and never repeats

Answer: B

Q219. Why are patterns important in computing?

- A. They only matter in artistic design projects
- B. They are unrelated to computing in any way
- C. They make computing more complex and slower
- D. They enable efficient algorithms and solutions

Answer: D

Q220. What does classification mean in pattern recognition?

- A. Sorting items into groups based on features
- B. Deleting all items that look similar to others
- C. Ignoring all features of the items completely
- D. Randomly mixing all items together in a pile

Answer: A

Q221. What is the result of 15 modulo 4?

- A. 3
- B. 2
- C. 4
- D. 5

Answer: A

Q222. What does the factorial of a number represent?

- A. Product of all positive integers up to that number
- B. Difference between the number and its predecessor
- C. Quotient when the number is divided by two here
- D. Sum of all positive integers up to that number here

Answer: A

Q223. What is the base of the decimal number system?

- A. Ten
- B. Two
- C. Eight
- D. Sixteen

Answer: A

Q224. What is a prime number?

- A. Any even number greater than two in the sequence
- B. A number divisible by any integer greater than one
- C. A number divisible only by one and by itself here
- D. A number that is always negative in value given

Answer: C

Q225. What is the order of operations in mathematics?

- A. Parentheses, exponents, multiply, divide, add, subtract
- B. Random order depending on the solver preference
- C. Addition then multiplication then parentheses used
- D. Subtraction first then all other operations after

Answer: A

Q226. What does the term average typically mean?

- A. The middle value when data is sorted in any order
- B. The smallest value in a given set of numbers here
- C. The sum of values divided by their total count
- D. The largest value in a given set of numbers here

Answer: C

Q227. What is a ratio?

- A. A comparison of two quantities using division here
- B. The product of multiplying two numbers together
- C. The sum of two different numbers added together
- D. The difference between two numbers subtracted here

Answer: A

Q228. What does percentage literally mean?

- A. Per hundred units in the total population here
- B. Per ten units in the total population being used
- C. Per thousand units in the total population here
- D. Per million units in the total population given

Answer: A

Q229. What is the greatest common divisor of 12 and 18?

- A. 4
- B. 6
- C. 3
- D. 9

Answer: B

Q230. What is an integer?

- A. A number used only in advanced calculus problems
- B. A number that can only be positive never negative
- C. A number that always has a decimal point value
- D. A whole number without any decimal part given

Answer: D

Q231. What is algorithmic thinking?

- A. Writing code without any planning or design thoughts
- B. Avoiding any structured approach to solving problems
- C. Memorizing code without understanding any of its logic
- D. Thinking about algorithms as step-by-step processes

Answer: D

Q232. What is the first step in designing an algorithm?

- A. Understanding the problem and its requirements here
- B. Writing the final code in a programming language
- C. Testing the code with random input data values given
- D. Optimizing the solution for maximum speed of running

Answer: A

Q233. What is a sorting algorithm used for?

- A. Deleting elements from a given data collection
- B. Searching for a specific element in the given data
- C. Adding new elements to an existing data structure
- D. Arranging elements in a specific defined order

Answer: D

Q234. What is a searching algorithm used for?

- A. Finding a specific element within a data collection
- B. Creating new data elements from existing data given
- C. Removing duplicate elements from the data structure
- D. Organizing data into a sorted specific order here

Answer: A

Q235. What does efficiency mean for an algorithm?

- A. Producing incorrect results as quickly as possible
- B. Ignoring resource constraints during the execution
- C. Using minimal resources to achieve correct results
- D. Using the maximum amount of time and memory here

Answer: C

Q236. What is a greedy algorithm?

- A. An algorithm that explores all possible solutions
- B. An algorithm that makes the best choice at each step
- C. An algorithm that randomly selects the next action
- D. An algorithm that always picks the worst choice made

Answer: B

Q237. What does it mean for an algorithm to terminate?

- A. The algorithm runs indefinitely without stopping here
- B. The algorithm eventually stops and produces a result
- C. The algorithm crashes due to a runtime error found
- D. The algorithm restarts itself from the beginning now

Answer: B

Q238. What is a brute force algorithm?

- A. An algorithm that only checks the first possibility
- B. An efficient algorithm that skips unnecessary steps
- C. An algorithm that guesses the answer without checking
- D. An algorithm that tries all possibilities for answer

Answer: D

Q239. What is the input to an algorithm?

- A. The steps that make up the algorithm procedure
- B. The data provided for the algorithm to process
- C. The result that the algorithm produces as output
- D. The language used to write the algorithm code in

Answer: B

Q240. What is the output of an algorithm?

- A. The result produced after executing the algorithm
- B. The initial data given before processing starts
- C. The errors encountered during algorithm execution
- D. The code written to implement the algorithm here

Answer: A

Q241. What is critical thinking?

- A. Criticizing others without constructive feedback
- B. Accepting all information without any questioning
- C. Avoiding making any decisions under circumstances
- D. Analyzing and evaluating information objectively

Answer: D

Q242. Why is evidence important in decision making?

- A. Evidence supports informed and rational decisions
- B. Evidence makes decisions slower and more confusing
- C. Evidence should be ignored for faster decisions
- D. Evidence is irrelevant to any decision process

Answer: A

Q243. What does it mean to evaluate an argument?

- A. Assessing the strength and validity of reasoning
- B. Accepting the argument without any analysis here
- C. Ignoring the argument and its conclusions entirely
- D. Repeating the argument without understanding it

Answer: A

Q244. What is bias in thinking?

- A. A balanced and objective viewpoint on all issues
- B. A tendency to favor one perspective over others
- C. An approach that considers every possible angle
- D. A guarantee of accurate reasoning in all cases

Answer: B

Q245. What is the purpose of asking why questions?

- A. To avoid having to think about the problem here
- B. To delay making any progress on the decision
- C. To annoy others by questioning their decisions
- D. To understand the reasoning behind a statement

Answer: D

Q246. What is a fact versus an opinion?

- A. A fact is verifiable while an opinion is subjective
- B. Facts cannot be proven while opinions always can be
- C. An opinion is always more reliable than any fact
- D. Facts and opinions are exactly the same thing here

Answer: A

Q247. Why should multiple sources be consulted?

- A. Only the first source found should ever be used
- B. Multiple sources help verify and cross-check facts
- C. One source always provides complete information
- D. Consulting multiple sources wastes valuable time

Answer: B

Q248. What does objectivity mean in critical thinking?

- A. Judging based on facts without personal prejudice
- B. Ignoring all evidence and trusting only intuition
- C. Always agreeing with the majority group opinion
- D. Being influenced by personal feelings and biases

Answer: A

Q249. What is a conclusion in an argument?

- A. The starting assumptions of the argument given
- B. A random unrelated statement at the argument end
- C. The statement that follows from the given premises
- D. The evidence used to support the argument here

Answer: C

Q250. Why is open-mindedness important?

- A. It means accepting every claim without question
- B. It means rejecting all traditional knowledge given
- C. It prevents any evaluation of competing ideas here
- D. It allows considering alternative perspectives fairly

Answer: D

Q251. What is debugging in software development?

- A. The process of designing user interfaces for apps
- B. The process of writing new features from scratch
- C. The process of finding and fixing errors in code
- D. The process of deploying code to production server

Answer: C

Q252. What is a syntax error?

- A. An error that only occurs during program execution
- B. A violation of the programming language grammar rules
- C. A warning that does not affect program execution
- D. An error in the logic of the program algorithm

Answer: B

Q253. What is a runtime error?

- A. A warning shown during code compilation process here
- B. An error that occurs while the program is executing
- C. An error detected before the program runs at all
- D. A style issue found during code review by the team

Answer: B

Q254. What is a logical error?

- A. An error caused by missing semicolons in the code
- B. An error in the spelling of variable names in code
- C. An error that prevents the code from compiling here
- D. An error where code runs but produces wrong results

Answer: D

Q255. What is a breakpoint in debugging?

- A. A marked point where execution pauses for inspection
- B. A location where new code should be added to fix bug
- C. A section of code that should be deleted from source
- D. A point where the program permanently crashes here

Answer: A

Q256. What does a stack trace show?

- A. The number of lines of code in the entire program
- B. The amount of memory used by the entire program
- C. The sequence of function calls leading to an error
- D. The total execution time of the program in seconds

Answer: C

Q257. What is the simplest debugging technique?

- A. Using machine learning to find all the bugs here
- B. Adding print statements to trace program execution
- C. Deleting code randomly until the error disappears
- D. Rewriting the entire program from scratch every time

Answer: B

Q258. What is a bug in software?

- A. A standard part of every software design document
- B. An improvement suggestion from the user community
- C. A feature that works exactly as it was intended
- D. An error or flaw that causes unexpected behavior

Answer: D

Q259. Why is reproducing a bug important for debugging?

- A. It makes the bug worse and harder to fix in the code
- B. Bugs should never be reproduced only described aloud
- C. It helps understand the exact conditions causing the bug
- D. Reproducing bugs is unnecessary for fixing them here

Answer: C

Q260. What is a debugger tool?

- A. A tool that automatically writes new code features
- B. A tool that designs user interfaces for the program
- C. A tool that manages project timelines and schedules
- D. A tool that helps step through code to find errors

Answer: D

Q261. What does time complexity measure?

- A. How execution time grows with input size increase
- B. How long it takes to write the program code here
- C. How long the computer takes to boot up for work
- D. The calendar time to complete the whole project

Answer: A

Q262. What does space complexity measure?

- A. The disk space needed to store the source code here
- B. The number of monitors needed to display the output
- C. The physical size of the computer used for running
- D. How memory usage grows with the input size given

Answer: D

Q263. What does $O(1)$ mean in complexity notation?

- A. The algorithm takes constant time regardless of input
- B. The algorithm grows linearly with input size given
- C. The algorithm takes infinite time for all input sizes
- D. The algorithm cannot run with any input whatsoever

Answer: A

Q264. What does $O(n)$ mean in complexity notation?

- A. The time decreases as input size increases further
- B. The time grows proportionally with input size here
- C. The time grows exponentially with input size given
- D. The time remains constant for all input sizes given

Answer: B

Q265. Which is more efficient for large inputs: $O(n)$ or $O(n^2)$?

- A. Neither can handle any large input sizes effectively
- B. Both are equally efficient for all input sizes given
- C. $O(n^2)$ is always more efficient for large inputs
- D. $O(n)$ is more efficient because it grows more slowly

Answer: D

Q266. What is the complexity of accessing an array by index?

- A. $O(n^2)$ quadratic time
- B. $O(1)$ constant time
- C. $O(\log n)$ logarithmic
- D. $O(n)$ linear time

Answer: B

Q267. What does scalability mean for an algorithm?

- A. The algorithm requires manual intervention to run now
- B. The algorithm only works with very small inputs here
- C. How well the algorithm handles growing input sizes
- D. The algorithm cannot be modified or changed at all

Answer: C

Q268. What is the complexity of a simple for loop over n items?

- A. $O(n)$ linear time
- B. $O(n^2)$ quadratic time
- C. $O(1)$ constant time
- D. $O(\log n)$ logarithmic

Answer: A

Q269. Why do we ignore constants in Big-O notation?

- A. Constants are always exactly one in every algorithm
- B. Constants make Big-O notation too easy to calculate
- C. Constants become insignificant for large input sizes
- D. Constants matter more than the growth rate itself

Answer: C

Q270. What does $O(n^2)$ typically indicate?

- A. It has a nested loop iterating over the input data
- B. It accesses only one element of the input per run
- C. It requires no iterations through the data at all
- D. It has a single pass through the input data given

Answer: A

Q271. What is the purpose of a function in programming?

- A. To store data permanently on the hard disk drive
- B. To prevent any code from being executed by program
- C. To organize reusable blocks of code for tasks here
- D. To make code longer and more repetitive overall

Answer: C

Q272. What is a variable used for in problem solving?

- A. Displaying error messages on the computer screen
- B. Making the program run slower than it needs to run
- C. Storing data values that can change during execution
- D. Preventing any data from being stored in the program

Answer: C

Q273. What is a conditional statement used for?

- A. Making decisions based on whether conditions are true
- B. Defining a new function to be called later in code
- C. Repeating the same code block multiple times in row
- D. Importing external libraries for use in the program

Answer: A

Q274. What is an array used for in programs?

- A. Connecting to the internet for data downloading now
- B. Storing a single value at one memory location here
- C. Displaying graphics and images on the screen here
- D. Storing a collection of related values in structure

Answer: D

Q275. What does a return statement do in a function?

- A. It sends a value back to the code that called it
- B. It restarts the entire program from the beginning
- C. It permanently deletes the function from memory now
- D. It prints a message directly to the user screen

Answer: A

Q276. What is a string in programming?

- A. A physical wire connecting hardware parts together
- B. A numerical value used for math calculations here
- C. A Boolean value that is either true or false only
- D. A sequence of characters representing text data

Answer: D

Q277. Why are comments important in program code?

- A. Comments increase file size for better performance
- B. Comments make the program execute faster overall
- C. Comments are executed by the computer as commands
- D. Comments explain code purpose for human readers

Answer: D

Q278. What is an operator in programming?

- A. A file stored on the computer hard disk drive now
- B. A type of comment added for documentation purpose
- C. A symbol that performs an operation on values here
- D. A person who operates the computer hardware device

Answer: C

Q279. What is input in a programming context?

- A. The speed at which the program code is executed
- B. The visual appearance of the program user interface
- C. Data sent out by the program to external devices
- D. Data received by the program from external sources

Answer: D

Q280. What is output in a programming context?

- A. The errors found during compilation process of code
- B. The source code written by the programmer initially
- C. Data the program produces and sends to destination
- D. Data the program receives from the user or file

Answer: C

Q281. What is data-driven problem solving?

- A. Making decisions based on data and evidence collected
- B. Ignoring all available data when making any decisions
- C. Collecting data without ever using it for decisions
- D. Making decisions based on personal feelings and guesses

Answer: A

Q282. What is a dataset?

- A. A structured collection of related data records stored
- B. A single data point with no related information here
- C. A type of programming language used for data analysis
- D. A hardware device for storing physical paper documents

Answer: A

Q283. What is data collection?

- A. The process of ignoring available information sources
- B. The process of deleting all stored information here
- C. The process of randomizing existing information order
- D. The process of gathering relevant information needed

Answer: D

Q284. What is a chart used for in data analysis?

- A. Writing source code for a program application here
- B. Visually representing data to identify trends shown
- C. Encrypting data for secure network transmission now
- D. Storing raw data in a database management system

Answer: B

Q285. What does filtering data mean?

- A. Selecting data that meets specific criteria from set
- B. Adding fake data to increase the total dataset size
- C. Randomly removing data without any specific reason
- D. Including all data regardless of its relevance here

Answer: A

Q286. What is a spreadsheet commonly used for?

- A. Writing essays and long research documents only here
- B. Playing computer games and entertainment activities
- C. Organizing and analyzing data in rows and columns
- D. Creating animated videos and graphic design artwork

Answer: C

Q287. What is the mean of a dataset?

- A. The average calculated by sum divided by count here
- B. The range between the highest and lowest values
- C. The middle value when data is sorted in an order
- D. The most frequently occurring value in the data set

Answer: A

Q288. What is data cleaning?

- A. Fixing errors and removing inconsistencies in data
- B. Printing data on paper for physical storage needs
- C. Physically cleaning the computer hardware surfaces
- D. Adding more data to make the dataset larger in size

Answer: A

Q289. Why is sample size important in data analysis?

- A. Sample size only matters for qualitative research here
- B. Smaller samples always give the most accurate results
- C. Sample size has no effect on analysis reliability
- D. Larger samples generally give more reliable results

Answer: D

Q290. What is a survey used for in data collection?

- A. Writing computer programs for data processing tasks
- B. Gathering responses and opinions from participants
- C. Testing hardware components for defects and faults
- D. Designing user interfaces for software applications

Answer: B

Q291. What is creative thinking in problem solving?

- A. Avoiding all new ideas and sticking to tradition
- B. Following only established conventional methods here
- C. Generating novel and innovative solution approaches
- D. Copying solutions exactly from existing examples now

Answer: C

Q292. What is brainstorming?

- A. Avoiding all group discussion about the problem
- B. Selecting only the first idea that comes to mind
- C. Evaluating and criticizing every idea immediately
- D. Generating many ideas without initial judgment here

Answer: D

Q293. What is thinking outside the box?

- A. Considering unconventional approaches to problems
- B. Rejecting all ideas that are different or unusual
- C. Following the standard procedure without variation
- D. Staying within conventional boundaries and limits

Answer: A

Q294. What role does curiosity play in creative solving?

- A. Curiosity only matters in scientific research alone
- B. Curiosity drives exploration and new discoveries
- C. Curiosity has no effect on any problem solving work
- D. Curiosity hinders creative problem solving efforts

Answer: B

Q295. What is lateral thinking?

- A. Following the most obvious solution path directly
- B. Approaching problems from indirect creative angles
- C. Avoiding creative approaches to problem solving now
- D. Thinking in a straight linear logical sequence only

Answer: B

Q296. Why is it important to consider multiple solutions?

- A. The first solution found is always the best one here
- B. Considering options wastes time and adds confusion
- C. Multiple solutions provide options for best selection
- D. Only one solution can ever exist for any problem

Answer: C

Q297. What is an analogy in creative thinking?

- A. A mathematical formula for calculating precise values
- B. An exact copy of an existing solution without changes
- C. A list of rules that must be followed without question
- D. A comparison between two similar things for insight

Answer: D

Q298. What does innovation mean?

- A. Avoiding change and maintaining the current approach
- B. Copying competitors without adding any improvement
- C. Repeating the same old methods without any change
- D. Introducing new ideas or methods to solve problems

Answer: D

Q299. What is mind mapping used for?

- A. Performing complex mathematical calculations quickly
- B. Writing sequential linear text documents for reports
- C. Visually organizing ideas and their connections here
- D. Managing computer hardware and network connections

Answer: C

Q300. Why is failure important in the creative process?

- A. Failure provides learning opportunities for improvement
- B. Failure should be avoided at all costs in every case
- C. Failure indicates no more attempts should be tried now
- D. Failure means the creative process has ended forever

Answer: A

Q301. What makes real-world problems different from textbook?

- A. Textbook problems are always harder to solve properly
- B. Real-world problems are messy with multiple factors
- C. Real-world problems have one clear correct answer
- D. Real-world problems require no practical considerations

Answer: B

Q302. Why is stakeholder input important?

- A. Stakeholders have no interest in the solution at all
- B. Stakeholders provide requirements and perspective here
- C. Only developers should make all decisions on solutions
- D. Stakeholder input always complicates the solution work

Answer: B

Q303. What is a trade-off in real-world problem solving?

- A. A solution that satisfies every requirement completely
- B. An approach that has no disadvantages at all given
- C. Giving up something to gain something else of value
- D. A situation where all options are equally perfect here

Answer: C

Q304. Why is time management important in projects?

- A. Real-world projects have unlimited time for completion
- B. Time management only matters for personal life tasks
- C. Deadlines require efficient prioritization of tasks
- D. Time has no impact on real-world project outcomes

Answer: C

Q305. What is a prototype in real-world development?

- A. A financial budget plan for the entire project cost
- B. A document listing all project requirements in detail
- C. An early model built to test ideas and get feedback
- D. The final finished product ready for market launch

Answer: C

Q306. Why is user feedback important?

- A. User feedback is irrelevant to solution quality here
- B. Users never provide useful information about solutions
- C. Feedback should only be gathered after project ends now
- D. It ensures the solution meets actual user needs given

Answer: D

Q307. What is project scope?

- A. The boundaries of what the project will deliver here
- B. The budget allocated for team lunches and activities
- C. The number of meetings held during the project now
- D. The personal goals of the project team manager now

Answer: A

Q308. Why are deadlines important in real-world projects?

- A. Projects should never have any deadlines set for them
- B. Deadlines have no effect on project delivery at all
- C. They create urgency and help prioritize work efforts
- D. Deadlines always reduce the quality of work produced

Answer: C

Q309. What is teamwork in real-world problem solving?

- A. Working alone without any help from other people here
- B. Competing against team members for individual credit
- C. Avoiding all communication with colleagues on team
- D. Collaborating with others to achieve a shared goal now

Answer: D

Q310. What is a requirement in a real-world project?

- A. A personal preference of the lead developer on team
- B. A suggestion that has no impact on the final product
- C. A condition the solution must meet to be acceptable
- D. An optional nice-to-have feature that can be skipped

Answer: C

Q311. Why is clear communication important in solving?

- A. Communication has no role in problem solving at all
- B. It ensures everyone understands problem and solution
- C. Only technical people need to communicate solutions
- D. Clear communication slows down problem solving work

Answer: B

Q312. What is the purpose of documentation in a project?

- A. To waste time that could be spent writing more code
- B. To prevent anyone from understanding project work
- C. To make projects more complicated than they need be
- D. To record decisions and provide reference for team

Answer: D

Q313. What should a good bug report include?

- A. The entire source code of the application in report
- B. Steps to reproduce, expected and actual behavior
- C. Only the statement that something is broken wrong
- D. A complaint about the development team performance

Answer: B

Q314. What is a README file typically used for?

- A. Listing personal information about team members now
- B. Recording all meetings held during project lifetime
- C. Storing the database passwords for the application
- D. Providing overview and setup instructions for project

Answer: D

Q315. Why is audience awareness important in documenting?

- A. Audience does not matter when writing documentation
- B. Technical details should always be included for all
- C. Different audiences need different levels of detail
- D. All audiences need exactly the same level of detail

Answer: C

Q316. What is the purpose of code comments?

- A. To replace the need for well-named variables entirely
- B. To make the code file unnecessarily larger in size
- C. To explain the intent and reasoning behind the code
- D. To prevent the code from compiling and running here

Answer: C

Q317. What is a technical specification document?

- A. A list of cafes near office for team lunch options
- B. A detailed description of how a system should be built
- C. A personal diary of developer daily activities here
- D. A marketing brochure for selling product to users

Answer: B

Q318. Why is version history important in documentation?

- A. It makes documents harder to read and understand now
- B. Version history should be deleted before sharing files
- C. It tracks changes and evolution of document over time
- D. Version history is never useful for any purpose here

Answer: C

Q319. What is a user manual?

- A. A guide for developers to write source code for app
- B. A guide that explains how to use a product or system
- C. A financial report about company revenue numbers here
- D. A legal contract between two companies for services

Answer: B

Q320. What makes writing concise and effective?

- A. Repeating the same information multiple times through
- B. Conveying meaning clearly with minimal extra words
- C. Using as many words as possible for every point made
- D. Including every tangential thought in documentation

Answer: B

Q321. What is the role of defining a problem clearly before solving it?

- A. It wastes time
- B. It ensures effort is directed at the right issue
- C. It limits creative solutions
- D. It is only needed for simple problems

Answer: B

Q322. Which of these is a real-world example of problem solving?

- A. Memorizing a textbook
- B. Debugging a program that crashes on startup
- C. Copying someone else's solution
- D. Ignoring error messages

Answer: B

Q323. What does the term 'solution space' refer to?

- A. The room where you work
- B. The set of all possible solutions to a problem
- C. A specific programming language
- D. The final answer only

Answer: B

Q324. Why is breaking a problem into smaller parts helpful?

- A. It makes the problem more complex
- B. Smaller parts are easier to understand and solve individually
- C. It increases the number of problems
- D. It is only useful for math problems

Answer: B

Q325. What is a common first reaction when encountering a difficult problem?

- A. Immediately writing code
- B. Understanding and restating the problem
- C. Skipping it entirely
- D. Asking someone else to do it

Answer: B

Q326. Which of the following best describes a well-structured problem?

- A. A problem with no clear answer
- B. A problem with clear goals, initial state, and known operations
- C. A problem that changes over time
- D. A problem that requires guessing

Answer: B

Q327. What does 'trial and error' mean as a problem-solving approach?

- A. Giving up after one attempt
- B. Trying different solutions until one works
- C. Only using proven methods
- D. Avoiding all mistakes

Answer: B

Q328. Which quality is most important for effective problem solving?

- A. Speed only
- B. Persistence and patience
- C. Memorizing all formulas
- D. Working alone always

Answer: B

Q329. What is the purpose of verifying a solution after implementing it?

- A. To waste time
- B. To ensure the solution actually solves the original problem
- C. To make the problem harder
- D. Verification is unnecessary

Answer: B

Q330. Which term describes learning from solving one problem to solve another?

- A. Randomization
- B. Transfer of knowledge
- C. Avoidance
- D. Memorization

Answer: B

Q331. What is the purpose of gathering requirements before solving a problem?

- A. To delay the project
- B. To understand exactly what the solution must achieve
- C. To create more work
- D. Requirements are unnecessary

Answer: B

Q332. What is a constraint in problem analysis?

- A. A suggestion that can be ignored
- B. A limitation or restriction that the solution must respect
- C. A feature request
- D. A type of programming language

Answer: B

Q333. Why is identifying the expected output important during problem analysis?

- A. It is not important
- B. It defines what the solution should produce, guiding the design
- C. It makes problems harder
- D. Output is always obvious

Answer: B

Q334. What does it mean to identify stakeholders in problem analysis?

- A. Finding who will write the code
- B. Determining who is affected by or interested in the problem and its solution
- C. Counting team members
- D. Stakeholders are irrelevant

Answer: B

Q335. What is the difference between inputs and outputs in a problem?

- A. They are the same thing
- B. Inputs are data given to the solution; outputs are results produced by it
- C. Inputs are always numbers
- D. Outputs come before inputs

Answer: B

Q336. Why should you identify assumptions when analyzing a problem?

- A. Assumptions are always correct
- B. Unstated assumptions can lead to incorrect solutions if they prove false
- C. Assumptions make problems easier
- D. They are only for documentation

Answer: B

Q337. What is a use case in problem analysis?

- A. A storage container
- B. A description of how a user interacts with a system to achieve a goal
- C. A type of error
- D. A programming pattern

Answer: B

Q338. What role do examples play in understanding a problem?

- A. Examples are distracting
- B. They illustrate expected behavior and help clarify the problem
- C. They are only for beginners
- D. Examples always oversimplify

Answer: B

Q339. What is the benefit of restating a problem in your own words?

- A. It wastes time
- B. It helps verify understanding and can reveal misunderstandings
- C. It changes the problem
- D. It is only done in exams

Answer: B

Q340. What does prioritizing requirements mean in problem analysis?

- A. Ignoring some requirements
- B. Ranking requirements by importance to focus on the most critical ones first
- C. Making all requirements equal
- D. Removing constraints

Answer: B

Q341. What is a conclusion in a logical argument?

- A. The first statement
- B. A random guess
- C. The statement that follows logically from the premises
- D. An assumption without evidence

Answer: C

Q342. What does the logical implication 'If A then B' mean?

- A. A and B are the same
- B. Whenever A is true, B must also be true
- C. B causes A
- D. A is always false

Answer: B

Q343. Which logical operator returns true only when both inputs are true?

- A. OR
- B. NOT
- C. AND
- D. XOR

Answer: C

Q344. What is an argument in the context of logic?

- A. A heated disagreement
- B. A set of premises followed by a conclusion
- C. A programming variable
- D. A mathematical equation

Answer: B

Q345. If 'All dogs are mammals' and 'Rex is a dog,' what can we conclude?

- A. Rex is a cat
- B. Rex is a mammal
- C. Rex is not an animal
- D. No conclusion can be drawn

Answer: B

Q346. What is a valid argument in logic?

- A. An argument that is always wrong
- B. An argument where the conclusion necessarily follows from the premises
- C. An argument with true premises only
- D. An argument everyone agrees with

Answer: B

Q347. What does the OR operator do in Boolean logic?

- A. Returns true only when both inputs are true
- B. Returns true when at least one input is true
- C. Returns the opposite of the input
- D. Returns false always

Answer: B

Q348. What is a counterexample used for in reasoning?

- A. To prove a statement is always true
- B. To disprove a universal claim by showing one case where it fails
- C. To add confusion
- D. To support an argument

Answer: B

Q349. In a truth table, how many rows are needed for three variables?

- A. 3
- B. 6
- C. 8
- D. 12

Answer: C

Q350. What does it mean for two statements to be logically equivalent?

- A. They have the same words
- B. They have the same truth value in all possible situations
- C. They always contradict each other
- D. They are written in the same language

Answer: B

Q351. What does the parallelogram symbol represent in a flowchart?

- A. A decision
- B. An input or output operation
- C. The start of the program
- D. A loop

Answer: B

Q352. What is the purpose of an arrow in a flowchart?

- A. To decorate the chart
- B. To show the direction of flow from one step to the next
- C. To represent data storage
- D. To indicate an error

Answer: B

Q353. What happens when a condition in an if-statement is false?

- A. The program crashes
- B. The code in the else branch (if it exists) executes, or the if block is skipped
- C. The condition becomes true
- D. The program restarts

Answer: B

Q354. What is the difference between a flowchart and pseudocode?

- A. They are identical
- B. A flowchart uses visual symbols; pseudocode uses text resembling programming language
- C. Flowcharts are only for hardware
- D. Pseudocode requires a compiler

Answer: B

Q355. What is a counter variable commonly used for in loops?

- A. Storing user names
- B. Keeping track of how many times a loop has executed
- C. Measuring memory
- D. Defining functions

Answer: B

Q356. What is the result of executing a sequence of three assignment statements in order?

- A. Only the first one runs
- B. All three execute one after another in order
- C. They execute simultaneously
- D. Only the last one runs

Answer: B

Q357. What does a diamond shape represent in a flowchart?

- A. A process step
- B. An input operation
- C. A decision point with yes/no branches
- D. The end of the program

Answer: C

Q358. Why is indentation important in pseudocode?

- A. It is purely decorative
- B. It shows the structure and nesting level of control statements
- C. It affects execution speed
- D. Indentation is incorrect in pseudocode

Answer: B

Q359. What is the simplest type of control flow in programming?

- A. Recursion
- B. Sequential execution where statements run one after another
- C. Multi-threading
- D. Event-driven programming

Answer: B

Q360. What does a loop termination condition determine?

- A. How fast the loop runs
- B. When the loop should stop repeating
- C. The loop's color in a flowchart
- D. Which variables to declare

Answer: B

Q361. What is the purpose of a data type in programming?

- A. To make code longer
- B. To define what kind of values a variable can hold and what operations are valid
- C. To slow down execution
- D. Data types are optional

Answer: B

Q362. What is the difference between an integer and a floating-point number?

- A. They are the same
- B. An integer has no decimal part while a floating-point number can represent fractions
- C. Integers are always larger
- D. Floating-point is always exact

Answer: B

Q363. What does a character encoding like ASCII define?

- A. How to draw characters
- B. A mapping between numeric codes and text characters
- C. The speed of text processing
- D. The size of a monitor

Answer: B

Q364. What is a record or struct used for in data representation?

- A. Playing music
- B. Grouping related data of different types into a single unit
- C. Sorting numbers only
- D. Replacing arrays

Answer: B

Q365. Why is binary the fundamental number system for computers?

- A. It was chosen arbitrarily
- B. Computer hardware uses two states (on/off) which map directly to binary digits 0 and 1
- C. Binary is easier for humans
- D. Decimal does not exist

Answer: B

Q366. What is a list data structure?

- A. A type of database
- B. An ordered collection of elements that can grow or shrink in size
- C. A fixed-size array only
- D. A type of loop

Answer: B

Q367. What does the term 'bit' stand for?

- A. Big integer type
- B. Binary digit
- C. Byte in transfer
- D. Boolean input test

Answer: B

Q368. What is the difference between a constant and a variable?

- A. They are the same
- B. A constant's value cannot change after being set; a variable's value can be modified
- C. Constants are always zero
- D. Variables never hold numbers

Answer: B

Q369. How many values can a single byte represent?

- A. 8
- B. 16
- C. 256
- D. 1024

Answer: C

Q370. What is a dictionary or map data structure used for?

- A. Looking up words only
- B. Storing key-value pairs for efficient lookup by key
- C. Sequential processing only
- D. Replacing all other structures

Answer: B

Q371. What does it mean to identify a pattern in a sequence of numbers?

- A. Guessing randomly
- B. Finding a consistent rule that describes how the numbers relate to each other
- C. Memorizing all numbers
- D. It is not possible

Answer: B

Q372. Why is pattern recognition useful in everyday life?

- A. It is not useful
- B. It helps predict outcomes and make decisions based on past experiences
- C. It only works in computers
- D. It slows down thinking

Answer: B

Q373. What is an example of a visual pattern?

- A. A random arrangement of shapes
- B. A wallpaper design that repeats the same motif at regular intervals
- C. A single dot on a page
- D. An empty canvas

Answer: B

Q374. What is the next number in the pattern 3, 6, 9, 12, ...?

- A. 13
- B. 14
- C. 15
- D. 18

Answer: C

Q375. How does recognizing a pattern in code help a programmer?

- A. It does not help
- B. It allows reuse of known solutions and avoidance of reinventing approaches
- C. It makes coding harder
- D. It only helps with debugging

Answer: B

Q376. What is a growing pattern?

- A. A pattern that decreases
- B. A pattern where each element increases by a consistent amount or rule
- C. A pattern with no change
- D. A pattern that stops

Answer: B

Q377. What does sorting data help reveal?

- A. Nothing useful
- B. Patterns, duplicates, and outliers that are hidden in unsorted data
- C. Random errors
- D. The data's creation date

Answer: B

Q378. What type of pattern does alternating between two values represent?

- A. A growing pattern
- B. An alternating or cyclic pattern
- C. A random sequence
- D. No pattern at all

Answer: B

Q379. How can a table help in recognizing patterns in data?

- A. Tables hide patterns
- B. Tables organize data into rows and columns, making relationships and trends easier to see
- C. Tables only store text
- D. Tables are too complex to read

Answer: B

Q380. What is the value of finding similarities between two different problems?

- A. It has no value
- B. Similar problems often have similar solutions, saving time and effort
- C. It creates confusion
- D. It only works in mathematics

Answer: B

Q381. What is the result of 3 raised to the power of 4?

- A. 12
- B. 64
- C. 81
- D. 27

Answer: C

Q382. What does the modulo operation return?

- A. The quotient of division
- B. The remainder after integer division
- C. The product of two numbers
- D. The absolute value

Answer: B

Q383. What is the sum of the first 5 even numbers (2, 4, 6, 8, 10)?

- A. 20
- B. 25
- C. 30
- D. 35

Answer: C

Q384. What is the least common multiple (LCM) of 3 and 5?

- A. 8
- B. 10
- C. 15
- D. 30

Answer: C

Q385. What does 'n choose 2' or $C(n,2)$ count?

- A. The number of ways to arrange n items
- B. The number of ways to pick 2 items from n without regard to order
- C. The product of n and 2
- D. The square of n

Answer: B

Q386. What is the value of log base 10 of 1000?

- A. 2
- B. 3
- C. 10
- D. 100

Answer: B

Q387. What does it mean for a number to be composite?

- A. It has no factors
- B. It has factors other than 1 and itself
- C. It is always even
- D. It is the same as prime

Answer: B

Q388. What is the formula for the area of a rectangle?

- A. Length + Width
- B. Length x Width
- C. Length / Width
- D. $2 \times (\text{Length} + \text{Width})$

Answer: B

Q389. What is the result of 0 factorial (0!)?

- A. 0
- B. 1
- C. Undefined
- D. Infinity

Answer: B

Q390. If a set has 4 elements, how many subsets does it have?

- A. 4
- B. 8
- C. 16
- D. 24

Answer: C

Q391. Why is it important for an algorithm to have a clear stopping condition?

- A. It is not important
- B. Without a stopping condition, the algorithm may run forever and never produce a result
- C. It makes the algorithm slower
- D. Stopping conditions are only for loops

Answer: B

Q392. What does it mean to trace through an algorithm?

- A. Drawing the algorithm
- B. Manually following each step with specific input values to verify the result
- C. Deleting the algorithm
- D. Translating to a different language

Answer: B

Q393. What is the purpose of selection sort?

- A. To search for an element
- B. To sort by repeatedly finding the minimum element and placing it in its correct position
- C. To merge two arrays
- D. To compress data

Answer: B

Q394. What does it mean for an algorithm to be correct?

- A. It runs fast
- B. It produces the expected output for every valid input
- C. It uses the least memory
- D. It has no comments

Answer: B

Q395. What is a swap operation in sorting algorithms?

- A. Deleting two elements
- B. Exchanging the positions of two elements in a collection
- C. Adding new elements
- D. Comparing without moving

Answer: B

Q396. Why is it useful to consider the best, worst, and average cases of an algorithm?

- A. Only worst case matters
- B. Different cases reveal how the algorithm performs under various conditions, guiding practical decisions
- C. Best case is always the same as worst case
- D. Average case is meaningless

Answer: B

Q397. What is the key idea behind insertion sort?

- A. Dividing the array in half
- B. Building a sorted portion one element at a time by inserting each new element into its correct position
- C. Finding the maximum first
- D. Using recursion

Answer: B

Q398. What is a step-by-step procedure for solving a problem called?

- A. A guess
- B. An algorithm
- C. A hypothesis
- D. A theorem

Answer: B

Q399. What does 'divide and conquer' mean in algorithm design?

- A. Giving up on the problem
- B. Breaking a problem into smaller sub-problems, solving each, then combining the results
- C. Only solving half the problem
- D. Dividing the team into groups

Answer: B

Q400. Why is choosing the right data structure important for algorithm efficiency?

- A. Data structures do not affect efficiency
- B. The right data structure can dramatically reduce the time and space an algorithm requires
- C. All data structures perform the same
- D. It only matters for large companies

Answer: B

Q401. Why is it important to question assumptions in problem solving?

- A. Assumptions should never be questioned
- B. Unexamined assumptions may be wrong and lead to incorrect solutions
- C. Questioning slows down work
- D. Only experts question assumptions

Answer: B

Q402. What does it mean to think critically about a source of information?

- A. Believing everything you read
- B. Evaluating the credibility, bias, and evidence behind the information
- C. Ignoring all sources
- D. Only trusting one source

Answer: B

Q403. What is the difference between a fact and an opinion?

- A. They are the same
- B. A fact is verifiable and objective; an opinion is a personal judgment or interpretation
- C. Facts are always wrong
- D. Opinions are always right

Answer: B

Q404. Why should you consider alternative explanations when solving a problem?

- A. Alternatives waste time
- B. The first explanation may not be correct; considering alternatives helps find the true cause
- C. There is always only one explanation
- D. Alternatives create confusion

Answer: B

Q405. What is skepticism in the context of critical thinking?

- A. Rejecting everything
- B. A healthy questioning attitude that demands evidence before accepting claims
- C. Being negative about everything
- D. Never making decisions

Answer: B

Q406. Why is it important to consider the consequences of a decision?

- A. Consequences do not matter
- B. Understanding potential outcomes helps make better decisions and avoid unintended harm
- C. Only immediate results matter
- D. Consequences are unpredictable so should be ignored

Answer: B

Q407. What does it mean to evaluate the strength of an argument?

- A. Measuring physical strength
- B. Assessing whether the premises adequately support the conclusion
- C. Counting the number of words
- D. Checking the author's credentials only

Answer: B

Q408. Why is considering different perspectives valuable in critical thinking?

- A. Different perspectives are confusing
- B. They can reveal blind spots, biases, and aspects of the problem you might have missed
- C. Only one perspective is ever correct
- D. It wastes time

Answer: B

Q409. What role does evidence play in critical thinking?

- A. Evidence is unnecessary
- B. Evidence provides the factual basis for claims and helps distinguish justified beliefs from unjustified ones
- C. Only opinions matter
- D. Evidence always proves things completely

Answer: B

Q410. What is a hasty generalization?

- A. A well-supported conclusion
- B. Drawing a broad conclusion from too few examples or insufficient evidence
- C. A type of algorithm
- D. A fast way to solve problems

Answer: B

Q411. What is the first thing you should do when you encounter a bug?

- A. Rewrite the entire program
- B. Try to reproduce the bug consistently
- C. Ignore it and hope it goes away
- D. Delete the code

Answer: B

Q412. What information does an error message typically provide?

- A. Nothing useful
- B. The type of error, the location in code, and often a description of what went wrong
- C. The solution to the bug
- D. The name of the developer who caused it

Answer: B

Q413. Why is reading the error message carefully the best first step in debugging?

- A. Error messages are always misleading
- B. They often tell you exactly what went wrong and where, saving significant debugging time
- C. They are too technical to understand
- D. Only experts can read them

Answer: B

Q414. What is the difference between a compile-time error and a runtime error?

- A. They are the same
- B. Compile-time errors are caught before the program runs; runtime errors occur during execution
- C. Runtime errors are easier to find
- D. Compile-time errors are more dangerous

Answer: B

Q415. Why should you test your fix after correcting a bug?

- A. Testing is unnecessary
- B. To verify the fix actually resolves the bug without introducing new problems
- C. Testing wastes time
- D. The fix is always correct

Answer: B

Q416. What is the purpose of adding print statements for debugging?

- A. To make the program faster
- B. To display variable values and execution flow at specific points to understand program behavior
- C. To make the code permanent
- D. Print statements cannot help with debugging

Answer: B

Q417. What does 'isolating a bug' mean?

- A. Deleting the bug
- B. Narrowing down the specific section of code that causes the problem
- C. Moving the bug to another file
- D. Isolating the computer from the network

Answer: B

Q418. Why is version control helpful for debugging?

- A. It has nothing to do with debugging
- B. It lets you compare current code with working versions to identify what change introduced the bug
- C. It makes code slower
- D. It only helps with team coordination

Answer: B

Q419. What is a common cause of an 'undefined variable' error?

- A. Using too much memory
- B. Trying to use a variable that has not been declared or initialized before use
- C. The variable is too large
- D. The computer is too slow

Answer: B

Q420. Why is it helpful to simplify the test case when debugging?

- A. Simpler tests are meaningless
- B. A minimal test case removes irrelevant complexity, making it easier to identify the exact cause of the bug
- C. Complex tests always find more bugs
- D. Simplification hides bugs

Answer: B

Q421. Why is it important to understand how an algorithm's performance changes with larger inputs?

- A. It does not matter
- B. Poor scaling means an algorithm that works for small inputs may become impractically slow for large ones
- C. All algorithms scale the same
- D. Performance never changes with input size

Answer: B

Q422. What does $O(n)$ mean in simple terms?

- A. The algorithm is broken
- B. The running time grows proportionally to the input size
- C. The algorithm always takes n seconds
- D. It means the algorithm is optimal

Answer: B

Q423. Which is faster for large n : $O(n)$ or $O(n^2)$?

- A. $O(n^2)$
- B. $O(n)$
- C. They are equal
- D. It depends on the computer

Answer: B

Q424. What does it mean when an algorithm has $O(1)$ space complexity?

- A. It uses one byte
- B. It uses a fixed amount of memory regardless of input size
- C. It uses no memory
- D. It is the fastest algorithm

Answer: B

Q425. Why does doubling the input size quadruple the time for an $O(n^2)$ algorithm?

- A. It does not quadruple
- B. Because $(2n)^2 = 4n^2$, so doubling input leads to four times the operations
- C. All algorithms behave this way
- D. This only applies to sorting

Answer: B

Q426. What is the complexity of looking up a value in a hash table on average?

- A. $O(n)$
- B. $O(\log n)$
- C. $O(1)$
- D. $O(n^2)$

Answer: C

Q427. What does it mean for an algorithm to be 'efficient'?

- A. It uses the most resources
- B. It solves the problem using reasonable time and space resources relative to the input size
- C. It always runs in $O(1)$
- D. Efficiency is subjective and unmeasurable

Answer: B

Q428. What is the time complexity of iterating through each element of an array once?

- A. $O(1)$
- B. $O(\log n)$
- C. $O(n)$
- D. $O(n^2)$

Answer: C

Q429. Why does Big-O notation focus on the fastest-growing term?

- A. Other terms are incorrect
- B. For large inputs, the fastest-growing term dominates the running time and other terms become negligible
- C. It is just a convention with no meaning
- D. All terms are equally important

Answer: B

Q430. How many comparisons does binary search need in the worst case for n elements?

- A. n comparisons
- B. $n/2$ comparisons
- C. $\log_2(n)$ comparisons
- D. 1 comparison

Answer: C

Q431. Why is understanding the problem important before writing code?

- A. It wastes development time
- B. Writing code for a misunderstood problem leads to wrong solutions that must be rewritten
- C. Understanding is optional
- D. Code can always be fixed later easily

Answer: B

Q432. What is the purpose of a loop in solving programming problems?

- A. To make code longer
- B. To repeat a block of code a specific number of times or until a condition is met
- C. To define a function
- D. To declare variables

Answer: B

Q433. What is the benefit of breaking code into functions?

- A. It makes programs slower
- B. Functions make code reusable, readable, and easier to test and debug independently
- C. Functions are unnecessary
- D. They increase the file size

Answer: B

Q434. What does it mean to 'hardcode' a value in a program?

- A. To write code that is difficult to understand
- B. To embed a specific value directly in the code rather than using a variable or configuration
- C. To compress the code
- D. To encrypt the value

Answer: B

Q435. Why is choosing meaningful variable names important?

- A. It does not matter what names you use
- B. Meaningful names make the code self-documenting and easier for others to understand
- C. Short names are always better
- D. Only the computer reads variable names

Answer: B

Q436. What is the role of indentation in code readability?

- A. It affects program execution
- B. It visually shows the structure and nesting of code blocks, making logic easier to follow
- C. It is purely decorative
- D. Indentation slows down typing

Answer: B

Q437. What is the difference between = and == in most programming languages?

- A. They are the same
- B. = assigns a value to a variable; == compares two values for equality
- C. == assigns a value
- D. = compares values

Answer: B

Q438. What is an index in the context of arrays?

- A. A table of contents
- B. A number that specifies the position of an element in an array
- C. The size of the array
- D. The data type of the array

Answer: B

Q439. Why is testing code with different inputs important?

- A. Testing is unnecessary
- B. Different inputs may reveal bugs that only appear under specific conditions
- C. One test case is always sufficient
- D. Testing slows down development

Answer: B

Q440. What is a syntax error in programming?

- A. A logical mistake
- B. A violation of the programming language's grammar rules that prevents compilation or interpretation
- C. A runtime crash
- D. An incorrect algorithm

Answer: B

Q441. Why is data important for making good decisions?

- A. Data is always irrelevant
- B. Data provides objective evidence that helps make informed decisions rather than relying on guesswork
- C. Data only confuses people
- D. Intuition is always better than data

Answer: B

Q442. What is the difference between qualitative and quantitative data?

- A. They are the same
- B. Quantitative data is numerical and measurable; qualitative data describes qualities and characteristics
- C. Qualitative data is always better
- D. Quantitative data is not useful

Answer: B

Q443. What is the mode of a dataset?

- A. The average value
- B. The middle value
- C. The most frequently occurring value
- D. The range of values

Answer: C

Q444. What is a line chart typically used to show?

- A. Categorical comparisons
- B. Changes in data over time or continuous trends
- C. Part-to-whole relationships
- D. Frequency distributions

Answer: B

Q445. Why should you check data for errors before analyzing it?

- A. Errors do not affect analysis
- B. Errors in data can lead to incorrect conclusions and bad decisions
- C. Data never has errors
- D. Checking wastes time

Answer: B

Q446. What does 'range' mean in the context of a dataset?

- A. The number of data points
- B. The difference between the maximum and minimum values in the dataset
- C. The average of all values
- D. The most common value

Answer: B

Q447. What is a histogram used for?

- A. Showing relationships between two variables
- B. Showing the frequency distribution of a single continuous variable across intervals
- C. Displaying text data
- D. Showing part-to-whole relationships

Answer: B

Q448. What does it mean to aggregate data?

- A. To delete data
- B. To combine multiple data points into summary values like totals, averages, or counts
- C. To split data into individual records
- D. To encrypt data

Answer: B

Q449. Why is comparing data to a baseline useful?

- A. Baselines are meaningless
- B. A baseline provides a reference point to measure whether current performance is better, worse, or unchanged
- C. Baselines always change
- D. Comparison is not useful

Answer: B

Q450. What is the purpose of labeling axes on a chart?

- A. It is optional decoration
- B. Labels tell the reader what each axis represents and the units of measurement, making the chart interpretable
- C. Labels make charts look cluttered
- D. Only one axis needs a label

Answer: B

Q451. What does 'thinking outside the box' mean in problem solving?

- A. Literally stepping outside a box
- B. Approaching a problem from an unconventional or unexpected angle
- C. Ignoring the problem
- D. Only using standard methods

Answer: B

Q452. Why is asking 'what if' questions valuable in creative thinking?

- A. It wastes time
- B. It opens up new possibilities by imagining different scenarios and challenging current assumptions
- C. It creates confusion
- D. 'What if' questions have no answers

Answer: B

Q453. What is the benefit of generating many ideas before selecting one?

- A. More ideas means more confusion
- B. A larger pool of ideas increases the chance of finding a truly creative and effective solution
- C. Only one idea is ever worth considering
- D. Generating many ideas wastes time

Answer: B

Q454. How does looking at problems from different perspectives help creativity?

- A. It does not help
- B. Different viewpoints reveal aspects of the problem that a single perspective might miss
- C. Only one perspective is ever correct
- D. It complicates the problem

Answer: B

Q455. What is the role of experimentation in creative problem solving?

- A. Experimentation is always risky
- B. Trying out ideas and learning from both successes and failures leads to better solutions
- C. Only proven methods should be used
- D. Experimentation guarantees success

Answer: B

Q456. Why is collaboration important for creative thinking?

- A. Collaboration kills creativity
- B. Different people bring different experiences and ideas, leading to more diverse and innovative solutions
- C. Creativity only comes from working alone
- D. Collaboration slows down thinking

Answer: B

Q457. What does it mean to build on someone else's idea?

- A. Copying their idea exactly
- B. Taking an existing idea and extending, combining, or adapting it to create something new
- C. Stealing intellectual property
- D. Rejecting the idea

Answer: B

Q458. What is the value of sketching or drawing when solving problems creatively?

- A. Drawing is unrelated to problem solving
- B. Visual representations can reveal patterns and relationships that are hard to see in text or numbers alone
- C. Only artists should draw
- D. Sketching wastes time

Answer: B

Q459. Why should you avoid judging ideas too early during brainstorming?

- A. Judging is always helpful
- B. Premature criticism stops the flow of ideas and prevents creative possibilities from being explored
- C. All ideas should be judged immediately
- D. Only good ideas should be shared

Answer: B

Q460. What is an example of creative problem solving in everyday technology?

- A. Using a phone only for calls
- B. Using a smartphone's camera as a document scanner
- C. Using technology exactly as intended
- D. Avoiding all new technology

Answer: B

Q461. Why do real-world problems often have multiple valid solutions?

- A. Real-world problems always have one solution
- B. Different trade-offs between cost, time, quality, and other factors lead to different acceptable solutions
- C. Multiple solutions indicate a poorly defined problem
- D. Only textbook problems have multiple solutions

Answer: B

Q462. What is the importance of understanding the end user in real-world problem solving?

- A. End users are irrelevant
- B. Understanding who will use the solution ensures it meets their actual needs and abilities
- C. Only developers matter
- D. End users always want the same thing

Answer: B

Q463. Why is budget an important constraint in real-world projects?

- A. Budget does not affect projects
- B. Limited resources mean trade-offs must be made about what features and quality levels are feasible
- C. Budget is unlimited in real projects
- D. Only non-profit projects have budgets

Answer: B

Q464. What is the role of testing in real-world software projects?

- A. Testing is unnecessary in the real world
- B. Testing verifies that the software works correctly under various conditions before reaching users
- C. Only developers test software
- D. Testing guarantees zero bugs

Answer: B

Q465. Why is communication important in real-world problem-solving teams?

- A. Communication slows down work
- B. Clear communication ensures everyone understands the problem, plan, and their responsibilities
- C. Only managers need to communicate
- D. Written communication is unnecessary

Answer: B

Q466. What does the term 'deploy' mean in software development?

- A. To delete software
- B. To release and make the software available for actual use by end users
- C. To write documentation
- D. To start coding

Answer: B

Q467. Why might a simpler solution be preferred over a complex one in real-world projects?

- A. Complex solutions are always better
- B. Simpler solutions are easier to build, test, maintain, and less likely to contain bugs
- C. Simplicity means low quality
- D. Complex solutions are always more reliable

Answer: B

Q468. What is user acceptance testing?

- A. Testing by the development team
- B. Testing performed by actual end users to verify the solution meets their real-world needs
- C. Automatic testing by computers
- D. Testing only at the end of the project

Answer: B

Q469. Why is iterating on a solution important in real-world development?

- A. The first version is always perfect
- B. Real-world feedback reveals issues and improvements that cannot be anticipated in advance
- C. Iteration wastes resources
- D. Customers prefer first versions

Answer: B

Q470. What does 'backwards compatibility' mean in software?

- A. Software that runs backward
- B. New versions continue to work with data, interfaces, or systems from older versions
- C. All old features are removed
- D. Only backward-looking analysis

Answer: B

Q471. Why is it important to write clearly when documenting a solution?

- A. Clear writing is unnecessary
- B. Clear documentation helps others understand and maintain the solution without needing to ask the original author
- C. Only the author needs to understand the docs
- D. Documentation is never read

Answer: B

Q472. What should a good code comment explain?

- A. What the code does line by line
- B. Why the code exists or why a particular approach was chosen
- C. The author's name only
- D. Nothing; comments are unnecessary

Answer: B

Q473. What is the purpose of a project README file?

- A. To list all bugs
- B. To provide an overview of the project, how to set it up, and how to use it
- C. To store passwords
- D. READMEs have no purpose

Answer: B

Q474. Why should error messages be written in clear, plain language?

- A. Error messages do not matter
- B. Clear error messages help users understand what went wrong and what they can do to fix the issue
- C. Technical jargon is always better
- D. Error messages should be hidden

Answer: B

Q475. What is the benefit of using consistent formatting in documentation?

- A. Formatting is irrelevant
- B. Consistent formatting makes documents easier to scan, read, and navigate, helping readers find information quickly
- C. It makes documents longer
- D. Only printed documents need formatting

Answer: B

Q476. What should you include when reporting a bug to another developer?

- A. Just say 'it is broken'
- B. Steps to reproduce, expected behavior, actual behavior, and environment details
- C. Only the error message
- D. Your opinion about the code quality

Answer: B

Q477. Why is keeping documentation up to date important?

- A. Outdated docs are fine
- B. Outdated documentation is worse than no documentation because it can mislead users into incorrect actions
- C. Documentation never needs updating
- D. Only initial documentation matters

Answer: B

Q478. What is the difference between internal and external documentation?

- A. They are identical
- B. Internal docs are for the development team; external docs are for end users or clients
- C. Internal docs are secret
- D. External docs are optional

Answer: B

Q479. Why is using examples helpful in documentation?

- A. Examples make docs longer without benefit
- B. Examples show exactly how to use something, making abstract descriptions concrete and actionable
- C. Examples are always misleading
- D. Only API docs need examples

Answer: B

Q480. What is the value of documenting lessons learned from a project?

- A. Past lessons have no value
- B. Documented lessons help future projects avoid the same mistakes and apply successful strategies
- C. Only failures should be documented
- D. Lessons learned are always obvious

Answer: B

Problem Solving & Analytical Skills — Medium

480 questions

Q481. In Polya's method, what comes after devising a plan?

- A. Looking back at the work
- B. Understanding the problem
- C. Carrying out the plan
- D. Restating the problem

Answer: C

Q482. What is divide and conquer?

- A. Only solving the first half of the problem
- B. Solving the problem without breaking it down
- C. Ignoring the most difficult parts of a task
- D. Dividing into smaller independent sub-problems

Answer: D

Q483. Which strategy works backward from the goal?

- A. Trial and error method
- B. Forward chaining approach
- C. Brute force technique
- D. Working backward strategy

Answer: D

Q484. What distinguishes computational thinking?

- A. It cannot involve algorithmic processes
- B. It requires absolutely no human input
- C. It is applicable only to mathematics
- D. Formulating problems for computer solutions

Answer: D

Q485. What is abstraction in problem solving?

- A. Ignoring the problem entirely to move on
- B. Removing unnecessary details to focus
- C. Adding more details to the problem
- D. Making the problem significantly harder

Answer: B

Q486. Which uses previously learned solutions?

- A. Analogical reasoning method
- B. Brute force technique
- C. Random guessing approach
- D. Creative problem solving

Answer: A

Q487. What is a constraint?

- A. An algorithm type to implement
- B. A ready-made solution to use
- C. An optional feature to consider
- D. A limitation that must be satisfied

Answer: D

Q488. What role does evaluation play?

- A. Assesses whether the solution is correct
- B. It only checks for syntax errors
- C. It fully replaces all testing
- D. It is completely unnecessary

Answer: A

Q489. Which technique tries all possible solutions?

- A. Dynamic programming
- B. Greedy algorithm
- C. Heuristic approach
- D. Brute force search

Answer: D

Q490. What is iterative refinement?

- A. Repeatedly improving through cycles
- B. Copying work from other sources
- C. Deleting the solution entirely
- D. Solving the problem only once

Answer: A

Q491. Most important sorting constraint?

- A. Array variable name
- B. Output color format
- C. Input data set size
- D. Code indentation style

Answer: C

Q492. What is root cause analysis?

- A. Addressing surface symptoms only
- B. Ignoring all errors and warnings
- C. Guessing the source of the bug
- D. Identifying the fundamental cause

Answer: D

Q493. Maximum sum subarray is what type?

- A. Optimization problem type
- B. String matching problem
- C. Graph traversal problem
- D. Sorting algorithm problem

Answer: A

Q494. What is a precondition?

- A. Must be true before process begins
- B. An error message shown to users
- C. Checked only after the execution
- D. A specific type of loop construct

Answer: A

Q495. What is a postcondition?

- A. Checked before execution begins
- B. A particular variable type in code
- C. A syntax rule for the compiler
- D. Must be true after process completes

Answer: D

Q496. What does feasibility analysis determine?

- A. Whether a solution is practical and achievable
- B. The color scheme for the interface
- C. The marketing strategy to implement
- D. The number of programmers to hire

Answer: A

Q497. Functional vs non-functional requirements?

- A. Functional: what it does; non-functional: how well
- B. There is no meaningful difference
- C. Functional requirements are always optional
- D. Non-functional requirements are more important

Answer: A

Q498. What is stakeholder identification?

- A. Finding all parties affected by it
- B. Identifying who wrote the code
- C. Counting the team members involved
- D. Listing programming languages used

Answer: A

Q499. What is a use case?

- A. A user-system interaction for a goal
- B. A specific test case for the code
- C. A bug report filed by the team
- D. A comment added to source code

Answer: A

Q500. Role of domain knowledge?

- A. Helps understand context and terms
- B. It only serves for documentation
- C. It slows down the development
- D. It is completely irrelevant

Answer: A

Q501. Contrapositive of If P then Q?

- A. If P then not Q
- B. If not Q then not P
- C. If Q then P
- D. If not P then not Q

Answer: B

Q502. Which is De Morgan's Law?

- A. $\text{NOT } A = A$
- B. $\text{NOT}(A \text{ AND } B) = \text{NOT } A \text{ OR } \text{NOT } B$
- C. $A \text{ AND } B = A \text{ OR } B$
- D. $A \text{ XOR } B = A \text{ AND } B$

Answer: B

Q503. What is modus ponens?

- A. $P \rightarrow Q$, P false, therefore Q false
- B. $P \rightarrow Q$, Q false, therefore P false
- C. $P \rightarrow Q$, Q true, therefore P
- D. $P \rightarrow Q$, P true, therefore Q

Answer: D

Q504. What is modus tollens?

- A. $P \rightarrow Q$, P true, therefore Q
- B. $P \rightarrow Q$, Q true, therefore P
- C. $P \rightarrow Q$, Q false, therefore P false
- D. $P \rightarrow Q$, P false, therefore Q false

Answer: C

Q505. Truth table with 4 variables?

- A. 16 rows
- B. 32 rows
- C. 8 rows
- D. 4 rows

Answer: A

Q506. What is a tautology?

- A. Always true regardless of values
- B. A logical contradiction in reasoning
- C. A statement that is always false
- D. A statement that is sometimes true

Answer: A

Q507. What is a contradiction?

- A. Always false regardless of values
- B. A tautology under all conditions
- C. A valid argument with true premises
- D. A statement that is always true

Answer: A

Q508. What does $P \leftrightarrow Q$ mean?

- A. P implies Q in one direction only
- B. P is true if and only if Q is true
- C. Q implies P in one direction only
- D. Both P and Q are always true together

Answer: B

Q509. What is proof by contradiction?

- A. A type of mathematical induction
- B. Proving the statement is false
- C. A direct proof from axioms
- D. Assume opposite, show contradiction

Answer: D

Q510. Affirming the consequent fallacy?

- A. Rain $\rightarrow$ wet. No rain. Not wet.
- B. Rain $\rightarrow$ wet. Dry. No rain.
- C. Rain $\rightarrow$ wet. Wet. Therefore rained.
- D. Rain $\rightarrow$ wet. Rained. Wet.

Answer: C

Q511. While vs do-while?

- A. While only runs the loop body once
- B. Do-while never actually executes code
- C. While checks first; do-while runs once
- D. They are exactly the same construct

Answer: C

Q512. Connector symbol indicates?

- A. A terminal or endpoint state
- B. Continuation on another page
- C. A decision point in the flow
- D. An input or output operation

Answer: B

Q513. What is a nested loop?

- A. A loop inside another loop
- B. A loop that runs just once
- C. A loop with no condition
- D. A loop that never ends

Answer: A

Q514. for i=1 to 5 executes how many times?

- A. 5 times
- B. 6 times
- C. 0 times
- D. 4 times

Answer: A

Q515. What is an infinite loop?

- A. Condition never becomes false
- B. A very fast loop execution
- C. A loop that runs just twice
- D. A loop with an empty body

Answer: A

Q516. What is a sentinel value?

- A. A random value in the data
- B. Special value signaling end of input
- C. An error code returned by functions
- D. A variable name in the program

Answer: B

Q517. What does break do?

- A. Skips to the next iteration
- B. Restarts the loop from start
- C. Pauses the loop temporarily
- D. Immediately exits the loop

Answer: D

Q518. What does continue do?

- A. Breaks out of all nested loops
- B. Skips rest, goes to next iteration
- C. Restarts the entire program run
- D. Exits the loop entirely now

Answer: B

Q519. What is switch-case for?

- A. Handling errors and exceptions
- B. Selecting blocks based on a value
- C. Iterating through a collection
- D. Managing memory allocation tasks

Answer: B

Q520. $x=0$; while $x<3$: print x ; $x=x+1$. Output?

- A. 1 2 3
- B. 1 2
- C. 0 1 2 3
- D. 0 1 2

Answer: D

Q521. Hex of 255?

- A. 1F
- B. FE
- C. EF
- D. FF

Answer: D

Q522. Two's complement is for?

- A. Floating point values
- B. Signed integers in binary
- C. Data compression formats
- D. Text character encoding

Answer: B

Q523. 8-bit two's complement of -1?

- A. 00000001
- B. 10000001
- C. 11111110
- D. 11111111

Answer: D

Q524. Data abstraction in OOP?

- A. Hide implementation, expose interface
- B. Write longer and more verbose code
- C. Expose all internal details fully
- D. Avoid using functions or methods

Answer: A

Q525. ADT vs data structure?

- A. Data structure is always better to use
- B. ADT: interface; structure: implementation
- C. ADT can never be implemented in code
- D. They are exactly the same thing

Answer: B

Q526. IEEE 754 representation?

- A. Character encoding standard
- B. Sign, exponent, and mantissa
- C. Only the mantissa component
- D. Simple integer format only

Answer: B

Q527. Unicode vs ASCII?

- A. Unicode only supports English letters
- B. They are exactly the same standard
- C. ASCII has more characters than Unicode
- D. Unicode: all scripts; ASCII: 128 chars

Answer: D

Q528. What is a struct?

- A. Composite grouping different types
- B. A single primitive variable only
- C. A file format for data storage
- D. A database table definition format

Answer: A

Q529. What is serialization?

- A. Encrypting sensitive data
- B. Deleting data permanently
- C. Converting to storable format
- D. Sorting data in an order

Answer: C

Q530. Unsigned 8-bit range?

- A. 1-256
- B. 0-127
- C. -128 to 127
- D. 0-255

Answer: D

Q531. What are design patterns?

- A. Random unstructured coding practices
- B. User interface visual elements
- C. Reusable solutions to common problems
- D. Compiler-level code optimizations

Answer: C

Q532. Next: 1,1,2,3,5,8,13,?

- A. 20
- B. 18
- C. 26
- D. 21

Answer: D

Q533. BFS and DFS both being traversals is?

- A. A data entry verification task
- B. A debugging analysis approach
- C. Problem decomposition technique
- D. Pattern recognition and generalization

Answer: D

Q534. Merge sort, quicksort, binary search share?

- A. $O(n)$ time complexity bound
- B. Divide and conquer pattern
- C. Linked list usage only
- D. Hash table data structure

Answer: B

Q535. What is template method pattern?

- A. A C++ template for generics
- B. Base defines skeleton, subclasses override
- C. Writing documentation for the system
- D. Creating document templates for use

Answer: B

Q536. Stacks, queues, priority queues share?

- A. Array-based implementation only
- B. Specific insertion/removal ordering
- C. Random access to all elements
- D. Tree-based structure internally

Answer: B

Q537. What is observer pattern?

- A. Objects notified on state changes
- B. A debugging analysis approach
- C. A sorting algorithm technique
- D. A testing framework structure

Answer: A

Q538. Top-k problems commonly use?

- A. Full array sorting first
- B. Recursion-only based approach
- C. Linked list traversal method
- D. Heap or priority queue structure

Answer: D

Q539. What is sliding window?

- A. Moving window across data efficiently
- B. A sorting algorithm for arrays
- C. A database query optimization trick
- D. A graphical user interface widget

Answer: A

Q540. What is two-pointer technique?

- A. A linked list creation method only
- B. Using two mouse cursors at once
- C. Two references moving through data
- D. A memory management allocation tool

Answer: C

Q541. Euclidean GCD complexity?

- A. $O(a+b)$
- B. $O(a*b)$
- C. $O(\log(\min(a,b)))$
- D. $O(n^2)$

Answer: C

Q542. Subsets of n elements?

- A. 2^n
- B. n^2
- C. n
- D. $n!$

Answer: A

Q543. Pigeonhole Principle?

- A. n items in m slots ($n > m$): one has >1
- B. A specific data structure type
- C. A search algorithm for arrays
- D. A sorting algorithm technique

Answer: A

Q544. Geometric series sum?

- A. $a \cdot n$
- B. $a(r^n - 1)/(r - 1)$
- C. $n \cdot r$
- D. $a + r \cdot n$

Answer: B

Q545. Permutations of n distinct objects?

- A. 2^n
- B. $n!$
- C. n
- D. n^2

Answer: B

Q546. Modular arithmetic useful for?

- A. Only for division operations needed
- B. Cyclic ops, hashing, overflow prevention
- C. Only formatting output display
- D. It has no practical use in CS

Answer: B

Q547. P(heads) fair coin?

- A. 0.5
- B. 1
- C. 0
- D. 0.25

Answer: A

Q548. P(5,3)?

- A. 15
- B. 120
- C. 10
- D. 60

Answer: D

Q549. C(10,3)?

- A. 120
- B. 1000
- C. 720
- D. 30

Answer: A

Q550. Base case importance?

- A. Stops infinite recursion as a condition
- B. It is used only for display output
- C. It makes the formula longer overall
- D. It is optional and unneeded

Answer: A

Q551. Binary search complexity?

- A. $O(1)$
- B. $O(\log n)$
- C. $O(n)$
- D. $O(n^2)$

Answer: B

Q552. Merge sort complexity?

- A. $O(n^2)$
- B. $O(n \log n)$
- C. $O(n)$
- D. $O(\log n)$

Answer: B

Q553. Quicksort worst case?

- A. $O(n \log n)$
- B. $O(n^2)$
- C. $O(\log n)$
- D. $O(n)$

Answer: B

Q554. When greedy over DP?

- A. Always use greedy by default
- B. When the problem is NP-hard
- C. Never use greedy over DP
- D. Greedy-choice property plus optimal substructure

Answer: D

Q555. What is backtracking?

- A. A data compression approach
- B. Never going back on choices
- C. A sorting algorithm technique
- D. Try solutions, undo failed ones

Answer: D

Q556. N-Queens uses?

- A. Greedy approach
- B. Backtracking method
- C. Linear search scan
- D. Dynamic programming

Answer: B

Q557. Dijkstra's is for?

- A. String pattern matching
- B. Balancing binary trees
- C. Shortest paths, non-negative weights
- D. Sorting array elements

Answer: C

Q558. Hash function purpose?

- A. Compressing files for storage
- B. Mapping keys to indices for lookup
- C. Encrypting sensitive user data
- D. Sorting data in an array

Answer: B

Q559. What is stable sort?

- A. Maintains equal elements relative order
- B. A sort using constant space only
- C. A sort that never crashes at all
- D. The fastest sorting algorithm known

Answer: A

Q560. Which sort is not comparison-based?

- A. Heap sort
- B. Counting sort
- C. Quick sort
- D. Merge sort

Answer: B

Q561. What is premature optimization?

- A. Optimizing only after profiling done
- B. A good development practice overall
- C. Optimizing before knowing bottlenecks
- D. Never optimizing the code at all

Answer: C

Q562. What is sunk cost fallacy?

- A. Never investing resources into any project
- B. Continuing failed effort due to past investment
- C. A financial concept only used in accounting
- D. Stopping all projects at the first issue

Answer: B

Q563. What is decision matrix?

- A. A database query optimization method
- B. A matrix multiplication technique
- C. A compilation step in the process
- D. Options scored against weighted criteria

Answer: D

Q564. Root cause vs symptoms?

- A. Fixing symptoms is always better
- B. Root cause fixes underlying problem
- C. Distinguishing them is unnecessary
- D. They are essentially the same thing

Answer: B

Q565. What is cost-benefit analysis?

- A. Comparing expected costs and benefits
- B. A marketing strategy for products
- C. Ignoring all costs of the project
- D. Analyzing only the project costs

Answer: A

Q566. What is SWOT?

- A. Strengths, Weaknesses, Opportunities, Threats
- B. A testing framework for software
- C. A debugging technique for code
- D. A sorting algorithm for datasets

Answer: A

Q567. What is anchoring bias?

- A. Over-relying on first information
- B. An HTML anchor element tag
- C. Avoiding all first impressions
- D. A secure authentication method

Answer: A

Q568. What is risk assessment?

- A. Taking every risk available
- B. Identify, analyze, evaluate risks
- C. Ignoring all potential risks
- D. Avoiding all risks completely

Answer: B

Q569. Why worst cases?

- A. Prepare contingency plans for them
- B. Considering them wastes all time
- C. Just for pessimistic outlook
- D. They never actually happen at all

Answer: A

Q570. What is peer review?

- A. A formal software testing phase
- B. Colleagues examine work for quality
- C. A social media interaction type
- D. A completely unnecessary activity

Answer: B

Q571. What is rubber duck debugging?

- A. A rubber duck shaped mouse device
- B. A testing framework for web apps
- C. A waterfowl-based sorting algorithm
- D. Explaining code aloud to find errors

Answer: D

Q572. What is binary search debugging?

- A. The binary search algorithm itself
- B. Eliminate half the code each step
- C. Debugging binary executable files
- D. Searching binary number patterns

Answer: B

Q573. What is null pointer exception?

- A. Accessing a null reference value
- B. A logic error in computation
- C. A warning from the compiler
- D. A syntax error in the code

Answer: A

Q574. What is off-by-one error?

- A. Loop iterates one too many or few
- B. A network connectivity issue seen
- C. A rounding error in calculations
- D. A compilation error in the build

Answer: A

Q575. What is memory leak?

- A. Physical damage to memory chips
- B. Allocated memory that is never freed
- C. A security hack on the system
- D. Fast access to memory locations

Answer: B

Q576. What is regression testing?

- A. The first test ever written for code
- B. User interface design testing only
- C. Re-testing after changes are made
- D. Performance benchmark testing only

Answer: C

Q577. What is race condition?

- A. A type of loop in the code
- B. A competition between two programs
- C. A fast algorithm execution speed
- D. Depends on concurrent event timing

Answer: D

Q578. What is deadlock?

- A. A locked database table in use
- B. Processes waiting for each other
- C. A system crash due to hardware
- D. An infinite loop in the code

Answer: B

Q579. What is step-through debugging?

- A. Running the entire program at once
- B. Walking through the office halls
- C. One line at a time in debugger
- D. Deleting steps from the program

Answer: C

Q580. What is defensive programming?

- A. Being defensive about your code
- B. Avoiding writing any error handlers
- C. Anticipating and handling errors well
- D. Programming without any planning

Answer: C

Q581. Two nested n-loops?

- A. $O(n^2)$
- B. $O(n)$
- C. $O(2n)$
- D. $O(n \log n)$

Answer: A

Q582. Merge sort space?

- A. $O(n)$
- B. $O(1)$
- C. $O(\log n)$
- D. $O(n^2)$

Answer: A

Q583. Bubble sort best case?

- A. $O(n)$
- B. $O(n^2)$
- C. $O(1)$
- D. $O(n \log n)$

Answer: A

Q584. Big-O vs Omega vs Theta?

- A. O describes the lower bound only
- B. Theta describes the upper bound
- C. They are all the same thing
- D. O: upper, Omega: lower, Theta: tight

Answer: D

Q585. Insert at array beginning?

- A. $O(n^2)$
- B. $O(\log n)$
- C. $O(1)$
- D. $O(n)$

Answer: D

Q586. What is in-place?

- A. A very slow sorting algorithm
- B. $O(1)$ or $O(\log n)$ extra space only
- C. An algorithm that does not work
- D. Uses lots of extra memory space

Answer: B

Q587. Hash table average lookup?

- A. $O(n)$
- B. $O(\log n)$
- C. $O(1)$
- D. $O(n^2)$

Answer: C

Q588. Heap insert/extract?

- A. $O(1)$
- B. $O(n \log n)$
- C. $O(\log n)$
- D. $O(n)$

Answer: C

Q589. Why $O(n \log n)$ efficient for sorting?

- A. It is slower than $O(n^2)$ sorting
- B. Theoretical lower bound for comparisons
- C. It is the slowest sort possible
- D. It only works on small inputs

Answer: B

Q590. Adjacency matrix space?

- A. $O(V^2)$
- B. $O(V+E)$
- C. $O(V)$
- D. $O(E)$

Answer: A

Q591. Modular code benefit?

- A. Smaller testable reusable pieces made
- B. It makes the program run slower
- C. It uses significantly more memory
- D. It makes code much longer overall

Answer: A

Q592. Find max in unsorted array?

- A. Use binary search on the array
- B. Sort the entire array first
- C. Iterate tracking the maximum value
- D. Use a hash table for lookup

Answer: C

Q593. What is variable scope?

- A. The data type of the variable
- B. The current value of the variable
- C. The name given to the variable
- D. The region where it is accessible

Answer: D

Q594. Recursion useful for?

- A. Self-similar sub-problem structures
- B. Replacing all loops in the code
- C. Only Fibonacci number computation
- D. It is always faster than loops

Answer: A

Q595. Count char frequency efficiently?

- A. Use nested loops over the string
- B. Use hash map in one single pass
- C. Sort the string first then count
- D. Use recursion to count each char

Answer: B

Q596. What is exception handling?

- A. A testing framework for the app
- B. Ignoring all errors in the code
- C. Handle runtime errors without crashing
- D. A design pattern for structure

Answer: C

Q597. Reverse linked list?

- A. Sort the list then reverse order
- B. Change next pointers to point back
- C. Copy all elements to an array first
- D. Delete and recreate the whole list

Answer: B

Q598. Check balanced parens?

- A. Hash table use
- B. Stack structure
- C. Array scanning
- D. Queue structure

Answer: B

Q599. Detect cycle in linked list?

- A. Count all the nodes in the list
- B. Sort the list nodes by value
- C. Check all pairs of nodes found
- D. Floyd's slow and fast pointers

Answer: D

Q600. What is memoization?

- A. A debugging technique for errors
- B. Writing memos about the code
- C. Caching results to avoid recomputation
- D. A type of loop construct used

Answer: C

Q601. Data preprocessing?

- A. It is entirely unnecessary to do
- B. Only formatting the output display
- C. Deleting data that is not needed
- D. Clean, transform, and prepare data

Answer: D

Q602. Standard deviation?

- A. The average of the dataset values
- B. The range of the data values
- C. Measure of spread from the mean
- D. The most common value in data

Answer: C

Q603. Outlier detection?

- A. Points significantly different from rest
- B. Sorting the data in ascending order
- C. Counting all the elements in data
- D. Finding the average of all values

Answer: A

Q604. Data normalization?

- A. Deleting abnormal data entries
- B. Making all values exactly equal
- C. Sorting data in a specific order
- D. Scaling values to standard range

Answer: D

Q605. Scatter plot for?

- A. Displaying hierarchical data structures
- B. Showing only time-based trend data
- C. Relationship between two numeric variables
- D. Comparing different categories shown

Answer: C

Q606. Hypothesis testing?

- A. Ignoring all evidence and results
- B. Making assumptions without any data
- C. Statistical method to check a claim
- D. Guessing the answer randomly given

Answer: C

Q607. Feature selection?

- A. Remove all features from the data
- B. Add random features to the dataset
- C. Choose the most relevant attributes
- D. Select all features without filtering

Answer: C

Q608. A/B testing?

- A. Alphabetical sorting of data values
- B. Compare two versions using real data
- C. A debugging technique for finding bugs
- D. Testing all possible combinations made

Answer: B

Q609. Data aggregation?

- A. Splitting data into smaller pieces
- B. Formatting data for the display
- C. Combining data into summary values
- D. Deleting data from the database

Answer: C

Q610. Cross-validation?

- A. Validating data only twice total
- B. Checking data types for correctness
- C. Removing all invalid data entries
- D. Multiple train/test splits for assessment

Answer: D

Q611. What is design thinking?

- A. Testing software for correctness
- B. Only visual design of interfaces
- C. Writing code without any planning
- D. Human-centered iterative problem process

Answer: D

Q612. What is SCAMPER?

- A. A debugging technique for code errors
- B. A testing methodology for the app
- C. Substitute, Combine, Adapt, Modify, etc.
- D. Running away from problems fast

Answer: C

Q613. What is reverse engineering?

- A. Manufacturing physical hardware parts
- B. Analyzing existing to understand it
- C. Destroying the existing system fully
- D. Writing code in reverse order now

Answer: B

Q614. Cross-pollination of ideas?

- A. Avoiding all ideas from other fields
- B. Copying solutions exactly from others
- C. A gardening technique for plants
- D. Applying concepts from one field to another

Answer: D

Q615. What is agile?

- A. Traditional waterfall development method
- B. Iterative cycles with feedback loops
- C. A one-time release without updates
- D. Quick development without any planning

Answer: B

Q616. Constraint-based creativity?

- A. Remove all constraints to create
- B. Limitations serve as creative catalysts
- C. Ignore all constraints in design
- D. Add unnecessary constraints to work

Answer: B

Q617. What is a hackathon?

- A. A security audit of the platform
- B. A hacking attack on the system
- C. A testing marathon for the app
- D. Time-limited creative solution event

Answer: D

Q618. What is biomimicry?

- A. Studying biology exclusively only
- B. Cloning organisms in a laboratory
- C. Medical research on human diseases
- D. Nature's designs for human problems

Answer: D

Q619. Six Thinking Hats?

- A. A fashion design methodology concept
- B. Different thinking perspectives via hats
- C. Wearing six physical hats at once
- D. Limiting team size to six people

Answer: B

Q620. What is rapid prototyping?

- A. Final development of the product
- B. A very slow development approach
- C. Quick rough models for fast testing
- D. Testing only without any building

Answer: C

Q621. What is load balancing?

- A. Managing the project budget funds
- B. Distributing traffic across servers
- C. Balancing items on a desk surface
- D. Reducing the size of source code

Answer: B

Q622. Cold start problem?

- A. Starting a server in cold weather
- B. Database initialization taking too long
- C. No history for new users or items
- D. A computer boot issue on startup

Answer: C

Q623. What is data privacy?

- A. Deleting all data after collection
- B. Only encrypting data at rest here
- C. Making all data publicly available
- D. Handling personal info per standards

Answer: D

Q624. What is fault tolerance?

- A. Ignoring all faults that occur
- B. Continuing operation despite failures
- C. Removing hardware that has faults
- D. Having no faults in the system

Answer: B

Q625. Technical feasibility?

- A. Whether the team is interested
- B. Can build with available resources
- C. Only the project budget matters
- D. Only the team size is relevant

Answer: B

Q626. Backward compatibility?

- A. Only old code is maintained still
- B. Only old hardware is supported now
- C. New version works with old systems
- D. Running programs in reverse order

Answer: C

Q627. MVP approach?

- A. Build all features before releasing
- B. Smallest core feature set to test
- C. Only planning without implementation
- D. Building nothing and just planning

Answer: B

Q628. What is interoperability?

- A. Avoiding all system integrations
- B. Systems operating independently only
- C. Only one system at a time runs
- D. Systems communicating and working together

Answer: D

Q629. What is CI/CD?

- A. A manual deployment process used
- B. Automate build, test, and deploy
- C. A programming language for web
- D. A database management system type

Answer: B

Q630. Disaster recovery?

- A. Only for natural disaster events
- B. Deleting all backups from storage
- C. Strategies to restore after disruption
- D. Ignoring all disruptions that occur

Answer: C

Q631. What is API documentation?

- A. A marketing brochure for the app
- B. Describes endpoints, params, examples
- C. Internal code comments only written
- D. A testing plan for quality checks

Answer: B

Q632. What is technical spec?

- A. A user guide for end-user features
- B. Detailed technical design and architecture
- C. A marketing brochure for clients
- D. A sales pitch for stakeholders

Answer: B

Q633. Inline vs block comments?

- A. Block comments are for single lines
- B. Inline: single line; block: section
- C. Inline comments are always longer
- D. They are exactly the same thing

Answer: B

Q634. What is RFC/design doc?

- A. Change proposal with rationale given
- B. A financial budget document for project
- C. A test plan for quality assurance
- D. A bug report filed by the team

Answer: A

Q635. Documentation-driven development?

- A. Only writing docs for bugs that are found
- B. Avoiding all documentation in the project
- C. Write docs before code to clarify design
- D. Writing docs only after the project ends

Answer: C

Q636. What is UML?

- A. Just drawing pictures for fun only
- B. Only for database schema diagrams
- C. Standardized software modeling notation
- D. Only for creating flowchart diagrams

Answer: C

Q637. Document the why not what?

- A. Never write any comments at all
- B. Comment every single line of code
- C. Explain reasoning, not restate code
- D. Only document variable names used

Answer: C

Q638. What is knowledge base?

- A. A collection of team passwords stored
- B. A textbook for learning to code
- C. A code repository on version control
- D. Repository of solutions for problems

Answer: D

Q639. Version control for docs?

- A. Track changes with revision history
- B. Deleting old versions of documents
- C. Counting the number of documents
- D. Numbering pages in the document

Answer: A

Q640. What is an ADR?

- A. A blueprint for building construction
- B. A server log of all system events
- C. Architectural decision with context noted
- D. A test result from quality checks

Answer: C

Q641. What distinguishes computational thinking from general problem solving?

- A. It focuses on solutions computers can execute
- B. It ignores data and pattern recognition
- C. It avoids using any logical reasoning skills
- D. It only applies to mathematical equations

Answer: A

Q642. In Polya's method, what does looking back involve?

- A. Reviewing and verifying the solution found
- B. Abandoning the solution without any check
- C. Forgetting all steps that were already taken
- D. Starting the entire problem over from scratch

Answer: A

Q643. How does analogical reasoning help in problem solving?

- A. By applying solutions from similar problems
- B. By avoiding any comparison between cases
- C. By creating entirely new unrelated methods
- D. By ignoring all previously solved problems

Answer: A

Q644. What is means-ends analysis in problem solving?

- A. Reducing difference between current and goal
- B. Randomly selecting any available next action
- C. Ignoring the goal and focusing on the start
- D. Solving unrelated problems simultaneously

Answer: A

Q645. Which best describes the role of feedback in solving?

- A. It replaces the need for initial planning
- B. It should always be completely ignored here
- C. It guides adjustments and improvements now
- D. It only matters after final deployment done

Answer: C

Q646. What happens when a problem is ill-defined?

- A. Every possible approach will definitely succeed
- B. The solution is immediately and easily obvious
- C. Goals and constraints are unclear or ambiguous
- D. No analysis or clarification is ever required

Answer: C

Q647. How does top-down design relate to problem solving?

- A. It starts with details then moves to overview
- B. It solves only the bottom-level sub-problems
- C. It avoids any form of hierarchical thinking
- D. It begins with overview then refines details

Answer: D

Q648. What is the benefit of using pseudocode in planning?

- A. It is only used for database query design
- B. It replaces the actual implementation fully
- C. It compiles and runs like real program code
- D. It describes logic without language specifics

Answer: D

Q649. Which cognitive bias can hinder effective problem solving?

- A. Confirmation bias limiting consideration
- B. Careful review of alternative approaches
- C. Systematic analysis of available evidence
- D. Open-minded evaluation of all options

Answer: A

Q650. What is the relationship between creativity and solving?

- A. Creativity replaces the need for all logic
- B. Creativity has no role in problem solving
- C. Creativity enables novel solution approaches
- D. Creativity only matters in artistic fields

Answer: C

Q651. How does stakeholder analysis contribute to understanding?

- A. It identifies who is affected and their needs
- B. It eliminates the need for requirements review
- C. It replaces technical analysis of the problem
- D. It only focuses on the developer preferences

Answer: A

Q652. What is the purpose of creating a problem taxonomy?

- A. Making problems more confusing to understand
- B. Classifying problems by type and characteristics
- C. Eliminating the need for any further analysis
- D. Reducing all problems to a single type only

Answer: B

Q653. Which technique uses 5 Whys for root cause analysis?

- A. Statistical regression analysis of the data
- B. Single question to find immediate surface cause
- C. Iterative questioning to find deeper causes
- D. Random hypothesis generation without any logic

Answer: C

Q654. What is a fishbone diagram used for in analysis?

- A. Mapping potential causes of a specific problem
- B. Displaying the project timeline and milestones
- C. Showing the organizational hierarchy structure
- D. Tracking daily progress of the development team

Answer: A

Q655. How does gap analysis help in problem understanding?

- A. It only measures the size of the code written
- B. It ignores differences between states entirely
- C. It focuses on past states and ignores the goal
- D. It compares current state to the desired state

Answer: D

Q656. What role do acceptance criteria play in analysis?

- A. They replace the need for testing the solution
- B. They are optional and can be safely left out
- C. They only apply to user interface design tasks
- D. They define when the solution is complete enough

Answer: D

Q657. Why distinguish symptoms from causes in analysis?

- A. Causes are never important in problem analysis
- B. Treating symptoms alone leads to recurring issues
- C. Symptoms and causes are always the same thing
- D. Symptoms always reveal the true root cause here

Answer: B

Q658. What is SWOT analysis in problem solving context?

- A. A method for encrypting sensitive user information
- B. Evaluating strengths, weaknesses, opportunities, threats
- C. A type of algorithm for sorting large data arrays
- D. A technique for compressing files to save space

Answer: B

Q659. How does domain knowledge affect analysis quality?

- A. Domain knowledge enables deeper issue understanding
- B. Domain knowledge only matters for simple problems
- C. Domain knowledge always leads to incorrect analysis
- D. Domain knowledge is irrelevant to analysis quality

Answer: A

Q660. What is the purpose of feasibility analysis?

- A. Determining if the proposed solution is achievable
- B. Skipping the analysis and starting implementation
- C. Eliminating all constraints from the problem scope
- D. Proving that every solution is equally feasible

Answer: A

Q661. What is the contrapositive of If it rains then the ground is wet?

- A. If it does not rain then ground is not wet
- B. If the ground is wet then it definitely rained
- C. If ground is not wet then it did not rain
- D. If the ground is wet then it did not rain

Answer: C

Q662. What is modus ponens as a rule of inference?

- A. If P implies Q and P is false then Q is true
- B. If P implies Q and P is true then Q is true
- C. If P implies Q and Q is false then P is true
- D. If P implies Q and Q is true then P is true

Answer: B

Q663. What is the difference between validity and soundness?

- A. Soundness only concerns form not actual truth
- B. Validity concerns form while soundness adds truth
- C. Neither concept is relevant to logical reasoning
- D. They mean exactly the same thing in all cases

Answer: B

Q664. Which logical equivalence is known as De Morgan's Law?

- A. $\text{NOT}(A \text{ OR } B)$ equals $(\text{NOT } A) \text{ OR } (\text{NOT } B)$
- B. $\text{NOT}(A \text{ AND } B)$ equals $(\text{NOT } A) \text{ OR } (\text{NOT } B)$
- C. $\text{NOT}(A \text{ AND } B)$ equals $(\text{NOT } A) \text{ AND } (\text{NOT } B)$
- D. $(A \text{ AND } B)$ equals $\text{NOT}(\text{NOT } A \text{ OR } \text{NOT } B \text{ AND } C)$

Answer: B

Q665. What is proof by contradiction?

- A. Providing a direct proof from given premises here
- B. Assuming opposite is true and showing a conflict
- C. Using examples to demonstrate the statement true
- D. Assuming the conclusion is true then verifying

Answer: B

Q666. What does the exclusive OR operation return?

- A. True only when both inputs are actually true
- B. True when both inputs have the same value
- C. True when exactly one input is true not both
- D. True only when both inputs are actually false

Answer: C

Q667. A valid argument with a false conclusion has what?

- A. At least one false premise in the reasoning
- B. An informal logical fallacy in the reasoning
- C. An inductive argument based on strong evidence
- D. A sound argument with reliable conclusions

Answer: A

Q668. How does propositional logic differ from predicate logic?

- A. Propositional logic uses quantifiers over domains
- B. Predicate logic adds quantifiers and predicates
- C. They are identical in expressive power and scope
- D. Propositional logic is more expressive than predicate

Answer: B

Q669. What is the law of excluded middle in classical logic?

- A. A statement is either true or false no third option
- B. Some statements have no truth value whatsoever
- C. A statement can be both true and false at once
- D. Logic only applies to mathematical statements made

Answer: A

Q670. What is a tautology in propositional logic?

- A. A statement that is always false for all values
- B. A statement that depends on context for truth
- C. A statement that is true for all possible values
- D. A statement that cannot be evaluated at all now

Answer: C

Q671. What is the key difference between while and do-while?

- A. A while loop always executes at least one time
- B. A do-while loop never executes the loop body
- C. A do-while executes body at least once first
- D. Both loops behave in exactly the same manner

Answer: C

Q672. What is a nested loop in algorithm design?

- A. A loop that runs exactly once then stops forever
- B. Two loops running in parallel at the same time
- C. A loop placed inside another loop structure here
- D. A loop that never terminates under any condition

Answer: C

Q673. What does a break statement do inside a loop?

- A. It pauses the loop for a specified time period
- B. It exits the loop immediately when it executes
- C. It restarts the loop from the very beginning
- D. It doubles the speed of the loop iterations

Answer: B

Q674. What is a sentinel value in loop control?

- A. A random value used to start the loop running
- B. A special value that signals end of input data
- C. An error value that crashes the program here
- D. The first value always processed in every loop

Answer: B

Q675. How does switch-case differ from if-else chains?

- A. If-else chains are never used in any programs
- B. Switch-case cannot handle any conditions at all
- C. Switch-case only supports two possible outcomes
- D. Switch-case matches a value against multiple cases

Answer: D

Q676. What is an infinite loop and when is it intentional?

- A. A loop that runs forever, used in event-driven code
- B. A loop that always runs exactly five times at most
- C. A loop that cannot be written in any programming code
- D. A compiler error that prevents program from running

Answer: A

Q677. What does the continue statement do in a loop?

- A. It stops the loop entirely and exits the code
- B. It reverses the direction of the loop execution
- C. It skips remaining body and goes to next iteration
- D. It duplicates the current iteration one extra time

Answer: C

Q678. What is a swim lane diagram used for in flowcharts?

- A. Showing which actor performs each process step
- B. Counting the number of variables in the program
- C. Displaying the color scheme of the user interface
- D. Measuring the execution speed of each algorithm

Answer: A

Q679. How does recursion relate to flow control?

- A. A function calling itself to solve smaller instances
- B. Recursion always runs faster than iterative loops
- C. Recursion never involves any function calls at all
- D. Recursion eliminates the need for base conditions

Answer: A

Q680. What is short-circuit evaluation in conditional flow?

- A. Randomly skipping some conditions during evaluation
- B. Evaluating all conditions regardless of early results
- C. Always evaluating every condition fully before acting
- D. Stopping evaluation when the outcome is determined

Answer: D

Q681. How does a hash table provide efficient data lookup?

- A. It stores data randomly without any structure
- B. It maps keys to indices using a hash function
- C. It searches through every element sequentially
- D. It sorts all data before every lookup operation

Answer: B

Q682. What is the difference between a stack and a queue?

- A. Stack is LIFO while queue follows the FIFO order
- B. Stack is FIFO while queue follows the LIFO order
- C. Both follow LIFO ordering for all their elements
- D. Both follow FIFO ordering for all their elements

Answer: A

Q683. Why is Unicode preferred over ASCII for text?

- A. Unicode only supports the English alphabet letters
- B. Unicode supports characters from all world languages
- C. ASCII already supports every language in the world
- D. Unicode uses fewer bytes than ASCII for every text

Answer: B

Q684. What is a linked list's advantage over an array?

- A. Linked lists use less memory for each stored item
- B. Linked lists provide faster random element access
- C. Linked lists allow efficient insertion and deletion
- D. Linked lists have better cache performance overall

Answer: C

Q685. How does floating-point handle decimal numbers?

- A. It converts all decimals to exact integer values
- B. It uses mantissa and exponent for approximation
- C. It stores decimals with exact infinite precision
- D. It only handles whole numbers with no decimals

Answer: B

Q686. What is data normalization in database design?

- A. Deleting all tables and recreating from scratch
- B. Making all data values uppercase text strings
- C. Organizing data to reduce redundancy and dependency
- D. Duplicating data across many tables for speed

Answer: C

Q687. What is the purpose of an abstract data type?

- A. To eliminate the need for any data structures here
- B. To expose all implementation details to the user
- C. To define behavior independently of implementation
- D. To force one specific implementation approach only

Answer: C

Q688. How does two's complement represent negative integers?

- A. By adding a minus sign before the binary number
- B. By inverting all bits and adding one to the result
- C. By storing the sign in a completely separate byte
- D. By using only positive numbers for all integers

Answer: B

Q689. What is a graph data structure used to represent?

- A. Relationships between entities as nodes and edges
- B. Only linear sequential data ordered by position
- C. A chart displayed on a computer monitor screen
- D. A single variable holding one value at a time

Answer: A

Q690. What is serialization in data representation?

- A. Running multiple processes at the exact same time
- B. Permanently deleting data from all storage systems
- C. Converting data to a format for storage or transfer
- D. Displaying data on a graphical user interface only

Answer: C

Q691. How do design patterns help software developers?

- A. They create unique solutions for every new problem
- B. They eliminate the need for any coding entirely
- C. They only apply to user interface design alone
- D. They provide tested reusable solution templates

Answer: D

Q692. What is the Observer pattern used for?

- A. Notifying multiple objects when a state changes
- B. Ensuring only one instance of a class exists
- C. Preventing any communication between components
- D. Sorting elements in a collection by value order

Answer: A

Q693. How does regular expression pattern matching work?

- A. It only works with numerical data not text at all
- B. It randomly selects text from a given document
- C. It matches text against a defined symbolic pattern
- D. It creates new text rather than matching existing

Answer: C

Q694. What distinguishes the Singleton from other patterns?

- A. It ensures a class has only one global instance
- B. It creates objects without specifying their class
- C. It creates unlimited instances of the given class
- D. It clones existing objects to create new instances

Answer: A

Q695. How does the Strategy pattern enable flexible algorithms?

- A. It hardcodes the algorithm directly into the client
- B. It encapsulates algorithms to make them swappable
- C. It prevents changing algorithms after initial setup
- D. It uses only one algorithm for all situations given

Answer: B

Q696. What role does feature extraction play in recognition?

- A. It adds noise to make pattern recognition more hard
- B. It removes all characteristics from the input data
- C. It identifies the most relevant characteristics used
- D. It ignores all features and uses random selection

Answer: C

Q697. How does MVC pattern separate concerns in design?

- A. It separates model, view, and controller components
- B. It only separates the database from the application
- C. It combines all code into one single monolithic file
- D. It merges user interface with business logic together

Answer: A

Q698. What is anomaly detection in pattern recognition?

- A. Deleting all data that follows expected patterns
- B. Finding patterns that match the expected baseline
- C. Ignoring unusual data to focus on common trends
- D. Identifying data points that deviate from normal

Answer: D

Q699. How does the Factory pattern help manage creation?

- A. It delegates creation to a method or subclass here
- B. It requires knowing the exact class to instantiate
- C. It only works with primitive data types not objects
- D. It prevents any new objects from being created now

Answer: A

Q700. What is clustering in pattern recognition?

- A. Assigning data to already known fixed categories here
- B. Deleting data points that are similar to each other
- C. Spreading data points as far apart as possible now
- D. Grouping similar data points without predefined labels

Answer: D

Q701. How is mathematical induction used to prove statements?

- A. By assuming the conclusion without any proof at all
- B. By testing every natural number one at a time here
- C. By using only examples without any formal reasoning
- D. By proving base case and that step k implies step $k+1$

Answer: D

Q702. What is the relationship between logarithms and exponents?

- A. Logarithms are the inverse operation of exponents
- B. Exponents are a special case of logarithm functions
- C. Logarithms and exponents give the same exact result
- D. They are completely unrelated math operations here

Answer: A

Q703. How does the pigeonhole principle apply to problem solving?

- A. It proves that distribution is always exactly even
- B. If items exceed containers at least one has two items
- C. It states every container will always remain empty
- D. It only applies to problems involving actual pigeons

Answer: B

Q704. What is the significance of base-2 logarithms in CS?

- A. They represent the halving steps in divide-and-conquer
- B. They are never used in algorithm time complexity work
- C. They measure the linear growth of algorithm runtime
- D. They only apply to algorithms with constant runtime

Answer: A

Q705. How do permutations differ from combinations in counting?

- A. Permutations ignore order while combinations respect it
- B. Permutations consider order while combinations do not
- C. Both permutations and combinations always ignore order
- D. Both permutations and combinations always respect order

Answer: B

Q706. What is the purpose of modular arithmetic in CS?

- A. It eliminates the need for division operations here
- B. It replaces all arithmetic with simpler string operations
- C. It handles cyclic computations like hashing and crypto
- D. It only works with even numbers in all calculations

Answer: C

Q707. How does a math function relate to a programming function?

- A. Programming functions never return any output values
- B. Math functions can only handle numerical input data
- C. Math functions have no relation to programming ones
- D. Both map inputs to outputs based on defined rules

Answer: D

Q708. What does Big-O notation describe mathematically?

- A. The minimum possible time any algorithm can achieve
- B. The average case performance of every algorithm run
- C. The exact running time of a given algorithm here
- D. The upper bound growth rate of a function or algorithm

Answer: D

Q709. How are sets useful for solving collection problems?

- A. Sets have no operations defined for them at all here
- B. Sets allow duplicate elements in the collection given
- C. Sets can only contain numerical values nothing more
- D. Sets model unique elements with union and intersection

Answer: D

Q710. What is the significance of Fibonacci in problem solving?

- A. It is a random sequence with no mathematical pattern
- B. It only appears in recreational math puzzles alone
- C. It has no application in computer science at all here
- D. It models growth patterns in nature and algorithms

Answer: D

Q711. How does binary search achieve $O(\log n)$ time?

- A. It halves the search space with each comparison made
- B. It checks every element in the array sequentially
- C. It uses extra memory to store all comparison results
- D. It sorts the array before performing the search here

Answer: A

Q712. What differs between stable and unstable sorting?

- A. Stable sorts are always faster than unstable ones here
- B. Unstable sorts always preserve order of equal elements
- C. Stable sorts preserve relative order of equal elements
- D. The distinction has no practical significance at all

Answer: C

Q713. Why is data structure choice important for efficiency?

- A. All data structures perform exactly the same in practice
- B. Data structures have no effect on algorithm performance
- C. The right structure optimizes time and space complexity
- D. Data structures only affect code readability not speed

Answer: C

Q714. What is the key idea behind dynamic programming?

- A. Storing solutions to subproblems to avoid recomputation
- B. Solving each subproblem repeatedly from scratch always
- C. Using random choices to find an approximate solution
- D. Solving the problem without breaking it into any parts

Answer: A

Q715. How does a hash function contribute to $O(1)$ lookup?

- A. It searches through all elements in sorted order here
- B. It uses binary search on the underlying array structure
- C. It maps keys directly to array indices for fast access
- D. It sorts all elements before every lookup operation run

Answer: C

Q716. What problem does the traveling salesman represent?

- A. Finding the shortest path between two specific nodes
- B. Sorting a list of city names in alphabetical order now
- C. Finding the minimum cost tour visiting all cities once
- D. Counting the number of cities in a given country here

Answer: C

Q717. What is backtracking in algorithmic problem solving?

- A. Always choosing the first available option and staying
- B. Exploring paths and reverting when they lead to failure
- C. Randomly jumping between different solution candidates
- D. Moving forward without ever reconsidering past choices

Answer: B

Q718. How does merge sort achieve guaranteed $O(n \log n)$?

- A. It divides the array and merges sorted halves repeatedly
- B. It randomly shuffles elements until they are sorted now
- C. It uses a single pass through the data without splitting
- D. It compares every pair of elements in the entire array

Answer: A

Q719. What is the purpose of a priority queue?

- A. Randomly selecting elements for processing each time
- B. Processing elements in the order they were inserted
- C. Always processing the highest priority element first
- D. Storing elements without any ordering consideration

Answer: C

Q720. What is the space-time tradeoff in algorithm design?

- A. Space and time always increase together simultaneously
- B. Using more space can reduce time and vice versa here
- C. Space and time are completely independent of each other
- D. Reducing space always also reduces the time required

Answer: B

Q721. How does the straw man fallacy misrepresent arguments?

- A. By distorting the argument to make it easier to attack
- B. By ignoring the argument and discussing another topic
- C. By accurately representing the original argument
- D. By strengthening the argument before responding to it

Answer: A

Q722. What is the appeal to authority fallacy?

- A. Consulting specialists before making important decisions
- B. Citing a relevant expert in the field of discussion
- C. Claiming truth because an authority said so not evidence
- D. Using evidence and data to support a logical argument

Answer: C

Q723. How does the sunk cost fallacy affect decisions?

- A. Past investments irrationally influence future choices
- B. It always leads to the most profitable decisions taken
- C. It encourages ignoring all previous investments made
- D. It helps make rational forward-looking decisions here

Answer: A

Q724. What differs between correlation and causation?

- A. Correlation and causation are the same concept exactly
- B. Correlation shows relationship while causation shows cause
- C. Correlation always proves one thing causes another
- D. Causation never requires any evidence to establish here

Answer: B

Q725. How does Occam's Razor guide critical thinking?

- A. It eliminates the need for any evidence or reasoning
- B. It says the most complex explanation is usually best
- C. The simplest adequate explanation is usually preferred
- D. It requires considering only one possible explanation

Answer: C

Q726. What is a false dilemma in argumentation?

- A. Offering too many options making the choice impossible
- B. Presenting all available options for fair evaluation
- C. Presenting only two options when more actually exist
- D. Accurately describing the limited available options here

Answer: C

Q727. How does the Dunning-Kruger effect impact solving?

- A. It ensures accurate self-assessment of all abilities
- B. It has no effect on problem solving or decision making
- C. Low ability individuals may overestimate their competence
- D. High ability individuals always overestimate their skills

Answer: C

Q728. What is the bandwagon fallacy?

- A. Rejecting an idea because too many people accept it
- B. Forming conclusions based on strong evidence and logic
- C. Assuming something is true because many people believe it
- D. Carefully evaluating popular opinions before acceptance

Answer: C

Q729. How does devil's advocate improve decision making?

- A. By deliberately arguing against to expose weaknesses
- B. By preventing any discussion of alternative viewpoints
- C. By always agreeing with the group without questioning
- D. By accepting the first proposed solution immediately

Answer: A

Q730. What differs between deductive and inductive strength?

- A. Deductive arguments are certain while inductive probable
- B. Inductive are always certain while deductive are not
- C. Neither type of argument provides any certainty at all
- D. Both provide the same level of certainty in all cases

Answer: A

Q731. What is rubber duck debugging?

- A. A proprietary tool made by a company called RubberDuck
- B. Randomly changing code until the error goes away now
- C. Using an actual rubber duck to fix code errors here
- D. Explaining code line-by-line to find errors by talking

Answer: D

Q732. How does binary search debugging help isolate a bug?

- A. It halves the suspicious code region with each test
- B. It only works when there is exactly one bug present
- C. It requires sorting the code alphabetically first here
- D. It tests every line of code one at a time in order

Answer: A

Q733. What is a regression bug?

- A. A bug that fixes itself without any developer action
- B. A bug introduced by changes that broke working features
- C. A bug that only appears on the first run of program
- D. A bug that exists in the very first version of code

Answer: B

Q734. What is the purpose of logging in error analysis?

- A. Logging replaces the need for any testing procedures
- B. Logging slows programs and provides no real value
- C. Logging is only useful during initial development here
- D. Logging records events and data for post-mortem analysis

Answer: D

Q735. How do assertions help catch bugs early?

- A. They only work in production not development environments
- B. Assertions have no effect on catching bugs at all here
- C. They validate assumptions and fail fast when violated
- D. Assertions slow the program without providing any value

Answer: C

Q736. What is a memory leak and how is it detected?

- A. A fast and efficient use of all available system memory
- B. Unreleased memory growing over time detected by profiling
- C. A type of syntax error caught during code compilation
- D. A physical leak of liquid from computer hardware here

Answer: B

Q737. How do unit and integration testing differ for bugs?

- A. Both test exactly the same scope of functionality here
- B. Unit tests are always more comprehensive than integration
- C. Unit tests check components while integration checks interactions
- D. Integration tests replace the need for any unit testing

Answer: C

Q738. How does version control help in debugging?

- A. Version control only stores the latest version of code
- B. It automatically fixes all bugs without developer action
- C. Version control has no relationship to debugging at all
- D. It allows comparing changes to identify when bugs appeared

Answer: D

Q739. What is a null pointer exception?

- A. A valid operation that always succeeds without any issue
- B. An error from accessing a reference that points to nothing
- C. An efficient way to manage memory in any program here
- D. A type of optimization that improves program performance

Answer: B

Q740. What role does code review play in preventing bugs?

- A. Reviews should only happen after bugs are found in code
- B. Code review only focuses on formatting not logic errors
- C. Peers examine code to catch errors before deployment now
- D. Code review is a waste of time and finds no bugs here

Answer: C

Q741. What is the time complexity of binary search?

- A. $O(1)$ constant time regardless of the array size
- B. $O(\log n)$ logarithmic time by halving search space
- C. $O(n)$ linear time for the search operation given
- D. $O(n^2)$ quadratic time for comparing element pairs

Answer: B

Q742. What is the best-case time complexity of quicksort?

- A. $O(1)$ when the array has only a single element in
- B. $O(n \log n)$ when pivots divide the array evenly
- C. $O(n^2)$ when the pivot is always the worst choice
- D. $O(n)$ when the array is already sorted in order

Answer: B

Q743. How do hash table operations degrade with many collisions?

- A. Collisions have no effect on the performance at all
- B. Performance improves significantly with collisions
- C. Hash tables never experience any collisions at all
- D. Operations degrade from $O(1)$ average to $O(n)$ worst

Answer: D

Q744. What is the space complexity of recursion with depth d ?

- A. Space complexity is unaffected by recursion depth
- B. $O(1)$ constant space regardless of recursion depth
- C. $O(d)$ due to the call stack frames stored in memory
- D. $O(n^2)$ quadratic space for all recursive algorithms

Answer: C

Q745. Why is $O(2^n)$ considered impractical for large inputs?

- A. $O(2^n)$ grows slower than $O(n)$ for all input sizes
- B. Exponential algorithms are the most efficient overall
- C. Exponential growth makes computation time infeasible
- D. It is actually faster than linear for large inputs

Answer: C

Q746. What is the complexity of inserting at linked list head?

- A. $O(n^2)$ because the entire list must be reorganized
- B. $O(n)$ because all elements must be shifted forward
- C. $O(\log n)$ because binary search finds insertion point
- D. $O(1)$ because only pointer adjustments are needed

Answer: D

Q747. How does matrix multiplication complexity impact design?

- A. Standard matrix multiplication is $O(n^3)$ which is costly
- B. Matrix multiplication has $O(1)$ constant time complexity
- C. Matrix multiplication complexity has no practical impact
- D. Standard matrix multiplication is always $O(n)$ time

Answer: A

Q748. What does amortized analysis tell us about operations?

- A. Every single operation has the exact same cost always
- B. The average cost per operation over many operations here
- C. Amortized analysis ignores all expensive operations now
- D. Only the most expensive operation matters for analysis

Answer: B

Q749. What is the relationship between $O(n \log n)$ and $O(n^2)$?

- A. $O(n^2)$ grows slower than $O(n \log n)$ for large inputs
- B. Both grow at exactly the same rate for all input sizes
- C. $O(n \log n)$ is always worse than $O(n^2)$ in every case
- D. $O(n \log n)$ grows slower than $O(n^2)$ for large inputs

Answer: D

Q750. How does choosing list versus set affect complexity?

- A. Lists always have better performance than sets entirely
- B. Sets offer $O(1)$ lookup while lists require $O(n)$ search
- C. Lists and sets have identical performance for all tasks
- D. The choice between list and set never affects performance

Answer: B

Q751. How does modular programming improve problem solving?

- A. It prevents any code reuse across different programs
- B. It puts all code in one file for simplicity of access
- C. It breaks solutions into independent reusable modules
- D. It eliminates the need for functions in the code base

Answer: C

Q752. What advantage does recursion have for tree problems?

- A. Recursion always uses less memory than iterative ways
- B. Tree structures naturally match recursive decomposition
- C. Recursion cannot handle any tree structure at all
- D. Trees can only be traversed using iterative approaches

Answer: B

Q753. How does exception handling improve robustness?

- A. It prevents all errors from occurring in the program
- B. It makes programs run faster by skipping error checks
- C. It gracefully manages errors without crashing program
- D. It eliminates the need for any input validation checks

Answer: C

Q754. What benefit do dictionaries provide for lookup problems?

- A. Dictionaries enable fast $O(1)$ average key-value lookups
- B. Dictionaries require sorting before any lookup is done
- C. Dictionaries can only store numerical data values now
- D. Dictionaries provide slow key-value pair access speed

Answer: A

Q755. How does the DRY principle improve maintainability?

- A. By making code as long and verbose as possible always
- B. By avoiding all abstraction to keep code very simple
- C. By duplicating code in many places for redundancy
- D. By eliminating repetition through abstraction and reuse

Answer: D

Q756. What role does algorithm selection play in solving?

- A. Any algorithm works equally well for every problem
- B. Choosing the right algorithm significantly impacts speed
- C. The fastest algorithm is the one with the most code
- D. Algorithm selection has no effect on execution at all

Answer: B

Q757. How does test-driven development improve quality?

- A. Writing tests first clarifies requirements before coding
- B. TDD eliminates the need for any code review process
- C. Tests are written after the code is fully complete
- D. TDD makes programs run faster than other approaches

Answer: A

Q758. What advantage do list comprehensions provide?

- A. They express transformations concisely and readably
- B. They prevent any filtering of data from the lists
- C. They make code longer and harder to understand here
- D. They always run slower than traditional loop methods

Answer: A

Q759. How do enumerations improve code clarity?

- A. They make code harder to understand and maintain now
- B. They define named constants for a set of related values
- C. They prevent any constants from being used in code
- D. They replace named constants with magic numbers here

Answer: B

Q760. What is the purpose of input validation?

- A. To accept all input without any checking whatsoever
- B. To reject all input regardless of its correctness here
- C. To slow down the program by adding unnecessary checks
- D. To ensure input meets expected format before processing

Answer: D

Q761. How does hypothesis testing help in data-driven solving?

- A. It replaces the need for any analysis of the data given
- B. It always confirms whatever hypothesis was proposed here
- C. It eliminates the need for any data collection at all
- D. It validates or rejects claims based on statistical evidence

Answer: D

Q762. What is the purpose of data visualization?

- A. To make data harder to understand and interpret here
- B. To store data more efficiently on the computer disk now
- C. To replace the need for any numerical analysis entirely
- D. To reveal patterns and insights through visual representation

Answer: D

Q763. What differs between qualitative and quantitative data?

- A. Quantitative data describes qualities not numerical values
- B. Qualitative describes qualities while quantitative measures
- C. Qualitative data can only be collected through surveys
- D. They are exactly the same type of data in all cases

Answer: B

Q764. How does A/B testing help in data-driven decisions?

- A. It eliminates the need for any control group in testing
- B. It compares two variants to determine which performs better
- C. It randomly assigns outcomes without measuring results
- D. It tests only one version without any comparison done

Answer: B

Q765. What is a key performance indicator used for?

- A. Tracking the number of emails received by the team
- B. Counting the number of lines of code written daily
- C. Measuring the physical weight of computer hardware
- D. Measuring progress toward specific business objectives

Answer: D

Q766. What is the purpose of exploratory data analysis?

- A. To skip all analysis and go directly to implementation
- B. To replace the need for any hypothesis testing process
- C. To discover patterns and anomalies before formal analysis
- D. To confirm predetermined conclusions without question

Answer: C

Q767. How does data normalization help when combining sources?

- A. It deletes data that does not match a specific format
- B. It scales data to a common range for fair comparison
- C. It prevents any data from multiple sources combined
- D. It makes data from different sources incompatible here

Answer: B

Q768. What role does statistical significance play in decisions?

- A. It has no bearing on the reliability of analysis results
- B. It guarantees the result will apply to all future cases
- C. It indicates results are likely due to chance variation
- D. It indicates results are unlikely due to chance alone

Answer: D

Q769. What is a data pipeline?

- A. A type of chart used for data visualization purposes
- B. A one-time manual data entry process for single use
- C. An automated flow for collecting and processing data
- D. A physical pipe that carries data through wires here

Answer: C

Q770. How does sampling bias affect conclusion validity?

- A. Sampling bias has no effect on conclusion validity here
- B. Sampling bias improves the accuracy of all conclusions
- C. Non-representative samples lead to misleading conclusions
- D. Only very large samples can have any sampling bias now

Answer: C

Q771. How does SCAMPER stimulate creative problem solving?

- A. It only works for product design not software solutions
- B. It provides a single fixed solution for every problem
- C. It uses prompts like substitute combine adapt to explore
- D. It eliminates the need for any creative thinking effort

Answer: C

Q772. What is design thinking's approach to creative solving?

- A. It skips understanding users and goes to building
- B. It eliminates the need for any user research or input
- C. It centers on empathy ideation prototyping and testing
- D. It follows a rigid linear process without iteration

Answer: C

Q773. How does cross-domain thinking enhance solutions?

- A. It only works when both domains are exactly identical
- B. It applies concepts from other fields to the problem
- C. It limits thinking to only one specific domain area
- D. It prevents any knowledge transfer between domains

Answer: B

Q774. What is rapid prototyping in creative problem solving?

- A. Avoiding prototypes and going to full production now
- B. Building a final polished product on the first attempt
- C. Spending years perfecting a prototype before testing
- D. Quickly creating a rough model to test ideas and learn

Answer: D

Q775. How does fail fast support creative innovation?

- A. Quick failures provide early learning to redirect effort
- B. Failing fast always results in permanent project failure
- C. It encourages building perfect solutions on first try
- D. It means giving up after the first minor setback here

Answer: A

Q776. What is divergent thinking and when is it most useful?

- A. It replaces the need for any convergent thinking here
- B. Narrowing down to a single solution immediately given
- C. Generating many different ideas during idea exploration
- D. It is only useful at the very end of problem solving

Answer: C

Q777. How does constraint-based creativity use limitations?

- A. Removing all constraints is the only path to creativity
- B. Constraints make every solution worse without exception
- C. Limitations force exploration of non-obvious solutions
- D. Constraints always prevent any creative thinking at all

Answer: C

Q778. What is the role of incubation in creative solving?

- A. It requires constant active focus without any breaks
- B. Incubation wastes time that should be spent working now
- C. Stepping away allows subconscious processing to happen
- D. It only works for artistic problems not technical ones

Answer: C

Q779. How does yes and from improv enhance creative collaboration?

- A. It requires saying no to all suggestions made by others
- B. It encourages rejecting every idea from team members
- C. It prevents any collaboration between team members now
- D. It builds on others ideas rather than dismissing them

Answer: D

Q780. What is reverse brainstorming?

- A. It prevents any new ideas from being generated at all
- B. It generates ideas in the usual forward direction only
- C. It asks how to cause the problem to find solutions
- D. It only works for problems that have been solved before

Answer: C

Q781. How does agile address uncertainty in projects?

- A. It eliminates all uncertainty through upfront planning
- B. It embraces change through iterative incremental delivery
- C. It prevents any changes once the project has started here
- D. It requires complete specifications before any work starts

Answer: B

Q782. What is technical debt and why does it matter?

- A. Shortcuts now that increase future maintenance effort
- B. It is a financial obligation to purchase new hardware
- C. Technical debt always improves long-term code quality
- D. It has no effect on future development speed at all

Answer: A

Q783. How does risk management improve project outcomes?

- A. It identifies and mitigates potential problems proactively
- B. It replaces the need for any contingency planning done
- C. Risk management is only needed for very large projects
- D. It eliminates all risks from the project completely

Answer: A

Q784. What is the minimum viable product approach?

- A. Launching with core features to validate and iterate
- B. Waiting until every feature is perfect before release
- C. Delivering a product with no features at all to users
- D. Building the most feature-complete product on first try

Answer: A

Q785. How does scope creep affect project delivery?

- A. Scope creep always improves project quality and timeline
- B. Adding features always speeds up project completion time
- C. Scope creep has no negative effects on any project work
- D. Uncontrolled additions can delay delivery and increase cost

Answer: D

Q786. Why is documentation important in projects?

- A. Documentation is never read by anyone on the team here
- B. Only large enterprise projects need any documentation now
- C. Documentation slows development without any benefits here
- D. It preserves knowledge and enables maintenance onboarding

Answer: D

Q787. How does continuous integration improve quality?

- A. By preventing any changes to the code after initial release
- B. By frequently merging and testing changes automatically
- C. By requiring manual testing of every change by the team
- D. By integrating code changes only at the end of project

Answer: B

Q788. What role does prioritization play in problem solving?

- A. Prioritization prevents any tasks from being completed
- B. All tasks are equally important and should be done at once
- C. It focuses resources on the most impactful tasks first
- D. Only the easiest tasks should be prioritized for completion

Answer: C

Q789. How does UX design relate to real-world problem solving?

- A. It replaces the need for any functional requirements here
- B. It ensures solutions are usable and meet user needs well
- C. UX design only matters for games not business applications
- D. User experience is unrelated to problem solving entirely

Answer: B

Q790. What is the purpose of a retrospective?

- A. To plan the next five years of the project in detail
- B. To assign blame for any mistakes made during sprint
- C. To celebrate success without identifying improvements
- D. To reflect on what worked and what needs improvement

Answer: D

Q791. How do architecture decision records improve decisions?

- A. They document context, decision, and consequences here
- B. They prevent any future decisions from being made
- C. They replace the need for any team discussions given
- D. They are only useful for very large enterprise projects

Answer: A

Q792. What is the benefit of diagrams in technical communication?

- A. They convey complex relationships visually and concisely
- B. Diagrams make documentation more confusing than text
- C. They replace the need for any written explanation text
- D. Diagrams are only useful for non-technical audiences

Answer: A

Q793. How does a runbook improve operational problem solving?

- A. It provides step-by-step procedures for common scenarios
- B. Runbooks are only needed for new team members starting
- C. It replaces the need for any operational knowledge
- D. It eliminates the possibility of any operational errors

Answer: A

Q794. What is the purpose of a project status report?

- A. To communicate progress, risks, and blockers transparently
- B. To replace all other forms of project communication
- C. To document only successes and never mention issues
- D. To hide problems from stakeholders and management

Answer: A

Q795. How does API documentation help developers?

- A. API documentation is only needed for internal use here
- B. They describe endpoints, parameters, and expected behaviors
- C. They replace the need for any testing of integrations
- D. API docs prevent any integration from being possible

Answer: B

Q796. What is the value of post-mortem documents?

- A. They are only required by policy not for any learning
- B. To capture lessons learned and prevent future recurrence
- C. Post-mortems have no value for improving reliability
- D. To assign blame to the person who caused incident

Answer: B

Q797. How does a changelog help track software evolution?

- A. It summarizes notable changes between software versions
- B. They only list bug fixes never mentioning new features
- C. It replaces the need for any version control system
- D. Changelogs make software harder to use and understand

Answer: A

Q798. What is the purpose of a design document?

- A. To propose and discuss a technical approach before building
- B. Design documents are never read by anyone on the team
- C. To record the final code implementation line by line
- D. To list all the technologies the developer knows well

Answer: A

Q799. How does structured commit formatting improve communication?

- A. It makes commit messages longer without adding value
- B. Only the first line of commit messages matters at all
- C. It enables automated changelog generation and tracking
- D. Commit format has no effect on project communication

Answer: C

Q800. What role does technical writing play for non-experts?

- A. It eliminates the need for any diagrams or visual aids
- B. Technical writing should only target expert audiences
- C. It translates complex concepts into accessible explanations
- D. Technical writing makes content harder to understand

Answer: C

Q801. How does the generate-and-test strategy work in problem solving?

- A. It avoids generating any solutions
- B. It creates possible solutions and tests each against the requirements
- C. It only uses the first solution found
- D. It eliminates testing entirely

Answer: B

Q802. What is the difference between a problem and a puzzle in computer science?

- A. They are identical concepts
- B. A problem has practical significance while a puzzle is typically recreational with a known solution
- C. Puzzles are harder than problems
- D. Problems have no solutions

Answer: B

Q803. Which problem-solving strategy involves starting from the desired outcome and working backward?

- A. Forward chaining
- B. Working backward
- C. Random search
- D. Brute force

Answer: B

Q804. How does collaboration improve the problem-solving process?

- A. It slows everything down
- B. Different perspectives and expertise can lead to better solutions
- C. It creates more problems
- D. Collaboration is never useful

Answer: B

Q805. What is the role of mental models in problem solving?

- A. They make thinking slower
- B. They provide simplified representations of how things work to guide reasoning
- C. They are only useful in psychology
- D. They prevent creative solutions

Answer: B

Q806. Why might simplifying assumptions be useful when solving complex problems?

- A. They always lead to wrong answers
- B. They reduce complexity to make the problem tractable while preserving key features
- C. They eliminate the need for analysis
- D. They are only used by beginners

Answer: B

Q807. What distinguishes algorithmic problem solving from heuristic problem solving?

- A. They are the same thing
- B. Algorithmic guarantees a correct solution; heuristic provides a good-enough solution faster
- C. Heuristic always gives the best answer
- D. Algorithmic is always faster

Answer: B

Q808. What is the importance of recognizing problem types in computer science?

- A. It has no importance
- B. Knowing the problem type helps select appropriate solving strategies and algorithms
- C. It only matters for exams
- D. It makes problems harder

Answer: B

Q809. How does the concept of incremental development relate to problem solving?

- A. It means solving the entire problem at once
- B. It involves building a solution step by step, testing at each stage
- C. It skips the planning phase
- D. It only applies to large teams

Answer: B

Q810. What is the relationship between problem representation and solution efficiency?

- A. Representation has no effect
- B. A good representation can make the solution much easier to find
- C. Only one representation exists for any problem
- D. Representation only affects readability

Answer: B

Q811. How does impact analysis help in understanding a proposed change?

- A. It has no value
- B. It identifies what parts of the system will be affected by a change and assesses the risks
- C. It only measures performance
- D. It replaces testing

Answer: B

Q812. What is the purpose of creating a problem decomposition tree?

- A. To make the problem look bigger
- B. To break a complex problem into a hierarchy of smaller, manageable sub-problems
- C. To avoid solving the problem
- D. It is only decorative

Answer: B

Q813. How do non-functional requirements differ from functional requirements in analysis?

- A. They are identical
- B. Functional requirements define what the system does; non-functional define how well it performs
- C. Non-functional requirements are optional
- D. Functional requirements are less important

Answer: B

Q814. What is the purpose of creating a context diagram during analysis?

- A. To draw pretty pictures
- B. To show the system boundary and its interactions with external entities
- C. To write code
- D. To replace all documentation

Answer: B

Q815. Why is it important to distinguish between causes and symptoms in problem analysis?

- A. They are always the same
- B. Fixing symptoms without addressing root causes leads to recurring problems
- C. Symptoms do not exist
- D. Causes are always obvious

Answer: B

Q816. How does benchmarking contribute to problem analysis?

- A. It has no contribution
- B. It provides reference points by comparing against known standards or competitors
- C. It slows down analysis
- D. It only applies to hardware

Answer: B

Q817. What is a decision table used for in problem analysis?

- A. Tracking expenses
- B. Mapping combinations of conditions to their corresponding actions systematically
- C. Storing data
- D. Scheduling tasks

Answer: B

Q818. How does Pareto analysis help prioritize problems?

- A. It solves all problems equally
- B. It identifies the vital few causes that produce the majority of effects
- C. It ignores important issues
- D. It only works for manufacturing

Answer: B

Q819. What is the role of traceability in requirements analysis?

- A. It is unnecessary overhead
- B. It links requirements to their sources, tests, and implementations for accountability
- C. It slows development
- D. It only matters for documentation

Answer: B

Q820. How does an affinity diagram help organize analysis findings?

- A. It replaces all other diagrams
- B. It groups related ideas and findings into clusters to reveal themes and patterns
- C. It is only for brainstorming
- D. It creates confusion

Answer: B

Q821. What is the law of contraposition and how is it applied?

- A. If P then Q is equivalent to If Q then P
- B. If P then Q is logically equivalent to If not Q then not P
- C. It reverses all logical operators
- D. It only applies to mathematics

Answer: B

Q822. What distinguishes deductive reasoning from abductive reasoning?

- A. They are identical
- B. Deductive derives certain conclusions from premises; abductive infers the best explanation from observations
- C. Abductive is always wrong
- D. Deductive is faster

Answer: B

Q823. What is the difference between necessary and sufficient conditions?

- A. They mean the same thing
- B. A necessary condition must be true for the conclusion; a sufficient condition guarantees the conclusion
- C. Necessary conditions are optional
- D. Sufficient conditions are always necessary

Answer: B

Q824. What is the fallacy of false cause (post hoc ergo propter hoc)?

- A. Correctly identifying causation
- B. Assuming that because one event followed another, the first caused the second
- C. A valid form of reasoning
- D. A mathematical proof technique

Answer: B

Q825. How does reductio ad absurdum work as a proof technique?

- A. It proves something by example
- B. It assumes the opposite of what you want to prove and shows this leads to a contradiction
- C. It avoids contradictions
- D. It only works for simple statements

Answer: B

Q826. What is the closed-world assumption in logical reasoning?

- A. The world is physically closed
- B. Anything not known or stated to be true is assumed to be false
- C. Everything is assumed to be true
- D. It applies only to databases

Answer: B

Q827. What is the difference between strong and weak inductive arguments?

- A. There is no difference
- B. A strong inductive argument makes the conclusion very probable; a weak one does not
- C. Strong arguments are always deductive
- D. Weak arguments are always fallacious

Answer: B

Q828. How does the principle of explosion apply in inconsistent logical systems?

- A. It prevents contradictions
- B. From a contradiction, any statement can be derived as true
- C. It destroys the computer
- D. It only applies to explosives

Answer: B

Q829. What is the relationship between logical equivalence and biconditional?

- A. They are unrelated
- B. Two statements are logically equivalent if and only if their biconditional is a tautology
- C. Biconditional is always false
- D. Equivalence means one implies the other

Answer: B

Q830. What is the difference between material and strict implication?

- A. They are identical
- B. Material implication is truth-functional; strict implication requires a necessary connection
- C. Strict implication is weaker
- D. Material implication requires causation

Answer: B

Q831. What is the difference between pre-test and post-test loops?

- A. There is no difference
- B. Pre-test loops check the condition before each iteration; post-test loops check after, guaranteeing at least one execution
- C. Pre-test loops are faster
- D. Post-test loops never terminate

Answer: B

Q832. How does a flag variable control flow in an algorithm?

- A. It draws a flag on screen
- B. It is a Boolean variable that signals when a specific condition has been met, altering subsequent flow
- C. It stores numerical data
- D. It replaces all loops

Answer: B

Q833. What is an accumulator pattern in algorithm design?

- A. A hardware component
- B. A variable that collects a running total or result through repeated updates in a loop
- C. A type of flowchart symbol
- D. A debugging technique

Answer: B

Q834. How does exception handling affect the normal flow of a program?

- A. It has no effect
- B. When an exception occurs, normal flow is interrupted and control transfers to an exception handler
- C. It speeds up execution
- D. It prevents all errors

Answer: B

Q835. What is the purpose of a guard clause in function flow?

- A. To protect against hackers
- B. To handle edge cases early by returning or throwing before the main logic executes
- C. To make code longer
- D. To add complexity

Answer: B

Q836. How does a state machine manage complex flow control?

- A. It uses random states
- B. It defines a finite set of states and transitions between them based on inputs or events
- C. It eliminates all states
- D. It only works for games

Answer: B

Q837. What is the difference between iteration and recursion for solving the same problem?

- A. They always produce different results
- B. Iteration uses loops while recursion uses self-calling functions; both can solve the same problems
- C. Recursion is always better
- D. Iteration cannot solve recursive problems

Answer: B

Q838. What is the role of an initialization step before a loop?

- A. It is unnecessary
- B. It sets up variables to their starting values so the loop operates correctly from the first iteration
- C. It makes loops slower
- D. It only applies to for loops

Answer: B

Q839. How does early termination improve algorithm efficiency?

- A. It does not help
- B. It stops processing as soon as the answer is found, avoiding unnecessary work on remaining data
- C. It makes algorithms incorrect
- D. It only works for sorting

Answer: B

Q840. What is a structured programming theorem and what does it guarantee?

- A. It is a math theorem only
- B. Any computable function can be implemented using only sequence, selection, and iteration
- C. It requires goto statements
- D. It is outdated and irrelevant

Answer: B

Q841. How does a priority queue differ from a regular queue?

- A. They are identical
- B. In a priority queue, elements are dequeued based on priority rather than insertion order
- C. Priority queues are always slower
- D. Regular queues use priorities too

Answer: B

Q842. What is the trade-off between using an array versus a linked list?

- A. There is no trade-off
- B. Arrays provide $O(1)$ random access but expensive insertion; linked lists provide $O(1)$ insertion but $O(n)$ access
- C. Arrays are always better
- D. Linked lists are always better

Answer: B

Q843. What problem does data compression solve?

- A. It makes data larger
- B. It reduces the size of data for efficient storage and transmission while preserving information
- C. It deletes data permanently
- D. It only works on images

Answer: B

Q844. How does a tree data structure organize hierarchical data?

- A. In a flat list
- B. With a root node and child nodes forming a branching hierarchy where each node has at most one parent
- C. In a circular pattern
- D. Randomly without structure

Answer: B

Q845. What is the purpose of a hash function in data representation?

- A. To encrypt data
- B. To map data of arbitrary size to fixed-size values for efficient lookup and comparison
- C. To sort data alphabetically
- D. To compress data

Answer: B

Q846. How does JSON represent structured data?

- A. In binary format
- B. As human-readable text using key-value pairs, arrays, and nested objects
- C. As machine code
- D. In XML format only

Answer: B

Q847. What is the difference between a set and a list as data structures?

- A. They are identical
- B. A set contains unique unordered elements while a list maintains order and allows duplicates
- C. Sets are always larger
- D. Lists cannot store numbers

Answer: B

Q848. Why might a circular buffer be preferred over a regular array for streaming data?

- A. It is always slower
- B. It efficiently reuses fixed memory by overwriting the oldest data when full, ideal for constant-size windows
- C. It uses infinite memory
- D. It cannot store data

Answer: B

Q849. How does an enum (enumeration) type improve data representation?

- A. It makes code slower
- B. It defines a fixed set of named values, making code more readable and preventing invalid values
- C. It replaces all data types
- D. It is only for numbers

Answer: B

Q850. What is the purpose of a schema in data representation?

- A. To make data larger
- B. To define the structure, types, and constraints of data, ensuring consistency and validation
- C. It is only decorative
- D. Schemas are unnecessary

Answer: B

Q851. How does the Decorator pattern extend object behavior without modifying existing code?

- A. It replaces the original class
- B. It wraps objects with additional functionality layers while maintaining the same interface
- C. It deletes existing behavior
- D. It only works with strings

Answer: B

Q852. What is the difference between structural and behavioral design patterns?

- A. There is no difference
- B. Structural patterns deal with object composition; behavioral patterns deal with communication between objects
- C. Structural patterns are always better
- D. Behavioral patterns are outdated

Answer: B

Q853. How does the sliding window pattern optimize problems involving contiguous subarrays?

- A. It processes all subarrays independently
- B. It maintains a window that slides through the data, updating results incrementally rather than recomputing from scratch
- C. It only works for sorting
- D. It makes problems slower

Answer: B

Q854. What common pattern connects depth-first search, backtracking, and recursive tree traversal?

- A. They all use queues
- B. They all explore paths fully before backtracking, using a stack-like mechanism
- C. They are unrelated algorithms
- D. They all require sorted data

Answer: B

Q855. How does the Iterator pattern decouple traversal from data structure implementation?

- A. It does not decouple anything
- B. It provides a uniform interface for accessing elements sequentially regardless of the underlying data structure
- C. It only works with arrays
- D. It replaces all loops

Answer: B

Q856. What pattern do problems like finding duplicates, counting frequencies, and grouping share?

- A. Sorting pattern
- B. Hash-based lookup pattern using dictionaries or hash maps
- C. Recursive pattern
- D. Divide and conquer

Answer: B

Q857. How does the Adapter pattern solve interface incompatibility?

- A. It deletes one interface
- B. It wraps an existing interface to make it compatible with a different expected interface
- C. It requires rewriting both interfaces
- D. It only works in Java

Answer: B

Q858. What is the two-pointer pattern and when is it typically applied?

- A. Using two mouse pointers
- B. Using two indices that traverse a sorted array from different positions to find pairs or subarrays
- C. It only works for strings
- D. It requires extra memory

Answer: B

Q859. How does the publish-subscribe pattern differ from the observer pattern?

- A. They are identical in every way
- B. Publish-subscribe uses a message broker for decoupling while observer has direct subject-observer coupling
- C. Observer is always better
- D. Publish-subscribe cannot handle events

Answer: B

Q860. What pattern connects binary search, bisection method, and git bisect?

- A. Linear search pattern
- B. Divide-and-eliminate pattern that halves the search space at each step
- C. Random search pattern
- D. Brute force pattern

Answer: B

Q861. How does the inclusion-exclusion principle count the size of a union of sets?

- A. It adds all set sizes
- B. It adds individual sizes, subtracts pairwise intersections, adds triple intersections, and so on
- C. It only counts the largest set
- D. It ignores overlaps

Answer: B

Q862. What is the sum formula for an arithmetic series with n terms, first term a and last term l ?

- A. $n * a$
- B. $n * (a + l) / 2$
- C. $a * l / n$
- D. $(a + l) * n$

Answer: B

Q863. Why is the floor function useful in algorithm analysis?

- A. It rounds numbers up
- B. It gives the largest integer less than or equal to a value, useful for integer division in index calculations
- C. It is only for architecture
- D. It has no use in CS

Answer: B

Q864. What does it mean for a function to grow exponentially?

- A. It grows by adding a constant
- B. Its growth rate is proportional to its current value, resulting in doubling over fixed intervals
- C. It stays the same
- D. It decreases over time

Answer: B

Q865. How is the binomial coefficient $C(n,k)$ computed using factorials?

- A. $n! / k!$
- B. $n! / (k! * (n-k)!)$
- C. $n * k$
- D. $(n+k)!$

Answer: B

Q866. What is the significance of the number e (Euler's number) in mathematics?

- A. It is just a random constant
- B. It is the base of natural logarithms and appears in growth, decay, probability, and calculus
- C. It equals exactly 3
- D. It is only used in physics

Answer: B

Q867. How does proof by strong induction differ from regular mathematical induction?

- A. They are identical
- B. Strong induction assumes the statement holds for all values up to k , not just k , when proving $k+1$
- C. Strong induction is weaker
- D. Regular induction is more powerful

Answer: B

Q868. What is the mathematical principle behind hashing with modular arithmetic?

- A. It has no mathematical basis
- B. Mapping keys to indices using $h(k) = k \bmod m$ distributes keys across m buckets, where prime m reduces clustering
- C. It requires floating-point math
- D. It only works for strings

Answer: B

Q869. Why are logarithms important in analyzing divide-and-conquer algorithms?

- A. They are not important
- B. Each halving of the problem adds one step, so algorithms that halve the input have $O(\log n)$ depth
- C. Logarithms make algorithms slower
- D. They only apply to sorting

Answer: B

Q870. What is a recurrence relation in algorithm analysis?

- A. A loop in code
- B. An equation that defines a function in terms of its value at smaller inputs, used to describe recursive algorithm complexity
- C. A fixed formula
- D. A type of data structure

Answer: B

Q871. How does memoization improve the performance of recursive algorithms?

- A. It makes recursion deeper
- B. It caches results of expensive function calls and returns the cached result when the same inputs recur
- C. It eliminates recursion
- D. It increases memory but not speed

Answer: B

Q872. What is the key difference between BFS and DFS in terms of the data structure used?

- A. Both use stacks
- B. BFS uses a queue (FIFO) while DFS uses a stack (LIFO)
- C. Both use queues
- D. They use the same structure

Answer: B

Q873. What makes a greedy algorithm 'greedy'?

- A. It uses the most memory
- B. It makes the locally optimal choice at each step without reconsidering previous choices
- C. It always finds the global optimum
- D. It tries every possibility

Answer: B

Q874. How does counting sort achieve $O(n)$ time complexity?

- A. It compares elements pairwise
- B. It counts occurrences of each value and uses the counts to place elements directly, avoiding comparisons
- C. It divides the array in half
- D. It is only $O(n)$ for small inputs

Answer: B

Q875. What is the advantage of using a heap for finding the k smallest elements in a stream?

- A. Heaps are slower than sorting
- B. A max-heap of size k maintains the k smallest elements seen so far with $O(\log k)$ per insertion
- C. Heaps require sorting first
- D. There is no advantage

Answer: B

Q876. What problem does Dijkstra's algorithm solve and what is its key limitation?

- A. It sorts arrays but is slow
- B. It finds shortest paths from a source vertex but requires all edge weights to be non-negative
- C. It works with negative weights
- D. It finds the longest path

Answer: B

Q877. How does the concept of invariants help prove algorithm correctness?

- A. Invariants are not useful
- B. An invariant is a condition that holds before and after each iteration; proving it holds throughout proves correctness
- C. They only apply to mathematics
- D. They make algorithms slower

Answer: B

Q878. What is the difference between a deterministic and a non-deterministic algorithm?

- A. They are the same
- B. A deterministic algorithm always produces the same output for the same input; a non-deterministic one may explore multiple paths simultaneously
- C. Non-deterministic is always wrong
- D. Deterministic is always faster

Answer: B

Q879. How does the two-phase approach of map-reduce enable distributed algorithm design?

- A. It only works on single machines
- B. Map distributes work by applying a function to each element; reduce aggregates results, enabling parallelism
- C. Map and reduce must be sequential
- D. It replaces all algorithms

Answer: B

Q880. What is the significance of algorithm stability in sorting?

- A. It means the algorithm never crashes
- B. A stable sort preserves the relative order of elements with equal keys, which matters when sorting by multiple criteria
- C. Stability is irrelevant
- D. All sorts are stable

Answer: B

Q881. How does the Socratic method promote critical thinking?

- A. It provides answers directly
- B. It uses systematic questioning to challenge assumptions, expose contradictions, and deepen understanding
- C. It avoids all questions
- D. It is only for philosophy

Answer: B

Q882. What is the difference between deductive and inductive reasoning in critical analysis?

- A. They are identical
- B. Deductive moves from general principles to specific conclusions with certainty; inductive moves from specific observations to general conclusions with probability
- C. Inductive is always correct
- D. Deductive is always wrong

Answer: B

Q883. What is the appeal to emotion fallacy?

- A. Using emotions to enhance understanding
- B. Manipulating emotional responses instead of providing logical arguments to support a claim
- C. All emotions are fallacious
- D. Emotions and logic are the same

Answer: B

Q884. How does steel-manning differ from straw-manning an argument?

- A. They are the same technique
- B. Steel-manning presents the strongest version of an opponent's argument; straw-manning misrepresents it as weaker
- C. Steel-manning is always wrong
- D. Straw-manning is more honest

Answer: B

Q885. Why is the distinction between correlation and causation important in data interpretation?

- A. They mean the same thing
- B. Two things occurring together (correlation) does not mean one causes the other; confusing them leads to false conclusions
- C. Correlation always implies causation
- D. Causation never involves correlation

Answer: B

Q886. What is the role of falsifiability in evaluating claims?

- A. It makes claims false
- B. A claim that cannot be tested or potentially disproven is not scientifically meaningful
- C. All claims are falsifiable
- D. Falsifiability weakens claims

Answer: B

Q887. How does the principle of charity improve critical discourse?

- A. It involves donating money
- B. It interprets others' arguments in the strongest reasonable way before critiquing, leading to more productive dialogue
- C. It weakens your position
- D. It is only for debates

Answer: B

Q888. What is the red herring fallacy in argumentation?

- A. A type of fish
- B. Introducing an irrelevant topic to divert attention from the original argument
- C. A strong counterargument
- D. A valid reasoning technique

Answer: B

Q889. How does systems thinking enhance critical analysis of complex problems?

- A. It simplifies everything to one cause
- B. It examines how components interact within a larger system, revealing feedback loops and unintended consequences
- C. It ignores relationships between parts
- D. It only applies to engineering

Answer: B

Q890. What is the difference between necessary and sufficient evidence?

- A. They are identical
- B. Necessary evidence must be present but may not be enough alone; sufficient evidence is enough to support the conclusion
- C. Necessary evidence is optional
- D. Sufficient evidence is always necessary

Answer: B

Q891. How does git bisect help locate the commit that introduced a bug?

- A. It searches every commit linearly
- B. It uses binary search through commit history, testing the midpoint to halve the search space each time
- C. It only works for merge conflicts
- D. It requires manual inspection of every commit

Answer: B

Q892. What is the difference between a bug caused by incorrect logic versus incorrect data?

- A. They are the same thing
- B. Logic bugs are errors in the algorithm or control flow; data bugs are caused by unexpected or corrupt input data
- C. Data bugs do not exist
- D. Logic bugs are always harder

Answer: B

Q893. How do assertions help catch bugs during development?

- A. They slow programs down without benefit
- B. They check that expected conditions hold at runtime, failing immediately with a clear message when violated
- C. They replace all testing
- D. They only work in production

Answer: B

Q894. What is the watchpoint feature in a debugger used for?

- A. Watching a clock
- B. It pauses execution when a specific variable's value changes, helping track unexpected modifications
- C. It watches network traffic
- D. It monitors CPU usage

Answer: B

Q895. Why is logging often preferred over print debugging in production systems?

- A. Logging is slower
- B. Logging provides structured, level-based output that can be filtered, persisted, and analyzed without code changes
- C. Print statements are always better
- D. Logging adds no value over print

Answer: B

Q896. What is a stack overflow error and what typically causes it?

- A. Too much data in an array
- B. Excessive recursion or very deep call chains that exhaust the available call stack memory
- C. Writing too many functions
- D. Using too many variables

Answer: B

Q897. How does defensive programming help prevent bugs?

- A. It makes code defensive against hackers only
- B. It validates inputs, checks preconditions, and handles edge cases explicitly, catching errors before they propagate
- C. It slows development without benefit
- D. It replaces all testing

Answer: B

Q898. What is the purpose of a core dump in debugging?

- A. To delete core files
- B. It captures the memory state of a program at the time of a crash, allowing post-mortem analysis
- C. To improve performance
- D. To backup the database

Answer: B

Q899. How does pair debugging improve bug-finding efficiency?

- A. It doubles the number of bugs
- B. Two people bring different perspectives and can catch errors that a single person might overlook due to familiarity
- C. It always slows things down
- D. It only works for beginners

Answer: B

Q900. What is the significance of boundary value analysis in testing and debugging?

- A. Boundaries are not important
- B. Bugs frequently occur at the edges of input ranges, so testing boundary values (0, max, min, empty) catches common errors
- C. It only applies to numerical data
- D. Boundary values are always correct

Answer: B

Q901. What is the difference between Big-O and Big-Theta notation?

- A. They are the same
- B. Big-O gives an upper bound; Big-Theta gives a tight bound that is both upper and lower
- C. Big-O is always tighter
- D. Big-Theta is less precise

Answer: B

Q902. Why is $O(n \log n)$ considered an efficient complexity class for sorting?

- A. It is the slowest possible
- B. It is the theoretical lower bound for comparison-based sorting, so no comparison sort can do better
- C. It is the same as $O(n)$
- D. It is worse than $O(n^2)$

Answer: B

Q903. How does space complexity differ from auxiliary space complexity?

- A. They are identical
- B. Space complexity includes input space; auxiliary space measures only the extra memory the algorithm uses beyond the input
- C. Auxiliary space is always larger
- D. They measure different things entirely

Answer: B

Q904. What is the time complexity of searching an unsorted array for a specific element?

- A. $O(1)$
- B. $O(\log n)$
- C. $O(n)$
- D. $O(n^2)$

Answer: C

Q905. How does the time complexity of building a heap compare to sorting?

- A. Building a heap is $O(n \log n)$
- B. Building a heap is $O(n)$ which is faster than the $O(n \log n)$ required for comparison-based sorting
- C. They are both $O(n)$
- D. Building a heap is slower than sorting

Answer: B

Q906. What is the worst-case complexity of a hash table lookup?

- A. $O(1)$ always
- B. $O(n)$ when all elements hash to the same bucket
- C. $O(\log n)$
- D. $O(n^2)$

Answer: B

Q907. Why is $O(2^n)$ considered intractable for even moderate input sizes?

- A. It is actually very fast
- B. Exponential growth means even $n=40$ requires over a trillion operations, far exceeding practical limits
- C. It only seems slow on old hardware
- D. It is the same as $O(n^2)$

Answer: B

Q908. How does memoization change the complexity of the naive recursive Fibonacci?

- A. It has no effect
- B. It reduces the complexity from $O(2^n)$ to $O(n)$ by caching previously computed results
- C. It increases complexity
- D. It only affects space, not time

Answer: B

Q909. What is the relationship between recursion depth and stack space complexity?

- A. They are unrelated
- B. Each recursive call adds a frame to the call stack, so space complexity is at least $O(\text{depth})$
- C. Recursion uses no stack space
- D. Stack depth is always $O(1)$

Answer: B

Q910. What is the time complexity of matrix multiplication for two $n \times n$ matrices using the standard algorithm?

- A. $O(n)$
- B. $O(n^2)$
- C. $O(n^3)$
- D. $O(n \log n)$

Answer: C

Q911. How does the strategy of pre-sorting input data simplify certain programming problems?

- A. Pre-sorting always wastes time
- B. Sorted data enables binary search, simplifies duplicate detection, and allows efficient two-pointer techniques
- C. Sorting has no benefit for problem solving
- D. Pre-sorting makes problems harder

Answer: B

Q912. What is the advantage of using a stack for matching brackets in an expression?

- A. Stacks are the only data structure
- B. A stack naturally matches each opening bracket with its corresponding closing bracket in correct nesting order
- C. Arrays cannot solve this problem
- D. Stacks make the problem harder

Answer: B

Q913. How does the concept of 'separation of concerns' improve code quality?

- A. It has no effect
- B. Each module handles a single responsibility, making code easier to understand, test, modify, and reuse
- C. It makes code more complex
- D. It only applies to large projects

Answer: B

Q914. What is the benefit of using a hash map for counting frequency of elements?

- A. Hash maps are slower than arrays for counting
- B. Hash maps provide $O(1)$ average-time lookup and update, making frequency counting $O(n)$ for any hashable element type
- C. Arrays are always better for counting
- D. Hash maps cannot count frequencies

Answer: B

Q915. How does refactoring improve code without changing its behavior?

- A. Refactoring changes program output
- B. It restructures existing code for better readability, maintainability, or performance while preserving the same functionality
- C. It is the same as rewriting from scratch
- D. Refactoring introduces bugs

Answer: B

Q916. What is the purpose of edge case testing in programming?

- A. Edge cases are unimportant
- B. It tests boundary conditions and unusual inputs where bugs commonly hide, such as empty inputs, single elements, or maximum values
- C. It only tests happy paths
- D. Edge cases never cause bugs

Answer: B

Q917. How does the builder pattern solve the problem of complex object construction?

- A. It deletes complex objects
- B. It separates the construction of a complex object from its representation, allowing step-by-step construction with validation
- C. It only works for strings
- D. It makes construction harder

Answer: B

Q918. What advantage does recursion provide for tree and graph traversal?

- A. Recursion cannot traverse trees
- B. Recursion naturally maps to the hierarchical structure of trees and graphs, where each recursive call handles a subtree or neighbor
- C. It is always slower than iteration
- D. It only works for binary trees

Answer: B

Q919. How does using a set data structure help solve problems involving duplicates?

- A. Sets allow duplicates
- B. Sets automatically ensure uniqueness, making duplicate detection and removal trivially efficient
- C. Sets are slower than lists for this
- D. Sets cannot store data

Answer: B

Q920. What is the importance of writing clean code for future problem solving?

- A. Clean code is unnecessary
- B. Clean code is easier to understand, modify, and extend, reducing the time needed to solve future problems in the same codebase
- C. Only compilers care about code quality
- D. Clean code is slower

Answer: B

Q921. How does stratified sampling improve the representativeness of a sample?

- A. It uses random selection only
- B. It divides the population into subgroups and samples from each proportionally, ensuring all groups are represented
- C. It selects only from the largest group
- D. It eliminates sampling entirely

Answer: B

Q922. What is the purpose of a box plot (box and whisker diagram)?

- A. To show exact data values
- B. To visualize the distribution's center, spread, skewness, and outliers through quartiles
- C. To display trends over time
- D. To show part-to-whole relationships

Answer: B

Q923. How does the concept of standard deviation complement the mean?

- A. It replaces the mean
- B. The mean shows the center of data while standard deviation shows how spread out the data is around that center
- C. They measure the same thing
- D. Standard deviation is always larger than mean

Answer: B

Q924. What is selection bias and how does it affect data analysis?

- A. It is not a real problem
- B. It occurs when the sample is not representative of the population, leading to conclusions that do not generalize
- C. It only affects large datasets
- D. It improves data quality

Answer: B

Q925. How do dashboards help in data-driven problem solving?

- A. They are purely decorative
- B. They consolidate key metrics and visualizations in one place for quick monitoring, trend identification, and decision support
- C. They replace all data analysis
- D. They only show historical data

Answer: B

Q926. What is the purpose of a control group in an experiment?

- A. To control the experiment's cost
- B. To provide a baseline comparison by not receiving the treatment, allowing measurement of the treatment's true effect
- C. Control groups are unnecessary
- D. It receives all treatments simultaneously

Answer: B

Q927. How does data interpolation differ from data extrapolation?

- A. They are identical
- B. Interpolation estimates values within the range of known data; extrapolation estimates beyond it, with higher uncertainty
- C. Extrapolation is more reliable
- D. Interpolation ignores existing data

Answer: B

Q928. What is the difference between precision and recall in data-driven evaluation?

- A. They measure the same thing
- B. Precision measures accuracy of positive predictions; recall measures completeness of finding all actual positives
- C. Precision is always more important
- D. Recall is always more important

Answer: B

Q929. How does cohort analysis provide insights that aggregate analysis misses?

- A. It is identical to aggregate analysis
- B. It groups users by shared characteristics or time periods and tracks them separately, revealing patterns hidden in aggregate data
- C. It only works for medical data
- D. It requires larger samples

Answer: B

Q930. What role does data provenance play in trusting analysis results?

- A. It is irrelevant
- B. Knowing where data came from, how it was collected, and how it was transformed helps assess its reliability and potential biases
- C. Only the final result matters
- D. Provenance makes data less trustworthy

Answer: B

Q931. How does the technique of random stimulation promote creative thinking?

- A. It produces random solutions
- B. It introduces unrelated concepts or words as stimuli to force new associations and break fixation on familiar approaches
- C. It replaces structured thinking
- D. Random input always confuses

Answer: B

Q932. What is the concept of 'creative constraints' and how do they help?

- A. Constraints always limit creativity
- B. Deliberate limitations force innovative thinking by requiring novel approaches within boundaries
- C. Constraints eliminate all creativity
- D. They only apply to art

Answer: B

Q933. How does the concept of bisociation explain creative breakthroughs?

- A. It is about bicycle associations
- B. It describes creativity as connecting two previously unrelated frames of reference or mental matrices
- C. It only applies to science
- D. It prevents creative thinking

Answer: B

Q934. How does the 'worst possible idea' technique lead to creative solutions?

- A. It produces only bad ideas
- B. Deliberately generating terrible ideas removes the fear of failure and often reveals the inverse as a viable creative solution
- C. It wastes time
- D. It is never used in practice

Answer: B

Q935. What is the difference between incremental and radical innovation?

- A. They are identical
- B. Incremental innovation makes small improvements to existing solutions; radical innovation creates fundamentally new approaches
- C. Incremental is always better
- D. Radical innovation is impossible

Answer: B

Q936. How does role-playing or perspective-taking enhance creative problem solving?

- A. It is only for actors
- B. Adopting different roles forces you to see the problem from unfamiliar viewpoints, generating ideas you would not otherwise consider
- C. It has no cognitive benefit
- D. It always produces wrong answers

Answer: B

Q937. What is the purpose of an innovation sprint or design sprint?

- A. To run fast
- B. To compress the creative process into a short time-boxed period with structured activities for rapid prototyping and validation
- C. Sprints eliminate the need for planning
- D. They replace long-term development

Answer: B

Q938. How does the concept of serendipity contribute to creative discoveries?

- A. It is pure luck with no preparation
- B. Prepared minds recognize unexpected connections or accidental findings as opportunities for creative solutions
- C. It cannot be facilitated
- D. It only happens in laboratories

Answer: B

Q939. How does analogical reasoning across domains drive innovation?

- A. Analogies are always misleading
- B. Solutions from one domain can be adapted to solve problems in a completely different domain through structural similarity
- C. Analogies only work within the same field
- D. They always produce identical solutions

Answer: B

Q940. What is the role of constructive criticism in refining creative ideas?

- A. All criticism destroys creativity
- B. Constructive feedback identifies weaknesses and suggests improvements, helping transform raw creative ideas into viable solutions
- C. Criticism should always be avoided
- D. Only praise is productive

Answer: B

Q941. How does feature flagging help manage risk in real-world deployments?

- A. It adds unnecessary complexity
- B. It allows new features to be deployed but disabled by default, enabling gradual rollout and quick rollback without redeployment
- C. It prevents all deployments
- D. It only works for web applications

Answer: B

Q942. What is the significance of Service Level Agreements (SLAs) in real-world systems?

- A. They are just legal documents with no practical impact
- B. They define measurable commitments for availability, performance, and support, setting clear expectations between providers and consumers
- C. SLAs are only for large companies
- D. They guarantee perfect uptime

Answer: B

Q943. How does the concept of technical debt affect long-term project health?

- A. Technical debt has no long-term effect
- B. Shortcuts taken now accumulate and make future changes increasingly expensive and risky if not addressed
- C. It is always worth taking on
- D. It can be completely avoided

Answer: B

Q944. Why is monitoring important after deploying a solution to production?

- A. Monitoring is unnecessary after deployment
- B. It detects issues, performance degradation, and unexpected behavior before they significantly impact users
- C. Only broken systems need monitoring
- D. Monitoring slows down the system

Answer: B

Q945. How does the strangler pattern help modernize legacy systems?

- A. It shuts down the legacy system immediately
- B. It gradually replaces legacy components with new implementations while both coexist, reducing migration risk
- C. It requires a complete rewrite
- D. It only works for small systems

Answer: B

Q946. What role does capacity planning play in preventing real-world system failures?

- A. Capacity planning is unnecessary
- B. It forecasts future resource needs based on growth projections to ensure the system can handle expected load
- C. It only matters for new systems
- D. Current capacity is always sufficient

Answer: B

Q947. How does A/B testing help make better product decisions?

- A. It tests the alphabet
- B. It compares two versions with real users simultaneously, using statistical analysis to determine which performs better
- C. It requires two separate products
- D. It always gives clear answers

Answer: B

Q948. What is the purpose of an incident response plan?

- A. To assign blame for incidents
- B. To define clear procedures, roles, and communication channels for responding to production issues quickly and effectively
- C. Incidents do not need plans
- D. It replaces prevention measures

Answer: B

Q949. How does the concept of graceful degradation improve user experience during failures?

- A. It lets the system crash gracefully
- B. The system continues operating with reduced functionality rather than failing completely when components malfunction
- C. It prevents all failures
- D. It is the same as zero downtime

Answer: B

Q950. Why is stakeholder management important throughout a project lifecycle?

- A. Stakeholders are only relevant at the start
- B. Continuous stakeholder engagement ensures alignment, manages expectations, and prevents costly surprises from misaligned priorities
- C. Only technical stakeholders matter
- D. It creates unnecessary overhead

Answer: B

Q951. How do Architecture Decision Records (ADRs) improve long-term project understanding?

- A. They are unnecessary bureaucracy
- B. They document the context, options considered, and rationale behind architectural decisions, preserving decision-making context for future teams
- C. They replace all other documentation
- D. They only record final decisions

Answer: B

Q952. What is the difference between a tutorial, a how-to guide, and a reference document?

- A. They serve the same purpose
- B. Tutorials teach through examples; how-to guides provide steps for specific tasks; references provide comprehensive technical details
- C. Only references are needed
- D. Tutorials replace all other docs

Answer: B

Q953. How does the STAR method help communicate problem-solving achievements?

- A. It involves astronomy
- B. It structures communication as Situation, Task, Action, Result to clearly convey the context and impact of problem-solving work
- C. It is only for job interviews
- D. It replaces technical documentation

Answer: B

Q954. Why is audience awareness critical when writing technical documentation?

- A. All audiences need the same information
- B. Different audiences have different knowledge levels, needs, and goals requiring adapted content, depth, and terminology
- C. Technical writing has no audience
- D. Only experts read documentation

Answer: B

Q955. How do diagrams complement textual documentation?

- A. Diagrams are always unnecessary
- B. Diagrams convey relationships, architectures, and flows more clearly than text alone, especially for complex systems
- C. Text is always sufficient
- D. Diagrams replace all text

Answer: B

Q956. What is the purpose of semantic commit messages?

- A. They look more professional
- B. They categorize commits by type (feat, fix, docs, refactor) making history easier to search, understand, and automate changelog generation
- C. They are only for open-source projects
- D. They make commits longer for no reason

Answer: B

Q957. How does a glossary improve technical documentation for diverse audiences?

- A. Glossaries are outdated
- B. A glossary defines domain-specific terms in one place, ensuring consistent understanding across readers with different backgrounds
- C. Everyone knows all technical terms
- D. Glossaries make docs too long

Answer: B

Q958. What is the value of writing a pre-mortem document before starting a project?

- A. Pre-mortems are only for medical contexts
- B. It imagines the project has already failed and identifies potential causes, enabling proactive risk mitigation before problems occur
- C. It is too pessimistic
- D. Post-mortems are sufficient

Answer: B

Q959. How do inline code documentation tools like JSDoc or Javadoc improve developer productivity?

- A. They slow down development
- B. They generate navigable API documentation directly from annotated source code, keeping documentation close to the code it describes
- C. They are only for Java
- D. They replace all testing

Answer: B

Q960. Why is asynchronous communication important for distributed teams?

- A. Asynchronous communication is always inferior
- B. It allows team members in different time zones to contribute at their own pace, with documented discussions that anyone can reference later
- C. All communication must be synchronous
- D. It creates communication gaps

Answer: B

Problem Solving & Analytical Skills — Hard

480 questions

Q961. Shortest path with negative edges (no negative cycles)?

- A. Simple BFS traversal
- B. Bellman-Ford algorithm
- C. Dijkstra's algorithm
- D. DFS with backtracking

Answer: B

Q962. Which paradigm makes locally optimal choices at each step?

- A. Dynamic programming approach
- B. Backtracking technique
- C. Divide and conquer method
- D. Greedy algorithm approach

Answer: D

Q963. What is reduction in problem solving?

- A. Reducing the number of variables used
- B. Transforming into another known problem
- C. Simplifying the program output format
- D. Making smaller by removing input data

Answer: B

Q964. What characterizes an NP-hard problem?

- A. Has absolutely no solution at all
- B. At least as hard as hardest NP problems
- C. Solvable easily in polynomial time
- D. Always unsolvable by any algorithm

Answer: B

Q965. What is optimal substructure?

- A. Only greedy algorithms work on it
- B. Problem cannot be decomposed at all
- C. Sub-problems are fully independent
- D. Optimal solution uses optimal sub-solutions

Answer: D

Q966. Which concept stores sub-problem solutions?

- A. Iterative computation
- B. Memoization technique
- C. Simple recursion method
- D. Compilation process

Answer: B

Q967. Overlapping sub-problems + optimal substructure = which paradigm?

- A. Linear search method
- B. Brute force search
- C. Dynamic programming
- D. Simple recursion only

Answer: C

Q968. Decidable vs undecidable?

- A. An algorithm solves every decidable instance
- B. Decidable problems cannot use any recursion
- C. Undecidable problems have no instances at all
- D. Decidable problems are always harder to solve

Answer: A

Q969. What is state space search?

- A. Querying a database for records
- B. Allocating memory for data storage
- C. Searching through files on a disk
- D. Exploring all states to find a path

Answer: D

Q970. Which problem is undecidable?

- A. Binary search tree lookup
- B. Finding the shortest path
- C. Sorting a list of elements
- D. The Halting Problem itself

Answer: D

Q971. 10K concurrent users, <100ms response is what type?

- A. Usability requirement type
- B. Security requirement type
- C. Functional requirement type
- D. Performance non-functional type

Answer: D

Q972. Mapping to TSP implies?

- A. Likely NP-hard, needs heuristics
- B. Only sorting is required to solve
- C. It can be solved in linear time
- D. The problem is trivially solvable

Answer: A

Q973. Most critical real-time constraint?

- A. Deterministic response time guarantees
- B. Code readability standards
- C. Number of comments in code
- D. Descriptive variable naming style

Answer: A

Q974. Problem space vs solution space?

- A. They are exactly the same thing
- B. Neither concept is used in CS at all
- C. They refer to code versus documentation
- D. Problem: all states; solution: valid answers

Answer: D

Q975. Multiple conflicting constraints need?

- A. Multi-objective optimization approach
- B. Only brute force can work on them
- C. Ignoring some constraints entirely
- D. Removing all constraints from scope

Answer: A

Q976. Stream processing each element once?

- A. Static analysis technique
- B. Online streaming algorithm type
- C. Batch processing approach
- D. Compile-time optimization

Answer: B

Q977. What is invariant analysis?

- A. Measuring code coverage metrics
- B. Checking syntax for correctness
- C. Properties remaining true throughout
- D. Enforcing a code style guide

Answer: C

Q978. Scheduling with dependencies uses?

- A. Binary tree structure
- B. Linked list structure
- C. Hash table structure
- D. Directed acyclic graph

Answer: D

Q979. What is amortized analysis for?

- A. Measuring space complexity alone
- B. Analyzing the best case scenario
- C. Average cost per operation over time
- D. Analyzing single worst case only

Answer: C

Q980. Well-defined vs ill-defined problems?

- A. All CS problems are always well-defined
- B. Ill-defined problems have no solution at all
- C. Well-defined are always easier to solve
- D. Well-defined: clear states; ill-defined: ambiguous

Answer: D

Q981. Universal quantifier means?

- A. For exactly one element only
- B. For all elements in the domain
- C. For no elements in the domain
- D. There exists at least one element

Answer: B

Q982. Negation of forall x P(x)?

- A. exists x NOT P(x)
- B. forall x NOT P(x)
- C. NOT exists x P(x)
- D. forall x P(x) still

Answer: A

Q983. What is mathematical induction?

- A. Proving base case plus if n then n+1
- B. Testing all possible cases manually
- C. Guessing the answer to the problem
- D. Only proving the base case alone

Answer: A

Q984. What is a clause in resolution?

- A. A disjunction of logical literals
- B. A conjunction of several functions
- C. A complete executable program
- D. A variable assignment statement

Answer: A

Q985. Logical completeness means?

- A. The system runs very fast
- B. The system has no bugs at all
- C. Every valid formula is provable
- D. It uses all logical operators

Answer: C

Q986. Godel's first incompleteness theorem?

- A. All formal systems are complete
- B. All true statements are provable
- C. Consistent systems have unprovable truths
- D. Logic is unnecessary for proofs

Answer: C

Q987. What is satisfiability (SAT)?

- A. Whether an assignment makes it true
- B. A formula that is always true
- C. A formula with no variable bindings
- D. A formula that is always false

Answer: A

Q988. What is soundness?

- A. It uses auditory signal processing
- B. The system proves everything possible
- C. The system has zero error states
- D. Every provable statement is valid

Answer: D

Q989. SAT worst case?

- A. $O(n^2)$ quadratic time
- B. $O(n)$ linear time
- C. Exponential (NP-complete)
- D. $O(n \log n)$ linearithmic

Answer: C

Q990. What is semantic entailment?

- A. A programming concept for objects
- B. A specific data type for storage
- C. A syntax rule in formal languages
- D. Premises true implies conclusion true

Answer: D

Q991. for i=1..3, j=1..2: print(i*j). Output?

- A. 1 1 2 2 3 3
- B. 1 2 2 4 3 6
- C. 1 2 3 4 5 6
- D. 2 4 6 8 10 12

Answer: B

Q992. What is tail recursion?

- A. A recursion that never ends at all
- B. Last operation is the recursive call
- C. Recursion with multiple base cases
- D. A specific type of iterative loop

Answer: B

Q993. Recursion vs iteration relationship?

- A. They are completely unrelated concepts
- B. Only recursion works for any problem
- C. Every recursive can become iterative
- D. Iteration is always faster than recursion

Answer: C

Q994. Factorial base case?

- A. $n * \text{factorial}(n-1)$ recursive call
- B. $\text{factorial}(n)=n$ for all values
- C. $\text{factorial}(0)=1$ or $\text{factorial}(1)=1$
- D. No base case is ever needed

Answer: C

Q995. No base case means?

- A. The function runs much faster
- B. Stack overflow from infinite recursion
- C. The function gives correct results
- D. The function returns zero always

Answer: B

Q996. What is mutual recursion?

- A. A function calling only itself
- B. Functions calling each other cyclically
- C. A loop nested inside a function
- D. Recursion without any base case

Answer: B

Q997. Max depth for naive fibonacci(n)?

- A. $n/2$
- B. $\log n$
- C. 2^n
- D. n

Answer: D

Q998. What is a coroutine?

- A. A specific type of loop construct
- B. Can suspend and resume execution
- C. A normal sequential function call
- D. An error handler for exceptions

Answer: B

Q999. What is short-circuit evaluation?

- A. Evaluating expressions in reverse order
- B. A hardware electrical circuit fault
- C. Always evaluate all operands fully
- D. Stop evaluating when result is known

Answer: D

Q1000. What is a basic block?

- A. A class definition in a program
- B. A function definition in code
- C. One entry, one exit, no branches
- D. Any arbitrary piece of code

Answer: C

Q1001. Float problems vs integer?

- A. Rounding errors only
- B. Overflow errors only
- C. All of the above
- D. Underflow errors only

Answer: C

Q1002. What is endianness?

- A. A sorting algorithm technique
- B. A programming paradigm style
- C. A data compression approach
- D. Byte order stored in memory

Answer: D

Q1003. IEEE 754 single exponent bits?

- A. 32 bits
- B. 23 bits
- C. 8 bits
- D. 1 bit

Answer: C

Q1004. What is a tagged union?

- A. An array with mixed types
- B. A union with a type tag field
- C. A regular union without tags
- D. A sorting data structure type

Answer: B

Q1005. What is information hiding?

- A. Encrypting all data in storage
- B. Deleting data after processing
- C. Hiding files on the filesystem
- D. Restricting access, exposing interfaces

Answer: D

Q1006. What is data alignment?

- A. Alphabetical sorting of records
- B. Formatting text in a document
- C. Aligning elements on the screen
- D. Placing at boundaries for CPU access

Answer: D

Q1007. What is Liskov Substitution Principle?

- A. Subtypes substitutable for supertypes
- B. Making all class members public
- C. An encryption standard for data
- D. Functions without any parameters

Answer: A

Q1008. What is a phantom type?

- A. A null type with no useful value
- B. A runtime type for dynamic checks
- C. Type param with compile-time info only
- D. A type that does not exist at all

Answer: C

Q1009. IEEE 754 single max value?

- A. $3.4e38$
- B. 2^{64}
- C. $1.8e308$
- D. 2^{31}

Answer: A

Q1010. Structural vs nominal typing?

- A. They are exactly the same thing
- B. Nominal typing is not used in practice
- C. Structural is always better to use here
- D. Structural: shape match; nominal: name match

Answer: D

Q1011. LCS, edit distance, knapsack share?

- A. Linear search through all values
- B. Simple recursion without caching
- C. DP with optimal substructure pattern
- D. Greedy algorithm approach works

Answer: C

Q1012. Monotonic stack is for?

- A. Next greater/smaller element queries
- B. Graph traversal and pathfinding
- C. Sorting elements in an array
- D. String matching and searching

Answer: A

Q1013. Prefix sum is for?

- A. Searching for specific elements
- B. Sorting elements in an array
- C. Balancing a binary search tree
- D. Efficient range sum query answers

Answer: D

Q1014. Floyd's cycle detection?

- A. Two pointers at different speeds
- B. Graph coloring and partitioning
- C. Sorting elements in a list
- D. Tree traversal and path finding

Answer: A

Q1015. All permutations/combinations use?

- A. Hashing data structure
- B. Binary search technique
- C. Greedy algorithm approach
- D. Backtracking exploration

Answer: D

Q1016. What is meet in the middle?

- A. A compromise between two parties
- B. A social meeting of team members
- C. Split, process halves, then combine
- D. Finding the median of a dataset

Answer: C

Q1017. Topological sort is for?

- A. Sorting numbers in order
- B. Binary search in sorted data
- C. Sorting strings alphabetically
- D. Tasks ordered by dependencies

Answer: D

Q1018. Union-find is for?

- A. Sorting elements in arrays
- B. String matching and parsing
- C. Connected components and membership
- D. Queue operation management

Answer: C

Q1019. Binary lifting in trees?

- A. Sorting with binary comparisons
- B. Balancing a binary tree structure
- C. Ancestors and LCA in $O(\log n)$
- D. Binary search in an array

Answer: C

Q1020. Sweep line is for?

- A. Geometric event processing tasks
- B. Sorting points by coordinates
- C. Drawing shapes on screen
- D. Rendering graphics on display

Answer: A

Q1021. Master Theorem for?

- A. Finding prime numbers only
- B. Sorting arrays of elements
- C. Recurrences $T(n)=aT(n/b)+f(n)$
- D. Graph traversal algorithms

Answer: C

Q1022. $T(n)=2T(n/2)+n$?

- A. $O(n)$
- B. $O(\log n)$
- C. $O(n \log n)$
- D. $O(n^2)$

Answer: C

Q1023. Euler's totient $\phi(n)$?

- A. Count of numbers coprime with n
- B. The factorial of the number n
- C. The largest prime factor of n
- D. The sum of all divisors of n

Answer: A

Q1024. Catalan numbers count?

- A. Prime numbers in a range
- B. Greatest common divisor values
- C. Parenthesizations and binary trees
- D. Sorted permutations of arrays

Answer: C

Q1025. Fermat's Little Theorem?

- A. $a^p = a \pmod{p}$ for prime p
- B. All prime numbers are odd values
- C. $a^n = n^a$ for all values
- D. Every integer has unique factors

Answer: A

Q1026. Sieve of Eratosthenes complexity?

- A. $O(n \log n)$
- B. $O(n^2)$
- C. $O(n \log \log n)$
- D. $O(n)$

Answer: C

Q1027. Chinese Remainder Theorem for?

- A. Sorting elements in arrays
- B. Graph coloring and partitioning
- C. Factorization of large numbers
- D. Simultaneous modular congruences

Answer: D

Q1028. Stirling's approximation significance?

- A. Approximates $n!$ for growth analysis
- B. It sorts elements efficiently
- C. It finds prime numbers quickly
- D. It solves algebraic equations

Answer: A

Q1029. Handshaking Lemma?

- A. All vertices have even degree
- B. The graph is always fully connected
- C. Sum of degrees equals 2 times edges
- D. All edges have equal weight values

Answer: C

Q1030. Matrix exponentiation for?

- A. Sorting matrix rows and columns
- B. n th term of linear recurrences fast
- C. Displaying matrices on screen
- D. Reading matrix input from files

Answer: B

Q1031. Comparison sort lower bound?

- A. $O(\log n)$
- B. $O(n \log n)$
- C. $O(n)$
- D. $O(n^2)$

Answer: B

Q1032. What is approximation algorithm?

- A. Produces random unpredictable results
- B. Always gives exact solutions
- C. Near-optimal with provable guarantees
- D. Tries all solutions by brute force

Answer: C

Q1033. What is A*?

- A. A string matching technique
- B. A linear search method only
- C. Best-first search using heuristics
- D. A sorting algorithm for arrays

Answer: C

Q1034. What is branch and bound?

- A. A tree data structure type
- B. A sorting algorithm approach
- C. A compression technique used
- D. Prunes branches worse than best

Answer: D

Q1035. Floyd-Warshall complexity?

- A. $O(V^2)$
- B. $O(V^3)$
- C. $O(V \cdot E)$
- D. $O(E \log V)$

Answer: B

Q1036. What is randomized algorithm?

- A. An algorithm with no clear structure
- B. An algorithm that has bugs in it
- C. Uses random choices for performance
- D. An algorithm that is always wrong

Answer: C

Q1037. Las Vegas vs Monte Carlo?

- A. LV: always correct; MC: may err sometimes
- B. Monte Carlo is always correct output
- C. They are exactly the same algorithm
- D. Las Vegas is always faster in practice

Answer: A

Q1038. Amortized $O(1)$ dynamic array?

- A. The array never resizes at all
- B. Every single operation is $O(1)$
- C. Some $O(n)$ resizes, average $O(1)$
- D. All operations take $O(\log n)$

Answer: C

Q1039. Kruskal's key idea?

- A. Use DFS for tree construction
- B. Start from a single vertex first
- C. Sort edges, add non-cycle-forming
- D. Use BFS for tree construction

Answer: C

Q1040. KMP solves?

- A. Sorting array elements fast
- B. Pattern matching in $O(n+m)$
- C. Matrix multiplication task
- D. Graph traversal and search

Answer: B

Q1041. What is CAP theorem?

- A. A programming language feature
- B. At most 2 of C, A, P guaranteed
- C. All 3 properties always achievable
- D. A sorting algorithm for databases

Answer: B

Q1042. What is technical debt?

- A. Money owed to hardware vendors
- B. Software licensing fees to pay
- C. Future rework cost from shortcuts
- D. Physical hardware that is broken

Answer: C

Q1043. Pareto Principle in software?

- A. 80% of code should be comments
- B. All components contribute equally
- C. 80% of effects from 20% of causes
- D. 80% developers write 20% of code

Answer: C

Q1044. What is opportunity cost?

- A. The cost of running cloud servers
- B. The best alternative you gave up
- C. A pricing type for subscriptions
- D. The cost of buying software tools

Answer: B

Q1045. What is survivorship bias?

- A. A sorting algorithm technique used
- B. A programming paradigm approach
- C. A database concept for backups
- D. Focus on successes, ignore failures

Answer: D

Q1046. Correlation vs causation?

- A. Neither concept applies in CS work
- B. They are exactly the same concept
- C. Correlation: related; causation: causes
- D. Causation is weaker than correlation

Answer: C

Q1047. Conway's Law?

- A. A sorting algorithm for lists
- B. A complexity class in theory
- C. A data storage optimization law
- D. Systems mirror communication structure

Answer: D

Q1048. What is expected value?

- A. The maximum possible outcome
- B. The most common outcome seen
- C. Sum of outcome times probability
- D. The minimum possible outcome

Answer: C

Q1049. What is post-mortem?

- A. Structured review of what happened
- B. Only done after deployment phase
- C. Deleting all code from the project
- D. A security audit of the system

Answer: A

Q1050. Occam's Razor in software?

- A. All solutions are equally valid here
- B. Simplest adequate solution preferred
- C. Maximize complexity for robustness
- D. Complex solutions are always best

Answer: B

Q1051. What is a Heisenbug?

- A. A syntax error in the source code
- B. A bug that changes when observed
- C. A normal reproducible program bug
- D. A hardware component malfunction

Answer: B

Q1052. What is static analysis?

- A. Testing at runtime with real data
- B. Analyzing code without executing it
- C. User acceptance testing of the app
- D. Running the code to find bugs

Answer: B

Q1053. What is dynamic analysis?

- A. Code review by team members only
- B. Reading documentation about the code
- C. Analyzing behavior during execution
- D. Static code review without running

Answer: C

Q1054. What is buffer overflow?

- A. Writing beyond allocated memory area
- B. A database overflowing with data
- C. A network packet overflow error
- D. A full disk with no free space

Answer: A

Q1055. What is fault localization?

- A. Finding faulty hardware components
- B. Locating the server in a datacenter
- C. Finding the compiler installation path
- D. Automatically identifying bug locations

Answer: D

Q1056. What is delta debugging?

- A. Running multiple programs together
- B. Using version control for changes
- C. Minimizing failure-inducing input size
- D. Making small incremental changes

Answer: C

Q1057. What is a sanitizer?

- A. A cleaning product for hardware
- B. A code formatting utility tool
- C. An antivirus software application
- D. Runtime tool detecting memory errors

Answer: D

Q1058. What is symbolic execution?

- A. Using symbols in variable names
- B. Deploying to production environment
- C. Symbolic values exploring all paths
- D. Code review by senior developers

Answer: C

Q1059. What is a Mandelbug?

- A. A graphical rendering issue seen
- B. A simple reproducible error found
- C. A user interface display problem
- D. Complex chaotic causes, hard to fix

Answer: D

Q1060. What is time-travel debugging?

- A. Predicting future bugs in advance
- B. Record execution, step backward
- C. Debugging legacy system software
- D. Debugging very old legacy code

Answer: B

Q1061. Naive Fibonacci complexity?

- A. $O(n)$
- B. $O(n^2)$
- C. $O(2^n)$
- D. $O(n \log n)$

Answer: C

Q1062. Splay tree amortized?

- A. $O(n)$
- B. $O(n \log n)$
- C. $O(\log n)$
- D. $O(1)$

Answer: C

Q1063. Adjacency list vs matrix?

- A. No meaningful difference exists
- B. $O(V+E)$ vs $O(V^2)$, saves space sparse
- C. Matrix is always better to use here
- D. They use exactly the same space

Answer: B

Q1064. Dijkstra with binary heap?

- A. $O(V^3)$
- B. $O((V+E) \log V)$
- C. $O(V^2)$
- D. $O(E)$

Answer: B

Q1065. P=NP asks?

- A. Are primes equal to naturals
- B. Verifiable in poly time means solvable?
- C. Do programs always need parameters
- D. Are pointers always needed in code

Answer: B

Q1066. Naive matrix mult $n \times n$?

- A. $O(n^2)$
- B. $O(n^2 \log n)$
- C. $O(n^3)$
- D. $O(n)$

Answer: C

Q1067. Strassen significance?

- A. Same time as naive multiplication
- B. It is slower than naive approach
- C. Below $O(n^3)$ at about $O(n^{2.807})$
- D. It only works on 2×2 matrices

Answer: C

Q1068. What is NP-complete?

- A. An easy problem with fast solution
- B. In NP, all NP reduces to it in poly
- C. A problem with known $O(n)$ solution
- D. A problem that is unsolvable always

Answer: B

Q1069. Cache complexity?

- A. Same as hash table complexity analysis
- B. Cache misses affect real performance
- C. Only measures the cache memory size
- D. Same as space complexity of program

Answer: B

Q1070. PSPACE?

- A. Polynomial time class problems
- B. Only applies to sorting algorithms
- C. The simplest complexity class known
- D. Polynomial space, contains P and NP

Answer: D

Q1071. Find LCS?

- A. DP with a 2D table approach
- B. Hashing all subsequences found
- C. Brute force all subsequences
- D. Simple character matching scan

Answer: A

Q1072. 0/1 Knapsack optimal?

- A. Sort items by weight first
- B. DP over items and capacities
- C. Greedy algorithm approach
- D. Random selection of items

Answer: B

Q1073. LRU cache efficient implementation?

- A. Only an array data structure
- B. A stack-based data structure
- C. A binary search tree structure
- D. Hash map plus doubly linked list

Answer: D

Q1074. All connected components?

- A. BFS/DFS from unvisited vertices
- B. Use hash table for the lookup
- C. Check all pairs of vertex nodes
- D. Sort all the vertices first

Answer: A

Q1075. What is tree DP?

- A. Sorting the tree node values
- B. Balancing the tree structure
- C. Bottom-up from children results
- D. Simple tree traversal method

Answer: C

Q1076. Median in stream?

- A. Two heaps: max lower, min upper
- B. Use a single array for storage
- C. Sort the entire stream each time
- D. Use a linked list for ordering

Answer: A

Q1077. Bit manipulation for?

- A. Making code completely unreadable
- B. Graphics rendering operations only
- C. Low-level hardware drivers only
- D. Sets, flags, and $O(1)$ optimizations

Answer: D

Q1078. Shortest path unweighted?

- A. BFS from the source vertex
- B. Bellman-Ford algorithm approach
- C. Dijkstra's algorithm approach
- D. DFS from the source vertex

Answer: A

Q1079. What is segment tree?

- A. Sorting segments of an array
- B. String matching and pattern search
- C. Range queries and updates in $O(\log n)$
- D. Graph traversal and pathfinding

Answer: C

Q1080. Find articulation points?

- A. BFS from every vertex in the graph
- B. Sort the graph vertices by degree
- C. Tarjan's DFS with discovery/low values
- D. Check all vertices by removal

Answer: C

Q1081. Dimensionality reduction?

- A. Reduce features, preserve information
- B. Compressing files on the filesystem
- C. Deleting rows from the data table
- D. Reducing disk storage size used

Answer: A

Q1082. What is overfitting?

- A. Having too little training data
- B. A model achieving perfect accuracy
- C. Learns noise, fails on new data
- D. A model that is far too simple

Answer: C

Q1083. Bias-variance tradeoff?

- A. Balance simplicity and complexity
- B. A network optimization technique
- C. A database optimization approach
- D. A hiring process consideration

Answer: A

Q1084. What is MapReduce?

- A. A database type for storing data
- B. A visualization tool for charts
- C. Parallel map and reduce for big data
- D. A sorting algorithm for arrays

Answer: C

Q1085. Simpson's Paradox?

- A. Trend reverses when groups combined
- B. A sorting algorithm special case
- C. A code error in the implementation
- D. A normal expected statistical result

Answer: A

Q1086. Curse of dimensionality?

- A. Having too few data dimensions
- B. A graphics rendering limitation
- C. More features makes data sparser
- D. Simple linear scaling of the data

Answer: C

Q1087. Time series analysis?

- A. Sorting data by alphabetical name
- B. Counting occurrences in a dataset
- C. Analyzing static unchanging data only
- D. Temporal data for trends and forecasts

Answer: D

Q1088. Ensemble learning?

- A. A data cleaning preprocessing step
- B. Using a single model for predictions
- C. A database optimization technique used
- D. Combining multiple models for accuracy

Answer: D

Q1089. Data lineage?

- A. Formatting data for the display
- B. Deleting data after processing it
- C. The age of the data in storage
- D. Tracking origin and transformations

Answer: D

Q1090. Streaming vs batch?

- A. Streaming: real-time; batch: scheduled
- B. Streaming is always much slower
- C. They are exactly the same approach
- D. Batch processing is always real-time

Answer: A

Q1091. What is TRIZ?

- A. Systematic innovation from patents
- B. A testing framework for software
- C. A programming language for apps
- D. A database management system type

Answer: A

Q1092. What is first principles?

- A. Copying existing best practices always
- B. Following established procedures only
- C. Using templates without modification
- D. Reasoning up from fundamental truths

Answer: D

Q1093. Emergent vs planned design?

- A. Emergent evolves; planned defines upfront
- B. They are exactly the same approach
- C. Planned design never actually works
- D. Emergent design is always worse overall

Answer: A

Q1094. Innovator's dilemma?

- A. A database design issue to address
- B. A code problem in the application
- C. Established orgs struggle with disruption
- D. A sorting algorithm special case issue

Answer: C

Q1095. What is double diamond?

- A. A testing framework for the app
- B. A type of jewelry design pattern
- C. A data structure for binary trees
- D. Two diverge/converge phases in design

Answer: D

Q1096. What is technology transfer?

- A. Upgrading hardware on all machines
- B. Installing new software on systems
- C. Technology from one domain to another
- D. Moving servers to a new location

Answer: C

Q1097. What is morphological analysis?

- A. Analyzing word forms in linguistics
- B. All parameter variation combinations
- C. Parsing source code for syntax errors
- D. Static code analysis for quality

Answer: B

Q1098. Serendipity in innovation?

- A. It fully replaces structured research
- B. It always guarantees success achieved
- C. It plays no role in innovation
- D. Unexpected discoveries become breakthroughs

Answer: D

Q1099. What is disruptive innovation?

- A. Simpler solutions for overlooked segments
- B. Optimization of existing system performance
- C. Breaking things without any purpose
- D. Normal incremental improvement of product

Answer: A

Q1100. Platform thinking?

- A. Extensible platforms creating ecosystems
- B. Building stages for presentations
- C. Only using existing platforms given
- D. Targeting a specific operating system

Answer: A

Q1101. CAP for distributed DBs?

- A. Choose consistency or availability during partitions
- B. Only consistency matters in databases
- C. All 3 guarantees always achievable
- D. Only availability matters in databases

Answer: A

Q1102. Eventual consistency?

- A. System is never consistent at all
- B. Only works on very small systems
- C. Becomes consistent given enough time
- D. System is always immediately consistent

Answer: C

Q1103. Two generals problem?

- A. Reliable communication over unreliable channel
- B. A graph traversal algorithm technique
- C. A military history lesson studied
- D. A sorting algorithm for dual arrays

Answer: A

Q1104. Horizontal vs vertical scaling?

- A. Vertical is always the better choice
- B. They are exactly the same approach
- C. Horizontal: more machines; vertical: more power
- D. Horizontal is always the better choice

Answer: C

Q1105. Microservices?

- A. Very small source code files only
- B. Independent services communicating via APIs
- C. Micro-controller hardware programming
- D. A testing framework for small units

Answer: B

Q1106. Chaos engineering?

- A. Breaking systems permanently by intent
- B. General disorder in the codebase
- C. Creating chaos without any purpose
- D. Deliberate failures to test resilience

Answer: D

Q1107. Byzantine Generals?

- A. Consensus with faulty participants
- B. A graph coloring algorithm used
- C. A military strategy for battles
- D. A history lesson about generals

Answer: A

Q1108. GDPR compliance?

- A. A programming language standard
- B. Protecting EU personal data rights
- C. A database type for EU servers
- D. A testing methodology for quality

Answer: B

Q1109. Edge computing?

- A. A CSS styling technique for borders
- B. An optimization for rendering speed
- C. Processing near source for low latency
- D. Computing at the edge of a desk

Answer: C

Q1110. Zero-knowledge proof?

- A. A testing methodology for software
- B. An authentication method for users
- C. Knowing nothing about the subject
- D. Prove knowledge without revealing info

Answer: D

Q1111. What is literate programming?

- A. Adding many comments everywhere found
- B. Writing poetry about the algorithms
- C. Writing literary fiction about code
- D. Code embedded in documentation text

Answer: D

Q1112. What is docs-as-code?

- A. Using code tools to write docs
- B. Docs with version control and CI/CD
- C. Writing code inside documentation
- D. Replacing code with documentation

Answer: B

Q1113. What is C4 model?

- A. A C programming language standard
- B. A security classification framework
- C. Four levels: Context, Container, Component, Code
- D. A testing methodology for the app

Answer: C

Q1114. Technical writing for diverse audiences?

- A. Write only for experts in the field
- B. Always use maximum jargon in writing
- C. Adapt style and depth for audience
- D. Write the longest documents possible

Answer: C

Q1115. What are runbooks?

- A. Step-by-step operational procedure guides
- B. A marathon training plan document
- C. Code execution logs from the system
- D. Competitions between team members

Answer: A

Q1116. Diagramming in distributed systems?

- A. Designing logos for the application
- B. Diagrams for complex interactions shown
- C. Creating art for the office walls
- D. Drawing pictures for decoration only

Answer: B

Q1117. Retrospective documentation?

- A. Documenting very old legacy systems
- B. Record successes, failures, improvements
- C. Archiving old documents for storage
- D. Deleting outdated documentation files

Answer: B

Q1118. What is traceability?

- A. Link requirements to design and tests
- B. Numbering all the pages in documents
- C. Finding lost documentation on the server
- D. Tracking user access to the system

Answer: A

Q1119. What is communication debt?

- A. Outstanding phone bills to be paid
- B. Cost of poor communication accumulated
- C. Too many meetings in the schedule
- D. Email overload in the team inbox

Answer: B

Q1120. Post-incident review docs?

- A. Legal documents for court proceedings
- B. Documents assigning blame to someone
- C. Deleting all evidence of the incident
- D. Blameless timeline, root cause, actions

Answer: D

Q1121. A team repeatedly solves symptoms instead of root causes. Which problem-solving failure does this represent?

- A. Premature optimization of solution components
- B. Surface-level analysis without deep investigation
- C. Excessive decomposition of the original problem
- D. Functional fixedness limiting their perspective

Answer: B

Q1122. When solving a novel problem, which approach is most effective for generating initial solution ideas?

- A. Immediately coding the first idea that emerges
- B. Avoiding all brainstorming to save time needed
- C. Selecting the most popular known framework now
- D. Divergent thinking before convergent evaluation

Answer: D

Q1123. A developer uses only familiar tools despite better alternatives existing. What cognitive trap is this?

- A. The law of the instrument or Maslow's hammer
- B. The Dunning-Kruger effect on skill assessment
- C. Analysis paralysis from too many viable options
- D. Anchoring bias toward the first piece of data

Answer: A

Q1124. In complex system failures, why is linear cause-and-effect thinking often insufficient for analysis?

- A. Because complex problems only have single causes
- B. Because simple problems always have simple fixes
- C. Because linear thinking always finds root causes
- D. Because complex systems have interacting factors

Answer: D

Q1125. Which meta-cognitive strategy most improves problem-solving performance across different domains?

- A. Avoiding all unfamiliar problem types encountered
- B. Memorizing solutions to every known problem type
- C. Self-monitoring and reflective evaluation of approach
- D. Working faster without pausing to think or assess

Answer: C

Q1126. A problem has multiple valid solutions with different trade-offs. What is the best decision approach?

- A. Randomly select any solution since all are valid
- B. Always pick the solution with the lowest cost only
- C. Evaluate solutions against weighted criteria matrix
- D. Choose the first solution that seems acceptable

Answer: C

Q1127. How does the concept of problem space relate to search-based problem solving in AI research?

- A. It eliminates the need for any search algorithms
- B. It defines all possible states and transitions
- C. It only refers to physical workspace dimensions
- D. It is unrelated to state-based representations

Answer: B

Q1128. Why might an expert sometimes perform worse than a novice on certain restructured problems?

- A. Experts always think slower than beginners do
- B. Experts lack any knowledge of the problem domain
- C. Novices have inherently superior problem-solving skills
- D. Expert schemas can create fixation on known patterns

Answer: D

Q1129. In design thinking, why is reframing a problem often more valuable than directly solving it as stated?

- A. Because the original framing usually is always correct
- B. Because it eliminates the need for any user research
- C. Because reframing reveals hidden assumptions and needs
- D. Because it allows avoiding the problem entirely now

Answer: C

Q1130. Which approach best handles wicked problems that have no definitive formulation or stopping rule?

- A. Ignoring complexity and applying the simplest solution
- B. Applying a single fixed algorithm repeatedly here
- C. Iterative prototyping with stakeholder feedback loops
- D. Waiting until all information is completely available

Answer: C

Q1131. A system shows intermittent failures only under peak load. Which analysis approach best identifies the root cause?

- A. Reviewing the code during low traffic periods only
- B. Restarting the system each time a failure is noted
- C. Asking users to describe the error from their memory
- D. Load testing combined with systematic log analysis

Answer: D

Q1132. When requirements conflict with each other in a project, what is the most appropriate analysis strategy?

- A. Implementing all conflicting requirements simultaneously
- B. Ignoring the conflict and proceeding with work
- C. Prioritization and trade-off analysis with stakeholders
- D. Removing all conflicting requirements from the scope

Answer: C

Q1133. How does causal loop diagramming improve analysis of complex interdependent system problems?

- A. It focuses only on the most obvious direct causes
- B. It reveals feedback loops and systemic interactions
- C. It eliminates the need for any quantitative analysis
- D. It simplifies systems to single cause-effect chains

Answer: B

Q1134. A bug appears only when three specific modules interact simultaneously. What analysis type is most appropriate?

- A. Code style review of each module independently
- B. Performance benchmarking of individual module speed
- C. Unit testing each module in complete isolation
- D. Integration and interaction analysis across modules

Answer: D

Q1135. Why might quantitative data alone be insufficient for complete analysis in socio-technical systems?

- A. Statistical methods solve every analytical challenge
- B. Human factors require qualitative analysis as well
- C. Quantitative data always provides complete insights
- D. Numbers are always more important than any context

Answer: B

Q1136. When analyzing a legacy system for modernization, which approach best captures both explicit and tacit knowledge?

- A. Automated scanning without any human involvement
- B. Only reading the existing documentation available
- C. Code archaeology combined with expert interviews
- D. Rewriting everything from scratch without analysis

Answer: C

Q1137. How does Cynefin framework help classify problems for appropriate analysis strategies?

- A. It treats all problems as equally complex always
- B. It categorizes problems by their domain complexity
- C. It eliminates the need for any problem analysis
- D. It only works for simple and obvious problem types

Answer: B

Q1138. A data pipeline fails silently, producing incorrect results. What analysis method best detects this?

- A. Data validation with expected output range assertions
- B. Reviewing the pipeline code for syntax errors only
- C. Checking only if the pipeline completes without errors
- D. Monitoring only the system CPU and memory utilization

Answer: A

Q1139. When performing risk analysis for a critical system, why is fault tree analysis preferred over simple checklists?

- A. Fault trees are simpler to create than any checklist
- B. Fault trees model causal relationships between failures
- C. Checklists capture all possible failure combinations
- D. Checklists are always more thorough and comprehensive

Answer: B

Q1140. How does assumption mapping improve problem analysis in uncertain environments?

- A. It eliminates the need for any empirical validation
- B. It only works when all assumptions are already known
- C. It proves all assumptions are completely correct
- D. It identifies and ranks assumptions by risk and impact

Answer: D

Q1141. A programmer argues: The tests pass so the code is correct. What logical error is this?

- A. Affirming the consequent in conditional logic
- B. Correct use of biconditional logical equivalence
- C. Valid modus ponens application of the argument
- D. Denying the antecedent in conditional reasoning

Answer: A

Q1142. In formal verification, how does temporal logic extend propositional logic for system reasoning?

- A. It removes the need for any truth values used
- B. It only applies to static non-changing systems now
- C. It simplifies logic by removing all connectives
- D. It adds operators for reasoning about time states

Answer: D

Q1143. Why is the halting problem significant for the limits of logical reasoning in CS?

- A. It proves all programs can be verified completely
- B. It proves formal logic can solve every CS problem
- C. It demonstrates that all programs eventually halt
- D. It shows some properties are logically undecidable

Answer: D

Q1144. When building a logical argument for system requirements, why might closed-world assumption lead to errors?

- A. It ensures all possible scenarios are fully covered
- B. Unknown facts default to false which may be wrong
- C. It correctly handles all unknown facts in systems
- D. It makes reasoning more complete and more accurate

Answer: B

Q1145. How does non-monotonic reasoning differ from classical logic in practical problem solving?

- A. Both systems handle new evidence in the same way
- B. Classical logic allows retracting past conclusions
- C. Non-monotonic reasoning retracts with new evidence
- D. Non-monotonic reasoning never revises conclusions

Answer: C

Q1146. A spec states the system shall not fail more than once per year. What logical challenge arises in verification?

- A. It is trivially easy to verify this requirement
- B. Negative requirements need no special handling at all
- C. The requirement is too simple to need verification
- D. Proving a negative over time requires exhaustive testing

Answer: D

Q1147. In fuzzy logic, how is truth handled differently than in classical Boolean logic systems?

- A. Truth values range continuously between zero and one
- B. Truth values are restricted to only zero and one
- C. Fuzzy logic uses only three discrete truth values
- D. Fuzzy logic eliminates the concept of truth entirely

Answer: A

Q1148. Why is the principle of explosion problematic in inconsistent knowledge bases?

- A. It only affects mathematical proofs not software
- B. It makes consistent knowledge bases unreliable now
- C. From a contradiction any statement can be derived
- D. It prevents deriving any conclusions whatsoever

Answer: C

Q1149. How does abductive reasoning complement deductive and inductive reasoning in diagnostics?

- A. It provides mathematically certain proofs of causes
- B. It generates the best explanation for observations
- C. It replaces both deductive and inductive methods
- D. It ignores observed evidence during the reasoning

Answer: B

Q1150. When designing a rule-based system, why is the Rete algorithm important for efficient inference?

- A. It only works with exactly two rules at one time
- B. It optimizes pattern matching across many rules
- C. It eliminates the need for any rules in the system
- D. It replaces logical inference with random selection

Answer: B

Q1151. A recursive algorithm has overlapping subproblems. Which technique best optimizes its flow?

- A. Removing the base case to simplify the algorithm
- B. Adding more recursive calls to increase coverage
- C. Converting all recursion to deeply nested if-else
- D. Dynamic programming with memoization of results

Answer: D

Q1152. When converting recursive to iterative, which data structure commonly replaces the call stack?

- A. An explicit stack data structure for depth tracking
- B. A queue data structure for breadth-first processing
- C. A linked list for sequential data traversal only
- D. A hash table for random access to stored elements

Answer: A

Q1153. In concurrent programming, what flow control challenge does a race condition represent?

- A. Unpredictable behavior from unsynchronized shared access
- B. A guaranteed correct ordering of all operations
- C. A performance optimization technique for parallel code
- D. A standard debugging technique for multi-threaded code

Answer: A

Q1154. Why might tail-call optimization be important for recursive algorithms in functional programming?

- A. It increases the number of stack frames allocated
- B. It is only relevant to non-recursive function calls
- C. It prevents stack overflow by reusing stack frames
- D. It makes recursion slower but more memory intensive

Answer: C

Q1155. A state machine has unreachable states in its design. What does this indicate about algorithm flow?

- A. The design is optimal and completely correct here
- B. All states must be unreachable for a correct design
- C. Dead code or design flaws exist in the state model
- D. Unreachable states improve system fault tolerance

Answer: C

Q1156. How does coroutine-based flow control differ from traditional subroutine calls?

- A. Subroutines can pause execution at arbitrary points
- B. Coroutines cannot pause and resume their execution
- C. Coroutines and subroutines behave identically always
- D. Coroutines can suspend and resume at specific points

Answer: D

Q1157. What does cyclomatic complexity measure about code in a control flow graph?

- A. The number of independent paths through the program
- B. The execution speed of the algorithm in milliseconds
- C. The amount of memory the program needs to execute
- D. The total number of lines in the source code file

Answer: A

Q1158. In reactive programming, how does event-driven flow handle backpressure in data streams?

- A. By processing all events regardless of capacity now
- B. By shutting down the entire system under any load
- C. By dropping all incoming events without any notice
- D. By signaling the producer to slow down its rate

Answer: D

Q1159. Why are continuation-passing style transformations relevant to complex flow control?

- A. They make the flow of control explicit as a parameter
- B. They simplify all algorithms to a single instruction
- C. They eliminate the need for any function calls at all
- D. They only apply to algorithms with no branching logic

Answer: A

Q1160. In microservices, what flow pattern handles distributed transactions across services?

- A. Ignoring transaction failures across all the services
- B. Locking all services until one transaction completes
- C. The saga pattern with compensating transactions here
- D. Running all operations in a single local transaction

Answer: C

Q1161. For handling both structured and unstructured data, which representation provides the most flexibility?

- A. Schema-on-read approach using document-based storage
- B. Single flat file with comma separated value fields
- C. Rigid relational schema with strict type enforcement
- D. Fixed-width record format with no flexibility given

Answer: A

Q1162. Why might a trie be preferred over a hash table for storing a dictionary of words?

- A. Hash tables cannot store string data at all reliably
- B. Tries support efficient prefix-based search queries
- C. Tries have faster exact match lookup than hash tables
- D. Tries use less memory than hash tables for words

Answer: B

Q1163. In distributed systems, what challenge does the CAP theorem pose for data representation across nodes?

- A. The CAP theorem only applies to single-node systems
- B. Distributed systems never face any consistency issues
- C. All three properties can always be fully guaranteed
- D. Only two of consistency, availability, partition tolerance

Answer: D

Q1164. How does a bloom filter represent set membership differently from a traditional hash set?

- A. Bloom filters use more memory than standard hash sets
- B. Bloom filters guarantee no false positives ever occur
- C. Bloom filters store every element in full original form
- D. Bloom filters allow false positives never false negatives

Answer: D

Q1165. When representing time-series data, why might columnar storage outperform row-based storage?

- A. Columnar enables better compression and faster scans
- B. Columnar is worse for analytical aggregation queries
- C. Columnar storage cannot handle time-series data at all
- D. Row-based always compresses better than column based

Answer: A

Q1166. What is the significance of endianness in cross-platform data representation?

- A. Byte order differences can cause incorrect data reading
- B. Endianness has no effect on data interpretation ever
- C. All systems universally use the same byte ordering now
- D. Endianness only affects text data never numerical data

Answer: A

Q1167. Why might persistent data structures be preferred over mutable ones in concurrent programming?

- A. Persistent structures use less memory than mutable ones
- B. Mutable structures never cause concurrency problems here
- C. Persistent structures avoid race conditions by immutability
- D. Persistent structures are always faster to modify now

Answer: C

Q1168. How does a B-tree optimize data representation for disk-based storage systems?

- A. B-trees minimize disk reads with high branching factors
- B. B-trees cannot be used in any database system at all
- C. B-trees are binary trees optimized for memory only use
- D. B-trees store all data in memory avoiding disk reads

Answer: A

Q1169. When choosing between adjacency matrix and adjacency list for a sparse graph, what drives the decision?

- A. Adjacency matrices are always better for sparse graphs
- B. The choice has no impact on space or time performance
- C. Adjacency lists require more space than matrices always
- D. Adjacency lists use less space for graphs with few edges

Answer: D

Q1170. What challenge does IEEE 754 floating-point create for financial calculations?

- A. It represents all decimal values with perfect precision
- B. Rounding errors accumulate making exact comparison unsafe
- C. It is designed specifically for financial math operations
- D. Floating-point numbers never have any precision issues

Answer: B

Q1171. When applying CQRS, what does separating read and write models provide?

- A. It prevents any data from being written to the store
- B. It ensures read and write use the same model always
- C. It optimizes each model for its specific access pattern
- D. It eliminates the need for any database queries at all

Answer: C

Q1172. How does the circuit breaker pattern help in distributed systems?

- A. It stops calls to failing services preventing cascading
- B. It keeps retrying failed calls indefinitely always
- C. It ignores all failures and continues normal operation
- D. It speeds up calls to services that are already down

Answer: A

Q1173. In ML, why might overfitting be a failure of pattern generalization?

- A. The model learns general patterns that transfer well
- B. Overfitting and generalization are the same concept
- C. The model memorizes training data without generalizing
- D. Overfitting always improves performance on new data

Answer: C

Q1174. How does event sourcing differ from CRUD in representing state changes?

- A. It stores only the current state deleting all history
- B. It uses the same approach as traditional CRUD always
- C. It prevents any modification to the application state
- D. It stores all state changes as a sequence of events

Answer: D

Q1175. Why is the anti-corruption layer important when integrating with legacy systems?

- A. It translates between new and legacy domain models
- B. It forces legacy systems to adopt all new interfaces
- C. It eliminates the need for any integration testing
- D. It directly couples new code to the legacy interface

Answer: A

Q1176. How does hierarchical pattern recognition improve efficiency over flat matching?

- A. Hierarchical approaches ignore high-level pattern data
- B. It processes all patterns at the same level of detail
- C. Flat matching is always more efficient than hierarchical
- D. It recognizes patterns at multiple levels of abstraction

Answer: D

Q1177. When should the strangler fig pattern be used instead of a big-bang rewrite?

- A. When the system is small enough for instant rewrite
- B. When gradual incremental migration reduces overall risk
- C. When no part of the old system needs to keep running
- D. When all features must be replaced simultaneously now

Answer: B

Q1178. What challenge does concept drift pose for deployed recognition systems?

- A. Patterns remain perfectly stable over all time periods
- B. Concept drift improves model accuracy automatically
- C. Deployed models never need any retraining or updates
- D. Data distributions change making learned patterns stale

Answer: D

Q1179. How does the bulkhead pattern prevent total system failure?

- A. It shares all resources across every system component
- B. It eliminates redundancy to reduce system complexity
- C. It isolates components so failure in one stays contained
- D. It connects all components tightly for maximum speed

Answer: C

Q1180. Why is transfer learning effective when labeled data is scarce?

- A. It requires more labeled data than training from scratch
- B. It trains from scratch ignoring all prior knowledge
- C. It reuses patterns learned from related tasks or domains
- D. It only works when unlimited labeled data is available

Answer: C

Q1181. A problem requires counting paths in a grid. Which mathematical concept most efficiently solves this?

- A. Combinatorics using binomial coefficients for paths
- B. Simple addition of row and column dimensions here
- C. Subtraction of blocked cells from the total cell count
- D. Multiplication of the grid area by a constant factor

Answer: A

Q1182. Why is the master theorem important for analyzing divide-and-conquer complexity?

- A. It only applies to iterative non-recursive algorithms
- B. It provides closed-form solutions for recurrence relations
- C. It gives exact constant factors for algorithm run time
- D. It eliminates the need for any complexity analysis done

Answer: B

Q1183. How does probability theory help in analyzing randomized algorithms?

- A. Randomized algorithms are always worse than deterministic
- B. Probability is irrelevant to algorithm analysis work
- C. Expected value analysis predicts average case behavior
- D. Randomized algorithms always have fixed performance

Answer: C

Q1184. When solving optimization, why might linear programming be preferred over brute force?

- A. Brute force is always faster for optimization problems
- B. Linear programming efficiently handles continuous constraints
- C. Brute force uses less memory than linear programming does
- D. Linear programming cannot handle any constraints at all

Answer: B

Q1185. How does graph theory help model and solve network flow problems?

- A. Graph theory is unrelated to any network flow problems
- B. Graph theory only applies to social network analysis
- C. Network flows cannot be modeled using any graph theory
- D. Max-flow min-cut theorem optimizes network throughput

Answer: D

Q1186. Why is NP-completeness important for practical problem-solving decisions?

- A. NP-complete problems are always easy to solve quickly
- B. NP-completeness is only theoretical with no practical use
- C. It identifies problems unlikely to have efficient solutions
- D. All NP-complete problems have known polynomial solutions

Answer: C

Q1187. How does Bayes' theorem enable updating beliefs with new evidence?

- A. It updates posterior probability using prior and likelihood
- B. It ignores prior probabilities during all calculations
- C. It only works when all probabilities are exactly known
- D. It eliminates uncertainty from all decision-making tasks

Answer: A

Q1188. In cryptography, why is the discrete logarithm problem computationally significant?

- A. Computing discrete logarithms takes constant time only
- B. Its difficulty forms the basis for public key security
- C. Discrete logarithms are trivially easy to compute
- D. Discrete logarithms have no use in modern cryptography

Answer: B

Q1189. How does amortized analysis provide more accurate picture than worst-case analysis?

- A. It only applies to algorithms with constant time cost
- B. It always shows worse performance than worst-case does
- C. It ignores expensive operations in the total cost here
- D. It averages cost over a sequence of operations performed

Answer: D

Q1190. Why are Markov chains useful for modeling stochastic processes?

- A. They can only model deterministic non-random processes
- B. They require knowledge of entire history for prediction
- C. They model memoryless state transitions with probability
- D. They are unable to predict any future system states

Answer: C

Q1191. For real-time systems, why is worst-case complexity more critical than average-case?

- A. Average case guarantees meeting real-time deadlines
- B. Worst case determines if deadlines are always met here
- C. Real-time systems have no time constraints whatsoever
- D. Average and worst case are always the same value given

Answer: B

Q1192. How does A* search combine heuristic and actual cost for optimal paths?

- A. It always explores every node in the entire graph now
- B. It ignores heuristics and only uses actual path cost
- C. It uses $f(n) = g(n) + h(n)$ to guide the search path
- D. It randomly explores nodes without any cost function

Answer: C

Q1193. Why might approximation be preferred over exact for NP-hard problems?

- A. Approximation gives near-optimal results in polynomial time
- B. Approximation algorithms always find the exact best answer
- C. NP-hard problems are always easy to solve exactly in time
- D. Exact algorithms always run faster for NP-hard problems

Answer: A

Q1194. How does amortized $O(1)$ apply to dynamic array resizing?

- A. Every individual insertion takes $O(n)$ time to complete
- B. Resizing always takes constant time for every operation
- C. Expensive resizing is rare enough to average $O(1)$ cost
- D. Dynamic arrays never need to resize their storage space

Answer: C

Q1195. What advantage does randomized quicksort have over deterministic quicksort?

- A. Randomization makes the algorithm produce wrong results
- B. Deterministic quicksort handles all inputs optimally here
- C. It avoids worst-case $O(n^2)$ for adversarial input data
- D. Randomized quicksort is always slower in every case

Answer: C

Q1196. When should BFS be chosen over DFS for graph traversal?

- A. When finding any path regardless of its total length
- B. When memory usage must be minimized above all else
- C. When the shortest path in unweighted graph is required
- D. When the deepest node must be found first before rest

Answer: C

Q1197. How do algorithmic paradigms help select problem-solving strategies?

- A. They categorize problems to match with proven techniques
- B. They eliminate the need for understanding the problem
- C. Paradigms are only relevant for theoretical not practical
- D. Paradigms restrict thinking to only one possible approach

Answer: A

Q1198. Why is cache-oblivious algorithm design important for modern architectures?

- A. Modern computers have no memory hierarchy at all now
- B. It optimizes memory access patterns without cache knowledge
- C. Cache-oblivious algorithms are always slower in practice
- D. Memory hierarchy has no effect on algorithm performance

Answer: B

Q1199. How does minimax with alpha-beta pruning optimize game tree search?

- A. It prunes branches that cannot affect the final decision
- B. It only works for games with a single possible outcome
- C. It randomly selects moves without evaluating game state
- D. It explores every node in the entire game tree fully

Answer: A

Q1200. What is the significance of reduction in proving problem complexity?

- A. It proves all problems have polynomial-time solutions
- B. It shows one problem is at least as hard as another one
- C. Reduction makes all problems easier to solve quickly
- D. Reduction eliminates the need for complexity analysis

Answer: B

Q1201. A team unanimously agrees on a risky strategy without discussing alternatives. Which failure is this?

- A. Groupthink suppressing dissent and alternative views
- B. Effective consensus building through open discussion
- C. Healthy team alignment on a well-analyzed strategy
- D. Structured decision making with diverse perspectives

Answer: A

Q1202. How does survivorship bias distort analysis of success factors in startups?

- A. It equally considers both successful and failed startups
- B. It overweights successful cases ignoring failed examples
- C. It focuses specifically on understanding failure patterns
- D. It provides accurate insights by studying all startups

Answer: B

Q1203. Which critical thinking approach best reveals hidden assumptions in a technical proposal?

- A. Focusing only on cost without considering other factors
- B. Accepting the proposal based on presenter credentials
- C. Systematically questioning each premise and dependency
- D. Approving proposals that use the latest technology stack

Answer: C

Q1204. How does the availability heuristic lead to errors in risk assessment?

- A. It ensures accurate risk assessment in every situation
- B. It provides a mathematically rigorous probability model
- C. It eliminates all cognitive biases from risk assessment
- D. Recent or vivid events are overweighted in probability

Answer: D

Q1205. Why might a technically superior solution fail in practice despite logical analysis?

- A. Logical analysis always guarantees practical success here
- B. Technical superiority is the only factor that ever matters
- C. Organizational, cultural, or political factors override logic
- D. Rigorous analysis eliminates all implementation challenges

Answer: C

Q1206. How does the framing effect influence evaluation of equivalent options?

- A. People are immune to how information is framed and shown
- B. Presentation format changes perception of same information
- C. Statistically equivalent options always are seen the same
- D. Framing has no effect on rational decision making here

Answer: B

Q1207. In post-mortem analysis, why is just culture more effective than blame?

- A. Blame assignment encourages honest reporting of errors
- B. Just culture encourages reporting without fear of punishment
- C. Blame assignment always identifies the true root cause now
- D. Just culture prevents anyone from being held accountable

Answer: B

Q1208. How does confirmation bias affect the software testing process?

- A. It has no effect on how testers design their test cases
- B. It ensures all bugs are found during the testing phase
- C. Confirmation bias improves testing coverage automatically
- D. Testers may unconsciously write tests that confirm code works

Answer: D

Q1209. When stakeholders present competing priorities, which framework structures trade-off analysis best?

- A. Multi-criteria decision analysis with weighted scoring
- B. Ignoring all priorities and using developer intuition
- C. Choosing the priority of the most senior stakeholder
- D. Randomly selecting between the competing priorities here

Answer: A

Q1210. How does bounded rationality explain why optimal decisions are often impossible?

- A. Humans always make perfectly optimal rational decisions
- B. Cognitive limits and incomplete information constrain choices
- C. With enough time humans can always find optimal solutions
- D. Bounded rationality only applies to irrational individuals

Answer: B

Q1211. A distributed system produces inconsistent results but all services pass unit tests. What approach helps?

- A. Restarting all services simultaneously to clear the state
- B. Ignoring the inconsistency since all unit tests pass now
- C. Running more unit tests on each individual service here
- D. Distributed tracing to track requests across all services

Answer: D

Q1212. How does time-travel debugging improve analysis of intermittent bugs?

- A. It only works for bugs that occur at specific clock times
- B. It records execution history to replay and inspect states
- C. It speeds up program execution to find bugs more quickly
- D. It prevents bugs from ever occurring in the future here

Answer: B

Q1213. A heisenbug disappears under debugging conditions. What technique helps?

- A. Non-intrusive logging and statistical analysis of behavior
- B. Adding more breakpoints to catch the bug during debug
- C. Removing all error handling code to expose the real bug
- D. Attaching a debugger which will reliably reproduce the bug

Answer: A

Q1214. How does fuzzing complement traditional testing in discovering vulnerabilities?

- A. Fuzzing only tests the normal expected execution paths here
- B. It verifies that existing test cases are correctly written
- C. It provides random unexpected inputs to expose edge cases
- D. Fuzzing replaces the need for any traditional test cases

Answer: C

Q1215. Why might a bug be reproducible in dev but not staging?

- A. Development environments are always more strict than staging
- B. Bugs always behave identically across all environments here
- C. Environments are always identical in every possible way
- D. Configuration data or infrastructure differences mask bug

Answer: D

Q1216. How does static analysis complement dynamic testing?

- A. Static analysis replaces the need for runtime testing here
- B. Static finds code issues without running, dynamic tests behavior
- C. Dynamic testing makes static analysis completely unnecessary
- D. Both approaches find exactly the same types of bugs always

Answer: B

Q1217. Why is a data race in multi-threaded code hard to reproduce?

- A. Thread scheduling non-determinism changes execution order
- B. Data races are easily caught by standard unit testing here
- C. Data races always produce the same behavior every time
- D. Multi-threaded programs always execute in a fixed order

Answer: A

Q1218. How does chaos engineering improve reliability compared to reactive debugging?

- A. It deliberately injects failures to find weaknesses early
- B. It eliminates the need for any monitoring or logging here
- C. It waits for failures to occur then investigates cause
- D. It only applies to systems that have never had any bugs

Answer: A

Q1219. What is the significance of minimal reproducible examples for debugging?

- A. It should include the entire codebase for completeness now
- B. It is unnecessary when the full application code is given
- C. It includes as much code as possible for full context
- D. It isolates the essential conditions that trigger the bug

Answer: D

Q1220. How does observability differ from traditional monitoring for debugging?

- A. Observability and monitoring are exactly the same concept
- B. Observability enables exploring unknown unknowns from output
- C. Traditional monitoring covers all possible failure modes here
- D. Observability eliminates the need for any log collection now

Answer: B

Q1221. An algorithm has $O(n \log n)$ average but $O(n^2)$ worst case. What practical implications does this have?

- A. The algorithm should be rejected due to the worst case
- B. Average and worst case are always the same in practice
- C. The worst case never occurs in practice so ignore it
- D. Input distribution and adversarial scenarios must be considered

Answer: D

Q1222. How does the complexity class P differ from NP?

- A. P is solvable in polynomial time while NP is verifiable
- B. P problems are unsolvable while NP problems are solvable
- C. P and NP are exactly the same complexity class by definition
- D. All NP problems have been proven to be in P class now

Answer: A

Q1223. Why is the constant factor important despite Big-O ignoring it?

- A. For similar Big-O algorithms constants affect real performance
- B. The constant factor has no practical significance at all
- C. Big-O already accounts for all constant factors completely
- D. Constants only matter for theoretical analysis not practice

Answer: A

Q1224. How does the external memory model affect complexity analysis?

- A. Disk access patterns are irrelevant to algorithm performance
- B. I/O complexity measures disk accesses which dominate runtime
- C. Algorithms perform identically regardless of memory hierarchy
- D. External memory has no effect on algorithm running time now

Answer: B

Q1225. What implications does the preprocessing vs query time trade-off have?

- A. Heavy preprocessing can enable fast queries when many queries
- B. Preprocessing never provides any benefit for query speed
- C. The number of queries is irrelevant to this trade-off here
- D. Preprocessing always makes every query slower not faster

Answer: A

Q1226. How does parallel algorithm complexity differ from sequential?

- A. Parallel algorithms always achieve linear speedup exactly
- B. Parallel algorithms never provide any speedup over sequential
- C. Sequential algorithms are always faster than parallel versions
- D. Work and span define parallel complexity and potential speedup

Answer: D

Q1227. Why might worse theoretical complexity outperform better in practice?

- A. Input size never affects which algorithm performs better here
- B. Worse theoretical complexity always means worse real running
- C. Cache behavior and constants can dominate for real inputs
- D. Theoretical complexity perfectly predicts real performance

Answer: C

Q1228. How do output-sensitive algorithms challenge traditional analysis?

- A. Output-sensitive algorithms are always slower than others
- B. Output size is always the same as input size in all cases
- C. Traditional complexity already accounts for all output sizes
- D. Complexity depends on output size which may be much smaller

Answer: D

Q1229. What is the significance of time-space trade-off in constrained environments?

- A. Reducing memory may require more computation and vice versa
- B. Time and space requirements always decrease together here
- C. Memory constraints never affect algorithm design decisions
- D. More memory always means faster execution in every case

Answer: A

Q1230. How does parameterized complexity provide more nuanced analysis?

- A. It considers only input size for complexity analysis here
- B. It always gives the same result as worst-case analysis here
- C. It measures complexity in terms of problem-specific parameters
- D. It ignores all parameters except the total input size given

Answer: C

Q1231. For fast insertion and fast lookup, which data structure combination best addresses this?

- A. A singly linked list for all insertion and lookup tasks
- B. A balanced BST or hash table with appropriate trade-offs
- C. Using only an unsorted array for both operations here
- D. A stack that provides both fast insertion and fast lookup

Answer: B

Q1232. How does memoization transform the computational complexity of Fibonacci?

- A. It reduces complexity from $O(2^n)$ to $O(n)$ linear time
- B. It increases memory usage without any speed improvement
- C. It changes complexity from $O(n)$ to $O(2^n)$ exponential
- D. Memoization has no effect on Fibonacci computation time

Answer: A

Q1233. When should composition be chosen over inheritance for code reuse?

- A. Composition is preferred for flexible has-a relationships
- B. Inheritance and composition are the same thing entirely
- C. Composition prevents any code reuse between components
- D. Inheritance is always the better choice for code reuse

Answer: A

Q1234. How does lazy evaluation improve performance with potentially unused computations?

- A. It prevents any computation from ever being performed
- B. It defers computation until results are actually required
- C. Lazy evaluation always uses more memory than eager does
- D. It computes everything eagerly whether needed or not

Answer: B

Q1235. In concurrent solving, why might immutable structures be preferred?

- A. Mutable structures never cause any concurrency problems
- B. They eliminate synchronization issues between threads
- C. Immutable structures cannot be used in concurrent code
- D. Immutable structures are always slower to access here

Answer: B

Q1236. How does the sliding window technique optimize contiguous subarray problems?

- A. It maintains a window that moves incrementally through
- B. It requires sorting the array before any processing
- C. It examines every possible subarray from scratch each
- D. It only works on arrays with fewer than ten elements

Answer: A

Q1237. For complex state transitions, which programming paradigm best manages this?

- A. Unstructured code with global variables and goto jumps
- B. State machine pattern with explicit state and transitions
- C. Random conditional logic scattered throughout the codebase
- D. Ignoring state management and hoping for correct results

Answer: B

Q1238. How does the two-pointer technique improve efficiency for sorted arrays?

- A. It only works on unsorted arrays not on sorted ones
- B. It solves problems in $O(n)$ that naively require $O(n^2)$
- C. It always requires additional $O(n)$ extra memory to use
- D. It requires three passes through the array for queries

Answer: B

Q1239. Why is understanding constraints crucial before selecting an approach?

- A. Constraints should only be considered after implementing
- B. Constraints are irrelevant to the approach selection
- C. All approaches work regardless of any constraints given
- D. Constraints determine which approaches are feasible here

Answer: D

Q1240. How does Union-Find efficiently solve graph connectivity problems?

- A. It stores the entire adjacency matrix for the graph
- B. It tracks connected components with near-constant operations
- C. It requires $O(n^2)$ time for every connectivity query made
- D. It can only determine connectivity for two specific nodes

Answer: B

Q1241. When building a predictive model, why is train-test split essential?

- A. Training and testing on same data gives true accuracy
- B. Separate test data reveals how well the model generalizes
- C. Testing on training data always shows worst performance
- D. The split is unnecessary when the dataset is large enough

Answer: B

Q1242. How does Simpson's paradox challenge interpretation of aggregated data?

- A. Trends in subgroups can reverse when data is aggregated
- B. Aggregation always improves the accuracy of data trends
- C. Aggregated data always shows the same trend as subgroups
- D. Simpson's paradox only affects very small datasets here

Answer: A

Q1243. What challenges does multicollinearity pose for regression?

- A. Multicollinearity has no effect on regression analysis
- B. It improves reliability of all regression coefficients
- C. Correlated predictors make individual effects unreliable
- D. It always reduces the prediction error of the model now

Answer: C

Q1244. How does the curse of dimensionality affect high-dimensional data?

- A. Data becomes sparse making distance metrics less meaningful
- B. More dimensions always improve model accuracy and speed
- C. High-dimensional data is always easier to analyze than low
- D. The curse only affects two-dimensional data and nothing else

Answer: A

Q1245. Why are randomized controlled trials the gold standard for causal claims?

- A. Observational studies always provide stronger causal claims
- B. They are cheaper and faster than observational studies
- C. They require no statistical analysis of the results given
- D. Random assignment controls for confounding variables here

Answer: D

Q1246. How does feature engineering improve ML model performance?

- A. It only involves deleting features from the original data
- B. Feature engineering is unnecessary for all modern models
- C. It always reduces model accuracy by adding complexity
- D. It creates informative features that capture domain knowledge

Answer: D

Q1247. What ethical challenge does algorithmic bias present?

- A. Algorithmic bias only affects entertainment recommendations
- B. Algorithmic bias ensures perfectly fair decisions always
- C. Biased training data can perpetuate systemic discrimination
- D. Data-driven systems are always free from any human bias

Answer: C

Q1248. How does cross-validation provide more robust performance estimates?

- A. Cross-validation uses the same split every single time
- B. A single split always provides a more reliable estimate now
- C. It averages performance across multiple different data splits
- D. Cross-validation always uses more data for testing than training

Answer: C

Q1249. With imbalanced datasets, why might accuracy be misleading?

- A. Accuracy always reflects true model performance correctly
- B. A model predicting majority class gets high accuracy trivially
- C. Imbalanced datasets make accuracy harder to calculate here
- D. Accuracy is the only valid metric for classification tasks

Answer: B

Q1250. How does data lineage ensure trustworthiness of conclusions?

- A. Data lineage only matters for regulatory compliance now
- B. It eliminates the need for any data validation checks
- C. Data lineage is unrelated to data quality or trust here
- D. It tracks data origins transformations and dependencies

Answer: D

Q1251. How does biomimicry translate natural solutions to engineering challenges?

- A. Natural solutions never apply to technology or engineering
- B. It ignores nature and relies only on human engineering
- C. Biomimicry only applies to biological science research
- D. It studies natural systems to inspire engineered solutions

Answer: D

Q1252. When a team has creative block, which technique most likely generates breakthroughs?

- A. Restricting discussion to only the most obvious solutions
- B. Random stimulus technique introducing unrelated concepts
- C. Ending the creative process and accepting current results
- D. Repeating the same brainstorming approach more intensely

Answer: B

Q1253. How does TRIZ systematically approach creative problem solving unlike brainstorming?

- A. TRIZ relies on random idea generation without structure
- B. TRIZ uses patterns of invention from patent analysis here
- C. TRIZ only works for problems in mechanical engineering
- D. TRIZ eliminates the need for understanding the problem

Answer: B

Q1254. Why might psychological safety be more important than talent for team creativity?

- A. Safety enables risk-taking and sharing unconventional ideas
- B. Psychological safety has no effect on team creative output
- C. Teams perform best when members fear negative consequences
- D. Individual talent is always sufficient for team creativity

Answer: A

Q1255. How does creative destruction relate to innovative problem solving?

- A. It preserves all existing solutions without any changes
- B. Existing solutions should never be replaced or improved
- C. New innovations replace outdated approaches and systems
- D. Destruction of ideas always hinders creative progress now

Answer: C

Q1256. How does minimum viable product reduce innovation risk?

- A. MVP eliminates the need for any market validation work
- B. It tests core assumptions quickly with minimal investment
- C. It requires building the complete product before testing
- D. It always results in a product needing no improvements

Answer: B

Q1257. How does combinatorial creativity generate novel solutions?

- A. It requires discarding all prior knowledge before starting
- B. Combinatorial creativity only replicates existing solutions
- C. It creates ideas from nothing without any prior knowledge
- D. It connects existing concepts in unexpected combinations

Answer: D

Q1258. Why is the tension between exploitation and exploration important for innovation?

- A. Only exploration of new ideas matters for innovation now
- B. Exploitation and exploration are the same activity exactly
- C. Balancing refinement of known and exploration of new ideas
- D. Only exploitation of known solutions matters for success

Answer: C

Q1259. How does the adjacent possible explain emergence of innovations?

- A. Innovation builds incrementally on what currently exists
- B. The adjacent possible concept only applies to biology here
- C. All innovations appear randomly without any preconditions
- D. Innovation can jump to any possibility regardless of now

Answer: A

Q1260. What advantage does considering extreme scenarios provide in planning?

- A. They reveal vulnerabilities and opportunities others miss
- B. Extreme scenarios always lead to impractical solutions now
- C. Only the most likely scenario needs to be considered here
- D. Extreme scenarios are unrealistic and should be ignored

Answer: A

Q1261. When balancing reliability, performance, and cost, which approach best manages these concerns?

- A. Ignoring cost and focusing only on maximum performance
- B. Maximizing all three simultaneously without trade-off
- C. Service level objectives with explicit trade-off decisions
- D. Choosing the cheapest option regardless of reliability

Answer: C

Q1262. How does Conway's Law affect software architecture decisions?

- A. Organization structure has no effect on software design
- B. Conway's Law only applies to hardware not software design
- C. Small teams always produce more complex software systems
- D. System design tends to mirror organizational communication

Answer: D

Q1263. When migrating a legacy system, what approach minimizes business disruption?

- A. Shutting down old system before building new one from scratch
- B. Incremental strangler pattern migrating component by component
- C. Complete rewrite deployed all at once replacing old system
- D. Running both systems indefinitely without any migration plan

Answer: B

Q1264. How should a team handle deadline pressure versus code quality?

- A. Negotiate scope reduction while maintaining quality standards
- B. Ignore the deadline and deliver whenever team is ready now
- C. Always sacrifice quality to meet every deadline set here
- D. Deliver buggy code and fix it in a future release cycle

Answer: A

Q1265. In a production incident, what balances speed and thoroughness?

- A. Immediate mitigation followed by systematic root cause analysis
- B. Thorough root cause analysis before attempting any fix here
- C. Trying random fixes until something appears to work now
- D. Ignoring the incident until customers report it to support

Answer: A

Q1266. How does the build versus buy decision framework apply?

- A. Always buy commercial solutions and never build anything
- B. Evaluate total cost of ownership, fit, and strategic value
- C. Always build custom solutions for every requirement needed
- D. The decision has no long-term impact on the project here

Answer: B

Q1267. What challenge does data sovereignty present for global systems?

- A. Data sovereignty is only a concern for government systems
- B. All countries have identical data storage regulations here
- C. Data can be stored anywhere without any legal implications
- D. Local laws require data to remain within specific jurisdictions

Answer: D

Q1268. How does antifragility apply to designing resilient systems?

- A. Antifragile systems break under any stress or volatility
- B. Antifragile systems improve and strengthen under stress
- C. All systems should be designed to avoid all stressors now
- D. Antifragility means the system never changes or adapts

Answer: B

Q1269. When requirements evolve faster than development, which approach maintains alignment?

- A. Continuous discovery with short feedback loops and adaptation
- B. Freezing requirements at the start and never changing them
- C. Waiting for all requirements to stabilize before any coding
- D. Building features based on predicted future requirements only

Answer: A

Q1270. How does observability-driven development change problem solving?

- A. It eliminates the need for any pre-production testing here
- B. It replaces all monitoring with manual health checks only
- C. Observability is only useful for debugging not development
- D. It instruments systems to understand behavior in production

Answer: D

Q1271. How does documentation-as-code improve maintainability and accuracy?

- A. It keeps documentation separate from code in different repo
- B. It makes documentation harder to update than traditional
- C. It versions documentation alongside code using same workflow
- D. Documentation-as-code eliminates the need for any reviews

Answer: C

Q1272. When communicating architecture to mixed audiences, what approach is most effective?

- A. Layered communication with different detail levels per group
- B. Only communicating with the technical team and no one else
- C. Avoiding all diagrams and using only dense text paragraphs
- D. Using the same highly technical presentation for everyone

Answer: A

Q1273. How does the C4 model improve communication of software architecture?

- A. It describes architecture at only one level of detail here
- B. It replaces all other forms of technical documentation used
- C. The C4 model only works for microservices architecture now
- D. It provides context, container, component, and code views

Answer: D

Q1274. Why might a team adopt RFC documents for major technical decisions?

- A. RFCs prevent any discussion or debate about decisions
- B. RFCs are only used in internet standards not in companies
- C. They eliminate the need for any meetings about the decision
- D. RFCs enable structured async review and knowledge sharing

Answer: D

Q1275. How does living documentation address the problem of stale docs?

- A. Living docs are written once and never updated afterwards
- B. It eliminates the need for any human-written documentation
- C. Living docs are generated or verified from source of truth
- D. Living documentation is identical to traditional static docs

Answer: C

Q1276. When writing a post-mortem for a major outage, what balance must be struck?

- A. Only management should have access to post-mortem docs
- B. Post-mortems should avoid discussing what went wrong
- C. Focus entirely on blaming individuals who caused issue
- D. Blameless analysis focuses on systemic causes prevention

Answer: D

Q1277. How do structured communication frameworks like STAR help present work?

- A. Structured frameworks prevent any creative communication
- B. They make communication longer without adding clarity
- C. They organize information into situation task action result
- D. STAR only works for job interviews not project contexts

Answer: C

Q1278. What role do knowledge management systems play for problem-solving expertise?

- A. They capture and make searchable collective solving knowledge
- B. They replace the need for any experienced team members
- C. Knowledge management has no effect on team productivity
- D. They only store administrative documents not technical ones

Answer: A

Q1279. How does diagramming-as-code improve consistency and version control?

- A. It eliminates ability to include diagrams in any documents
- B. It generates diagrams from text making them versionable
- C. Diagramming-as-code produces lower quality diagrams always
- D. It requires manual drawing tools for all diagram creation

Answer: B

Q1280. When establishing documentation standards for cross-functional teams, how to reconcile competing needs?

- A. Eliminate all documentation standards let everyone decide
- B. Enforce one rigid format regardless of audience or purpose
- C. Create templates by documentation type with flexible sections
- D. Only the team lead should write all documentation for team

Answer: C

Q1281. How does the concept of computational irreducibility challenge traditional problem-solving approaches?

- A. It means all problems have shortcuts
- B. Some problems require running the full computation with no shortcut possible
- C. It only applies to simple problems
- D. It eliminates the need for algorithms

Answer: B

Q1282. When applying the Cynefin framework, how should a problem in the 'complex' domain be approached differently from one in the 'complicated' domain?

- A. Complex problems need expert analysis; complicated ones need experimentation
- B. Complex problems require probe-sense-respond; complicated ones require sense-analyze-respond
- C. Both domains use the same approach
- D. Complex problems are simpler than complicated ones

Answer: B

Q1283. Why is problem decomposition sometimes counterproductive for highly coupled systems?

- A. Decomposition always works perfectly
- B. Splitting tightly coupled components can obscure critical interactions between them
- C. Coupled systems are always simple
- D. Decomposition is only for hardware

Answer: B

Q1284. How does the concept of 'satisficing' differ from 'optimizing' in real-world problem solving?

- A. They are identical
- B. Satisficing seeks a good-enough solution that meets criteria; optimizing seeks the absolute best
- C. Satisficing always produces worse results
- D. Optimizing is always faster

Answer: B

Q1285. What role does epistemic humility play in effective problem solving?

- A. It has no role
- B. Recognizing the limits of one's knowledge prevents overconfidence and encourages seeking additional information
- C. It means never attempting difficult problems
- D. It only applies to beginners

Answer: B

Q1286. In systems thinking, why is understanding feedback loops essential for solving complex problems?

- A. Feedback loops only exist in electronics
- B. Feedback loops can amplify or dampen effects, creating non-linear behaviors that simple cause-effect analysis misses
- C. They make problems simpler
- D. They are irrelevant to software

Answer: B

Q1287. How does the 'no free lunch' theorem impact the selection of problem-solving strategies?

- A. One strategy works best for all problems
- B. No single strategy outperforms all others across every possible problem, so strategy must be matched to the problem
- C. It means all strategies are equally bad
- D. It only applies to machine learning

Answer: B

Q1288. What is the significance of problem isomorphism in problem solving?

- A. It means problems are always unique
- B. Two problems with the same underlying structure can be solved using the same approach despite different surface appearances
- C. It only applies to mathematics
- D. It makes problems harder

Answer: B

Q1289. Why is the frame problem significant in artificial intelligence and automated problem solving?

- A. It is about picture frames
- B. It concerns determining which aspects of a state change and which remain unchanged after an action
- C. It is a solved problem
- D. It only affects image processing

Answer: B

Q1290. How does emergence complicate reductionist problem-solving approaches?

- A. Emergence is not a real concept
- B. System-level behaviors can arise from component interactions that cannot be predicted by studying components in isolation
- C. It simplifies all problems
- D. It only applies to biology

Answer: B

Q1291. When analyzing a system with emergent behavior, why might traditional requirements gathering be insufficient?

- A. Traditional methods always work
- B. Emergent behaviors arise from component interactions and cannot be specified upfront from individual component analysis
- C. Emergent behavior does not exist
- D. Requirements gathering is never useful

Answer: B

Q1292. How does event storming improve problem analysis for complex business domains?

- A. It replaces all analysis techniques
- B. It maps domain events collaboratively, revealing process flows, boundaries, and hidden complexity
- C. It only applies to simple systems
- D. It eliminates the need for stakeholders

Answer: B

Q1293. Why is sensitivity analysis important when evaluating solutions to problems with uncertain parameters?

- A. It is only for finance
- B. It reveals which input variables most affect outcomes, guiding where to focus validation effort
- C. It adds unnecessary complexity
- D. It always gives definitive answers

Answer: B

Q1294. How does formal methods analysis differ from informal analysis in verifying system correctness?

- A. They are the same approach
- B. Formal methods use mathematical proofs to guarantee properties; informal methods rely on testing and review
- C. Informal methods are always better
- D. Formal methods are faster

Answer: B

Q1295. When a problem has multiple valid framings, how does frame analysis help select the most productive one?

- A. All framings are equally good
- B. Frame analysis evaluates how different perspectives highlight different aspects and solutions
- C. Only one framing ever exists
- D. Frame analysis is purely theoretical

Answer: B

Q1296. What is the significance of identifying invariants during problem analysis?

- A. Invariants are unimportant
- B. Invariants are properties that must always hold true, providing anchors for correctness verification
- C. They only apply to mathematics
- D. They change with every iteration

Answer: B

Q1297. How does theory of constraints analysis identify the most impactful improvement point?

- A. It improves everything simultaneously
- B. It identifies the single bottleneck that limits overall system throughput
- C. It only works for manufacturing
- D. It ignores system interactions

Answer: B

Q1298. Why might Bayesian analysis be preferred over frequentist analysis for problems with prior knowledge?

- A. It is always simpler
- B. Bayesian analysis formally incorporates prior knowledge and updates beliefs as new evidence is gathered
- C. It ignores prior knowledge
- D. Frequentist analysis is always wrong

Answer: B

Q1299. How does sociotechnical systems analysis improve over purely technical analysis?

- A. It ignores people
- B. It considers both human and technical factors, recognizing that system behavior depends on their interaction
- C. It replaces all technical analysis
- D. Social factors never matter in CS

Answer: B

Q1300. What challenge does Goodhart's Law pose when selecting metrics for problem analysis?

- A. It has no relevance
- B. When a measure becomes a target, it ceases to be a good measure because people optimize for it
- C. Metrics are always reliable
- D. It only applies to economics

Answer: B

Q1301. How does paraconsistent logic handle contradictions differently from classical logic?

- A. It ignores contradictions entirely
- B. It allows contradictions to exist without the entire system becoming trivially true
- C. It prevents all contradictions
- D. It is identical to classical logic

Answer: B

Q1302. What is the significance of the resolution inference rule in automated theorem proving?

- A. It is not used in practice
- B. It provides a single, complete inference rule for refutation-based proofs in propositional and first-order logic
- C. It replaces all other logic
- D. It only works for simple problems

Answer: B

Q1303. Why is the distinction between decidable and semi-decidable problems important for logical reasoning?

- A. It has no practical importance
- B. For decidable problems an algorithm always terminates with an answer; for semi-decidable it may run forever on false inputs
- C. All logical problems are decidable
- D. Semi-decidable problems cannot be computed at all

Answer: B

Q1304. How does intuitionistic logic differ from classical logic regarding the law of excluded middle?

- A. It accepts the law fully
- B. It rejects the law of excluded middle, requiring constructive proof of existence rather than proof by contradiction
- C. It is identical to classical logic
- D. It only applies to mathematics

Answer: B

Q1305. What is the role of Skolemization in converting first-order logic formulas for resolution?

- A. It adds more quantifiers
- B. It eliminates existential quantifiers by introducing Skolem functions, enabling clause-form conversion
- C. It removes all variables
- D. It simplifies propositional logic only

Answer: B

Q1306. How does model checking verify properties of finite-state systems?

- A. It checks physical models
- B. It exhaustively explores all possible states to verify whether a specification holds
- C. It only tests one scenario
- D. It replaces formal logic entirely

Answer: B

Q1307. What is the significance of Curry-Howard correspondence for logic and programming?

- A. It is a cooking technique
- B. It establishes that proofs in logic correspond to programs in computation, connecting these fields deeply
- C. It only applies to functional programming
- D. It has been disproven

Answer: B

Q1308. How do many-valued logics extend classical two-valued logic for practical reasoning?

- A. They reduce the number of truth values
- B. They introduce additional truth values beyond true and false to handle uncertainty or partial truth
- C. They eliminate true and false
- D. They are purely theoretical

Answer: B

Q1309. What is the frame problem in the context of logical reasoning about actions?

- A. It is about framing pictures
- B. It concerns efficiently representing what does not change when an action is performed in a logical formalism
- C. It has been fully solved
- D. It only affects visual computing

Answer: B

Q1310. Why is the compactness theorem important for first-order logic?

- A. It compresses files
- B. If every finite subset of a set of sentences is satisfiable, then the entire set is satisfiable
- C. It only applies to propositional logic
- D. It makes proofs shorter

Answer: B

Q1311. How does continuation-passing style transform the flow control of a program?

- A. It has no effect on flow
- B. Each function receives an explicit continuation (callback) representing the rest of the computation, making control flow explicit
- C. It removes all functions
- D. It only applies to web programming

Answer: B

Q1312. What problem does backpressure solve in reactive flow control?

- A. It increases data speed
- B. It prevents fast producers from overwhelming slow consumers by providing feedback to regulate flow
- C. It eliminates all data flow
- D. It only applies to plumbing

Answer: B

Q1313. Why is the halting problem relevant to understanding the limits of flow analysis?

- A. It is not relevant
- B. It proves that no general algorithm can determine whether an arbitrary program will terminate or loop forever
- C. It only affects old computers
- D. It has been solved

Answer: B

Q1314. How does the actor model handle concurrent flow control differently from shared-memory threading?

- A. They are identical
- B. Actors communicate through message passing with no shared state, eliminating data races by design
- C. Actors share all memory
- D. Actors are slower in every case

Answer: B

Q1315. What is the significance of control flow integrity in security?

- A. It has no security relevance
- B. It ensures program execution follows only legitimate paths, preventing attacks like return-oriented programming
- C. It makes programs slower
- D. It only applies to web apps

Answer: B

Q1316. How does algebraic effects provide a structured alternative to exceptions and monads for control flow?

- A. It eliminates all errors
- B. It separates the declaration of effects from their handling, allowing composable and resumable control flow
- C. It is identical to try-catch
- D. It only works in Haskell

Answer: B

Q1317. What is the role of Petri nets in modeling concurrent and distributed flow control?

- A. They are fishing nets
- B. They formally model concurrent processes with places, transitions, and tokens, capturing synchronization and conflict
- C. They only model sequential flow
- D. They are obsolete

Answer: B

Q1318. How does speculative execution affect the intended flow of an algorithm?

- A. It strictly follows the algorithm
- B. The processor executes instructions ahead of confirmed branches, potentially needing to discard results if the prediction was wrong
- C. It never makes predictions
- D. It only applies to interpreted languages

Answer: B

Q1319. What is the difference between structured concurrency and unstructured concurrency?

- A. There is no difference
- B. Structured concurrency ties the lifetime of concurrent tasks to a lexical scope, ensuring all complete before the scope exits
- C. Unstructured is always better
- D. Structured concurrency prevents all parallelism

Answer: B

Q1320. How does dataflow programming differ from traditional control flow programming?

- A. They are the same
- B. In dataflow, execution is driven by data availability rather than a prescribed sequence of instructions
- C. Dataflow is always slower
- D. Dataflow eliminates all parallelism

Answer: B

Q1321. How does a skip list achieve probabilistic $O(\log n)$ search while maintaining simplicity over balanced trees?

- A. It uses sorted arrays
- B. It maintains multiple layers of linked lists with probabilistic promotion, providing expected $O(\log n)$ without complex rotations
- C. It is always $O(n)$
- D. It requires more memory than any tree

Answer: B

Q1322. What trade-offs does a content-addressable storage system make compared to location-addressable storage?

- A. They have no differences
- B. Content-addressable enables deduplication and integrity verification but requires computing hashes and complicates updates
- C. Content-addressable is always faster
- D. Location-addressable cannot store files

Answer: B

Q1323. How do CRDTs (Conflict-free Replicated Data Types) enable eventual consistency without coordination?

- A. They prevent all conflicts
- B. They are designed so that concurrent updates can always be merged deterministically without conflicts
- C. They require a central server
- D. They only work for text

Answer: B

Q1324. Why might a rope data structure be preferred over a string for text editing applications?

- A. Ropes are simpler
- B. Ropes represent strings as balanced binary trees of shorter strings, enabling $O(\log n)$ insertion and deletion
- C. Ropes are always slower
- D. Strings are better for all uses

Answer: B

Q1325. How does a Merkle tree enable efficient verification of large data structures?

- A. It stores data linearly
- B. Each non-leaf node contains the hash of its children, so any change propagates up, allowing efficient verification by checking only the path from root
- C. It prevents all data changes
- D. It only works for blockchains

Answer: B

Q1326. What is the significance of algebraic data types in representing domain models?

- A. They are only for algebra
- B. They combine sum types (tagged unions) and product types (structs) to precisely model domain concepts and make illegal states unrepresentable
- C. They replace all databases
- D. They are slower than primitives

Answer: B

Q1327. How does column-oriented storage improve analytical query performance compared to row-oriented storage?

- A. It does not improve performance
- B. Storing data by columns enables better compression and allows reading only the columns needed for a query
- C. Columns are always slower
- D. Row storage is always better

Answer: B

Q1328. What challenge does the object-relational impedance mismatch present?

- A. There is no mismatch
- B. Object-oriented models use inheritance, encapsulation, and identity while relational models use tables and joins, creating a fundamental tension in data mapping
- C. They map perfectly
- D. It only affects old systems

Answer: B

Q1329. How do probabilistic data structures like HyperLogLog achieve extreme space efficiency?

- A. They store all data precisely
- B. They trade exactness for space by using randomization to provide approximate answers with bounded error
- C. They use more space than exact structures
- D. They cannot provide useful answers

Answer: B

Q1330. Why is the choice between normalized and denormalized data representation critical for system design?

- A. It makes no difference
- B. Normalization reduces redundancy and anomalies but requires joins; denormalization improves read performance but risks inconsistency
- C. Normalization is always best
- D. Denormalization eliminates all problems

Answer: B

Q1331. How does the saga pattern manage distributed transactions across microservices?

- A. It uses a single database transaction
- B. It breaks a distributed transaction into a sequence of local transactions with compensating actions for rollback
- C. It prevents all failures
- D. It only works for monoliths

Answer: B

Q1332. How does recognizing the monotonic stack pattern help solve next-greater-element problems efficiently?

- A. It does not help
- B. A monotonic stack maintains elements in sorted order, allowing $O(n)$ computation of next-greater-element for all positions
- C. It requires $O(n^2)$ time
- D. It only works for sorted arrays

Answer: B

Q1333. Why is recognizing the interval scheduling pattern important for resource allocation problems?

- A. It is not important
- B. Many resource allocation problems reduce to interval scheduling, where greedy algorithms by end time produce optimal solutions
- C. It only applies to calendars
- D. It requires exponential time

Answer: B

Q1334. How does the sidecar pattern extend service functionality in microservices architectures?

- A. It replaces the main service
- B. It deploys auxiliary functionality alongside a service in a separate process, sharing the lifecycle but not the runtime
- C. It merges all services
- D. It only handles logging

Answer: B

Q1335. How does the reservoir sampling pattern enable fair sampling from streams of unknown length?

- A. It requires knowing the stream length
- B. It maintains a fixed-size sample where each element has an equal probability of being included regardless of stream length
- C. It only works for small streams
- D. It produces biased samples

Answer: B

Q1336. What is the significance of recognizing the topological sort pattern in dependency resolution?

- A. It is only for sorting numbers
- B. It identifies that a problem requires ordering items subject to precedence constraints, enabling linear-time solutions via DAG processing
- C. It requires exponential time
- D. It only works for trees

Answer: B

Q1337. How does the outbox pattern ensure reliable event publishing in database-centric systems?

- A. It skips events during failures
- B. It writes events to a database table atomically with state changes, then publishes them asynchronously, ensuring at-least-once delivery
- C. It guarantees exactly-once delivery
- D. It replaces the database

Answer: B

Q1338. How does recognizing anti-patterns help improve software quality?

- A. Anti-patterns are always good
- B. Identifying common negative patterns allows teams to refactor toward known better solutions before problems worsen
- C. Anti-patterns are rare in practice
- D. They only matter for new projects

Answer: B

Q1339. How does the space-partitioning pattern (like quad-trees or k-d trees) optimize spatial queries?

- A. It makes spatial queries slower
- B. It recursively divides space into regions, enabling logarithmic-time range and nearest-neighbor queries by pruning irrelevant regions
- C. It only works in 2D
- D. It requires sorting all points

Answer: B

Q1340. What is the relationship between the visitor pattern and double dispatch in object-oriented design?

- A. They are unrelated
- B. The visitor pattern uses double dispatch to select the correct method based on both the visitor and element types at runtime
- C. Double dispatch is unnecessary
- D. Visitor replaces all other patterns

Answer: B

Q1341. How does generating function methodology help solve counting problems in combinatorics?

- A. It generates random numbers
- B. It encodes a sequence as coefficients of a formal power series, enabling algebraic manipulation to derive closed-form or recursive solutions
- C. It only works for simple sequences
- D. It replaces all counting techniques

Answer: B

Q1342. Why is the extended Euclidean algorithm important for computing modular inverses?

- A. It only computes GCD
- B. It finds integers x and y such that $ax + by = \gcd(a,b)$, which yields the modular inverse when $\gcd$ is 1
- C. It requires floating-point arithmetic
- D. It is slower than brute force

Answer: B

Q1343. How does the probabilistic method prove existence of structures without constructing them?

- A. It always constructs the structure
- B. It shows that a random construction has nonzero probability of having the desired property, proving existence
- C. It disproves existence
- D. It only applies to small sets

Answer: B

Q1344. What is the significance of the fast Fourier transform for polynomial multiplication?

- A. It has no computational benefit
- B. It reduces polynomial multiplication from $O(n^2)$ to $O(n \log n)$ by converting to and from frequency domain
- C. It only works for integers
- D. It increases the complexity

Answer: B

Q1345. How does Ramsey theory relate to the inevitability of patterns in large structures?

- A. It eliminates all patterns
- B. It proves that sufficiently large structures must contain certain ordered substructures regardless of arrangement
- C. It only applies to graphs
- D. It is a programming technique

Answer: B

Q1346. Why is the Akra-Bazzi method more general than the Master Theorem for solving recurrences?

- A. It is less general
- B. It handles non-uniform sub-problem sizes and more general cost functions that the Master Theorem cannot accommodate
- C. It only works for linear recurrences
- D. It gives less accurate results

Answer: B

Q1347. How does information theory quantify the minimum number of bits needed to represent data?

- A. It always uses 8 bits
- B. Shannon's entropy $H = -\sum(p \cdot \log_2(p))$ gives the theoretical lower bound on average bits per symbol
- C. It requires fixed-length encoding
- D. It ignores probability distributions

Answer: B

Q1348. What is the role of linear algebra in solving systems of equations computationally?

- A. It has no role
- B. Matrix representations enable efficient algorithms like Gaussian elimination and LU decomposition for solving large systems
- C. It only works for 2 equations
- D. It is slower than substitution

Answer: B

Q1349. How do Galois fields underpin error-correcting codes like Reed-Solomon?

- A. They are unrelated
- B. Galois fields provide finite arithmetic structures where polynomial operations enable systematic encoding and decoding for error detection and correction
- C. They are only theoretical
- D. They increase errors

Answer: B

Q1350. What is the mathematical foundation of public-key cryptography based on?

- A. Simple addition
- B. The computational asymmetry of certain mathematical problems: easy to compute in one direction but infeasible to reverse without a key
- C. It has no mathematical basis
- D. Only random number generation

Answer: B

Q1351. How does the minimax algorithm with alpha-beta pruning reduce game tree exploration?

- A. It explores every node
- B. It maintains alpha and beta bounds to skip branches that cannot influence the final decision, reducing exploration exponentially
- C. It only works for tic-tac-toe
- D. It eliminates all pruning

Answer: B

Q1352. Why is the distinction between online and offline algorithms important for system design?

- A. It has no importance
- B. Online algorithms must make decisions as data arrives without seeing future input, while offline algorithms have access to all data upfront
- C. Online means internet-based
- D. Offline is always better

Answer: B

Q1353. How does the concept of algorithm reduction help prove computational complexity?

- A. It makes algorithms simpler
- B. If problem A can be reduced to problem B in polynomial time, then B is at least as hard as A
- C. Reductions always make problems easier
- D. It only applies to NP problems

Answer: B

Q1354. What advantage do Las Vegas algorithms have over Monte Carlo algorithms?

- A. They are always faster
- B. Las Vegas algorithms always produce correct results but with variable runtime; Monte Carlo algorithms always terminate but may produce incorrect results
- C. They are identical
- D. Monte Carlo is always better

Answer: B

Q1355. How does the external memory model change algorithm design compared to the RAM model?

- A. It has no effect
- B. It accounts for the cost of transferring data blocks between disk and memory, favoring algorithms that minimize I/O operations
- C. It only affects storage capacity
- D. The RAM model is always sufficient

Answer: B

Q1356. What is the significance of fixed-parameter tractability for NP-hard problems?

- A. It makes all NP-hard problems easy
- B. It shows that some NP-hard problems can be solved efficiently when a specific parameter is small, even if input size is large
- C. It disproves $P \neq NP$
- D. It only applies to graph problems

Answer: B

Q1357. How does the concept of work-stealing improve parallel algorithm scheduling?

- A. It prevents parallelism
- B. Idle processors steal tasks from busy processors' queues, dynamically balancing load without centralized coordination
- C. It requires a central scheduler
- D. It increases total work

Answer: B

Q1358. Why is the competitive ratio important for evaluating online algorithms?

- A. It is not important
- B. It measures the worst-case ratio between the online algorithm's cost and the optimal offline solution's cost
- C. It measures average performance
- D. It only applies to sorting

Answer: B

Q1359. How does the concept of self-stabilization ensure algorithm robustness in distributed systems?

- A. It prevents all errors
- B. A self-stabilizing algorithm eventually reaches a correct state from any arbitrary initial state, recovering from transient faults automatically
- C. It requires manual restart
- D. It only works for single machines

Answer: B

Q1360. What is the significance of communication complexity for understanding distributed algorithm limits?

- A. It measures code length
- B. It quantifies the minimum amount of communication needed between parties to compute a function, establishing lower bounds
- C. It only measures network speed
- D. It is identical to time complexity

Answer: B

Q1361. How does Kahneman's System 1 and System 2 thinking framework explain decision-making errors?

- A. All thinking is the same
- B. System 1 is fast and intuitive but error-prone; System 2 is slow and analytical but often defers to System 1, leading to cognitive biases
- C. System 1 is always correct
- D. System 2 makes more errors

Answer: B

Q1362. Why is the ecological fallacy problematic in data-driven critical thinking?

- A. It is not a real fallacy
- B. It incorrectly infers individual characteristics from aggregate group data, leading to false conclusions about individuals
- C. It only affects ecology
- D. It strengthens conclusions

Answer: B

Q1363. How does the concept of epistemic closure affect critical evaluation of belief systems?

- A. It has no effect
- B. A closed system of beliefs can be internally consistent yet disconnected from reality, resisting disconfirming evidence
- C. It ensures beliefs are correct
- D. It only applies to formal logic

Answer: B

Q1364. What is the significance of Bayesian reasoning for updating beliefs in the face of new evidence?

- A. It has no significance
- B. It provides a formal framework for rationally updating the probability of a hypothesis as new evidence is observed
- C. It always gives absolute answers
- D. It replaces all other reasoning

Answer: B

Q1365. How does the concept of incommensurability challenge rational comparison of competing theories?

- A. All theories are easily comparable
- B. When theories use different concepts and standards, direct comparison may be impossible since there is no neutral ground
- C. It makes all theories equal
- D. It only affects science

Answer: B

Q1366. Why is the base rate fallacy particularly dangerous in medical and security screening?

- A. It is not dangerous
- B. Ignoring the base rate of a condition causes overestimation of positive results' reliability when the condition is rare
- C. Base rates are always known
- D. It only affects medical testing

Answer: B

Q1367. How does the veil of ignorance thought experiment promote fairness in decision-making?

- A. It promotes ignorance
- B. By imagining you do not know your position in society, you design rules that are fair to everyone since you might be anyone
- C. It ignores fairness
- D. It only applies to politics

Answer: B

Q1368. What is the Dunning-Kruger effect's impact on team decision-making in technical contexts?

- A. It has no impact
- B. Less skilled members may overestimate their competence while highly skilled members underestimate theirs, distorting group decisions
- C. It makes everyone equally confident
- D. It only affects individuals

Answer: B

Q1369. How does motivated reasoning differ from rational analysis?

- A. They are the same
- B. Motivated reasoning starts with a desired conclusion and selects evidence to support it, while rational analysis follows evidence to a conclusion
- C. Motivated reasoning is more objective
- D. Rational analysis ignores evidence

Answer: B

Q1370. Why is understanding second-order effects crucial for critical analysis of policy decisions?

- A. Only first-order effects matter
- B. Policies create cascading consequences where the indirect effects may be larger and opposite to the intended first-order effects
- C. Second-order effects are always positive
- D. They are too complex to consider

Answer: B

Q1371. How does differential debugging help identify the cause of a regression?

- A. It compares two debuggers
- B. It systematically compares a working version with a broken version, isolating the specific changes that caused the regression
- C. It only works for mathematics
- D. It requires two computers

Answer: B

Q1372. Why are race conditions particularly difficult to debug in concurrent systems?

- A. They are easy to find
- B. Their occurrence depends on precise timing of thread execution, making them non-deterministic and often unreproducible under debugger conditions
- C. They always produce the same error
- D. They are caught by compilers

Answer: B

Q1373. How does taint analysis help track the propagation of untrusted data through a program?

- A. It checks for stains on hardware
- B. It marks untrusted inputs and tracks their flow through computation, alerting when tainted data reaches sensitive operations
- C. It only works for web applications
- D. It replaces input validation

Answer: B

Q1374. What advantage does record-and-replay debugging provide for non-deterministic bugs?

- A. It has no advantage
- B. It captures the exact execution including all non-deterministic events, allowing deterministic replay for investigation
- C. It only records inputs
- D. It makes programs slower permanently

Answer: B

Q1375. How does property-based testing complement example-based testing in finding bugs?

- A. It is identical to example-based testing
- B. It generates many random inputs based on declared properties, often discovering edge cases that manually written tests miss
- C. It replaces all other testing
- D. It only works for simple functions

Answer: B

Q1376. What is the challenge of debugging in a microservices architecture compared to a monolith?

- A. Microservices are easier to debug
- B. Bugs may span multiple services, requiring correlation of distributed logs, understanding of network interactions, and tracing across service boundaries
- C. There is no difference
- D. Monoliths have more bugs

Answer: B

Q1377. How does mutation testing evaluate the quality of a test suite?

- A. It mutates the tests themselves
- B. It introduces small changes (mutations) to the source code and checks whether the test suite detects them; undetected mutations indicate weak tests
- C. It only tests performance
- D. It replaces all other testing metrics

Answer: B

Q1378. What is the significance of distributed tracing in debugging latency issues across services?

- A. It is unnecessary for performance issues
- B. It correlates request flows across services with timing data, revealing which service or operation is the bottleneck
- C. It only measures network speed
- D. It replaces all logging

Answer: B

Q1379. How does symbolic execution differ from concrete execution in bug finding?

- A. They are identical
- B. Symbolic execution uses symbolic inputs to explore all possible execution paths simultaneously, finding bugs that specific concrete inputs might miss
- C. It is always slower and less useful
- D. It only works for toy programs

Answer: B

Q1380. What role does canary deployment play in debugging production issues?

- A. It involves actual birds
- B. It gradually rolls out changes to a small subset of users first, allowing detection of issues before they affect the entire user base
- C. It is identical to blue-green deployment
- D. It prevents all production bugs

Answer: B

Q1381. How does amortized analysis using the potential method work?

- A. It averages over all algorithms
- B. It assigns a potential function to the data structure state, distributing the cost of expensive operations across many cheap ones
- C. It only works for arrays
- D. It ignores worst-case performance

Answer: B

Q1382. What is the significance of the Strassen algorithm for matrix multiplication?

- A. It is slower than standard multiplication
- B. It was the first to show matrix multiplication can be done in less than $O(n^3)$ time, opening research into faster algorithms
- C. It only works for 2×2 matrices
- D. It has no practical importance

Answer: B

Q1383. Why is the complexity class PSPACE significant in computational theory?

- A. It is the same as P
- B. It contains all problems solvable with polynomial space, and PSPACE-complete problems are believed to be harder than NP-complete ones
- C. It only contains easy problems
- D. It has been proven equal to NP

Answer: B

Q1384. How does the gap between theoretical complexity and practical performance arise?

- A. Theory and practice always agree
- B. Constant factors, cache behavior, branch prediction, and input distribution can make theoretically slower algorithms faster in practice
- C. Theory is always wrong
- D. Practice cannot be measured

Answer: B

Q1385. What is the complexity-theoretic significance of the Cook-Levin theorem?

- A. It proves P equals NP
- B. It proves SAT is NP-complete, establishing the first NP-complete problem and enabling reduction-based proofs for others
- C. It only applies to SAT solvers
- D. It has been disproven

Answer: B

Q1386. How do sublinear algorithms achieve better than $O(n)$ performance?

- A. They are impossible
- B. They sample or process only a portion of the input, providing approximate or probabilistic answers without reading all data
- C. They read all data faster
- D. They compress the input first

Answer: B

Q1387. What is the practical significance of the smoothed analysis framework?

- A. It replaces worst-case analysis entirely
- B. It analyzes algorithms under slight random perturbation of worst-case inputs, often explaining why algorithms like simplex work well in practice despite poor worst-case bounds
- C. It is identical to average-case analysis
- D. It only applies to the simplex method

Answer: B

Q1388. How does the concept of fine-grained complexity refine classical complexity theory?

- A. It makes complexity simpler
- B. It establishes conditional lower bounds within P, showing that improving certain polynomial-time algorithms would require breakthroughs in long-standing open problems
- C. It only classifies NP problems
- D. It eliminates Big-O notation

Answer: B

Q1389. What is the relationship between circuit complexity and algorithm design?

- A. They are unrelated
- B. Circuit complexity lower bounds prove minimum resources needed for computation, and strong enough bounds would resolve P vs NP
- C. Circuits are only for hardware
- D. All circuits are polynomial

Answer: B

Q1390. How does quantum computing potentially change the complexity landscape?

- A. It solves all NP problems
- B. It offers exponential speedup for specific problems like factoring (Shor's) and quadratic speedup for search (Grover's), but not for all problems
- C. Quantum computers are infinitely fast
- D. It makes all problems $O(1)$

Answer: B

Q1391. How does the technique of coordinate compression simplify problems with large value ranges?

- A. It converts all values to zero
- B. It maps sparse values to a compact range preserving relative order, enabling use of arrays instead of hash maps
- C. It only works for coordinates
- D. It increases space usage

Answer: B

Q1392. What is the advantage of using persistent data structures in functional programming?

- A. They save data to disk
- B. They preserve previous versions when modified, enabling safe sharing across threads and efficient undo operations
- C. They are always slower
- D. They are identical to mutable structures

Answer: B

Q1393. How does the technique of small-to-large merging achieve efficient set union operations?

- A. It merges sets randomly
- B. It always merges the smaller set into the larger one, ensuring each element is moved $O(\log n)$ times total
- C. It merges the larger into the smaller
- D. It has no efficiency benefit

Answer: B

Q1394. What problem does the flyweight pattern solve in memory-intensive applications?

- A. It makes objects heavier
- B. It shares common state among many objects to reduce memory usage when many similar objects exist
- C. It eliminates all objects
- D. It only works for graphics

Answer: B

Q1395. How does the technique of bit manipulation solve subset enumeration problems efficiently?

- A. Bits cannot represent subsets
- B. Each bitmask of n bits represents a subset where bit i indicates inclusion of element i , allowing $O(2^n)$ enumeration with $O(1)$ per subset operation
- C. It only works for sets of 8 elements
- D. Bit manipulation is always slower

Answer: B

Q1396. Why is understanding memory layout important for performance-critical programming?

- A. Memory layout has no effect on performance
- B. Cache-friendly data layout exploiting spatial and temporal locality can yield orders-of-magnitude speedup over cache-unfriendly access patterns
- C. It only matters for assembly language
- D. Modern compilers handle all optimization

Answer: B

Q1397. How does the technique of loop unrolling improve performance at the cost of code size?

- A. It makes loops shorter
- B. It processes multiple loop iterations in one pass, reducing branch overhead and enabling instruction-level parallelism
- C. It has no performance benefit
- D. It always causes bugs

Answer: B

Q1398. What is the advantage of event-driven architecture for handling high concurrency?

- A. It requires more threads
- B. A single thread handles many connections by reacting to events via callbacks or async IO, avoiding the overhead of thread-per-connection
- C. It is always slower
- D. It only works for web servers

Answer: B

Q1399. How does the concept of zero-copy optimize data transfer in systems programming?

- A. It copies data zero times per second
- B. It eliminates unnecessary data copying between kernel and user space, reducing CPU and memory bandwidth usage
- C. It always reduces reliability
- D. It only works on Linux

Answer: B

Q1400. Why is understanding instruction-level parallelism important for optimizing tight loops?

- A. Modern CPUs execute one instruction at a time
- B. Modern CPUs can execute multiple independent instructions simultaneously, and code that exposes this parallelism runs faster
- C. It only applies to GPUs
- D. Compilers always optimize this perfectly

Answer: B

Q1401. How does the bootstrap method enable statistical inference without distributional assumptions?

- A. It requires normal distribution
- B. It resamples the observed data with replacement many times, using the distribution of the statistic across resamples to estimate confidence intervals
- C. It only works for large samples
- D. It replaces all statistical tests

Answer: B

Q1402. What challenges does concept drift pose for deployed machine learning models?

- A. It has no effect
- B. The statistical properties of the target variable change over time, degrading model performance and requiring continuous monitoring and retraining
- C. It only affects unsupervised models
- D. It makes models more accurate

Answer: B

Q1403. How does causal inference differ from predictive modeling in data science?

- A. They are the same approach
- B. Causal inference aims to determine if X causes Y using techniques like instrumental variables; predictive modeling only needs correlations for accurate predictions
- C. Predictive modeling always finds causes
- D. Causal inference is less rigorous

Answer: B

Q1404. Why might accuracy be a misleading metric for imbalanced classification problems?

- A. Accuracy is always the best metric
- B. A classifier that always predicts the majority class achieves high accuracy while completely failing on the minority class
- C. Imbalanced datasets do not exist
- D. Accuracy accounts for class imbalance

Answer: B

Q1405. How does differential privacy protect individual information while enabling useful data analysis?

- A. It encrypts all data
- B. It adds carefully calibrated noise to query results so that individual records cannot be distinguished, with mathematically proven privacy guarantees
- C. It deletes sensitive data
- D. It only works for small datasets

Answer: B

Q1406. What is the significance of the bias-variance decomposition for understanding model error?

- A. It is purely theoretical
- B. It decomposes prediction error into bias (systematic error from model simplicity), variance (sensitivity to training data), and irreducible noise
- C. It only applies to linear models
- D. It eliminates all model error

Answer: B

Q1407. How does SHAP (SHapley Additive exPlanations) improve interpretability of complex models?

- A. It simplifies the model itself
- B. It uses game theory's Shapley values to assign each feature a contribution to each prediction, providing consistent and locally accurate explanations
- C. It only works for linear models
- D. It replaces the model

Answer: B

Q1408. Why is the multiple comparisons problem important when testing many hypotheses simultaneously?

- A. Testing many hypotheses is fine without correction
- B. Testing many hypotheses increases the chance of false positives, requiring corrections like Bonferroni or FDR control
- C. It only affects small datasets
- D. It makes all results significant

Answer: B

Q1409. How does survival analysis handle censored data that traditional methods cannot?

- A. It ignores incomplete data
- B. It models time-to-event data where some subjects have not yet experienced the event, using techniques like Kaplan-Meier and Cox regression
- C. It only applies to medical data
- D. Censored data should always be removed

Answer: B

Q1410. What is the significance of the reproducibility crisis for data-driven problem solving?

- A. It is not a real issue
- B. Many published findings fail to replicate, highlighting the importance of rigorous methodology, pre-registration, and transparent analysis for trustworthy results
- C. It only affects academic research
- D. It means all data analysis is wrong

Answer: B

Q1411. How does the theory of inventive problem solving (TRIZ) differ from unstructured creative techniques?

- A. It is identical to brainstorming
- B. TRIZ uses a systematic methodology based on patterns of invention derived from patent analysis, providing structured inventive principles
- C. It eliminates creativity
- D. It only works in engineering

Answer: B

Q1412. Why might diverse teams outperform homogeneous expert teams in creative problem solving?

- A. Diversity always causes conflict
- B. Diverse perspectives prevent groupthink and bring varied mental models that combine in novel ways, offsetting the efficiency of expert consensus
- C. Homogeneous teams are always better
- D. Diversity has no effect on creativity

Answer: B

Q1413. How does the concept of 'exaptation' explain how existing capabilities lead to novel innovations?

- A. It is a misspelling of adaptation
- B. A feature evolved for one purpose is co-opted for a different use, driving innovation through repurposing rather than invention
- C. It only applies to biology
- D. It prevents innovation

Answer: B

Q1414. How does the paradox of expertise affect creative thinking in specialist fields?

- A. Expertise always enhances creativity
- B. Deep expertise can create mental ruts and functional fixedness that prevent seeing novel solutions available to less specialized thinkers
- C. Expertise has no downsides
- D. Only non-experts are creative

Answer: B

Q1415. How does design fiction as a creative technique help explore future possibilities?

- A. It writes science fiction novels
- B. It creates tangible artifacts from an imagined future to provoke discussion about the implications and desirability of emerging technologies
- C. It predicts the future precisely
- D. It has no practical value

Answer: B

Q1416. What is the role of abductive reasoning in creative hypothesis generation?

- A. It replaces all other reasoning
- B. It infers the most likely explanation from incomplete observations, generating creative hypotheses that can then be tested deductively
- C. It is unreliable and useless
- D. It only works in medicine

Answer: B

Q1417. How does the concept of boundary objects facilitate creative collaboration across disciplines?

- A. They are physical walls
- B. They are shared artifacts flexible enough to be meaningful to different communities while maintaining enough structure for productive collaboration
- C. They prevent collaboration
- D. They only exist in museums

Answer: B

Q1418. Why is the incubation effect important for solving problems that resist conscious analytical effort?

- A. It is a myth
- B. Stepping away from a problem allows unconscious processing to continue, often producing insights that conscious effort could not achieve
- C. It only works for artistic problems
- D. Continuous work is always more productive

Answer: B

Q1419. How does the concept of bricolage describe innovation under resource constraints?

- A. It is a type of brick construction
- B. It involves creating solutions by combining whatever resources are available in novel ways rather than seeking purpose-built tools
- C. It requires unlimited resources
- D. It always produces inferior solutions

Answer: B

Q1420. What is the significance of psychological safety for fostering creative innovation in teams?

- A. It makes teams complacent
- B. It enables team members to take creative risks, share unconventional ideas, and learn from failures without fear of punishment
- C. It prevents all conflict
- D. It only matters for new teams

Answer: B

Q1421. How does the concept of antifragility go beyond resilience in system design?

- A. It is the same as resilience
- B. Antifragile systems actually improve and strengthen from stressors and disorder, unlike resilient systems that merely withstand them
- C. It makes systems fragile
- D. It only applies to financial systems

Answer: B

Q1422. How does the CAP theorem constrain distributed system design decisions in practice?

- A. It allows all three guarantees simultaneously
- B. Under network partitions, a system must choose between consistency and availability, with different choices suited to different use cases
- C. It is only theoretical
- D. It has been disproven

Answer: B

Q1423. What is the significance of the 'two pizza rule' for team structure in complex problem solving?

- A. It is about ordering pizza
- B. Small teams (fed by two pizzas) communicate more effectively and make faster decisions than large teams
- C. Larger teams are always better
- D. It only applies to Amazon

Answer: B

Q1424. How does the concept of reversible versus irreversible decisions affect real-world decision-making speed?

- A. All decisions are equally important
- B. Reversible decisions should be made quickly since they can be undone; irreversible decisions warrant more analysis and caution
- C. Irreversible decisions should be made faster
- D. Decision speed does not matter

Answer: B

Q1425. Why is the concept of eventual consistency important for globally distributed systems?

- A. All systems must be immediately consistent
- B. Eventual consistency allows better availability and performance across regions, accepting temporary inconsistency that resolves over time
- C. It means data is never consistent
- D. It only applies to social media

Answer: B

Q1426. How does Site Reliability Engineering (SRE) balance reliability with feature velocity?

- A. It prioritizes reliability over all else
- B. It uses error budgets where excess reliability budget allows more risky feature releases, and exhausted budgets prioritize stability
- C. It ignores reliability
- D. It only measures uptime

Answer: B

Q1427. What are the practical challenges of implementing zero-trust security architecture?

- A. It has no challenges
- B. It requires verifying every request regardless of network location, demanding identity management, micro-segmentation, and continuous validation infrastructure
- C. It is identical to perimeter security
- D. It eliminates all security threats

Answer: B

Q1428. How does the concept of blast radius minimization improve system reliability?

- A. It is about controlling explosions
- B. It limits the scope of potential failures through isolation, ensuring that a single component failure cannot cascade and take down the entire system
- C. It prevents all failures
- D. It only applies to physical systems

Answer: B

Q1429. Why is the build-versus-buy decision critical for real-world engineering organizations?

- A. Always build everything custom
- B. Building provides customization but costs more time and maintenance; buying saves time but creates dependencies, requiring analysis of strategic value
- C. Always buy off-the-shelf solutions
- D. The decision has no impact

Answer: B

Q1430. How does the concept of toil reduction from SRE improve engineering productivity?

- A. Toil is unavoidable and should be accepted
- B. Identifying and automating repetitive manual operational work frees engineers to focus on improvements that compound in value
- C. Automation is always too expensive
- D. Toil reduction only applies to large companies

Answer: B

Q1431. How does the concept of documentation as executable specifications bridge the gap between docs and tests?

- A. Documentation cannot be executable
- B. Tools like Cucumber or doctest embed testable examples in documentation, ensuring docs stay accurate because they are verified by automated testing
- C. It replaces all manual documentation
- D. It only works for simple functions

Answer: B

Q1432. How does the Divio documentation framework classify documentation into four types?

- A. It uses only two types
- B. It separates docs into tutorials (learning), how-to guides (goals), explanation (understanding), and reference (information), each serving a distinct need
- C. All documentation is the same type
- D. It only applies to APIs

Answer: B

Q1433. What is the significance of the bus factor for knowledge documentation in teams?

- A. It relates to public transportation
- B. A low bus factor means few people hold critical knowledge; documentation mitigates this risk by distributing knowledge across the team
- C. High bus factor is dangerous
- D. It only applies to open-source projects

Answer: B

Q1434. How does the concept of documentation debt parallel technical debt?

- A. Documentation debt does not exist
- B. Like technical debt, documentation debt accumulates when documentation is neglected, increasing the cost of onboarding, maintenance, and knowledge transfer over time
- C. Documentation is never a priority
- D. It resolves itself automatically

Answer: B

Q1435. Why is the distinction between prescriptive and descriptive documentation important?

- A. There is no meaningful distinction
- B. Prescriptive documentation tells how things should be done (standards, guidelines); descriptive documents how things currently work (architecture, behavior)
- C. They serve the same purpose
- D. Only prescriptive docs are needed

Answer: B

Q1436. How does the practice of writing RFCs (Request for Comments) improve technical decision-making?

- A. RFCs are outdated internet documents only
- B. RFCs provide structured proposals that invite feedback from stakeholders before implementation, improving decisions through diverse input and creating permanent records
- C. They slow down all decisions
- D. Only protocol designers use RFCs

Answer: B

Q1437. What role does information architecture play in making large documentation sets navigable?

- A. It is the same as document formatting
- B. It organizes content through meaningful categorization, hierarchies, and navigation paths so users can efficiently find what they need
- C. It only applies to websites
- D. It makes documentation harder to find

Answer: B

Q1438. How does Conway's Law influence the relationship between team communication and system architecture?

- A. It has no influence
- B. Organizations design systems that mirror their communication structures, so documentation and communication patterns directly shape system architecture
- C. It only affects hardware
- D. It has been disproven

Answer: B

Q1439. What is the value of decision logs versus meeting minutes for preserving institutional knowledge?

- A. They are identical
- B. Decision logs capture decisions with context and rationale in a searchable format; meeting minutes capture discussions but often bury decisions in narrative
- C. Meeting minutes are always sufficient
- D. Neither has value

Answer: B

Q1440. How does the concept of progressive disclosure improve documentation usability?

- A. It hides all information
- B. It presents essential information first and reveals additional detail on demand, reducing cognitive overload while keeping comprehensive information accessible
- C. It removes detailed content
- D. It only works for UI design

Answer: B

SUBJECT

Programming (C++, Java, Python)

1,440 Multiple Choice Questions

Programming (C++, Java, Python) — Easy

480 questions

Q1. What is a computer program?

- A. A set of instructions for a computer to execute
- B. A system that manages hardware resources
- C. A physical component inside the computer case
- D. A protocol for transmitting network data

Answer: A

Q2. Which of the following is a low-level programming language?

- A. JavaScript
- B. Kotlin
- C. Assembly
- D. Python

Answer: C

Q3. What does a compiler do?

- A. Allocates and manages runtime memory
- B. Converts entire source code to machine code
- C. Identifies and removes bugs from programs
- D. Executes source code one line at a time

Answer: B

Q4. Which is NOT a programming paradigm?

- A. Procedural Programming
- B. Functional Programming
- C. Structural Networking
- D. Object-Oriented Programming

Answer: C

Q5. What is the purpose of pseudocode?

- A. To plan program logic in plain language
- B. To write actual machine-executable code
- C. To compile a program into bytecode
- D. To debug errors in existing code

Answer: A

Q6. What is a syntax error?

- A. A failure in the underlying hardware
- B. A violation of the language grammar rules
- C. A mistake in the program logic flow
- D. A crash that occurs during runtime only

Answer: B

Q7. What is the difference between a compiler and an interpreter?

- A. An interpreter generates standalone machine code files for later execution
- B. They are functionally identical with no meaningful differences at all
- C. A compiler translates all code at once; an interpreter translates line by line
- D. A compiler always executes faster than an interpreter in every case

Answer: C

Q8. What is an algorithm?

- A. A step-by-step problem-solving procedure
- B. A high-level programming language
- C. A standalone software application
- D. A structured data storage format

Answer: A

Q9. Which of the following is a high-level language?

- A. Python
- B. Assembly
- C. FORTRAN
- D. Haskell

Answer: A

Q10. What does IDE stand for?

- A. Integrated Development Environment
- B. Internal Data Engine
- C. Indexed Database Engine
- D. Internet Development Extension

Answer: A

Q11. Which of the following is a primitive data type in most languages?

- A. List
- B. Integer
- C. Object
- D. Array

Answer: B

Q12. What data type is used to store true or false values?

- A. string
- B. float
- C. boolean
- D. integer

Answer: B

Q13. What is a variable?

- A. A fixed constant value in the program
- B. A reusable block of executable code
- C. A named storage location in memory
- D. A control structure for iteration

Answer: C

Q14. What is the default value of an integer variable in Java?

- A. 1
- B. -1
- C. 0
- D. null

Answer: C

Q15. Which data type would you use to store a single letter?

- A. String
- B. double
- C. boolean
- D. char

Answer: D

Q16. What is the difference between float and double?

- A. They are exactly the same data type
- B. double has more precision than float
- C. double is used for storing integers
- D. float has more precision than double

Answer: B

Q17. What does declaring a variable mean?

- A. Specifying its name and type
- B. Assigning it an initial value
- C. Printing its current value
- D. Removing it from the memory

Answer: A

Q18. Which keyword is used to declare a constant in Java?

- A. static
- B. fixed
- C. final
- D. const

Answer: C

Q19. What is a string?

- A. A boolean value
- B. A single character
- C. A sequence of characters
- D. A numeric data type

Answer: C

Q20. What is type casting?

- A. Removing a variable from program memory
- B. Declaring a variable with a specific name
- C. Converting a value from one type to another
- D. Defining a brand new composite data type

Answer: C

Q21. What does the + operator do with two integers?

- A. Concatenates them together
- B. Compares their values
- C. Multiplies them together
- D. Adds them arithmetically

Answer: D

Q22. What is the result of 10 % 3?

- A. 10
- B. 0
- C. 3
- D. 1

Answer: D

Q23. Which operator is used for assignment in most programming languages?

- A. :=
- B. !=
- C. =
- D. ==

Answer: C

Q24. What does the == operator do?

- A. Declares a new variable in scope
- B. Assigns a value to a variable
- C. Performs addition of two values
- D. Checks if two values are equal

Answer: D

Q25. What is the result of 5 > 3?

- A. true
- B. 3
- C. false
- D. 5

Answer: A

Q26. Which operator performs integer division in Python 3?

- A. //
- B. %
- C. **
- D. /

Answer: A

Q27. What does the NOT (!) operator do?

- A. Reverses a boolean value
- B. Adds two values together
- C. Multiplies two operands
- D. Assigns a new value

Answer: A

Q28. What is the result of true AND false?

- A. false
- B. error
- C. true
- D. null

Answer: A

Q29. What does the ++ operator do?

- A. Adds two variables
- B. Doubles the value
- C. Increments the value by 1
- D. Concatenates strings

Answer: C

Q30. Which operator is used for string concatenation in Java?

- A. +
- B. &
- C. .
- D. ,

Answer: A

Q31. What does an if statement do?

- A. Declares a new variable in the scope
- B. Defines a reusable function or method
- C. Executes code only if condition is true
- D. Repeats a block of code in a loop

Answer: C

Q32. What is a loop?

- A. A conditional branching statement
- B. A construct that repeats code blocks
- C. A syntax for declaring variables
- D. A mechanism for calling functions

Answer: B

Q33. What is the purpose of the 'else' clause?

- A. To terminate the entire program execution
- B. To execute code when if condition is false
- C. To define a reusable function or procedure
- D. To start a new loop iteration immediately

Answer: B

Q34. Which loop checks the condition before executing the body?

- A. both of them
- B. do-while loop
- C. neither of them
- D. while loop

Answer: D

Q35. What does the 'break' statement do in a loop?

- A. Restarts the loop from start
- B. Terminates the loop entirely
- C. Pauses the loop temporarily
- D. Skips the current iteration only

Answer: B

Q36. How many times does a for loop 'for(int i=0; i<5; i++)' execute?

- A. 4
- B. 0
- C. 5
- D. 6

Answer: C

Q37. What is a switch statement?

- A. A syntax for declaring new functions
- B. A type of iterative loop construct
- C. A multi-way branch on expression value
- D. A mechanism for handling exceptions

Answer: C

Q38. What does the 'continue' statement do?

- A. Restarts the loop from the very beginning
- B. Skips rest of iteration and moves to next
- C. Pauses execution until user gives input
- D. Exits and terminates the entire loop

Answer: B

Q39. What is a nested loop?

- A. A loop that has been terminated
- B. A loop without any condition set
- C. A loop that runs only one time
- D. A loop placed inside another loop

Answer: D

Q40. What is the output of: if(false) print('A'); else print('B');?

- A. BA
- B. AB
- C. A
- D. B

Answer: D

Q41. What is a function in programming?

- A. A reusable block performing a specific task
- B. A loop that iterates over a collection
- C. A named variable storing data values
- D. A primitive or composite data type name

Answer: A

Q42. What is a function parameter?

- A. A global variable accessible from anywhere
- B. A variable that receives a passed-in value
- C. The return value that a function produces
- D. The identifier name given to the function

Answer: B

Q43. What does the return statement do?

- A. Declares a variable in the function scope
- B. Sends a value back and exits the function
- C. Terminates the entire running program
- D. Starts a new loop inside the function

Answer: B

Q44. What is the difference between a parameter and an argument?

- A. A parameter is the actual value passed during the function call
- B. A parameter is in the definition; an argument is the passed value
- C. They are exactly the same thing with no differences at all
- D. An argument is defined in the function formal declaration only

Answer: B

Q45. What is a void function?

- A. A function that takes no input parameters
- B. A function with an empty body and no code
- C. A function that does not return any value
- D. A function that is syntactically invalid

Answer: C

Q46. What is a function call?

- A. Writing the definition of a function
- B. Invoking a function to execute its code
- C. Removing a function from the program
- D. Adding a comment above a function

Answer: B

Q47. Can a function call another function?

- A. Only void functions can call other ones
- B. Only if they are defined in the same class
- C. No, functions cannot call other functions
- D. Yes, functions can call other functions

Answer: D

Q48. What is a built-in function?

- A. A function provided by the language or library
- B. A function that you must always write yourself
- C. A function that calls itself recursively
- D. A function defined inside a loop construct

Answer: A

Q49. What is the main function?

- A. Any function found anywhere in the program source
- B. The last function that gets called before exiting
- C. The entry point of execution in languages like C/Java
- D. A mathematical function for numeric calculations

Answer: C

Q50. What happens if a function is defined but never called?

- A. It causes an error
- B. It runs at the end
- C. It runs automatically
- D. Its code is never executed

Answer: D

Q51. What does I/O stand for in programming?

- A. Integer or Object
- B. Index with Order
- C. Internal Operations
- D. Input and Output

Answer: D

Q52. Which function is used to display output in Python?

- A. echo()
- B. write()
- C. print()
- D. printf()

Answer: C

Q53. What is standard input (stdin)?

- A. A persistent network data connection
- B. The default input source like keyboard
- C. A database for storing program data
- D. A file stored on the hard disk drive

Answer: B

Q54. What is standard output (stdout)?

- A. The default output destination like console
- B. A database for persisting program output
- C. A connected printer device only
- D. A file stored on persistent disk storage

Answer: A

Q55. Which function reads a line of input from the user in Python?

- A. input()
- B. scanf()
- C. getLine()
- D. readLine()

Answer: A

Q56. What is a newline character?

- A. A whitespace character for horizontal spacing
- B. A tab character for indenting text in documents
- C. A special character that moves cursor to next line
- D. A backspace character that deletes previous input

Answer: C

Q57. What does printf() do in C?

- A. Closes an open file descriptor handle
- B. Formats and prints output to the console
- C. Reads input from standard input stream
- D. Allocates dynamic memory on the heap

Answer: B

Q58. What is a format specifier?

- A. A variable name in source code
- B. A placeholder defining value display format
- C. A counter used inside loop bodies
- D. A file type extension on disk

Answer: B

Q59. What is console output?

- A. Storing structured data in a database table
- B. Sending formatted data to a connected printer
- C. Displaying information on screen or terminal
- D. Writing data to a file on disk storage

Answer: C

Q60. What does Scanner class do in Java?

- A. Compiles Java source code into bytecode
- B. Reads input from various sources like keyboard
- C. Scans the system for viruses and malware
- D. Writes formatted output to the console

Answer: B

Q61. How are strings typically stored in memory?

- A. As boolean value sequences
- B. As integer values in memory
- C. As an array of characters
- D. As floating-point numbers

Answer: C

Q62. What does the length/len function return for a string?

- A. The index of the last character
- B. The memory size in bytes
- C. The number of characters in the string
- D. The ASCII value of the first character

Answer: C

Q63. What is string concatenation?

- A. Splitting a string into parts
- B. Reversing a string order
- C. Joining two or more strings
- D. Deleting a string entirely

Answer: C

Q64. What is the index of the first character in a string in most languages?

- A. -1
- B. 2
- C. 0
- D. 1

Answer: C

Q65. What does the toUpperCase() method do?

- A. Converts all characters to lowercase
- B. Converts all characters to uppercase
- C. Reverses the character order fully
- D. Removes all whitespace characters

Answer: B

Q66. What is a substring?

- A. The name of a string variable only
- B. A string that contains no characters
- C. A contiguous portion of the string
- D. A complete copy of the entire string

Answer: C

Q67. What is an empty string?

- A. A string with spaces only inside it
- B. An uninitialized variable with no type
- C. A null value with no reference set
- D. A string with zero characters in it

Answer: D

Q68. What does the trim() method do?

- A. Removes leading and trailing whitespace
- B. Converts all characters to lowercase form
- C. Removes every character from the string
- D. Splits the string at every space character

Answer: A

Q69. How do you compare strings in Java?

- A. Using the == operator directly
- B. Using the + concatenation operator
- C. Using the > comparison operator
- D. Using the equals() method call

Answer: D

Q70. What is a character in programming?

- A. A string of exactly two characters
- B. A numeric whole number data type
- C. A single symbol like a letter or digit
- D. A boolean true or false data value

Answer: C

Q71. What is an array?

- A. A loop construct that iterates over data elements
- B. A collection of same-type elements in contiguous memory
- C. A reusable function that performs a specific task
- D. A single variable holding one data value

Answer: B

Q72. What is the index of the first element in an array in C/Java?

- A. 0
- B. -1
- C. 2
- D. 1

Answer: A

Q73. What is an ArrayList in Java?

- A. A doubly linked list of elements
- B. A key-value pair mapping structure
- C. A fixed-size array that cannot change
- D. A dynamic array that grows and shrinks

Answer: D

Q74. What is a list in Python?

- A. A dictionary mapping keys to their associated values
- B. An immutable sequence that cannot be modified at all
- C. A fixed-size array of uniform typed elements
- D. An ordered mutable collection of any element types

Answer: D

Q75. What happens when you access an array index out of bounds?

- A. It returns zero as the default value for any position
- B. It returns null as the default value for the position
- C. It throws an exception or causes undefined behavior
- D. It automatically resizes the array to fit the index

Answer: C

Q76. What is a dictionary/map?

- A. An ordered array with indexed elements
- B. A list of words in alphabetical order
- C. A string type for storing text values
- D. A collection of key-value pair entries

Answer: D

Q77. What is a stack?

- A. A LIFO (Last-In First-Out) structure
- B. A sorted list of ordered elements
- C. A random access data structure
- D. A FIFO (First-In First-Out) structure

Answer: A

Q78. What is a queue?

- A. A LIFO (Last-In First-Out) structure
- B. A sorted array with indexed elements
- C. A FIFO (First-In First-Out) structure
- D. A hierarchical tree data structure

Answer: C

Q79. What is the difference between an array and a linked list?

- A. Arrays can dynamically grow and shrink in size without any overhead
- B. Arrays use contiguous memory with $O(1)$ access; linked lists use nodes
- C. They are functionally identical data structures with no differences
- D. Linked lists are always faster than arrays for every type of operation

Answer: B

Q80. What is a set in programming?

- A. A sorted array of ordered elements
- B. A specialized type of key-value map
- C. A collection with only unique elements
- D. An array that allows duplicate elements

Answer: C

Q81. What is a class in OOP?

- A. A variable type for storing data
- B. A blueprint for creating objects
- C. A loop construct for iteration
- D. A reusable function for operations

Answer: B

Q82. What is an object in OOP?

- A. A standalone function
- B. A class definition in code
- C. A primitive data type
- D. An instance of a class

Answer: D

Q83. What is encapsulation?

- A. Inheriting properties from a parent class definition
- B. A sorting method that orders elements in a collection
- C. A type of iterative loop control structure in code
- D. Bundling data and methods while restricting direct access

Answer: D

Q84. What is inheritance in OOP?

- A. Acquiring properties and methods from a parent class
- B. Creating a new object instance from a class
- C. A form of encapsulation that hides internal details
- D. A design pattern for structuring application code

Answer: A

Q85. What is a constructor?

- A. A static method belonging to the class itself
- B. A getter method that returns a field value
- C. A destructor that cleans up object state
- D. A special method called when creating an object

Answer: D

Q86. What is a method in OOP?

- A. A standalone function not tied to any class
- B. A function defined inside a class for objects
- C. A variable that stores data in the program
- D. A data type that defines value constraints

Answer: B

Q87. What does the 'this' keyword refer to?

- A. The current object instance
- B. A static method reference
- C. The class definition itself
- D. The parent class reference

Answer: A

Q88. What is an access modifier?

- A. A keyword controlling member visibility level
- B. A data type for storing numeric values only
- C. A type of loop construct in the language
- D. A constructor for initializing object state

Answer: A

Q89. What is the 'new' keyword used for?

- A. Deleting an existing object from memory
- B. Declaring a new variable in the scope
- C. Defining a new method in the class
- D. Creating a new instance of a class

Answer: D

Q90. What is polymorphism?

- A. A class relationship defined by inheritance hierarchy
- B. A type of encapsulation hiding internal data details
- C. Different objects responding differently to same method
- D. Having only one single form of behavior always

Answer: C

Q91. What is RAM?

- A. Random Access Memory, volatile memory for programs
- B. A programming language for system-level development
- C. A type of hard drive for persistent data storage
- D. Read-Access Module for hardware communication

Answer: A

Q92. What is the stack in memory?

- A. A memory region storing locals and call info in LIFO
- B. The hard drive used for persistent data storage
- C. A data structure only with no memory role
- D. A type of processor cache for fast data access

Answer: A

Q93. What is the heap in memory?

- A. Read-only memory that cannot be written to
- B. Identical to the stack memory region exactly
- C. A sorted data structure for priority access
- D. A memory region for dynamic memory allocation

Answer: D

Q94. What is a memory leak?

- A. A security vulnerability in network protocols
- B. A hardware failure in the memory modules
- C. A type of buffer overflow in array operations
- D. Allocated memory no longer used but not freed

Answer: D

Q95. What is garbage collection?

- A. A sorting algorithm for ordering data values
- B. Deleting files from the disk storage system
- C. Running disk cleanup utilities on the system
- D. Automatic reclamation of unused object memory

Answer: D

Q96. What is a pointer?

- A. A composite data type definition
- B. A variable storing a memory address
- C. A reusable block of executable code
- D. A variable holding a computed value

Answer: B

Q97. What is the difference between stack and heap allocation?

- A. Stack is automatic and fast with limited size; heap is larger but slower
- B. Heap allocation is always faster than stack allocation in all scenarios
- C. There is no meaningful difference between the two allocation methods
- D. Stack memory is unlimited while heap memory has a very strict size cap

Answer: A

Q98. What is a null pointer?

- A. A pointer to the memory address zero
- B. A pointer that has been deleted already
- C. A pointer to the stack memory region
- D. A pointer to no valid memory location

Answer: D

Q99. What does 'free()' do in C?

- A. Frees up disk space on the storage drive
- B. Deallocates previously allocated heap memory
- C. Resets a variable back to its default value
- D. Allocates new memory on the heap for use

Answer: B

Q100. What is a variable's lifetime?

- A. The duration the variable exists in memory
- B. Its declared data type in the program
- C. Its identifier name in the source program
- D. Its accessibility scope in the source code

Answer: A

Q101. What is an exception?

- A. A function call to a subroutine in the program
- B. An abnormal event disrupting normal program flow
- C. A normal expected program event during execution
- D. A variable type for storing error code values

Answer: B

Q102. What is a try-catch block?

- A. A function definition with parameters and body
- B. A looping construct for repeated iteration
- C. A construct attempting code and handling exceptions
- D. A conditional statement for branching logic

Answer: C

Q103. What is the purpose of the 'finally' block?

- A. To define a new function with a specific name
- B. To terminate the program execution immediately
- C. To execute code regardless of exception outcome
- D. To finalize a variable and prevent changes

Answer: C

Q104. What is a NullPointerException?

- A. A compile-time syntax error in the code
- B. A valid and expected operation result
- C. A warning generated by the code linter
- D. An exception from using a null reference

Answer: D

Q105. What does 'throw' do in exception handling?

- A. Ignores an exception and continues running
- B. Explicitly raises and creates an exception
- C. Logs an exception message to the console
- D. Catches an exception from the try block

Answer: B

Q106. What is an ArithmeticException?

- A. A type mismatch between two variable types
- B. A file not found error during I/O access
- C. An exception for illegal arithmetic operations
- D. A syntax error in the source code text

Answer: C

Q107. What is the difference between compile-time and runtime errors?

- A. Compile-time errors cause the running program to crash unexpectedly
- B. Runtime errors are detected and caught by the compiler before running
- C. Compile-time errors are caught before execution; runtime during it
- D. There is no meaningful difference between the two error types

Answer: C

Q108. What is an ArrayIndexOutOfBoundsException?

- A. A type error from wrong array element types
- B. An exception for accessing an invalid index
- C. A memory leak from unreleased array memory
- D. A valid and expected array access result

Answer: B

Q109. Can you have multiple catch blocks for one try?

- A. Yes, to handle different exception types differently
- B. No, only one catch block is allowed per try statement
- C. Only for a maximum of exactly two exception types total
- D. Only in Python, not in any other programming language

Answer: A

Q110. What is an error message?

- A. A variable name used for storing error data
- B. A text description about what went wrong
- C. A function output returned to calling code
- D. A comment written inside the source code

Answer: B

Q111. What is a module in programming?

- A. A data structure for storing elements
- B. A variable type for storing data values
- C. A self-contained unit of related code
- D. A type of iterative loop construct

Answer: C

Q112. What is the purpose of the import statement?

- A. To include functionality from another module
- B. To delete a module from the project files
- C. To compile code into executable binaries
- D. To export code to other modules or files

Answer: A

Q113. What is a package in Python?

- A. A single Python source code file only
- B. A Python class with methods and attributes
- C. A directory containing modules and `__init__.py`
- D. A standalone Python function definition

Answer: C

Q114. What is a library in programming?

- A. A collection of pre-written reusable code modules
- B. A place for storing reference books and manuals
- C. An integrated development environment application
- D. A single function defined in the source code file

Answer: A

Q115. What is the difference between a module and a package?

- A. They are exactly the same thing with different names
- B. A module is a single file; a package is a collection
- C. A package is smaller in size compared to a module
- D. A module contains multiple packages within itself

Answer: B

Q116. What is pip in Python?

- A. A general-purpose programming language for development
- B. A compiler that converts Python to machine code binary
- C. A debugger tool for finding errors in Python programs
- D. The package installer for Python that downloads packages

Answer: D

Q117. What is npm?

- A. A Python-specific development tool only
- B. A database for storing application data values
- C. A compiler for JavaScript source code files
- D. Node Package Manager for JavaScript packages

Answer: D

Q118. What is an API?

- A. A general-purpose programming language for development
- B. A relational database for storing structured data records
- C. A set of definitions and protocols for software integration
- D. A compiler that translates source code to machine language

Answer: C

Q119. What is a namespace?

- A. A string type for storing text data values in memory
- B. A file extension that identifies the programming language
- C. A container for identifiers to avoid naming conflicts
- D. A variable name in the source code program

Answer: C

Q120. What does 'from math import sqrt' do in Python?

- A. Creates a new sqrt function in the current file
- B. Imports only the sqrt function from math module
- C. Deletes the sqrt function from the math module
- D. Imports the entire math module with all functions

Answer: B

Q121. What is a design pattern?

- A. A reusable solution template for common software problems
- B. A type of algorithm for data processing
- C. A user interface visual design layout
- D. A programming language for building applications

Answer: A

Q122. What is recursion?

- A. A function calling itself to solve problems
- B. An error in the program source code
- C. A type of iterative loop construct
- D. A data structure for storing elements

Answer: A

Q123. What is a data structure?

- A. A way of organizing data for efficient access
- B. A programming language for writing software
- C. A function that performs a specific operation
- D. A database management system for storage

Answer: A

Q124. What is an event in event-driven programming?

- A. An action or occurrence that a program can respond to
- B. A function that performs a specific task when called
- C. A syntax error in the source code of the program
- D. A variable that stores data values during execution

Answer: A

Q125. What is a generic/template in programming?

- A. A design pattern for building complex object hierarchies
- B. Code that works with any type specified at compile time
- C. A default value assigned to variables automatically
- D. A runtime check that validates data types dynamically

Answer: B

Q126. What is the Singleton pattern?

- A. Ensuring a class has only one global instance
- B. Creating many instances of a class at once
- C. An error pattern for handling exceptions
- D. A function pattern for recursive operations

Answer: A

Q127. What is abstraction in programming?

- A. A primitive data type for storing numeric values only
- B. Making code unnecessarily complex and hard to follow
- C. Hiding details while exposing essential features only
- D. A type of error that occurs during program execution

Answer: C

Q128. What is an annotation/decorator?

- A. Metadata or wrapper adding behavior without modifying code
- B. A loop construct for iterating over collection elements
- C. A comment in the source code for documentation
- D. A variable that stores a single data value in memory

Answer: A

Q129. What is refactoring?

- A. Adding new features and functionality to the existing codebase
- B. Restructuring code to improve structure without changing behavior
- C. Deleting unused code and removing deprecated function calls
- D. Rewriting the entire codebase from scratch completely

Answer: B

Q130. What is the DRY principle?

- A. A testing principle for validating dry run output
- B. A moisture management rule for server rooms
- C. Avoid duplicating code by abstracting common logic
- D. A drying algorithm for removing duplicate data

Answer: C

Q131. What is concurrency in programming?

- A. Running only one task at a time sequentially
- B. Handling multiple tasks in overlapping time periods
- C. Exactly the same concept as true parallelism
- D. Strictly sequential execution of all operations

Answer: B

Q132. What is a thread?

- A. A reusable function definition block
- B. A type of string data structure
- C. The smallest unit of execution within a process
- D. A data structure for storing elements

Answer: C

Q133. What is a process?

- A. A function call that executes and then returns
- B. An independent program with its own memory space
- C. A variable that stores data during execution
- D. The same thing as a thread in all respects

Answer: B

Q134. What is the difference between concurrency and parallelism?

- A. Concurrency always requires multiple CPU cores to function properly
- B. Parallelism is always slower than concurrency in every scenario
- C. They are exactly the same concept with no differences at all
- D. Concurrency manages overlapping tasks; parallelism runs simultaneously

Answer: D

Q135. What is multithreading?

- A. A type of recursion with multiple recursive calls
- B. Using multiple programming languages in one project
- C. Using multiple processors in a distributed system
- D. Running multiple threads within a single process

Answer: D

Q136. What is a race condition?

- A. A debugging tool for analyzing multi-threaded applications
- B. A bug where outcome depends on non-deterministic thread timing
- C. A performance competition between different threads
- D. A type of exception thrown during concurrent data access

Answer: B

Q137. What is synchronization?

- A. A file operation for reading and writing data safely
- B. Running code statements in sequential order always
- C. A network protocol for data transfer between systems
- D. Coordinating shared resource access to prevent races

Answer: D

Q138. What is a mutex?

- A. An algorithm for sorting elements in a collection
- B. A primitive data type for storing values
- C. A lock ensuring exclusive resource access for one thread
- D. A function for performing calculations on data values

Answer: C

Q139. What is a deadlock?

- A. A memory leak from unreleased dynamically allocated data
- B. A type of infinite loop caused by incorrect conditions
- C. Threads blocked forever waiting for each other resources
- D. A program crash caused by a runtime error in code

Answer: C

Q140. What is the main benefit of concurrency?

- A. Improved performance and responsiveness of programs
- B. Simpler code that is easier to read and understand
- C. Smaller program binaries with reduced memory footprint
- D. Fewer bugs and errors during program execution overall

Answer: A

Q141. What is debugging?

- A. Compiling code into machine language
- B. Writing new source code from scratch
- C. Finding and fixing errors in a program
- D. Running automated tests on the code

Answer: C

Q142. What is a breakpoint?

- A. An error in the source code of the program
- B. A marker where the debugger pauses for inspection
- C. An exit point from a loop during execution flow
- D. The end of a function where control is returned

Answer: B

Q143. What is a unit test?

- A. A user acceptance test for final validation
- B. Testing the entire application end to end
- C. A test verifying a single isolated code unit
- D. A performance test measuring execution speed

Answer: C

Q144. What is a bug?

- A. A type of variable storing debug data
- B. A feature that works as intended
- C. A function used for error handling
- D. An error or flaw causing wrong behavior

Answer: D

Q145. What is print/console debugging?

- A. Using print statements to display values and flow
- B. An automated process for finding and fixing bugs
- C. A graphical user interface tool for visual debugging
- D. Professional debugging using specialized IDE tools

Answer: A

Q146. What is a test case?

- A. A debug tool for stepping through code line by line
- B. A set of inputs and expected results for verification
- C. A code review process performed by team members
- D. A programming language for writing test scripts

Answer: B

Q147. What does the term 'step over' mean in debugging?

- A. Execute current line and move to next without entering calls
- B. Skip the entire program execution completely
- C. Jump directly to the end of the current function body
- D. Restart the debugging session from the very beginning

Answer: A

Q148. What is a syntax error?

- A. A logic mistake producing incorrect program results
- B. A runtime error that occurs during program execution
- C. A design flaw in the overall application architecture
- D. A violation of language grammar rules at compile time

Answer: D

Q149. What is the purpose of an assertion?

- A. To create a new variable in the current scope of code
- B. To import a module from an external library into the file
- C. To define a new function with parameters and return type
- D. To verify a condition is true and fail with message if not

Answer: D

Q150. What is a logic error?

- A. An error where program runs but produces incorrect results
- B. A syntax mistake caught during code compilation by compiler
- C. A type error from assigning incompatible values to variables
- D. A compile error that prevents the program from being built

Answer: A

Q151. What is version control?

- A. Tracking and managing code changes over time
- B. Controlling which program version users access
- C. A testing tool for validating software behavior
- D. A deployment tool for releasing code to servers

Answer: A

Q152. What is Git?

- A. A distributed version control system for code
- B. A general-purpose programming language
- C. A database for storing application data
- D. A compiler for translating source to binaries

Answer: A

Q153. What is a code review?

- A. Deleting unused code from the codebase
- B. Having developers examine code before merging
- C. Writing documentation for the application
- D. Running automated tests on the source code

Answer: B

Q154. What is agile development?

- A. A programming language designed for rapid application development tasks
- B. An iterative approach emphasizing collaboration and short delivery cycles
- C. A testing method for validating software behavior against requirements
- D. Developing software as fast as possible without any planning at all

Answer: B

Q155. What is a commit in version control?

- A. A deployment to the production environment
- B. A branch created for feature development
- C. Deleting code from the shared repository
- D. A snapshot of changes saved with a message

Answer: D

Q156. What is a branch in Git?

- A. An independent line of development from main
- B. A variable used to store temporary values
- C. A copy of a file in the project directory
- D. A function defined in the source code file

Answer: A

Q157. What is a merge in Git?

- A. Creating a new commit in the repository
- B. Splitting code into separate source files
- C. Combining changes from one branch into another
- D. Deleting a branch from the repository

Answer: C

Q158. What is documentation in software development?

- A. Only code comments written inside source files
- B. Written descriptions of software design and usage
- C. A testing tool for validating software correctness
- D. A deployment process for releasing to production

Answer: B

Q159. What is the purpose of code comments?

- A. To define variables inside the source file
- B. To make the code execute faster at runtime
- C. To handle errors during program execution
- D. To explain the intent and reasoning in code

Answer: D

Q160. What is a pull request?

- A. A data request sent to a server for information
- B. A build request to compile and package the application
- C. Pulling code packages from the internet registry
- D. A request to merge changes with review and discussion

Answer: D

Q161. What is source code in programming?

- A. Human-readable instructions for a program
- B. The memory address of a running process
- C. A hardware signal sent to the processor
- D. The binary output after compilation

Answer: A

Q162. What does a linker do in program building?

- A. Checks source code for syntax errors
- B. Allocates memory for runtime variables
- C. Converts code into an intermediate format
- D. Combines object files into an executable

Answer: D

Q163. Which term describes a fixed value written directly in code?

- A. Keyword
- B. Operator
- C. Literal
- D. Variable

Answer: C

Q164. What is the main purpose of a flowchart?

- A. To encrypt data before transmission
- B. To allocate memory during execution
- C. To compile source code into bytecode
- D. To visually represent program logic flow

Answer: D

Q165. What does the term debugging mean?

- A. Designing the user interface layout
- B. Compiling code into machine language
- C. Writing new features for software
- D. Finding and fixing errors in code

Answer: D

Q166. Which of these is an example of a markup language?

- A. HTML
- B. Python
- C. Java
- D. Ruby

Answer: A

Q167. What is a runtime environment?

- A. A database for storing application data
- B. A text editor for writing source code files
- C. A network protocol for data transmission
- D. Software that provides execution support

Answer: D

Q168. What is machine code composed of?

- A. English keywords and symbols
- B. Flowchart diagram elements
- C. Pseudocode planning statements
- D. Binary digits zeros and ones

Answer: D

Q169. What is the purpose of comments in code?

- A. To speed up program execution time
- B. To declare variables automatically
- C. To handle errors during runtime
- D. To explain code for human readers

Answer: D

Q170. What does open source mean for software?

- A. Software that runs only on the internet
- B. Software with publicly available source code
- C. Software that requires no installation
- D. Software created by a single developer

Answer: B

Q171. What is a Boolean data type used to represent?

- A. Text strings
- B. Memory addresses
- C. Decimal numbers
- D. True or false

Answer: D

Q172. What is the default value of an uninitialized integer in Java?

- A. Undefined
- B. Null
- C. Zero
- D. One

Answer: C

Q173. Which data type stores a single character?

- A. char
- B. integer
- C. float
- D. string

Answer: A

Q174. What does the term type casting mean?

- A. Creating a new custom data type
- B. Assigning a name to a constant
- C. Deleting a variable from memory
- D. Converting a value to another type

Answer: D

Q175. Which data type is used for numbers with decimal points?

- A. char
- B. integer
- C. float
- D. boolean

Answer: C

Q176. What is the difference between a variable and a constant?

- A. A variable is global; a constant is always local
- B. A constant uses more memory than a variable
- C. A variable is faster than a constant in speed
- D. A variable can change; a constant cannot change

Answer: D

Q177. What does null represent in programming?

- A. A Boolean false state
- B. The absence of a value
- C. The number zero value
- D. An empty text string

Answer: B

Q178. Which keyword declares a constant in many languages?

- A. dim
- B. let
- C. const
- D. var

Answer: C

Q179. What is an unsigned integer?

- A. An integer with unlimited storage capacity
- B. An integer storing only non-negative values
- C. An integer stored as a floating-point value
- D. An integer that can hold negative values only

Answer: B

Q180. What is a string in programming?

- A. A sequence of characters
- B. A type of loop structure
- C. A memory address pointer
- D. A numeric calculation result

Answer: A

Q181. What does the modulus operator (%) return?

- A. The quotient of division
- B. The remainder of division
- C. The absolute difference
- D. The product of two values

Answer: B

Q182. Which operator is used for logical AND in most languages?

- A. ||
- B. ^^
- C. &&
- D. !!

Answer: C

Q183. What does the increment operator (++) do?

- A. Multiplies a value by two
- B. Divides a value by one
- C. Decreases a value by two
- D. Increases a value by one

Answer: D

Q184. What is the assignment operator in most programming languages?

- A. ==
- B. =
- C. !=
- D. <=

Answer: B

Q185. Which operator checks if two values are NOT equal?

- A. ==
- B. >=
- C. <=
- D. !=

Answer: D

Q186. What does the logical OR operator (||) return?

- A. True if at least one operand is true
- B. False only when first operand is false
- C. The sum of two Boolean operand values
- D. True only if both operands are true

Answer: A

Q187. What type of operator is the greater-than sign (>)?

- A. Comparison
- B. Logical
- C. Assignment
- D. Arithmetic

Answer: A

Q188. What does the NOT operator (!) do to a Boolean value?

- A. Doubles the value
- B. Deletes the value
- C. Inverts the value
- D. Converts to string

Answer: C

Q189. Which symbol represents the multiplication operator?

- A. asterisk
- B. forward slash
- C. minus sign
- D. plus sign

Answer: A

Q190. What is the result of 10 / 3 in integer division?

- A. 0
- B. 3.33
- C. 4
- D. 3

Answer: D

Q191. What does a while loop do in programming?

- A. Executes code a fixed number of times
- B. Skips code blocks based on values
- C. Repeats code while condition remains true
- D. Runs code exactly once then stops

Answer: C

Q192. What is the purpose of the break statement?

- A. It restarts the loop from beginning
- B. It pauses execution for specified time
- C. It skips to the next function in code
- D. It exits the current loop or switch

Answer: D

Q193. Which control structure selects between multiple code paths?

- A. for loop
- B. while loop
- C. if-else statement
- D. do-while loop

Answer: C

Q194. What does the continue statement do in a loop?

- A. Ends the loop entirely and exits immediately
- B. Skips the rest of current iteration only
- C. Restarts the loop counter to initial value
- D. Pauses the loop until user provides input

Answer: B

Q195. What is a do-while loop?

- A. A loop needing no ending condition at all
- B. A loop that checks condition after executing
- C. A loop that runs exactly twice every time
- D. A loop that might never execute its body

Answer: B

Q196. What does a switch statement do?

- A. Switches between two different program files
- B. Transfers control to another function call
- C. Toggles a Boolean variable between states
- D. Compares a value against multiple case labels

Answer: D

Q197. What is an infinite loop?

- A. A loop whose condition never becomes false
- B. A loop that processes infinite data items
- C. A loop iterating exactly one million times
- D. A loop that runs fast but eventually stops

Answer: A

Q198. What is the purpose of the else clause?

- A. It always executes regardless of condition
- B. It executes when the if condition is false
- C. It executes when the if condition is true
- D. It ends the program after the if statement

Answer: B

Q199. What are the three parts of a standard for loop?

- A. Try, catch, and finally block sections
- B. Input, process, and output statements
- C. Initialization, condition, and update
- D. Declaration, assignment, and return

Answer: C

Q200. What is nesting in control structures?

- A. Converting loops into conditional checks
- B. Placing one structure inside another one
- C. Removing unused code from the program
- D. Connecting two programs externally together

Answer: B

Q201. What is a function in programming?

- A. A variable that holds numeric values
- B. A data type that stores text values
- C. A reusable block of code with a name
- D. A control structure for looping code

Answer: C

Q202. What is a parameter in a function?

- A. A variable that receives input to function
- B. The name given to the function itself
- C. The data type of the function output
- D. The return value of the function call

Answer: A

Q203. What does the return statement do?

- A. It repeats the function from the start
- B. It sends a value back to the caller
- C. It declares a new variable in the scope
- D. It prints output to the screen display

Answer: B

Q204. What is a function call?

- A. Defining a new function in the program
- B. Renaming a function to a new identifier
- C. Executing a function by using its name
- D. Deleting a function from the program

Answer: C

Q205. What is the difference between an argument and a parameter?

- A. Arguments are types; parameters are values only
- B. They are the same thing with no difference
- C. Arguments are in definition; parameters in calls
- D. Parameters are in definition; arguments in calls

Answer: D

Q206. What is a void function?

- A. A function that cannot be called by others
- B. A function that takes no parameters at all
- C. A function that returns no value to caller
- D. A function that runs only one time ever

Answer: C

Q207. What does the term function signature mean?

- A. The comments written above function body
- B. The digital signature used to secure code
- C. The total number of lines in the function
- D. The name, parameters, and return type info

Answer: D

Q208. What is a built-in function?

- A. A function that cannot be modified at all
- B. A function that runs before main function
- C. A function provided by the language itself
- D. A function written by the programmer only

Answer: C

Q209. What is a recursive function?

- A. A function that takes no arguments at all
- B. A function that runs in parallel threads
- C. A function that calls itself to solve tasks
- D. A function that never returns any values

Answer: C

Q210. What is the scope of a local variable inside a function?

- A. It is shared across all program functions
- B. It is accessible from anywhere in program
- C. It persists after the function has ended
- D. It exists only within that function body

Answer: D

Q211. What does standard input (stdin) refer to?

- A. The file system storage on disk
- B. The network connection to internet
- C. The default input source like keyboard
- D. The screen display output device

Answer: C

Q212. What does standard output (stdout) refer to?

- A. The printer attached to the computer
- B. The default output destination like screen
- C. The database server for data storage
- D. The keyboard input device

Answer: B

Q213. What is a file in the context of I/O?

- A. A temporary variable in program memory
- B. A named collection of data stored on disk
- C. A network connection between two machines
- D. A function that handles user inputs

Answer: B

Q214. What does reading from a file mean?

- A. Renaming the file to a new file name
- B. Retrieving data from a file into program
- C. Deleting the file from the file system
- D. Writing new data to the file on disk

Answer: B

Q215. What is the purpose of the print function?

- A. To create a new variable in the program
- B. To display output to the screen or console
- C. To save data permanently to a file on disk
- D. To read input from the keyboard device

Answer: B

Q216. What is a file path?

- A. The content stored inside a file itself
- B. The location address of a file on system
- C. The date when a file was last modified
- D. The size of a file measured in bytes

Answer: B

Q217. What does writing to a file mean?

- A. Storing data from program into a file
- B. Deleting all contents of the file only
- C. Moving a file to a different directory
- D. Reading data from the file into memory

Answer: A

Q218. What is a console in programming?

- A. A database management system for data
- B. A text-based interface for program I/O
- C. A gaming device for entertainment use
- D. A graphical design tool for layouts

Answer: B

Q219. What does the input function do in Python?

- A. It writes data to a file on the disk
- B. It creates a new variable in the code
- C. It displays formatted output on screen
- D. It reads a line of text from the user

Answer: D

Q220. What is the difference between text and binary file modes?

- A. Text mode is faster; binary mode is slower always
- B. Text mode is for images; binary mode is for documents
- C. Text mode is read only; binary mode is write only
- D. Text mode handles characters; binary handles raw bytes

Answer: D

Q221. What is string concatenation?

- A. Converting string to number
- B. Joining two strings together
- C. Splitting a string into parts
- D. Reversing a string backward

Answer: B

Q222. What does the length property of a string return?

- A. The number of words inside
- B. The index of last character
- C. The memory size in bytes
- D. The number of characters

Answer: D

Q223. What is a substring?

- A. A string used as a variable
- B. A string with no characters
- C. A portion of another string
- D. A string stored in memory

Answer: C

Q224. What is an empty string?

- A. A string that equals null
- B. A string with zero characters
- C. A string containing spaces only
- D. A string holding the digit zero

Answer: B

Q225. What does converting a string to uppercase mean?

- A. Removing all spaces from the string text
- B. Adding exclamation marks after each word
- C. Changing all letters to capital letters
- D. Increasing the font size of the string

Answer: C

Q226. What is a character index in a string?

- A. The total count of characters in string
- B. The numeric position of a character
- C. The ASCII value of the character
- D. The memory address of the character

Answer: B

Q227. What does the trim function do to a string?

- A. Converts the string to lowercase only
- B. Removes leading and trailing whitespace
- C. Shortens the string to ten characters
- D. Removes all characters from the string

Answer: B

Q228. What is an escape character in a string?

- A. A character that exits the program when found
- B. A character that deletes the previous character
- C. A character that converts text into numeric data
- D. A character starting a special sequence in text

Answer: D

Q229. What does string comparison do?

- A. Counts how many strings exist in program
- B. Joins two strings into one combined string
- C. Converts strings into integer numeric values
- D. Checks if two strings are equal or ordered

Answer: D

Q230. What is the newline character used for?

- A. To add a space between two words
- B. To underline the current text line
- C. To delete the current line of text
- D. To move text to the next line down

Answer: D

Q231. What is an array in programming?

- A. A collection of elements stored in order
- B. A function that processes numeric data
- C. A single variable that holds one value
- D. A loop that iterates over user input

Answer: A

Q232. What is an array index?

- A. The data type of array elements
- B. The position number of an element
- C. The memory size of the full array
- D. The total number of elements in array

Answer: B

Q233. What does it mean to traverse an array?

- A. To delete all elements from the array
- B. To visit each element of the array once
- C. To sort the array in ascending order
- D. To copy the array to another variable

Answer: B

Q234. What is a two-dimensional array?

- A. An array organized in rows and columns
- B. An array that stores only two elements
- C. An array with two different index types
- D. An array that can hold two data types

Answer: A

Q235. What is a list in Python?

- A. A read-only sequence of characters
- B. A function that returns multiple values
- C. A fixed-size array of integers only
- D. An ordered mutable collection of items

Answer: D

Q236. What does the push operation do on a stack?

- A. Searches for an element in the stack
- B. Adds an element to the top of stack
- C. Sorts all elements within the stack
- D. Removes the top element from stack

Answer: B

Q237. What is the difference between a stack and a queue?

- A. Both use the same ordering of elements
- B. Stack is LIFO; queue is FIFO ordering
- C. Stack stores numbers; queue stores text
- D. Stack is FIFO; queue is LIFO ordering

Answer: B

Q238. What is the purpose of a dictionary or map?

- A. To store data as key-value pair entries
- B. To count the number of elements in list
- C. To sort data in alphabetical order always
- D. To convert arrays into string format

Answer: A

Q239. What happens when you access an array out of bounds?

- A. The array automatically grows in size
- B. An error or undefined behavior occurs
- C. The array resets to its initial state
- D. The first element is returned instead

Answer: B

Q240. What is a set in programming?

- A. A collection of unique elements only
- B. A sorted sequence of text characters
- C. A fixed-size array of numeric values
- D. An ordered list allowing duplicate values

Answer: A

Q241. What is a class in object-oriented programming?

- A. A function that processes data values
- B. A loop structure for iteration
- C. A variable storing numeric data
- D. A blueprint for creating objects

Answer: D

Q242. What is an object in programming?

- A. A type of loop control structure
- B. A keyword reserved by the language
- C. An instance of a class in memory
- D. A comment in the source code text

Answer: C

Q243. What is inheritance in OOP?

- A. Copying code from one file to another
- B. A class acquiring properties from parent
- C. Renaming a class to a new identifier
- D. Deleting unused classes from program

Answer: B

Q244. What is encapsulation?

- A. Running multiple objects at the same time
- B. Converting objects into primitive data type values
- C. Splitting a class into multiple smaller classes
- D. Bundling data and methods while restricting access

Answer: D

Q245. What is a constructor in a class?

- A. A function that converts objects to strings
- B. A variable that stores the class name text
- C. A special method called when object is created
- D. A method that deletes the object from memory

Answer: C

Q246. What is a method in OOP?

- A. A comment describing the class purpose
- B. A variable declared inside a class body
- C. A function defined within a class body
- D. A data type used only in class files

Answer: C

Q247. What does the keyword 'this' refer to in a class?

- A. The first parameter of any method
- B. The current object instance itself
- C. The main method of the application
- D. The parent class of the current class

Answer: B

Q248. What is polymorphism in OOP?

- A. Creating multiple constructors in a class
- B. Deleting objects after they are used once
- C. Storing different data types in one array
- D. Objects taking many forms via shared interface

Answer: D

Q249. What is an attribute or property of a class?

- A. A variable that holds data within a class
- B. A method that performs calculations on data
- C. A class that inherits from another class
- D. A keyword that defines the class access

Answer: A

Q250. What is an interface in OOP?

- A. A tool for debugging object-oriented programs
- B. A type of variable used in class constructors
- C. A contract specifying methods a class must implement
- D. A graphical user interface for the application

Answer: C

Q251. What is RAM in a computer?

- A. Volatile memory for temporary data storage
- B. A processor that executes program instructions
- C. A permanent storage device like hard drive
- D. A network component for internet connection

Answer: A

Q252. What is the stack in memory management?

- A. A region for storing program source code
- B. A region for network data buffering use
- C. A region for dynamically allocated objects
- D. A region for function calls and local vars

Answer: D

Q253. What is the heap in memory management?

- A. Memory for dynamic allocation at runtime
- B. Memory for function call stack frames
- C. Memory for operating system kernel data
- D. Memory for storing constant values only

Answer: A

Q254. What is garbage collection?

- A. Deleting old files from the file system
- B. Defragmenting the hard drive storage
- C. Automatic reclamation of unused memory
- D. Manually freeing unused memory in code

Answer: C

Q255. What is a memory leak?

- A. When a program uses too much CPU power
- B. When a program writes to read-only memory
- C. When data is lost during a power outage
- D. When allocated memory is never freed up

Answer: D

Q256. What does memory allocation mean?

- A. Reserving a block of memory for program use
- B. Compressing data to use less memory space
- C. Transferring memory between two computers
- D. Deleting data from computer memory entirely

Answer: A

Q257. What is a pointer in programming?

- A. A variable storing a memory address value
- B. A function that points to another function
- C. A graphical cursor on the screen display
- D. A keyword used to create new variables

Answer: A

Q258. What is the difference between static and dynamic memory allocation?

- A. Static is for objects; dynamic is for integers
- B. Static is at compile time; dynamic is at runtime
- C. Static is slower; dynamic is faster always
- D. Static uses the heap; dynamic uses the stack

Answer: B

Q259. What happens when a program runs out of memory?

- A. The program automatically frees old memory up
- B. The program switches to using the CPU cache only
- C. The operating system adds more physical RAM chips
- D. The program crashes or throws an out-of-memory error

Answer: D

Q260. What is a null pointer?

- A. A pointer that stores a numeric zero value
- B. A pointer to the largest memory address
- C. A pointer that does not point to anything
- D. A pointer to the first element of array

Answer: C

Q261. What is an exception in programming?

- A. A normal expected program execution result
- B. An event disrupting normal program flow
- C. A type of variable for storing text
- D. A function for mathematical calculation

Answer: B

Q262. What does the try block do?

- A. It tests the program for performance speed
- B. It retries failed operations multiple times
- C. It wraps code that might throw exceptions
- D. It attempts to fix errors automatically

Answer: C

Q263. What does the catch block do?

- A. It prevents all exceptions from occurring
- B. It handles a specific type of exception
- C. It catches user input from the keyboard
- D. It stops the program from running further

Answer: B

Q264. What does the finally block do?

- A. It runs only when no exception was thrown
- B. It finalizes variable values permanently
- C. It ends the program after error handling
- D. It executes code regardless of exceptions

Answer: D

Q265. What does throwing an exception mean?

- A. Preventing the program from starting at all
- B. Deleting an error from the program memory
- C. Signaling that an error condition has occurred
- D. Converting an error into a warning message

Answer: C

Q266. What is a NullPointerException?

- A. An error from dividing a number by zero
- B. An error from accessing a null reference
- C. An error from invalid string formatting
- D. An error from array index being negative

Answer: B

Q267. What is an ArrayIndexOutOfBoundsException?

- A. Deleting an element from an empty array
- B. Creating an array with too many elements
- C. Accessing an array with invalid index value
- D. Accessing an array with valid index number

Answer: C

Q268. What is error handling?

- A. The process of creating custom error messages
- B. The process of hiding errors from the user
- C. The process of writing error-free code always
- D. The process of managing and responding to errors

Answer: D

Q269. What is the difference between an error and an exception?

- A. Errors are recoverable; exceptions are not recoverable
- B. Errors are serious system issues; exceptions are recoverable
- C. Both errors and exceptions are exactly the same thing
- D. Exceptions are more serious than errors in all cases

Answer: B

Q270. What is a stack trace?

- A. A list of variables currently stored on stack
- B. A report showing the sequence of method calls
- C. A diagram of the program memory layout
- D. A log of all user inputs during execution

Answer: B

Q271. What is a module in programming?

- A. A type of variable for storing numbers
- B. A comment explaining program functionality
- C. A self-contained unit of reusable code
- D. A single line of executable source code

Answer: C

Q272. What is a package in programming?

- A. A collection of related modules grouped together
- B. A compressed file containing program source code
- C. A user interface component for displaying data
- D. A single function within a larger program

Answer: A

Q273. What does the import statement do?

- A. It exports code to another program file
- B. It deletes unused variables from memory
- C. It compiles the program into machine code
- D. It brings external code into current scope

Answer: D

Q274. What is a library in programming?

- A. A network protocol for data communication
- B. A tool for designing graphical user interfaces
- C. A database for storing user information data
- D. A collection of pre-written reusable code

Answer: D

Q275. What is a namespace?

- A. A file that lists all variables in the program
- B. A container preventing naming conflicts in code
- C. A variable that stores the program name string
- D. A function that generates random names for vars

Answer: B

Q276. What is a dependency in software?

- A. A function that runs only in debug build mode
- B. A bug that depends on specific input to occur
- C. A variable that depends on user input value
- D. An external library that a program requires

Answer: D

Q277. What is the purpose of a package manager?

- A. To test programs for bugs and performance
- B. To compile source code into executable files
- C. To install and manage software dependencies
- D. To write new code automatically for developers

Answer: C

Q278. What is an API in the context of libraries?

- A. The documentation website for the library only
- B. The internal source code of the library files
- C. The set of functions and rules for using library
- D. The version number assigned to the library code

Answer: C

Q279. What does exporting a function from a module mean?

- A. Copying the function into a separate backup file
- B. Deleting the function from the current module
- C. Making the function available to other modules
- D. Converting the function into a different language

Answer: C

Q280. What is a framework in programming?

- A. A testing tool for verifying program correctness
- B. A type of database for storing application data
- C. A single function for processing data values
- D. A structure providing foundation for building apps

Answer: D

Q281. What is recursion in programming?

- A. A function that calls itself to solve tasks
- B. A variable that stores multiple data values
- C. A loop that runs a fixed number of times
- D. A technique for sorting array elements fast

Answer: A

Q282. What is a data structure?

- A. A way to organize and store data efficiently
- B. A network protocol for transmitting data
- C. A programming language for data analysis
- D. A tool for designing user interface layouts

Answer: A

Q283. What is a generic type in programming?

- A. A type that can only store integer values
- B. A type that cannot be used in functions
- C. A type specific to one programming language
- D. A type parameterized to work with any type

Answer: D

Q284. What is a design pattern?

- A. A reusable solution to a common coding problem
- B. A database schema for storing application data
- C. A testing method for finding software bugs
- D. A visual design for user interface screens

Answer: A

Q285. What is an iterator?

- A. A counter variable used inside for loops
- B. A class that stores a collection of objects
- C. A function that sorts elements in an array
- D. An object that traverses elements one by one

Answer: D

Q286. What is a linked list?

- A. A sorted array with binary search capability
- B. An array stored in non-contiguous memory blocks
- C. A list of all links on a web page document
- D. A sequence of nodes connected by pointers

Answer: D

Q287. What is a binary tree?

- A. A tree with exactly two levels of depth total
- B. A tree where each node has at most two children
- C. A tree that stores only binary data zero and one
- D. A tree that can only be traversed in one way

Answer: B

Q288. What is Big O notation used for?

- A. Counting the number of lines in source code
- B. Describing algorithm efficiency and scalability
- C. Determining the memory size of data types used
- D. Measuring the exact runtime in milliseconds

Answer: B

Q289. What is a graph data structure?

- A. A set of nodes connected by edges or links
- B. A chart displaying numerical data visually
- C. A database table with rows and column fields
- D. A type of array with two dimensions in it

Answer: A

Q290. What is an abstract class?

- A. A class that cannot be instantiated directly
- B. A class with no methods or properties defined
- C. A class that exists only in pseudocode form
- D. A class that is automatically garbage collected

Answer: A

Q291. What is concurrency in programming?

- A. Using multiple monitors for one computer
- B. Running one task at a time in sequence
- C. Managing multiple tasks overlapping in time
- D. Connecting multiple computers on a network

Answer: C

Q292. What is a thread in programming?

- A. A connection between client and server
- B. A lightweight unit of execution in process
- C. A complete independent running program
- D. A sequence of characters in a string

Answer: B

Q293. What is parallelism?

- A. Writing code in two programming languages
- B. Executing multiple tasks simultaneously on CPUs
- C. Connecting computers in a parallel network
- D. Running tasks one after another in order

Answer: B

Q294. What is a process in operating systems?

- A. A single instruction executed by the CPU
- B. An independent program instance in execution
- C. A file stored on the hard drive disk
- D. A connection between two network devices

Answer: B

Q295. What is a race condition?

- A. A competition between programmers writing code
- B. A performance test comparing two algorithms speed
- C. An error when threads access shared data unsafely
- D. A condition that causes programs to run too fast

Answer: C

Q296. What is a deadlock?

- A. When a program produces incorrect output data
- B. When the CPU overheats from excessive usage
- C. When threads wait for each other indefinitely
- D. When a program runs too slowly for the user

Answer: C

Q297. What is multithreading?

- A. Running multiple threads within one process
- B. Using multiple monitors for programming work
- C. Writing code in multiple languages at once
- D. Storing data in multiple databases together

Answer: A

Q298. What does synchronization mean in concurrent programming?

- A. Coordinating thread access to shared resources
- B. Converting synchronous code to asynchronous
- C. Copying data between two different computers
- D. Keeping system clocks at the same exact time

Answer: A

Q299. What is a mutex?

- A. A lock ensuring only one thread accesses resource
- B. A function for mathematical calculations only
- C. A type of data structure for sorting elements
- D. A variable that stores multiple values at once

Answer: A

Q300. What is the difference between concurrency and parallelism?

- A. Concurrency is faster than parallelism in all execution scenarios
- B. Concurrency runs tasks simultaneously; parallelism overlaps them
- C. Both terms mean exactly the same thing in all programming contexts
- D. Concurrency manages overlapping tasks; parallelism runs them at once

Answer: D

Q301. What is a unit test?

- A. A test of the entire application system
- B. A test of network connectivity speed
- C. A test of a single small code component
- D. A test of the user interface design layout

Answer: C

Q302. What is a breakpoint in debugging?

- A. A marker indicating the end of a code block
- B. A point where the code has a syntax error
- C. A point where the program terminates normally
- D. A marker where execution pauses for inspection

Answer: D

Q303. What is a bug in software?

- A. An error causing unexpected program behavior
- B. A security update for the application code
- C. A comment left in the source code text
- D. A feature that works exactly as intended

Answer: A

Q304. What does a debugger tool do?

- A. It compiles source code into machine language
- B. It designs the user interface for applications
- C. It helps find and fix errors in program code
- D. It automatically writes code for the developer

Answer: C

Q305. What is a test case?

- A. The outer casing of a testing hardware device
- B. A folder containing all source code files used
- C. A specific scenario to verify expected behavior
- D. A function that generates random test data sets

Answer: C

Q306. What is a syntax error?

- A. A violation of programming language grammar rules
- B. An error caused by network connection failure
- C. A mistake in the user interface visual design
- D. An error in the logic of the program algorithm

Answer: A

Q307. What does stepping through code mean?

- A. Executing code one line at a time to inspect flow
- B. Deleting unused lines of code from the program
- C. Running the entire program as fast as possible
- D. Skipping over sections of code during execution

Answer: A

Q308. What is a test suite?

- A. A software tool for writing source code files
- B. A single test case for one specific function
- C. A room where testing equipment is kept stored
- D. A collection of related test cases grouped together

Answer: D

Q309. What is a runtime error?

- A. An error that occurs while the program is running
- B. An error that prevents the IDE from opening up
- C. An error in the project configuration file text
- D. An error detected during code compilation phase

Answer: A

Q310. What is the purpose of print debugging?

- A. To print the source code on paper for review
- B. To display the final output to the end user
- C. To print test results in a formatted report doc
- D. To add print statements to trace program values

Answer: D

Q311. What is version control?

- A. Limiting the number of code files in project
- B. Controlling access permissions to the server
- C. Controlling which version of language is used
- D. Tracking and managing changes to source code

Answer: D

Q312. What is a code review?

- A. Examining code by peers for quality and issues
- B. An automated test that runs on every commit
- C. A process of deleting unused code from files
- D. A tool that formats source code automatically

Answer: A

Q313. What does DRY stand for in programming?

- A. Deploy Run Yield in application lifecycle
- B. Do Repeat Yourself for code clarity
- C. Debug Review Yell in team collaboration
- D. Don't Repeat Yourself in code writing

Answer: D

Q314. What is a branch in version control?

- A. A copy of the main code for independent work
- B. A tool for merging two programs into one
- C. A comment added to track code change history
- D. A permanent deletion of code from repository

Answer: A

Q315. What is refactoring?

- A. Adding new features to existing program code
- B. Restructuring code without changing its behavior
- C. Removing all comments from the source code
- D. Converting code to a different language entirely

Answer: B

Q316. What is a commit in version control?

- A. Running the program to check for any errors
- B. Deleting a file from the code repository
- C. A request to merge code into main branch
- D. A saved snapshot of changes to the codebase

Answer: D

Q317. What is agile software development?

- A. A single long development phase with no feedback
- B. An iterative approach with frequent delivery cycles
- C. A technique for debugging code more efficiently
- D. A method of writing code as fast as possible

Answer: B

Q318. What is documentation in software development?

- A. The user interface of the application system
- B. Written explanations of code and system design
- C. The compiled binary output of the program
- D. The source code of the program itself only

Answer: B

Q319. What is a pull request?

- A. A request to review and merge code into a branch
- B. A request to compile the code into executable
- C. A request to delete code from the repository
- D. A request to pull data from a database server

Answer: A

Q320. What is the purpose of coding standards?

- A. To limit the number of lines per source file
- B. To make all code run faster on every machine
- C. To prevent any bugs from occurring in programs
- D. To ensure consistent readable code across team

Answer: D

Q321. What is a variable in programming?

- A. A named storage location in memory
- B. A hardware component
- C. A type of loop
- D. A function that returns nothing

Answer: A

Q322. What is the purpose of indentation in Python?

- A. It improves readability only
- B. It defines code blocks and scope
- C. It is used for commenting
- D. It speeds up execution

Answer: B

Q323. Which of the following is an example of an interpreted language?

- A. Assembly
- B. Rust
- C. C
- D. Python

Answer: D

Q324. What is the output of a program?

- A. The compiler error messages
- B. The result produced by executing the program
- C. The input given by the user
- D. The source code written by the programmer

Answer: B

Q325. What does the term 'syntax' refer to in programming?

- A. The rules governing the structure of code
- B. The speed of program execution
- C. The process of finding bugs
- D. The meaning behind the code logic

Answer: A

Q326. What is a binary number system?

- A. A system that uses hexadecimal digits
- B. A system based on the number 8
- C. A system that uses digits 0 through 9
- D. A system that uses only 0 and 1

Answer: D

Q327. What is a logic error in a program?

- A. An error that crashes the program immediately
- B. An error caught by the compiler
- C. An error where the program runs but produces wrong results
- D. An error caused by missing semicolons

Answer: C

Q328. What is the role of an operating system for running programs?

- A. It designs the user interface of programs
- B. It writes the source code automatically
- C. It compiles all programs into machine code
- D. It manages hardware resources and provides services to programs

Answer: D

Q329. What is a keyword in a programming language?

- A. Any word used as a variable name
- B. A reserved word with a special meaning in the language
- C. A comment in the source code
- D. A user-defined function name

Answer: B

Q330. What does the term 'execute' mean in programming?

- A. To delete a program from memory
- B. To save a file to disk
- C. To run a program and carry out its instructions
- D. To write source code in an editor

Answer: C

Q331. What is an integer data type used for?

- A. Storing true or false values
- B. Storing whole numbers without decimal points
- C. Storing text values
- D. Storing images and multimedia

Answer: B

Q332. What is a floating-point number?

- A. A number used only for counting
- B. A number stored as text
- C. A number that is always positive
- D. A number that includes a decimal point

Answer: D

Q333. What is the byte data type typically used for?

- A. Storing long text strings
- B. Storing very large decimal numbers
- C. Storing date and time values
- D. Storing small integer values in the range 0 to 255

Answer: D

Q334. What does initializing a variable mean?

- A. Deleting the variable from memory
- B. Assigning a value to a variable for the first time
- C. Changing the data type of a variable
- D. Declaring the variable without a value

Answer: B

Q335. What is a long data type?

- A. A type used exclusively for loop counters
- B. An integer type with a larger range than int
- C. A type that stores only positive numbers
- D. A data type for storing very short strings

Answer: B

Q336. What is the difference between a literal and a variable?

- A. A variable cannot change its value while a literal can
- B. Variables are only used in loops
- C. A literal is a fixed value written directly in code while a variable is a named storage location
- D. Literals are only used for strings

Answer: C

Q337. What is an array data type?

- A. A collection of elements of the same type stored in contiguous memory
- B. A single value stored in memory
- C. A type that stores only Boolean values
- D. A special type of function

Answer: A

Q338. What does the short data type represent?

- A. A text string with few characters
- B. A fractional number with limited precision
- C. An integer type with a smaller range than int
- D. A Boolean value that is always false

Answer: C

Q339. What is a double data type?

- A. An integer type with twice the range of long
- B. A type that doubles the value stored automatically
- C. A type that stores two variables at once
- D. A floating-point type with double the precision of float

Answer: D

Q340. What is a reference type in programming?

- A. A type that can only hold numeric values
- B. A type used only in functional programming
- C. A type that references external files
- D. A type that stores the memory address of data rather than the data itself

Answer: D

Q341. What is an arithmetic operator?

- A. An operator that combines Boolean values
- B. An operator that assigns values to variables
- C. An operator that performs mathematical calculations like addition and subtraction
- D. An operator that compares two values

Answer: C

Q342. What is the result of 15 / 4 in integer arithmetic?

- A. 3
- B. 3.75
- C. 4
- D. 15

Answer: A

Q343. What does the less-than-or-equal operator (<=) check?

- A. Whether two values are exactly equal
- B. Whether the left value is less than or equal to the right value
- C. Whether the left value is greater than the right value
- D. Whether the left value is strictly less than the right value

Answer: B

Q344. What does the subtraction operator (-) do?

- A. It adds two numbers together
- B. It multiplies two numbers
- C. It divides one number by another
- D. It calculates the difference between two numbers

Answer: D

Q345. What is the decrement operator (--) used for?

- A. Increasing a variable's value by one
- B. Decreasing a variable's value by one
- C. Setting a variable to zero
- D. Doubling a variable's value

Answer: B

Q346. What is a comparison operator?

- A. An operator that performs string concatenation
- B. An operator that allocates memory
- C. An operator that compares two values and returns a Boolean result
- D. An operator that modifies the value of a variable

Answer: C

Q347. What does the division operator (/) do?

- A. It compares two numbers
- B. It finds the remainder of division
- C. It divides the left operand by the right operand
- D. It multiplies two numbers together

Answer: C

Q348. What is the compound assignment operator += used for?

- A. Comparing whether two values are equal
- B. Adding a value to a variable and storing the result back
- C. Creating a new variable with an initial value
- D. Concatenating two strings permanently

Answer: B

Q349. What does the equality operator (==) return?

- A. A new variable with the combined value
- B. True if two values are equal, false otherwise
- C. The sum of two values
- D. The larger of two values

Answer: B

Q350. What is the logical AND operator (&&) used for?

- A. Returning true only when both operands are true
- B. Negating a Boolean value
- C. Adding two numbers together
- D. Returning true when at least one operand is true

Answer: A

Q351. What is a conditional statement?

- A. A statement that repeats code multiple times
- B. A statement that terminates the program
- C. A statement that executes code based on whether a condition is true or false
- D. A statement that always executes regardless of conditions

Answer: C

Q352. What is the purpose of the 'elif' or 'else if' clause?

- A. To define a new function
- B. To declare a new variable
- C. To check an additional condition when the previous if condition was false
- D. To end a loop immediately

Answer: C

Q353. How many times does a for loop 'for(int i=1; i<=10; i++)' execute its body?

- A. 11 times
- B. 10 times
- C. 9 times
- D. 1 time

Answer: B

Q354. What is the default behavior when no break is used in a switch case?

- A. The program terminates immediately
- B. Execution falls through to the next case
- C. Only the matched case executes
- D. An error is thrown

Answer: B

Q355. What is a counter variable in a loop?

- A. A variable that stores the exit code
- B. A variable that holds the loop condition
- C. A variable that stores the loop result
- D. A variable that keeps track of how many iterations have occurred

Answer: D

Q356. What happens when a while loop condition is false from the start?

- A. The loop body never executes
- B. The loop body executes once before stopping
- C. The loop runs infinitely
- D. The program crashes with an error

Answer: A

Q357. What is the difference between a for loop and a while loop?

- A. A for loop is used when the number of iterations is known while a while loop is used when it depends on a condition
- B. A for loop cannot use counter variables
- C. A for loop can only iterate 10 times
- D. A while loop always runs at least once

Answer: A

Q358. What is a Boolean condition in an if statement?

- A. A string that names the condition
- B. A numeric value used for calculation
- C. An expression that evaluates to true or false
- D. A reference to another function

Answer: C

Q359. What does the 'return' statement do inside a loop?

- A. It exits only the current loop
- B. It skips to the next iteration of the loop
- C. It restarts the loop from the beginning
- D. It exits the entire function, ending the loop as well

Answer: D

Q360. What is the purpose of curly braces {} in control structures in languages like Java?

- A. They are used for arithmetic grouping
- B. They declare new variables
- C. They define the block of code that belongs to the control structure
- D. They indicate comments in the code

Answer: C

Q361. What is the purpose of a return value in a function?

- A. To end the program immediately
- B. To send a result back to the code that called the function
- C. To declare a new variable
- D. To print output to the console

Answer: B

Q362. What is a function prototype or declaration?

- A. A comment describing what the function does
- B. A statement that specifies the function name, return type, and parameters without providing the body
- C. The complete implementation of a function
- D. A test case for the function

Answer: B

Q363. What is the difference between a local variable and a global variable?

- A. A local variable is accessible only within its function while a global variable is accessible throughout the program
- B. Global variables can only store numbers
- C. Local variables are always faster than global variables
- D. Local variables persist longer than global variables

Answer: A

Q364. Can a function have multiple parameters?

- A. Functions cannot have any parameters
- B. No, functions can only have one parameter
- C. Only built-in functions can have multiple parameters
- D. Yes, functions can have multiple parameters separated by commas

Answer: D

Q365. What happens when a function calls itself?

- A. It creates a new copy of the function in memory permanently
- B. It creates an infinite loop that always runs forever
- C. It performs recursion, which requires a base case to stop
- D. The program crashes immediately

Answer: C

Q366. What is a method in the context of object-oriented programming?

- A. A variable that stores a function reference
- B. A module that contains multiple functions
- C. A function that is defined inside a class and operates on its data
- D. A standalone function not associated with any class

Answer: C

Q367. What is the difference between calling and defining a function?

- A. Calling creates the function while defining runs it
- B. Defining executes the function while calling creates it
- C. There is no difference between them
- D. Defining specifies what the function does while calling executes it

Answer: D

Q368. What is a nested function?

- A. A function that never returns a value
- B. A function defined inside another function
- C. A function that only calls built-in functions
- D. A function that takes no parameters

Answer: B

Q369. What is the purpose of the 'def' keyword in Python?

- A. It deletes a variable from memory
- B. It defines a new function
- C. It defers execution to another thread
- D. It declares a constant value

Answer: B

Q370. What is function reusability?

- A. Copying a function from one language to another
- B. Writing the same code in multiple places
- C. Restarting the program to run functions again
- D. Writing a function once and calling it from multiple places in the program

Answer: D

Q371. What is a file extension used for?

- A. To encrypt the file contents
- B. To increase the file size
- C. To indicate the type or format of a file
- D. To compress the file automatically

Answer: C

Q372. What does opening a file in write mode do?

- A. It creates a new file or overwrites the existing file with new data
- B. It locks the file from other users
- C. It reads the contents of the file
- D. It appends data to the end of the file

Answer: A

Q373. What does opening a file in append mode do?

- A. It adds new data to the end of the file without removing existing content
- B. It opens the file in read-only mode
- C. It deletes all existing content before writing
- D. It creates a backup of the file first

Answer: A

Q374. What is the purpose of closing a file after reading or writing?

- A. To delete the file from disk
- B. To convert the file to a different format
- C. To encrypt the file contents
- D. To release system resources and ensure all data is written

Answer: D

Q375. What is the difference between reading and writing a file?

- A. Writing can only handle text files
- B. Reading retrieves data from a file while writing sends data to a file
- C. Reading modifies the file while writing preserves it
- D. Reading is faster than writing in all cases

Answer: B

Q376. What is the `System.out.println()` method used for in Java?

- A. Printing a line of text to the console with a newline at the end
- B. Reading input from the keyboard
- C. Creating a new string variable
- D. Opening a file for writing

Answer: A

Q377. What is user input in programming?

- A. The source code written by the programmer
- B. The output displayed on the screen
- C. Data provided by the user to the program during execution
- D. Data that the program generates internally

Answer: C

Q378. What is the difference between `stdout` and `stderr`?

- A. `stdout` handles errors while `stderr` handles normal output
- B. `stdout` is for normal program output while `stderr` is for error messages
- C. `stderr` is faster than `stdout`
- D. They are the same output stream with different names

Answer: B

Q379. What does the `readline()` method do?

- A. It deletes a line from a file
- B. It counts the number of lines in a file
- C. It reads a single line of text from input
- D. It reads all lines from a file at once

Answer: C

Q380. What is a file handle or file descriptor?

- A. A reference that the program uses to access an open file
- B. The name of the file on disk
- C. The size of the file in bytes
- D. A physical part of the hard drive

Answer: A

Q381. What is a string literal?

- A. A sequence of characters enclosed in quotation marks written directly in code
- B. A variable that holds text
- C. A number converted to text
- D. A function that returns text

Answer: A

Q382. What does the replace() method do on a string?

- A. It deletes the entire string
- B. It converts the string to a number
- C. It substitutes occurrences of a specified substring with another substring
- D. It reverses the order of characters

Answer: C

Q383. What is the null terminator in C strings?

- A. The character '\0' that marks the end of a string in C
- B. A special character that marks the beginning of a string
- C. A method that deletes a string from memory
- D. A character used to separate words in a string

Answer: A

Q384. What does the indexOf() or find() method do on a string?

- A. It removes a character at a specified position
- B. It counts the total characters in the string
- C. It inserts a character at a specified position
- D. It returns the position of the first occurrence of a substring

Answer: D

Q385. What is the difference between single quotes and double quotes for strings?

- A. Single quotes create mutable strings and double quotes create immutable strings
- B. There is no difference in any programming language
- C. Single quotes are for characters and double quotes for strings in languages like Java and C
- D. Single quotes are faster to process than double quotes

Answer: C

Q386. What does the toLowerCase() method do?

- A. It converts all characters in a string to lowercase letters
- B. It counts the number of lowercase letters
- C. It removes all lowercase letters from the string
- D. It moves lowercase letters to the end of the string

Answer: A

Q387. What does string slicing allow you to do?

- A. Extract a portion of a string by specifying start and end indices
- B. Convert a string to an array
- C. Delete parts of a string permanently
- D. Join two strings together

Answer: A

Q388. What is the startsWith() method used for?

- A. It starts the execution of a string function
- B. It moves the first character to the end of the string
- C. It creates a new string starting with a given character
- D. It checks whether a string begins with a specified prefix

Answer: D

Q389. What does the contains() method check on a string?

- A. Whether the string is empty
- B. Whether the string is stored in memory
- C. Whether the string contains a specified substring anywhere within it
- D. Whether the string contains only numbers

Answer: C

Q390. What is string reversal?

- A. Moving the first character to the end
- B. Creating a new string with characters in reverse order
- C. Deleting the last character of a string
- D. Converting a string to uppercase

Answer: B

Q391. What is the length or size of an array?

- A. The number of elements it contains
- B. The amount of memory it uses in bytes
- C. The maximum value stored in the array
- D. The index of the last element

Answer: A

Q392. What is a dynamic array?

- A. An array that sorts itself automatically
- B. An array that changes its data type automatically
- C. An array that only holds dynamic data
- D. An array that can grow or shrink in size as elements are added or removed

Answer: D

Q393. What does the pop operation do on a stack?

- A. It removes and returns the element from the top of the stack
- B. It adds an element to the top of the stack
- C. It clears all elements from the stack
- D. It returns the bottom element without removing it

Answer: A

Q394. What is the enqueue operation in a queue?

- A. Checking if the queue is empty
- B. Adding an element to the back of the queue
- C. Removing an element from the front
- D. Sorting the elements in the queue

Answer: B

Q395. What is the dequeue operation in a queue?

- A. Removing and returning the element from the front of the queue
- B. Doubling the size of the queue
- C. Adding an element to the front
- D. Reversing the order of elements

Answer: A

Q396. What is a key-value pair?

- A. A data structure that associates a unique key with a corresponding value
- B. Two arrays linked together
- C. Two functions that call each other
- D. A pair of variables with the same data type

Answer: A

Q397. What is the difference between a sorted array and an unsorted array?

- A. Sorted arrays use more memory than unsorted arrays
- B. Sorted arrays can only store numbers
- C. Unsorted arrays cannot be searched
- D. A sorted array has elements arranged in order while an unsorted array has elements in arbitrary order

Answer: D

Q398. What is an empty collection?

- A. A collection that only holds null values
- B. A collection that has not been initialized
- C. A collection that contains no elements
- D. A collection that has been deleted from memory

Answer: C

Q399. What is the difference between a list and a set?

- A. A list cannot be sorted
- B. A set maintains insertion order while a list does not
- C. A set can hold more elements than a list
- D. A list allows duplicate elements and maintains order while a set contains only unique elements

Answer: D

Q400. What is the peek operation on a stack or queue?

- A. It counts the number of elements
- B. It adds a new element to the collection
- C. It returns the top or front element without removing it
- D. It removes the top or front element

Answer: C

Q401. What is a destructor in OOP?

- A. A method that copies one object to another
- B. A method that modifies all object properties
- C. A special method called when an object is being destroyed to clean up resources
- D. A method that creates new objects

Answer: C

Q402. What is an instance variable?

- A. A variable that belongs to a specific object instance of a class
- B. A variable shared by all objects of a class
- C. A variable that cannot change its value
- D. A variable declared outside any class

Answer: A

Q403. What is method invocation in OOP?

- A. Calling a method on an object to execute its code
- B. Deleting a method from a class
- C. Defining a new method inside a class
- D. Copying a method from one class to another

Answer: A

Q404. What is a subclass or child class?

- A. A class that cannot be instantiated
- B. A class that contains only static methods
- C. A class that inherits properties and methods from another class
- D. A class that has no methods

Answer: C

Q405. What is the 'private' access modifier?

- A. It restricts access to only within the same class
- B. It makes the member accessible from the same package
- C. It allows access from any class in the program
- D. It allows access from subclasses only

Answer: A

Q406. What is an abstract method?

- A. A method declared without implementation that must be overridden by subclasses
- B. A method that cannot be called
- C. A method that returns abstract data types
- D. A method with a complete implementation

Answer: A

Q407. What is object instantiation?

- A. Defining a class in source code
- B. Creating a new object from a class using a constructor
- C. Deleting an object from memory
- D. Copying all methods from one class to another

Answer: B

Q408. What is a getter method?

- A. A method that retrieves the value of a private field
- B. A method that sets the value of a private field
- C. A method that creates a new field in the class
- D. A method that deletes a field from the object

Answer: A

Q409. What is a setter method?

- A. A method that sets or updates the value of a private field with optional validation
- B. A method that removes a field from the object
- C. A method that creates a new class
- D. A method that returns the value of a field

Answer: A

Q410. What is the 'public' access modifier?

- A. It allows access only from subclasses
- B. It restricts access to the same class only
- C. It makes the member read-only
- D. It allows access from any class in the program

Answer: D

Q411. What is memory deallocation?

- A. Allocating additional memory for a program
- B. Compressing memory to save space
- C. Moving data from RAM to the hard drive
- D. Releasing memory that is no longer needed back to the system

Answer: D

Q412. What is the stack used for in program execution?

- A. Storing function call information, local variables, and return addresses
- B. Storing all program files on disk
- C. Running the operating system
- D. Permanently storing user data

Answer: A

Q413. Why is the heap used for dynamic memory allocation?

- A. Because it automatically cleans up all allocations
- B. Because it is the only available memory region
- C. Because it allows allocating memory of any size at runtime that persists until explicitly freed
- D. Because it is faster than the stack

Answer: C

Q414. What is a stack overflow?

- A. When the heap runs out of space
- B. When the stack exceeds its maximum size, typically due to deep or infinite recursion
- C. When a variable stores a value that is too large
- D. When the CPU overheats

Answer: B

Q415. What does the 'new' keyword do in languages like Java and C++?

- A. It allocates memory on the heap and creates a new object
- B. It deletes an existing object from memory
- C. It declares a new variable on the stack
- D. It opens a new file for writing

Answer: A

Q416. What is automatic memory management?

- A. Memory that never needs to be allocated
- B. A system where the runtime environment automatically handles memory allocation and deallocation
- C. Memory that is managed by an external program
- D. A tool for manually tracking memory usage

Answer: B

Q417. What is a reference in memory management?

- A. A copy of an object stored in a different location
- B. A comment in the source code
- C. A value that points to or identifies the location of an object in memory
- D. A backup of the original memory

Answer: C

Q418. What happens to local variables when a function returns?

- A. They are moved to the heap
- B. Their memory on the stack is automatically reclaimed
- C. They become global variables
- D. They are saved to disk permanently

Answer: B

Q419. What is the difference between manual and automatic memory management?

- A. Manual is always faster than automatic
- B. Automatic management cannot handle large programs
- C. Manual requires the programmer to allocate and free memory while automatic uses garbage collection
- D. Manual management is only used in interpreted languages

Answer: C

Q420. What is the purpose of the delete keyword in C++?

- A. It deletes an entry from a database
- B. It removes a variable declaration from the code
- C. It frees memory that was allocated with new and calls the destructor
- D. It deletes a source code file

Answer: C

Q421. What is the purpose of error handling in programming?

- A. To speed up program execution
- B. To detect and respond to errors gracefully without crashing the program
- C. To prevent users from running the program
- D. To add more features to the program

Answer: B

Q422. What is a FileNotFoundException?

- A. An exception thrown when a program tries to access a file that does not exist
- B. An error when the file has the wrong extension
- C. An exception when the file system is full
- D. An error when a file is too large to open

Answer: A

Q423. What is an IllegalArgumentException?

- A. An exception when the compiler finds syntax errors
- B. An error when the program has too many arguments
- C. An error when the program runs out of memory
- D. An exception thrown when a method receives an argument that is not valid

Answer: D

Q424. What happens if an exception is not caught?

- A. The exception is saved to a file for later review
- B. The exception is automatically fixed
- C. The program terminates abnormally and typically displays an error message
- D. The exception is ignored and the program continues

Answer: C

Q425. What is a ClassCastException?

- A. An error when a class has too many methods
- B. An exception when a class has duplicate names
- C. An error when a class cannot be loaded
- D. An exception thrown when trying to cast an object to a type it is not an instance of

Answer: D

Q426. Can a try block have no catch block?

- A. Yes, a try block can exist alone without catch or finally
- B. No, it must have both catch and finally
- C. Yes, if it has a finally block instead
- D. No, every try must have at least one catch

Answer: C

Q427. What is a NumberFormatException?

- A. An exception when arithmetic operations fail
- B. An exception thrown when trying to convert a non-numeric string to a number
- C. An error when numbers exceed the integer range
- D. An error when formatting numbers for display

Answer: B

Q428. What is the purpose of the 'catch' keyword?

- A. It throws a new exception
- B. It defines a block that handles a specific type of exception
- C. It stops the program from running
- D. It creates a new error object

Answer: B

Q429. What is an IndexOutOfBoundsException?

- A. An exception when too many indices are created
- B. An exception thrown when accessing a collection with an invalid index
- C. An error when the file index is corrupted
- D. An error when a database index is missing

Answer: B

Q430. What is the purpose of the 'raise' keyword in Python?

- A. It raises the value of a variable
- B. It increases the stack size
- C. It throws or raises an exception manually
- D. It increases the priority of a process

Answer: C

Q431. What is the purpose of importing a module?

- A. To compile the module into machine code
- B. To delete the module from the system
- C. To make the functions and classes defined in the module available in the current file
- D. To create a backup copy of the module

Answer: C

Q432. What is the `__init__.py` file used for in Python packages?

- A. It creates a backup of the package
- B. It stores initialization data for the program
- C. It marks a directory as a Python package and can contain initialization code
- D. It initializes all variables to zero

Answer: C

Q433. What is a third-party library?

- A. A library created by someone other than the language developers or the project team
- B. A library that requires three imports
- C. A library stored in a third location
- D. A library built into the programming language

Answer: A

Q434. What is the purpose of the 'export' keyword in JavaScript modules?

- A. It sends data to an external server
- B. It removes a variable from the module
- C. It converts the module to a different format
- D. It makes functions, classes, or variables available for other modules to import

Answer: D

Q435. What is a standard library?

- A. A collection of modules and packages included with the programming language by default
- B. A library that can only store standard data types
- C. A physical library where programmers study
- D. A library that follows a specific coding standard

Answer: A

Q436. What is a Maven repository in Java development?

- A. A version control system for Java code
- B. A storage location for Java libraries and dependencies that Maven can download automatically
- C. A testing framework for Java modules
- D. A database for storing Java application data

Answer: B

Q437. What does the 'require' function do in Node.js?

- A. It imports and loads a module, returning its exported functionality
- B. It checks if a condition is true
- C. It validates that a file exists
- D. It requires user authentication

Answer: A

Q438. What is the purpose of a .gitignore file in relation to packages?

- A. It ignores all errors in the package
- B. It specifies files and directories (like node_modules) that should not be tracked by version control
- C. It hides packages from other developers
- D. It imports Git-related packages

Answer: B

Q439. What is a framework in software development?

- A. A single function imported from a library
- B. A debugging tool for tracing program execution
- C. A comprehensive package that provides a structure and set of tools for building applications
- D. A physical frame used to hold a computer

Answer: C

Q440. What is the difference between a default export and a named export in JavaScript?

- A. Named exports can only export functions
- B. Default exports are faster than named exports
- C. A default export allows importing without curly braces while named exports require them
- D. Default exports must be constants

Answer: C

Q441. What is the Observer pattern used for?

- A. Notifying multiple objects when the state of another object changes
- B. Watching for syntax errors in code
- C. Monitoring memory usage in real time
- D. Observing CPU performance metrics

Answer: A

Q442. What is a hash function?

- A. A function that validates user passwords
- B. A function that maps input data to a fixed-size output value used for efficient data lookup
- C. A function that encrypts all data permanently
- D. A function that removes duplicate data

Answer: B

Q443. What is an event listener?

- A. A hardware device that detects sound
- B. A tool that monitors network events
- C. A function that waits for and responds to a specific event in event-driven programming
- D. A log file that records all system events

Answer: C

Q444. What is the MVC pattern?

- A. A method of compressing video files
- B. A security protocol for web applications
- C. A version control system for managing code
- D. A pattern that separates an application into Model (data), View (presentation), and Controller (logic)

Answer: D

Q445. What is time complexity?

- A. The time needed to compile the program
- B. A measure of how the running time of an algorithm grows relative to input size
- C. The time a program takes to load from disk
- D. The time taken to write the code

Answer: B

Q446. What is the main advantage of concurrent programming?

- A. It reduces the amount of memory needed
- B. It makes source code shorter
- C. It allows programs to perform multiple tasks simultaneously, improving responsiveness and throughput
- D. It eliminates all bugs from the program

Answer: C

Q447. What is a thread-safe data structure?

- A. A data structure that can be safely accessed by multiple threads concurrently without data corruption
- B. A data structure stored in thread-local memory
- C. A data structure that can only be used in single-threaded programs
- D. A data structure that is locked permanently

Answer: A

Q448. What does the 'join' method do on a thread?

- A. It makes the calling thread wait until the target thread finishes execution
- B. It joins the thread to a thread pool
- C. It combines two threads into one
- D. It creates a connection between two threads

Answer: A

Q449. What is a shared resource in concurrent programming?

- A. A library used by multiple projects
- B. A file that is shared on a network drive
- C. A CPU core shared between programs
- D. Data or resource that can be accessed by multiple threads or processes simultaneously

Answer: D

Q450. What is the start() method used for with threads?

- A. It restarts a completed thread
- B. It initializes thread-local variables
- C. It starts the main program execution
- D. It begins the execution of a new thread by calling its run method in a separate thread of execution

Answer: D

Q451. What is starvation in concurrent programming?

- A. When a thread is perpetually denied access to a resource because other threads keep acquiring it first
- B. When a program runs out of memory
- C. When a CPU has no tasks to execute
- D. When a thread pool runs out of threads

Answer: A

Q452. What is the run() method in a thread?

- A. It contains the code that the thread will execute when started
- B. It measures the thread's execution speed
- C. It runs the entire program from the beginning
- D. It runs diagnostics on the thread

Answer: A

Q453. What is a critical section in concurrent programming?

- A. The most important part of the source code
- B. A section of code that accesses shared resources and must not be executed by more than one thread at a time
- C. A section of code that runs with highest priority
- D. A code section that handles critical errors

Answer: B

Q454. What is the difference between a daemon thread and a user thread?

- A. User threads cannot access shared resources
- B. Daemon threads have higher priority than user threads
- C. Daemon threads run faster than user threads
- D. Daemon threads are background threads that do not prevent the program from exiting, while user threads keep the program alive

Answer: D

Q455. What is the sleep() method used for with threads?

- A. It releases all thread resources
- B. It pauses the thread for a specified duration before it resumes execution
- C. It permanently stops the thread
- D. It puts the entire program to sleep

Answer: B

Q456. What is the purpose of a test framework?

- A. A physical frame for holding test devices
- B. A tool that provides structure and utilities for writing and running automated tests
- C. A framework for testing hardware components
- D. A document that describes test procedures

Answer: B

Q457. What is the difference between a bug and an error?

- A. Errors are more severe than bugs
- B. A bug is a flaw in the code that causes incorrect behavior while an error is a mistake made by the programmer
- C. They are exactly the same thing
- D. Bugs only occur in production while errors occur in development

Answer: B

Q458. What does the 'step into' command do in a debugger?

- A. It executes the current line and enters the called function to debug it line by line
- B. It skips all remaining code in the function
- C. It restarts the debugging session
- D. It steps over the current line

Answer: A

Q459. What is a test assertion?

- A. A legal assertion about software ownership
- B. A claim that the software is bug-free
- C. A statement that checks whether an expected condition is true, failing the test if it is not
- D. A developer's confidence in their code

Answer: C

Q460. What is a watch expression in a debugger?

- A. A log message that watches for errors
- B. A timer that measures code execution speed
- C. A scheduled check that runs periodically
- D. An expression whose value is continuously monitored and displayed as you step through code

Answer: D

Q461. What is a test fixture?

- A. A fixture that holds test equipment
- B. A fixed baseline state or set of conditions set up before running tests
- C. A hardware device used for testing
- D. A test that runs automatically every day

Answer: B

Q462. What is the call stack in a debugger?

- A. A list of all breakpoints set in the program
- B. A display showing the chain of function calls that led to the current point of execution
- C. A list of all variables in the program
- D. A stack of error messages

Answer: B

Q463. What does a test pass or fail indicate?

- A. A pass means the code compiled successfully, a fail means it did not
- B. A pass means the code behaves as expected, a fail means the actual result differs from expected
- C. A pass means the test ran quickly, a fail means it was too slow
- D. A pass means no exceptions occurred regardless of correctness

Answer: B

Q464. What is the purpose of logging in debugging?

- A. Recording program events, state, and data flow to help diagnose issues without using a debugger
- B. Creating log files for regulatory compliance
- C. Recording user login attempts
- D. Logging the time spent on each function

Answer: A

Q465. What is a test runner?

- A. A thread dedicated to running tests in the background
- B. A tool that automatically discovers, executes, and reports the results of test cases
- C. A person who manually runs all tests
- D. A script that runs the program in test mode

Answer: B

Q466. What is a merge in version control?

- A. Reverting all changes to the original state
- B. Creating a new repository from scratch
- C. Deleting a branch permanently
- D. Combining changes from one branch into another

Answer: D

Q467. What is the purpose of code commenting?

- A. To make the code run faster
- B. To increase the file size for storage
- C. To prevent code from being compiled
- D. To add explanations that help developers understand the code's purpose and logic

Answer: D

Q468. What is a repository in version control?

- A. A physical location where servers are stored
- B. A collection of programming tutorials
- C. A database for user information
- D. A storage location containing the project files and their complete change history

Answer: D

Q469. What is the purpose of a README file?

- A. To list all errors found in the code
- B. To provide project documentation including description, setup instructions, and usage information
- C. To read data from a file at runtime
- D. To mark files as read-only

Answer: B

Q470. What is a coding convention?

- A. A conference about programming
- B. A legal contract for software development
- C. An agreed-upon set of guidelines for writing code consistently within a team or project
- D. A mathematical formula used in algorithms

Answer: C

Q471. What does 'git clone' do?

- A. It creates a new empty repository
- B. It copies a remote repository to the local machine including all history
- C. It merges two repositories together
- D. It deletes a remote repository

Answer: B

Q472. What is the purpose of an issue tracker?

- A. Recording and managing bugs, feature requests, and tasks for a software project
- B. Tracking performance issues in the CPU
- C. Monitoring server uptime issues
- D. Tracking internet connectivity issues

Answer: A

Q473. What is a deployment in software development?

- A. The process of making a software application available for use in a specific environment
- B. The process of hiring new developers
- C. The process of designing the user interface
- D. The process of writing new code

Answer: A

Q474. What is a staging environment?

- A. An environment that mirrors production, used for final testing before releasing to users
- B. A temporary development environment for writing code
- C. An environment used exclusively for database operations
- D. A backup environment for disaster recovery

Answer: A

Q475. What is the purpose of a .env file?

- A. Configuring the text editor environment
- B. Storing user session data
- C. Encrypting the entire project
- D. Storing environment-specific configuration variables like API keys and database URLs separately from code

Answer: D

Q476. What is the difference between a stack and a heap data structure?

- A. A heap follows First-In-First-Out order
- B. A heap is always faster than a stack
- C. A stack can hold more elements than a heap
- D. A stack follows Last-In-First-Out order while a heap is a tree-based structure that maintains a priority order

Answer: D

Q477. What is a queue data structure used for?

- A. Randomly accessing elements by index
- B. Processing elements in First-In-First-Out order where the first added element is removed first
- C. Storing elements in sorted order
- D. Storing elements in reverse order

Answer: B

Q478. What is the difference between linear search and binary search?

- A. Binary search works on unsorted data while linear search requires sorted data
- B. Linear search is always faster than binary search
- C. Linear search checks each element sequentially while binary search repeatedly divides a sorted collection in half
- D. Binary search checks every element one by one

Answer: C

Q479. What is a sorting algorithm?

- A. An algorithm that arranges elements in a specific order such as ascending or descending
- B. An algorithm that searches for a specific element
- C. An algorithm that deletes duplicate elements
- D. An algorithm that compresses data for storage

Answer: A

Q480. What is a callback function in event-driven programming?

- A. A function passed as an argument to be executed when an event or operation completes
- B. A function that always returns to the caller immediately
- C. A function that calls the main function repeatedly
- D. A function that calls back the operating system

Answer: A

Programming (C++, Java, Python) — Medium

480 questions

Q481. What is the difference between compilation and interpretation in terms of error detection?

- A. Compilation detects syntax errors before execution; interpretation at runtime
- B. Interpretation detects every error before execution even begins
- C. There is no practical difference between the two in error detection
- D. Compilation detects all errors only at runtime during execution

Answer: A

Q482. Which generation of programming languages includes languages like C and Java?

- A. First generation
- B. Second generation
- C. Fifth generation
- D. Third generation

Answer: D

Q483. What is Just-In-Time (JIT) compilation?

- A. Compiling code during execution for performance
- B. Handling runtime errors through catch blocks
- C. Compiling all code before shipping to users
- D. Writing source code at the very last minute

Answer: A

Q484. What is the role of a linker in program compilation?

- A. It combines object files into an executable
- B. It interprets bytecode during execution
- C. It validates syntax in the source code
- D. It manages runtime memory allocation

Answer: A

Q485. What is the difference between strongly typed and weakly typed languages?

- A. There is no real difference between the two type systems
- B. Strongly typed languages allow implicit type conversions freely
- C. Strongly typed enforces strict rules and restricts conversions
- D. Weakly typed languages enforce very strict type conversion rules

Answer: C

Q486. What is bytecode?

- A. Source code that is written using bytes
- B. Encrypted form of the original source
- C. Intermediate code run on a virtual machine
- D. Binary code identical to machine code

Answer: C

Q487. What is the purpose of a preprocessor in C/C++?

- A. To optimize generated machine code output
- B. To process directives like `#include` and `#define`
- C. To execute the compiled program directly
- D. To link external libraries to the binary

Answer: B

Q488. Which of the following best describes a runtime error?

- A. An error detected during the compilation phase
- B. An error that occurs while the program executes
- C. A syntax mistake caught before program runs
- D. A warning generated by the development IDE

Answer: B

Q489. What is type inference?

- A. A runtime error caused by wrong type assignment
- B. The compiler automatically determining variable type
- C. A debugging technique to inspect type values
- D. Manually specifying every variable data type

Answer: B

Q490. What is the call stack used for in program execution?

- A. Tracking function calls and local variables
- B. Handling incoming network data requests
- C. Storing all global variables and constants
- D. Managing all file input/output operations

Answer: A

Q491. What is the difference between implicit and explicit type casting?

- A. Implicit casting always causes errors while explicit casting never does
- B. Implicit is done automatically by the compiler; explicit is done manually
- C. Explicit casting executes much faster than implicit casting at runtime
- D. They are exactly the same type conversion process with no differences

Answer: B

Q492. What is the size of an int in Java?

- A. 32 bits
- B. 16 bits
- C. 8 bits
- D. 64 bits

Answer: A

Q493. What is the difference between value types and reference types?

- A. Value types store data directly; reference types store a memory address
- B. They are completely identical in behavior and memory representation
- C. Value types are immutable and cannot be modified after creation
- D. Reference types always execute faster than value types at all times

Answer: A

Q494. What happens when you assign a float value to an int variable without casting?

- A. It works fine with no issues at all
- B. Compilation error in strongly typed languages
- C. The decimal part is automatically rounded up
- D. The program crashes with a runtime exception

Answer: B

Q495. What is an enumeration (enum)?

- A. A looping construct for iteration
- B. A dynamically resizable array type
- C. A function for string manipulation
- D. A set of named constant values

Answer: D

Q496. What is the difference between signed and unsigned integers?

- A. There is no meaningful difference between signed and unsigned types
- B. Signed holds only positive values; unsigned holds negative values
- C. They have different byte sizes in memory but the same value range
- D. Signed holds negative and positive; unsigned holds non-negative only

Answer: D

Q497. What is a nullable type?

- A. A type that is always null
- B. A type with no default value
- C. A type that cannot be null
- D. A type that can hold a value or null

Answer: D

Q498. What is the result of integer overflow?

- A. The value automatically becomes null
- B. An exception is always thrown by runtime
- C. The program always crashes immediately
- D. The value wraps around in most languages

Answer: D

Q499. What is a union type in TypeScript?

- A. A built-in type for numeric calculations
- B. A type that can be one of several types
- C. A type that combines properties of two types
- D. A class that extends multiple parent classes

Answer: B

Q500. What is type aliasing?

- A. Hiding a type from external modules
- B. Renaming a declared variable in scope
- C. Converting between different data types
- D. Creating a new name for existing type

Answer: D

Q501. What is the difference between ++i and i++?

- A. ++i always executes faster than i++ does overall
- B. ++i increments before use; i++ increments after
- C. There is no difference between them at all
- D. i++ increments the value by two instead of one

Answer: B

Q502. What is short-circuit evaluation?

- A. Stopping evaluation once result is known
- B. A hardware feature of modern processors
- C. Evaluating all conditions in every case
- D. A compiler optimization for loop bodies

Answer: A

Q503. What is the ternary operator?

- A. An operator that takes a single operand only
- B. A conditional operator with three operands
- C. A bitwise operator for binary manipulation
- D. An operator that takes exactly two operands

Answer: B

Q504. What does the << operator do in C/Java?

- A. Reads input from the console
- B. Performs a left bitwise shift
- C. Compares two values for equality
- D. Performs string stream insertion

Answer: B

Q505. What is operator precedence?

- A. The order in which operators are defined in code
- B. The rules determining operator evaluation order
- C. The execution speed of a particular operator
- D. The data type returned by a given operator

Answer: B

Q506. What is the result of 5 & 3 (bitwise AND)?

- A. 5
- B. 15
- C. 8
- D. 1

Answer: D

Q507. What is the difference between == and === in JavaScript?

- A. There is no difference between them
- B. == checks value; === checks value and type
- C. === is used for variable assignment
- D. == checks type without value comparison

Answer: B

Q508. What does the XOR (^) operator return?

- A. True when the operands are different
- B. Always false regardless of the operands
- C. True when both operands are the same
- D. Always true regardless of the operands

Answer: A

Q509. What is the null coalescing operator (??) in JavaScript?

- A. It checks the code for syntax errors at compile time
- B. It returns the right operand if the left is null
- C. It always returns a null value to the caller
- D. It performs integer division on two operands

Answer: B

Q510. What is operator associativity?

- A. How operators are stored internally within the memory
- B. The direction operators of same precedence are evaluated
- C. The underlying data type of an operator in the language
- D. The execution speed of evaluating a given expression

Answer: B

Q511. What is the difference between while and do-while loops?

- A. do-while is faster in execution speed compared to while loop
- B. while always runs the loop body at least one time guaranteed
- C. while checks condition first; do-while runs body at least once
- D. There is no meaningful difference between them at all

Answer: C

Q512. What is an infinite loop?

- A. A syntax error detected at compile time
- B. A loop that has an empty body
- C. A loop whose condition never becomes false
- D. A loop that executes very fast

Answer: C

Q513. What is the purpose of the 'default' case in a switch statement?

- A. It executes when no other case matches
- B. It is required in every switch
- C. It must always be first
- D. It resets the switch value

Answer: A

Q514. What is fall-through in a switch statement?

- A. Execution continues to next case without break
- B. A compilation error in the switch block
- C. The switch statement fails and throws an error
- D. The default case gets skipped during execution

Answer: A

Q515. What is a sentinel value in loop control?

- A. A special value signaling loop termination
- B. A constant defined in the loop body
- C. A randomly generated input value
- D. A variable used as a loop counter

Answer: A

Q516. What is the time complexity of a nested loop where both run n times?

- A. $O(n)$
- B. $O(2^n)$
- C. $O(\log n)$
- D. $O(n^2)$

Answer: D

Q517. What is a for-each loop?

- A. A loop that counts its own iterations explicitly
- B. A loop that runs indefinitely until interrupted
- C. A loop that skips every other element in turn
- D. A loop iterating over each element without index

Answer: D

Q518. What happens if you forget the break in a switch case in Java?

- A. A compilation error is raised by the compiler immediately
- B. Nothing happens and the switch works as normally expected
- C. The program crashes with an unrecoverable runtime error
- D. Fall-through: subsequent cases execute until a break is hit

Answer: D

Q519. What is a labeled break statement?

- A. A break that exits a specific labeled outer loop
- B. A break statement followed by a comment
- C. A syntax error that the compiler will reject
- D. A break statement that logs a debug message

Answer: A

Q520. What is loop unrolling?

- A. Making loops intentionally execute more slowly for debugging
- B. Reducing loop overhead by executing multiple iterations per pass
- C. Adding additional iterations to increase the loop precision
- D. Removing a loop from the code entirely

Answer: B

Q521. What is function overloading?

- A. Calling a function too many times in a row
- B. Same name functions with different parameter lists
- C. A runtime error from excessive function calls
- D. Calling a function from within itself recursively

Answer: B

Q522. What is the difference between pass by value and pass by reference?

- A. Pass by value copies data; pass by reference passes address
- B. Pass by value uses pointers to access the original variable
- C. Pass by reference is always slower than pass by value mode
- D. They are exactly the same thing with different names only

Answer: A

Q523. What is recursion?

- A. A function calling itself
- B. A loop construct
- C. An error in programming
- D. A type of variable

Answer: A

Q524. What is a base case in recursion?

- A. The condition that stops recursion
- B. The recursive call itself in code
- C. An optional function parameter
- D. The very first function call made

Answer: A

Q525. What is a lambda function?

- A. An anonymous short inline function expression
- B. A function that returns no value at all
- C. A function that calls itself recursively
- D. A function with an explicitly declared name

Answer: A

Q526. What is a callback function?

- A. A function that calls itself in a recursive way
- B. A function passed as an argument for later call
- C. A constructor that initializes object properties
- D. A function that contacts the user directly

Answer: B

Q527. What is variable scope in the context of functions?

- A. The memory size of a variable in bytes
- B. The declared data type of the variable
- C. The current assigned value of a variable
- D. The region where a variable is accessible

Answer: D

Q528. What are default parameters?

- A. Parameters that are always initialized to zero
- B. Parameters whose values cannot ever be changed
- C. Parameters with values used when none is passed
- D. Parameters that the caller must always provide

Answer: C

Q529. What is a higher-order function?

- A. A function with high computational time complexity
- B. A function defined in a higher-level parent class
- C. A function that accepts many input parameters
- D. A function that takes or returns other functions

Answer: D

Q530. What is method signature?

- A. Only the access modifier of the declared method
- B. The complete body of the method implementation
- C. Only the return type of the declared method
- D. The method name combined with its parameter list

Answer: D

Q531. What is buffered I/O?

- A. I/O that applies encryption to all data before transferring
- B. I/O that exclusively works with pre-allocated buffer arrays
- C. I/O that collects data in a buffer to reduce system calls
- D. I/O operations performed without any processing delays

Answer: C

Q532. What is the difference between print() and println() in Java?

- A. There is no difference between them at all
- B. print() executes faster than println() does
- C. println() only prints numeric values to output
- D. println() adds a newline; print() does not

Answer: D

Q533. What is file I/O?

- A. Reading from and writing to disk files
- B. Reading data from the keyboard input
- C. Allocating memory for program variables
- D. Communicating over the internet network

Answer: A

Q534. What is serialization?

- A. Encrypting data for secure network transmission
- B. Sorting data in ascending or descending order
- C. Converting an object into a byte stream for storage
- D. Compressing files to reduce their storage size

Answer: C

Q535. What does the flush() method do in I/O?

- A. Deletes the entire output buffer contents
- B. Forces buffered output to be written immediately
- C. Clears the entire console screen display
- D. Reads all contents from the input buffer

Answer: B

Q536. What is the difference between text mode and binary mode in file I/O?

- A. Text mode processes files faster than binary mode does
- B. There is no functional difference between the two modes
- C. Text mode translates newlines; binary reads raw bytes
- D. Binary mode is exclusively used for image file formats

Answer: C

Q537. What is standard error (stderr)?

- A. A specialized tool used for debugging programs
- B. A separate output stream for error messages only
- C. An error that occurs in the standard input stream
- D. It is the same stream as standard output stdout

Answer: B

Q538. What is the purpose of the 'with' statement for file I/O in Python?

- A. It loops through all files in a given directory
- B. It reads binary files into memory as byte arrays
- C. It writes data to multiple files simultaneously
- D. It ensures the file is properly closed after use

Answer: D

Q539. What is CSV and how is it used in I/O?

- A. Comma-Separated Values, a text format for tabular data
- B. A relational database management system for queries
- C. A compiled programming language for systems work
- D. A binary file format for storing structured records

Answer: A

Q540. What is piping in command-line I/O?

- A. A specialized type of network I/O communication
- B. Writing output data to a physical pipe file
- C. Redirecting one program output as another input
- D. Processing multiple data streams in parallel mode

Answer: C

Q541. What is the difference between mutable and immutable strings?

- A. Mutable strings always execute faster than immutable string types
- B. Immutable strings can dynamically grow in size without any limit
- C. There is no meaningful difference between the two string types
- D. Mutable strings can be modified in place; immutable ones cannot

Answer: D

Q542. What is a regular expression?

- A. A standard normal string literal value
- B. An algorithm for sorting string collections
- C. A pattern for matching and manipulating text
- D. A mathematical expression using operators

Answer: C

Q543. What is string interpolation?

- A. Converting strings to their numeric equivalents
- B. Splitting strings by a delimiter character
- C. Comparing two strings for equality or ordering
- D. Embedding variables directly within a string

Answer: D

Q544. What is the String Pool in Java?

- A. A thread pool dedicated to string processing
- B. A mutable buffer for building string values
- C. A collection of string utility methods
- D. A heap area that caches string literals for reuse

Answer: D

Q545. What is the difference between String, StringBuilder, and StringBuffer in Java?

- A. StringBuilder is immutable while String and StringBuffer are both mutable
- B. They are all identical in function and performance with no differences
- C. String is immutable; StringBuilder is mutable; StringBuffer is thread-safe
- D. StringBuffer is always the fastest option among all three implementations

Answer: C

Q546. What does the split() method do?

- A. Divides a string into an array of parts
- B. Joins multiple strings into one
- C. Converts string to a character array
- D. Removes specified characters from string

Answer: A

Q547. What is Unicode?

- A. A type of data encryption algorithm
- B. An extension of the ASCII character set
- C. A universal character encoding standard
- D. A specific font type used in displays

Answer: C

Q548. What is the difference between ASCII and Unicode?

- A. ASCII supports many more characters than Unicode ever can
- B. Unicode is the older standard that was developed before ASCII
- C. They are exactly the same encoding standard with no differences
- D. ASCII uses 7 bits for 128 chars; Unicode supports over 143,000

Answer: D

Q549. What is string encoding (like UTF-8)?

- A. Compressing a string to reduce storage size
- B. Encrypting a string for secure transmission
- C. Converting a string value into an integer
- D. Representing Unicode characters as byte sequences

Answer: D

Q550. What is a string builder pattern and when should you use it?

- A. A design pattern for constructing user interface strings
- B. A utility tool for building complex regular expressions
- C. A parser that transforms strings into structured data types
- D. Using a mutable class for efficient repeated concatenation

Answer: D

Q551. What is the time complexity of accessing an element by index in an array?

- A. $O(n^2)$
- B. $O(1)$
- C. $O(n)$
- D. $O(\log n)$

Answer: B

Q552. What is a hash table?

- A. A balanced binary tree for hierarchical storage
- B. A sorted array of ordered elements in memory
- C. A structure mapping keys to values via hashing
- D. A stack variant that allows random element access

Answer: C

Q553. What is a collision in a hash table?

- A. A runtime error in the hash function code
- B. When a requested key is not found in map
- C. When the hash table is completely full up
- D. When two keys hash to the same index slot

Answer: D

Q554. What is the difference between a singly linked list and a doubly linked list?

- A. Singly linked lists always use more memory than doubly linked
- B. There is no meaningful difference between the two list types
- C. Singly linked has next pointer; doubly linked has next and prev
- D. Doubly linked lists are always slower than singly linked lists

Answer: C

Q555. What is a priority queue?

- A. A sorted array with indexed element access only
- B. A queue where higher-priority elements dequeue first
- C. A stack with last-in first-out element ordering
- D. A regular first-in first-out queue structure

Answer: B

Q556. What is the difference between HashMap and TreeMap in Java?

- A. They are exactly the same data structure implementation
- B. HashMap maintains the insertion order of all its entries
- C. TreeMap is always faster than HashMap for all operations
- D. HashMap is unordered $O(1)$; TreeMap is sorted $O(\log n)$

Answer: D

Q557. What is a tuple?

- A. A mutable list that can change its elements
- B. A key-value map for looking up data by key
- C. An immutable ordered collection of mixed types
- D. A sorted set that maintains element ordering

Answer: C

Q558. What is an iterator?

- A. An array index used for element access
- B. An object for traversing a collection sequentially
- C. A sorting algorithm for ordering collections
- D. A type of loop construct for repeating code

Answer: B

Q559. What is the difference between a shallow copy and a deep copy of a collection?

- A. They are the same operation with identical results in every case
- B. Shallow copy shares element references; deep copy duplicates all
- C. Shallow copy is always slower than deep copy for all collections
- D. Deep copy shares references while shallow copy duplicates objects

Answer: B

Q560. What is a deque (double-ended queue)?

- A. A regular first-in first-out queue
- B. A queue allowing insertion at both ends
- C. A sorted list with binary search access
- D. A stack variant for ordered elements

Answer: B

Q561. What is the difference between method overloading and method overriding?

- A. They are the same concept with different names in different languages
- B. Overloading requires an inheritance hierarchy between the classes involved
- C. Overriding changes the method parameter list in the parent class definition
- D. Overloading: same name different params; overriding: same signature in subclass

Answer: D

Q562. What is an abstract class?

- A. A fully implemented class with all methods defined
- B. A static class that only contains class-level members
- C. A final class that cannot be extended by other classes
- D. A class that cannot be instantiated with abstract methods

Answer: D

Q563. What is an interface?

- A. A class with full method implementations provided
- B. A contract defining method signatures without bodies
- C. An abstract class with some methods implemented
- D. A type of variable for storing object references

Answer: B

Q564. What is the difference between composition and inheritance?

- A. Composition is always slower than inheritance in performance
- B. Inheritance is an is-a relationship; composition is has-a
- C. They are exactly the same concept with no difference at all
- D. Inheritance is more flexible and loosely coupled overall

Answer: B

Q565. What is a static method?

- A. A method belonging to the class not any specific instance
- B. A constructor method for initializing new object state
- C. A method that can never be overridden by subclasses
- D. A private method accessible only within the same class

Answer: A

Q566. What are getter and setter methods?

- A. Static methods belonging to the class directly
- B. Methods for controlled field read and write access
- C. Loop constructs for iterating over collections
- D. Constructor types for initializing objects

Answer: B

Q567. What is the super keyword?

- A. A global variable accessible everywhere in code
- B. A reference to parent class for calling its members
- C. The main class that serves as program entry point
- D. A type of access modifier restricting member access

Answer: B

Q568. What is multiple inheritance and which languages support it?

- A. Creating multiple instances of the same class simultaneously
- B. Inheriting from a single parent class in a linear hierarchy
- C. Inheriting from multiple parents; supported by C++ and Python
- D. All programming languages support this feature universally

Answer: C

Q569. What is the diamond problem?

- A. An error related to geometric shapes in rendering code
- B. A memory leak issue caused by circular object references only
- C. An ambiguity when inheriting from classes with common ancestor
- D. A compilation warning about unused variables in the program

Answer: C

Q570. What is the Liskov Substitution Principle (LSP)?

- A. A design pattern for constructing complex object hierarchies
- B. A principle for sorting elements in a collection efficiently
- C. A rule for managing memory allocation in large applications
- D. Subclass objects should be substitutable for superclass objects

Answer: D

Q571. What is a dangling pointer?

- A. A pointer referencing already freed memory
- B. A pointer to a valid live object in memory
- C. A pointer with a null value assigned to it
- D. A smart pointer that manages object lifetime

Answer: A

Q572. What is reference counting as a garbage collection strategy?

- A. Counting variables in the current scope of code
- B. Tracking reference count and freeing at zero count
- C. A manual memory management technique using counters
- D. Counting the total number of memory blocks allocated

Answer: B

Q573. What is a buffer overflow?

- A. Slow memory access due to fragmentation issues
- B. Writing data beyond the allocated buffer boundary
- C. Running out of available system memory completely
- D. A type of memory leak from unreleased resources

Answer: B

Q574. What is the difference between malloc() and calloc() in C?

- A. They are exactly the same function with no differences
- B. malloc initializes all allocated memory bytes to zero
- C. malloc allocates uninitialized; calloc initializes to zero
- D. calloc is always faster than malloc in all scenarios

Answer: C

Q575. What is virtual memory?

- A. Shared memory accessible by multiple processes at once
- B. Cloud storage accessible over the internet network
- C. Using disk space to extend memory beyond physical RAM
- D. Encrypted memory for storing sensitive data securely

Answer: C

Q576. What is a memory pool?

- A. A pre-allocated block for smaller allocations reducing fragmentation
- B. A collection of physical memory hardware modules
- C. A memory leak caused by unreleased dynamically allocated resources
- D. A type of processor cache that stores frequently accessed data values

Answer: A

Q577. What is the difference between deep copy and shallow copy regarding memory?

- A. Deep copy allocates new memory for nested objects; shallow copies refs
- B. Shallow copy always uses more memory than deep copy for same structure
- C. There is no difference in memory usage between the two copy methods
- D. Deep copy uses less memory than shallow copy by sharing object refs

Answer: A

Q578. What are smart pointers in C++?

- A. Regular raw pointers with no extra features at all
- B. Objects that wrap pointers and manage memory automatically
- C. Pointers that point to function addresses in the program
- D. Pointers that find data faster through optimized algorithms

Answer: B

Q579. What is memory fragmentation?

- A. Memory that has been encrypted for data security
- B. Memory that has been compressed to save space
- C. Corrupted memory caused by hardware failures in RAM
- D. Free memory broken into small non-contiguous blocks

Answer: D

Q580. What is the RAII pattern?

- A. A design pattern for building user interfaces
- B. A type of garbage collection used in managed languages
- C. Tying resource lifetime to object lifetime via destructors
- D. A memory allocation algorithm for dynamic memory pools

Answer: C

Q581. What is the difference between checked and unchecked exceptions in Java?

- A. Unchecked exceptions are always more severe than checked exceptions
- B. Checked exceptions are always faster to handle than unchecked ones
- C. Checked must be declared or caught; unchecked RuntimeExceptions do not
- D. There is no meaningful difference between the two exception types

Answer: C

Q582. What is exception propagation?

- A. Unhandled exception moving up the call stack to find a handler
- B. Fixing an exception by catching and logging it
- C. Logging all exceptions to a file for later review and analysis
- D. Throwing multiple exceptions at the same time from one method

Answer: A

Q583. What is a custom exception?

- A. A user-defined exception class for domain-specific errors
- B. A syntax error detected during the compilation process
- C. A warning generated by the IDE during code development
- D. A built-in exception provided by the language runtime

Answer: A

Q584. What is the 'throws' keyword in Java?

- A. A keyword for suppressing exceptions so they are ignored
- B. A mechanism for catching and handling thrown exceptions
- C. It is exactly the same as the throw keyword in Java
- D. A declaration of which checked exceptions a method may throw

Answer: D

Q585. What is a stack trace?

- A. A log file containing all program output and error data
- B. The call stack data structure itself in memory
- C. A memory dump of the entire program state at crash
- D. A report showing the method call chain to an exception

Answer: D

Q586. What is the purpose of try-with-resources in Java?

- A. To automatically close resources when the try block finishes
- B. To try executing multiple methods in parallel simultaneously
- C. To retry the try block harder if the first attempt fails
- D. To allocate additional resources for the try block execution

Answer: A

Q587. What is the difference between Error and Exception in Java?

- A. Errors are serious JVM problems; Exceptions are recoverable issues
- B. Exceptions cannot be caught or handled by any catch block
- C. They are the same thing with no meaningful differences at all
- D. Errors are less serious and easier to handle than Exceptions

Answer: A

Q588. What is a try-except-else block in Python?

- A. An advanced loop construct for iterating with error handling
- B. try runs risky code; except handles; else runs if no exception
- C. Invalid syntax that causes a compilation error
- D. A triple try block for handling three exceptions at once

Answer: B

Q589. What is exception chaining?

- A. Throwing many exceptions at the same time from one method call
- B. Catching exceptions in a specific order from most to least severe
- C. Wrapping one exception inside another to preserve original cause
- D. A linked list of error objects stored in chronological order

Answer: C

Q590. What is the principle of 'fail fast'?

- A. Testing the speed of program execution under various load levels
- B. Detecting and reporting errors as soon as they occur in the program
- C. Making programs crash as quickly as possible for testing purposes
- D. Running programs as quickly as possible to improve the performance

Answer: B

Q591. What is dependency management?

- A. Managing memory allocation for program variables only
- B. Tracking and maintaining external libraries and versions
- C. Sorting functions in alphabetical order within modules
- D. Managing source code files in a project directory

Answer: B

Q592. What is a package.json file?

- A. A database schema definition for SQL table structures
- B. A log file recording application events and errors
- C. A Node.js config file listing metadata and dependencies
- D. A generic JSON data file for storing values

Answer: C

Q593. What is semantic versioning (SemVer)?

- A. Alphabetical versioning using letter-based identifiers
- B. Date-based versioning using release date as version
- C. MAJOR.MINOR.PATCH scheme indicating type of changes
- D. Random version numbers assigned without any rules

Answer: C

Q594. What is a virtual environment in Python?

- A. A Docker container with a bundled application inside
- B. An isolated Python environment with its own packages
- C. A cloud-hosted computing environment for deployment
- D. A full virtual machine running a complete OS image

Answer: B

Q595. What is the difference between static and dynamic linking of libraries?

- A. Static includes library at compile time; dynamic loads at runtime
- B. Dynamic linking is always slower than static linking in every case
- C. Static linking uses less memory than dynamic linking at all times
- D. There is no meaningful difference between the two linking methods

Answer: A

Q596. What is a circular dependency?

- A. When two or more modules depend on each other
- B. A circular buffer data structure for streaming
- C. A loop in the code that repeats indefinitely
- D. A recursive function calling itself repeatedly

Answer: A

Q597. What is tree shaking?

- A. A data structure operation on binary tree nodes only
- B. A sorting technique that organizes elements by weight
- C. Removing trees from the project code repository
- D. Removing unused dead code from the final build bundle

Answer: D

Q598. What is the purpose of a lock file (e.g., package-lock.json)?

- A. To lock the project and prevent any code changes
- B. To encrypt installed packages for secure data storage use
- C. To record exact dependency versions for reproducible builds
- D. To prevent anyone from making changes to the source code

Answer: C

Q599. What is the difference between devDependencies and dependencies in npm?

- A. devDependencies are completely optional and never actually installed
- B. dependencies are operating system packages not application libraries
- C. There is no meaningful difference between the two dependency types
- D. dependencies are for production; devDependencies for development only

Answer: D

Q600. What is a monorepo?

- A. A repository containing only a single source file
- B. A read-only repository that cannot accept any changes
- C. A private repository with restricted access permissions
- D. A single repository containing multiple related projects

Answer: D

Q601. What is the Observer pattern?

- A. A logging pattern for recording application events to a file
- B. A security pattern for authenticating users before granting access
- C. A subject notifies dependent observers of state changes automatically
- D. Watching code execute step by step in the debugger tool

Answer: C

Q602. What is the Factory pattern?

- A. A singleton variant that creates exactly one object instance
- B. A manufacturing system for producing physical goods
- C. An iterator pattern for traversing collections sequentially
- D. A pattern providing an interface for creating objects flexibly

Answer: D

Q603. What is the Strategy pattern?

- A. A sorting method that orders elements in a specific sequence
- B. A game strategy for winning competitions against opponents
- C. A pattern defining interchangeable algorithms varying independently
- D. A caching technique for storing frequently accessed data values

Answer: C

Q604. What is the difference between composition and inheritance as design strategies?

- A. Composition offers flexibility and loose coupling; inheritance is tighter
- B. They are identical design strategies with no practical differences
- C. Inheritance is always better than composition in every design scenario
- D. Composition is always slower in execution than inheritance-based designs

Answer: A

Q605. What is the Model-View-Controller (MVC) pattern?

- A. A database pattern for organizing tables and their relationships
- B. A testing framework for writing and running automated test suites
- C. Separating application into Model (data), View (UI), Controller (logic)
- D. A data structure for organizing elements in memory efficiently

Answer: C

Q606. What is lazy loading?

- A. Loading resources in reverse chronological order
- B. Deferring initialization until actually needed
- C. A slow loading technique for large file downloads
- D. Loading all data and resources at startup time

Answer: B

Q607. What is the difference between aggregation and composition in UML?

- A. Composition implies weaker ownership than aggregation in all designs
- B. Aggregation implies stronger ownership than composition in all cases
- C. Aggregation implies weak ownership; composition implies strong ownership
- D. They are identical concepts with no practical differences at all

Answer: C

Q608. What is the Iterator pattern?

- A. Accessing collection elements sequentially without exposing internals
- B. A search technique for finding elements in a sorted collection
- C. A sorting algorithm for ordering elements in ascending sequence
- D. A loop construct for iterating over elements in arrays

Answer: A

Q609. What is the Decorator pattern?

- A. A user interface pattern for visual element styling
- B. A logging tool for recording method execution traces
- C. Adding annotations to class methods for documentation
- D. Dynamically adding responsibilities to objects through wrapping

Answer: D

Q610. What is the Adapter pattern?

- A. A data converter for transforming file format encodings
- B. A sorting adapter for converting between sort algorithms
- C. A pattern wrapping one interface to match another interface
- D. A physical power adapter for electrical devices

Answer: C

Q611. What is a semaphore?

- A. A counter-based primitive controlling shared resource access
- B. A thread type for running background tasks in parallel
- C. A simple flag variable for boolean state tracking
- D. A type of lock for exclusive thread synchronization

Answer: A

Q612. What is thread safety?

- A. A testing concept for validating thread performance
- B. Code functioning correctly under multi-thread access
- C. A debugging feature for analyzing thread behavior
- D. Threads that are guaranteed to never crash or fail

Answer: B

Q613. What is a thread pool?

- A. Pre-created worker threads reused for tasks
- B. A memory pool for thread-local storage data
- C. A collection of all threads in the system
- D. A connection pool for database access threads

Answer: A

Q614. What is the producer-consumer problem?

- A. A database issue with concurrent read and write transaction conflicts
- B. A business problem related to supply chain management
- C. A network problem caused by packet loss during data transmission
- D. A synchronization problem with shared buffer between producers and consumers

Answer: D

Q615. What is an atomic operation?

- A. An indivisible operation completing without interruption
- B. A nuclear physics operation on atomic particles
- C. A file operation that reads and writes data to disk
- D. A database transaction that groups multiple queries

Answer: A

Q616. What is the volatile keyword in Java?

- A. A final variable that cannot be modified after assignment
- B. A keyword ensuring reads always come from main memory
- C. A variable that changes randomly without any cause
- D. A type of exception thrown during concurrent operations

Answer: B

Q617. What is a Future/Promise?

- A. An object representing a value available in the future
- B. A type of thread that executes tasks asynchronously
- C. A variable that is declared for use at a later time
- D. A timer that triggers a callback after a set delay

Answer: A

Q618. What is context switching?

- A. Saving one thread state and loading another for CPU switching
- B. Changing the scope of a variable from local to global
- C. Changing between different programming languages in a project
- D. Switching between different source code files in an editor

Answer: A

Q619. What is the dining philosophers problem?

- A. A sorting problem for ordering elements in a collection
- B. A classic problem illustrating resource allocation and deadlock
- C. A database problem with concurrent transaction conflicts
- D. A food-related algorithm for meal planning optimization

Answer: B

Q620. What is async/await?

- A. A type of loop for iterating over collection elements
- B. A testing pattern for validating asynchronous behavior
- C. A syntax for writing async code that looks synchronous
- D. A threading model for managing parallel execution of tasks

Answer: C

Q621. What is Test-Driven Development (TDD)?

- A. A deployment strategy for releasing code to production
- B. Testing code only after all development is complete
- C. Writing tests before code to drive the design process
- D. A debugging technique for finding runtime exceptions

Answer: C

Q622. What is code coverage?

- A. How much source code has been written in total
- B. The percentage of code lines executed during testing
- C. A code review tool for examining code quality metrics
- D. A security measure for protecting source code files

Answer: B

Q623. What is the difference between black-box and white-box testing?

- A. There is no meaningful difference between the two testing approaches
- B. Black-box tests functionality without code; white-box tests internals
- C. White-box testing is fully automated while black-box is manual
- D. Black-box testing is always better than white-box in every case

Answer: B

Q624. What is a mock object in testing?

- A. A simulated object mimicking real behavior for tests
- B. A real database connection used in production
- C. A randomly generated object with unpredictable state
- D. A deleted object that has been garbage collected

Answer: A

Q625. What is integration testing?

- A. Testing interaction between multiple components together
- B. Testing the user interface for visual correctness
- C. Testing individual functions in complete isolation
- D. Testing performance under heavy load conditions

Answer: A

Q626. What is regression testing?

- A. Testing for performance regression and slower execution
- B. Re-running tests after changes to ensure nothing broke
- C. Testing only brand new features and functionality
- D. Testing backward compatibility with previous versions

Answer: B

Q627. What is profiling in software optimization?

- A. Writing documentation for the application source code
- B. Analyzing execution to identify performance bottlenecks
- C. Formatting code according to style guide conventions
- D. Creating user profiles for application personalization

Answer: B

Q628. What is the difference between stubs and mocks?

- A. Mocks are simpler and provide less functionality than stubs do
- B. They are identical testing concepts with no differences at all
- C. Stubs are more complex to implement than mocks in every case
- D. Stubs provide predefined responses; mocks also verify interactions

Answer: D

Q629. What is a memory profiler?

- A. A physical RAM module installed in the computer
- B. A tool monitoring memory usage and detecting leaks
- C. A compiler for translating source code to binaries
- D. A debugger for stepping through code line by line

Answer: B

Q630. What is the difference between manual and automated testing?

- A. There is no practical difference between the two testing methods
- B. Manual testing is done by humans; automated uses scripts and tools
- C. Automated testing is less reliable than manual in every case
- D. Manual testing is faster than automated testing for all suites

Answer: B

Q631. What is Continuous Integration (CI)?

- A. A deployment method for releasing code directly to production servers
- B. Integrating code changes only once at the very end of development
- C. A branching strategy for organizing feature development work in Git
- D. Frequently merging code with automated builds and tests for feedback

Answer: D

Q632. What is Continuous Deployment/Delivery (CD)?

- A. Automatically deploying or making release-ready every passing change
- B. A testing strategy for validating software behavior before release
- C. A branching model for organizing development work into branches
- D. Manual deployment performed by operations team members only

Answer: A

Q633. What is the KISS principle?

- A. A security principle for protecting sensitive user data
- B. A greeting or expression of affection between people
- C. A testing method for validating software behavior fully
- D. Keep It Simple: avoid unnecessary complexity in design

Answer: D

Q634. What is technical debt?

- A. Financial debt from purchasing technology equipment
- B. Server costs for hosting applications in the cloud
- C. The cost of future rework from choosing quick solutions
- D. A type of bug caused by incorrect logic in the code

Answer: C

Q635. What is the YAGNI principle?

- A. A testing approach for validating code correctness
- B. A design pattern for creating reusable components
- C. A programming language for building applications
- D. Do not implement functionality until actually needed

Answer: D

Q636. What is pair programming?

- A. Programming alone without any collaboration at all
- B. Writing duplicate code for redundancy and safety nets
- C. Two developers at one workstation: driver and navigator
- D. Programming with matched pairs of source code files

Answer: C

Q637. What is a sprint in Scrum?

- A. A fixed-length iteration for completing prioritized tasks
- B. A deployment window for releasing features to production
- C. A programming technique for optimizing code performance
- D. Running fast to meet project deadlines under pressure

Answer: A

Q638. What is a code smell?

- A. Malicious code injected by an attacker into the system
- B. A runtime bug that crashes the application unexpectedly
- C. A surface indication suggesting a deeper design problem
- D. A syntax error detected during the compilation process

Answer: C

Q639. What is the difference between waterfall and agile methodologies?

- A. Agile has no planning or structure and relies entirely on improvisation
- B. There is no meaningful difference between the two methodologies at all
- C. Waterfall is always faster than agile for delivering software to production
- D. Waterfall is sequential with fixed phases; agile is iterative and flexible

Answer: D

Q640. What is a merge conflict?

- A. A deployment error when releasing code to production servers
- B. A runtime error that occurs during program execution in prod
- C. A syntax error caused by incorrect code in the source file
- D. When Git cannot auto-merge because branches changed same code

Answer: D

Q641. What distinguishes a strongly typed language from a weakly typed one?

- A. Strongly typed languages do not support any string operations
- B. Weakly typed languages cannot handle any numeric data types
- C. Strongly typed languages enforce strict rules with fewer casts
- D. Strongly typed languages run faster than weakly typed ones always

Answer: C

Q642. What is the difference between a logical error and a runtime error?

- A. Logical errors crash the program; runtime errors give wrong results
- B. Both types always occur at compile time before program even starts
- C. Logical errors give wrong results; runtime errors halt execution
- D. Runtime errors give wrong output; logical errors are always harmless

Answer: C

Q643. What is the role of a preprocessor in languages like C?

- A. It processes directives like includes and macros before compiling
- B. It manages memory allocation and garbage collection automatically
- C. It links all object files together into one single final binary
- D. It executes the program line by line during the runtime phase

Answer: A

Q644. What is the key advantage of using an interpreted language?

- A. Interpreted languages cannot be used for web applications
- B. Interpreted languages always produce faster executables overall
- C. Interpreted languages enable platform-independent execution
- D. Interpreted languages require a separate linking step always

Answer: C

Q645. What does bytecode represent in languages like Java?

- A. The final machine code running directly on hardware
- B. An intermediate code executed by a virtual machine
- C. The original source code written by the programmer
- D. A binary format used only for storing databases

Answer: B

Q646. How does just-in-time (JIT) compilation improve performance?

- A. It compiles the entire program before any execution begins
- B. It translates bytecode to native code at runtime for speed
- C. It replaces the interpreter with a static linker entirely
- D. It removes all comments and whitespace from source code

Answer: B

Q647. What is the purpose of a symbol table during compilation?

- A. It maps identifiers to their types, scopes, and locations
- B. It stores the graphical icons used in the program interface
- C. It records user input received during program execution
- D. It contains the complete list of syntax errors in the code

Answer: A

Q648. What is the difference between static and dynamic typing?

- A. Static typing checks types at compile time; dynamic at runtime
- B. Static typing is only for scripting; dynamic for systems coding
- C. Static typing has no types; dynamic typing has strict types only
- D. Static typing is slower; dynamic typing is always faster overall

Answer: A

Q649. What does lexical analysis do during compilation?

- A. It links external libraries into the final binary
- B. It checks the logic correctness of the algorithm
- C. It optimizes the generated machine code for speed
- D. It breaks source code into tokens for the parser

Answer: D

Q650. What is a cross-compiler used for?

- A. It compiles code for a different target platform
- B. It converts machine code back into source code
- C. It interprets and compiles code at the same time
- D. It compiles two languages into one single program

Answer: A

Q651. What is the risk of implicit type conversion in weakly typed languages?

- A. It may produce unexpected results by silently converting
- B. It prevents the programmer from using numeric data types
- C. It always causes the program to crash immediately on start
- D. It makes the program significantly faster than explicit casts

Answer: A

Q652. Why might a developer choose a double over a float?

- A. A double provides greater precision for decimal values
- B. A double is faster to process than a float on all CPUs
- C. A double uses less memory than a float in all cases
- D. A double can store text while a float cannot do that

Answer: A

Q653. What is an enumeration (enum) data type used for?

- A. Creating dynamic arrays that resize automatically
- B. Defining a set of named constant integer values
- C. Encrypting sensitive string data in the program
- D. Storing floating-point numbers with high precision

Answer: B

Q654. What happens when integer overflow occurs?

- A. The program automatically switches to a larger type
- B. The OS allocates additional memory for the variable
- C. The compiler prevents the code from being compiled
- D. The value wraps around to the minimum of its range

Answer: D

Q655. What is the difference between value types and reference types?

- A. Value types require garbage collection; references do not
- B. Value types are slower to access than reference types
- C. Value types store data directly; references store addresses
- D. Reference types cannot be used inside any functions

Answer: C

Q656. Why are floating-point numbers sometimes imprecise?

- A. Computers cannot perform mathematical operations
- B. Binary representation cannot express all decimal values
- C. Languages intentionally introduce errors for speed
- D. Floating-point numbers are always rounded to integers

Answer: B

Q657. What is a union data type in languages like C?

- A. A type combining two arrays into one structure
- B. A type that converts between different types
- C. A type where all members share the same location
- D. A type that stores multiple values simultaneously

Answer: C

Q658. What is type narrowing in TypeScript?

- A. Restricting variables to only hold integer values
- B. Refining a variable type within a conditional
- C. Converting a large data type to a smaller one
- D. Removing type information from variables at runtime

Answer: B

Q659. What is the purpose of the void type in programming?

- A. To indicate a function returns no value at all
- B. To represent null pointer references in memory
- C. To declare variables with unspecified data type
- D. To store empty string values in memory locations

Answer: A

Q660. How does a BigInteger type differ from a standard integer?

- A. BigInteger uses less memory than a standard integer
- B. BigInteger stores decimal fractions unlike standard int
- C. BigInteger can represent arbitrarily large whole numbers
- D. BigInteger is faster but less accurate than standard

Answer: C

Q661. What is the difference between prefix and postfix increment?

- A. Prefix decrements and postfix increments the variable
- B. Prefix returns value after increment; postfix returns before
- C. Prefix works on integers only; postfix works on floats
- D. Prefix is faster and postfix is slower in all languages

Answer: B

Q662. What does the ternary operator (?:) do?

- A. It performs three separate arithmetic operations
- B. It provides a shorthand for if-else expression
- C. It creates three variables from one declaration
- D. It loops through a collection exactly three times

Answer: B

Q663. What does the bitwise left shift operator (<<) do?

- A. Moves bits left, effectively multiplying by powers of two
- B. Rotates bits in circular pattern within byte boundaries
- C. Moves bits right, effectively dividing by powers of two
- D. Compares two values bit by bit and returns differences

Answer: A

Q664. What is short-circuit evaluation in logical expressions?

- A. It skips the second operand when the result is determined
- B. It converts logical expressions into arithmetic for speed
- C. It evaluates all operands regardless of intermediate results
- D. It caches the results of all Boolean operations for reuse

Answer: A

Q665. What is operator overloading in object-oriented programming?

- A. Increasing operator precedence beyond default values
- B. Combining multiple operators into one new operator
- C. Defining custom behavior for operators with objects
- D. Using more operators than necessary in expressions

Answer: C

Q666. What does the null coalescing operator (??) do?

- A. It checks if a value is undefined and throws error
- B. It converts null values to empty strings always
- C. It prevents null values from being assigned anywhere
- D. It returns the right operand when left is null

Answer: D

Q667. How does bitwise XOR (^) differ from bitwise OR (|)?

- A. XOR returns one only when bits differ; OR either is one
- B. XOR returns one when both bits are one; OR returns zero
- C. XOR works on signed integers only; OR on unsigned ones
- D. XOR does logical operations; OR does arithmetic only

Answer: A

Q668. What is the purpose of the instanceof operator in Java?

- A. It creates a new instance of a class from constructor
- B. It deletes an instance and frees allocated memory
- C. It counts how many instances of a class exist now
- D. It checks if an object is instance of a specific class

Answer: D

Q669. What does the spaceship operator (<==>) return?

- A. Negative, zero, or positive based on comparison
- B. The average of two numeric values as float result
- C. A new object merging properties from two objects
- D. A Boolean true or false based on strict equality

Answer: A

Q670. What is the difference between == and === in JavaScript?

- A. Both perform the same comparison with no difference
- B. === checks value only; == checks value and type
- C. == checks value only; === checks value and type
- D. == is for numbers only; === for strings and objects

Answer: C

Q671. What is the advantage of a for-each loop over a traditional for loop?

- A. For-each loops can modify collection size in iteration
- B. For-each loops work only with arrays not other types
- C. For-each loops simplify iteration by abstracting index
- D. For-each loops execute faster than for loops always

Answer: C

Q672. What is a sentinel value in loop control?

- A. A special value that signals the loop should stop
- B. A flag tracking whether the loop has been started
- C. A variable counting the number of loop iterations
- D. A constant defining maximum loop iteration count

Answer: A

Q673. What is fall-through behavior in a switch statement?

- A. Execution continues into next case when break missing
- B. The switch exits immediately after matching first case
- C. The switch skips all cases and goes to default
- D. The default case always executes before other cases

Answer: A

Q674. What is tail recursion and why is it important?

- A. A technique to reverse order of recursive calls only
- B. Recursion at the end of an array for faster access
- C. A recursive call as last action enabling optimization
- D. Recursion that always uses the tail element of list

Answer: C

Q675. When should a developer use a do-while loop instead of while?

- A. When the loop condition depends on hardware signals
- B. When the loop needs to run a predetermined count
- C. When the loop should never execute if condition false
- D. When the loop body must execute at least one time

Answer: D

Q676. What is the purpose of labeled break statements in Java?

- A. They exit a specific outer loop from nested loops
- B. They add descriptive labels for code documentation
- C. They create named variables within loop block scope
- D. They mark loops for automatic parallel JVM execution

Answer: A

Q677. What is a guard clause in programming?

- A. A security check validating user authentication
- B. A try-catch block placed around every function call
- C. A mutex lock preventing concurrent shared data access
- D. An early return handling edge cases at function start

Answer: D

Q678. How does pattern matching differ from a traditional switch?

- A. Pattern matching is slower and less efficient than switch
- B. Pattern matching can destructure and match complex shapes
- C. Pattern matching cannot have a default case for unmatched
- D. Pattern matching only works with integer and string values

Answer: B

Q679. What is the risk of deeply nested control structures?

- A. They cause undefined behavior in all languages used
- B. They always cause the program to run out of memory
- C. They reduce readability and increase complexity greatly
- D. They prevent the compiler from generating optimized code

Answer: C

Q680. What is a loop invariant and why is it useful?

- A. A special loop type that cannot be modified in execution
- B. An optimization removing unnecessary variables from loop
- C. A condition true before and after each loop iteration
- D. A constant declared outside the loop that never changes

Answer: C

Q681. What is the difference between pass-by-value and pass-by-reference?

- A. Pass-by-value is faster than pass-by-reference in all cases
- B. Pass-by-value sends a copy; reference sends original location
- C. Both methods behave identically with no differences at all
- D. Pass-by-reference creates a copy of the argument each time

Answer: B

Q682. What is a lambda function?

- A. A small anonymous function defined inline in the code
- B. A function that runs automatically at program startup
- C. A function that cannot accept any parameters at all
- D. A function that can only return Boolean values always

Answer: A

Q683. What is function overloading?

- A. Defining multiple functions with same name different params
- B. Replacing an existing function with a new one entirely
- C. Calling a function more times than it can handle
- D. Loading a function into memory before it is needed

Answer: A

Q684. What is a callback function?

- A. A function that undoes the actions of another function
- B. A function that calls back the previous function always
- C. A function passed as an argument to another function
- D. A function that only executes during error handling

Answer: C

Q685. What is a closure in programming?

- A. A function that cannot be called more than once
- B. A function that retains access to its enclosing scope
- C. A function that closes the program when called
- D. A function that blocks all other functions from running

Answer: B

Q686. What is a default parameter in a function?

- A. A parameter with a preset value used if no arg given
- B. A parameter that overrides all other parameter values
- C. A parameter that must always be provided by caller
- D. A parameter that is automatically deleted after use

Answer: A

Q687. What is a higher-order function?

- A. A function that requires administrator access to run
- B. A function that takes or returns other functions
- C. A function with more than ten parameters defined
- D. A function that runs with higher priority than others

Answer: B

Q688. What is the purpose of the main function in C or Java?

- A. It handles all error messages in the entire program
- B. It defines all variables used throughout the program
- C. It serves as the entry point for program execution
- D. It manages memory allocation for all data structures

Answer: C

Q689. What is a variadic function?

- A. A function that can only be called once in program
- B. A function that accepts a variable number of args
- C. A function that changes its name during runtime
- D. A function that returns different types each time

Answer: B

Q690. What is a pure function in functional programming?

- A. A function written in a pure functional language only
- B. A function that only works with primitive data types
- C. A function with no side effects and deterministic output
- D. A function that has been optimized for maximum speed

Answer: C

Q691. What is buffered I/O and why is it used?

- A. It converts binary data to text format before file output
- B. It compresses files automatically to save disk space used
- C. It stores data in a buffer to reduce system call overhead
- D. It encrypts all data before writing to prevent data loss

Answer: C

Q692. What is the purpose of flushing a buffer?

- A. To convert buffer contents from binary to text encoding
- B. To clear the buffer and discard all data within it
- C. To force buffered data to be written to its destination
- D. To increase the size of the buffer for more data storage

Answer: C

Q693. What is serialization in the context of I/O?

- A. Converting objects into a storable or transmittable format
- B. Executing I/O operations in a strict serial sequence
- C. Assigning serial numbers to files for organization
- D. Sorting data in sequential order before processing

Answer: A

Q694. What is the difference between synchronous and asynchronous I/O?

- A. Synchronous blocks until complete; asynchronous does not block
- B. Synchronous only works with network; asynchronous with files
- C. Synchronous is always faster than asynchronous operations
- D. Asynchronous cannot read files only write to them on disk

Answer: A

Q695. What is a stream in I/O programming?

- A. A sequence of data elements available over time to read
- B. A database connection for querying stored information
- C. A graphical interface for displaying program output data
- D. A fixed-size block of memory for storing data values

Answer: A

Q696. What is the purpose of the with statement for file I/O in Python?

- A. It prevents other programs from accessing the same file
- B. It converts file contents to uppercase text characters
- C. It creates a new file with specific permissions on disk
- D. It ensures the file is properly closed after operations

Answer: D

Q697. What is CSV format used for?

- A. Compressing large files to reduce their disk size
- B. Storing executable program code in text files
- C. Encrypting sensitive data for secure transmission
- D. Storing tabular data with comma-separated values

Answer: D

Q698. What is the difference between absolute and relative file paths?

- A. Absolute paths can only reference files not directories
- B. Absolute starts from root; relative from current directory
- C. Absolute is shorter than relative paths in all cases
- D. Relative paths work across different operating systems only

Answer: B

Q699. What is JSON and why is it popular for data interchange?

- A. A binary format designed for storing image data on disk
- B. A protocol for encrypting data during network transfer
- C. A programming language for building web applications
- D. A lightweight text format that is easy to read and parse

Answer: D

Q700. What does the file seek operation do?

- A. It creates a backup copy of the file on another disk
- B. It moves the read/write position to a specific location
- C. It searches for a specific string within file content
- D. It locks the file to prevent concurrent access by others

Answer: B

Q701. What is the difference between mutable and immutable strings?

- A. Mutable strings use less memory; immutable use more always
- B. Mutable strings are faster; immutable are always slower
- C. Mutable strings can be changed; immutable cannot after creation
- D. Mutable strings have a fixed length; immutable can grow

Answer: C

Q702. What is string interpolation?

- A. Converting a string to a different character encoding
- B. Splitting a string by a specific delimiter character
- C. Embedding expressions directly inside a string literal
- D. Comparing two strings character by character position

Answer: C

Q703. What is a regular expression used for?

- A. Encrypting strings to protect sensitive data values
- B. Converting strings between different data type formats
- C. Defining patterns for searching and matching text
- D. Storing mathematical formulas as text strings

Answer: C

Q704. Why is string concatenation in a loop often inefficient?

- A. Because each concatenation creates a new string object
- B. Because the loop counter interferes with string length
- C. Because loops cannot process string data type values
- D. Because concatenation requires network access each time

Answer: A

Q705. What is Unicode and why is it important for strings?

- A. A standard encoding supporting characters from all languages
- B. A programming language designed for text processing
- C. A compression algorithm that reduces string memory usage
- D. A sorting algorithm designed specifically for text strings

Answer: A

Q706. What does the split function do to a string?

- A. It converts the string from uppercase to lowercase
- B. It duplicates the string into two identical copies
- C. It removes every other character from the string text
- D. It divides a string into an array using a delimiter

Answer: D

Q707. What is the difference between StringBuilder and String in Java?

- A. StringBuilder is slower than String for all text operations
- B. StringBuilder is mutable for efficient string modifications
- C. StringBuilder is immutable; String is mutable and changeable
- D. StringBuilder cannot store text only numeric data values

Answer: B

Q708. What is a string pool in Java?

- A. A sorting algorithm specifically designed for string arrays
- B. A cache of string literals reused to save memory space
- C. A collection of string utility functions in a library
- D. A data structure for storing strings in sorted order

Answer: B

Q709. What is the purpose of the replace function on strings?

- A. It substitutes occurrences of a pattern with new text
- B. It moves the string to a different memory location
- C. It renames the variable holding the string value
- D. It converts the string from text into binary format

Answer: A

Q710. What is URL encoding for strings?

- A. Replacing special characters with percent-encoded values
- B. Compressing URL strings to reduce their total length
- C. Encrypting URL strings to prevent unauthorized access
- D. Converting URLs into string variables in the program

Answer: A

Q711. What is the time complexity of accessing an element by index in an array?

- A. $O(1)$ because arrays provide direct index-based access
- B. $O(\log n)$ because arrays use binary search internally
- C. $O(n^2)$ because index calculation is quadratic
- D. $O(n)$ because it must search through all elements

Answer: A

Q712. What is a linked list and how does it differ from an array?

- A. A linked list cannot store more than one hundred elements
- B. A linked list is faster for index access than an array
- C. A linked list uses nodes with pointers allowing dynamic size
- D. A linked list stores elements contiguously like an array does

Answer: C

Q713. What is a hash map collision and how is it typically resolved?

- A. When two values are identical in different hash map entries
- B. When the hash map runs out of memory during insertion
- C. When the hash function produces negative index values
- D. When two keys map to the same bucket needing chain or probe

Answer: D

Q714. What is the difference between an ArrayList and a LinkedList in Java?

- A. ArrayList stores only strings; LinkedList stores only numbers
- B. ArrayList uses nodes with pointers; LinkedList uses an array
- C. ArrayList provides fast index access; LinkedList fast insertion
- D. ArrayList is immutable; LinkedList is mutable and changeable

Answer: C

Q715. What is a priority queue?

- A. A queue that can only store integer data type elements
- B. A queue where elements are processed in insertion order always
- C. A queue that automatically removes duplicate element entries
- D. A queue where elements are processed based on priority value

Answer: D

Q716. What is the purpose of the Collections.sort() method in Java?

- A. It sorts the elements of a list in their natural ordering
- B. It removes all duplicate elements from the given collection
- C. It converts a collection into an array of primitive values
- D. It reverses the order of elements in a collection object

Answer: A

Q717. What is a deque (double-ended queue)?

- A. A queue that only allows deletion not insertion operations
- B. A queue allowing insertion and removal at both front and back
- C. A queue that stores exactly two elements at any given time
- D. A queue that automatically sorts elements upon insertion

Answer: B

Q718. What is the advantage of a balanced binary search tree over a sorted array?

- A. Balanced BST uses less memory than a sorted array always
- B. Balanced BST has faster element access by index position
- C. Balanced BST provides efficient insertion and deletion both
- D. Balanced BST does not require any comparison operations

Answer: C

Q719. What is a sparse array?

- A. An array where most elements are default or zero values
- B. An array that stores elements in non-contiguous memory
- C. An array that only contains string type data elements
- D. An array with a fixed small number of total elements

Answer: A

Q720. What is the difference between a shallow copy and a deep copy of a collection?

- A. Deep copies share nested references; shallow copies clone all
- B. Shallow copies are slower to create than deep copies always
- C. Shallow copies share nested references; deep copies clone all
- D. Both shallow and deep copies produce identical independent results

Answer: C

Q721. What is the difference between an abstract class and an interface?

- A. Abstract classes can have implementation; interfaces cannot always
- B. Abstract classes cannot have fields; interfaces can have fields
- C. Interfaces support multiple inheritance; abstract classes do not
- D. Abstract classes and interfaces are functionally identical always

Answer: A

Q722. What is method overriding in OOP?

- A. Creating a method with a different number of parameters
- B. Deleting a parent method from the child class entirely
- C. Defining a method with the same name in a child class
- D. Calling a method from an unrelated class in the program

Answer: C

Q723. What is composition in OOP and how does it differ from inheritance?

- A. Composition is slower than inheritance in all execution scenarios
- B. Composition and inheritance produce identical code structure always
- C. Composition cannot reuse code from other classes unlike inheritance
- D. Composition uses has-a relationship; inheritance uses is-a relation

Answer: D

Q724. What is the purpose of access modifiers like public, private, protected?

- A. They determine the execution speed of methods in the class
- B. They control the visibility and accessibility of class members
- C. They define the order in which methods are called in class
- D. They specify the data type that a class member can store

Answer: B

Q725. What is the Liskov Substitution Principle?

- A. Objects should have only one reason to change at any time
- B. Subtypes must be substitutable for their base types safely
- C. A class should be open for extension closed for modification
- D. Classes should depend on abstractions not concrete classes

Answer: B

Q726. What is a static method in a class?

- A. A method that cannot accept any parameters at all
- B. A method that cannot be called from outside the class
- C. A method belonging to the class itself not instances
- D. A method that runs automatically when program starts

Answer: C

Q727. What is the diamond problem in multiple inheritance?

- A. A class has too many methods causing diamond-shaped UML
- B. A class cannot be instantiated due to circular references
- C. A class inherits from itself creating an infinite loop
- D. A class inherits from two classes sharing a common ancestor

Answer: D

Q728. What is dependency injection and why is it useful?

- A. Providing dependencies externally rather than creating inside
- B. Injecting code into a running program for hot patching
- C. Adding extra methods to a class during program execution
- D. Inserting breakpoints into code for debugging at runtime

Answer: A

Q729. What is the difference between aggregation and composition?

- A. Aggregation has strong ownership; composition has weak ownership
- B. Both aggregation and composition represent identical relationships
- C. Aggregation has weak ownership; composition has strong ownership
- D. Aggregation uses inheritance; composition uses interfaces only

Answer: C

Q730. What is the purpose of the final keyword in Java?

- A. It prevents classes from being inherited or methods overridden
- B. It specifies that a variable can only hold final data type
- C. It marks a method as the first one to execute in class
- D. It indicates the method will return the final calculated value

Answer: A

Q731. What is a dangling pointer and why is it dangerous?

- A. A pointer that takes too long to dereference in code
- B. A pointer that points to the beginning of an allocated block
- C. A pointer that is declared but never used in the program
- D. A pointer referencing freed memory causing undefined behavior

Answer: D

Q732. What is reference counting as a garbage collection strategy?

- A. Tracking how many references point to each object in memory
- B. Tracking how many times each function is called at runtime
- C. Counting the total amount of memory allocated by program
- D. Counting how many variables exist in the program scope

Answer: A

Q733. What is the difference between stack overflow and heap overflow?

- A. Both stack and heap overflow are caused by the same exact problem
- B. Stack overflow uses too much heap; heap overflow uses too much stack
- C. Stack overflow is harmless; heap overflow always crashes program
- D. Stack overflow is from deep recursion; heap from excess allocation

Answer: D

Q734. What is virtual memory and why do operating systems use it?

- A. A type of memory used only by virtual machines and emulators
- B. A memory encryption technique to protect sensitive program data
- C. A technique using disk to extend available memory beyond RAM
- D. Physical RAM that is faster than regular memory modules

Answer: C

Q735. What is a memory pool allocation strategy?

- A. Pooling memory from multiple computers on a network
- B. Storing all variables in a single continuous memory block
- C. Allocating memory from a single global heap always
- D. Pre-allocating fixed-size blocks for fast repeated allocation

Answer: D

Q736. What is the purpose of smart pointers in C++?

- A. They convert pointers into integers for easier comparison
- B. They automatically manage memory ownership and deallocation
- C. They make pointer arithmetic faster on modern CPUs
- D. They prevent pointers from being passed to any functions

Answer: B

Q737. What is memory fragmentation?

- A. When free memory exists but is split into small unusable chunks
- B. When memory is completely full with no free space left
- C. When memory hardware develops physical defects over time
- D. When a program accesses memory belonging to another process

Answer: A

Q738. What is a buffer overflow?

- A. Writing more data than a buffer can hold overwriting adjacent memory
- B. Reading data from a buffer that has already been fully emptied
- C. Allocating a buffer that is larger than available memory space
- D. Writing less data than a buffer can hold in its capacity

Answer: A

Q739. What is the RAII pattern in C++?

- A. A pattern for parallel processing of arrays using multiple threads
- B. A technique for reducing memory usage of string operations
- C. A pattern for implementing recursive algorithms efficiently
- D. Resource acquisition tied to object lifetime via constructor and destructor

Answer: D

Q740. What is the difference between malloc and new in C++?

- A. malloc is faster than new in all performance benchmarks
- B. malloc is for C++ only; new is for C language only
- C. new cannot allocate memory for arrays unlike malloc can
- D. malloc allocates raw memory; new also calls the constructor

Answer: D

Q741. What is the difference between checked and unchecked exceptions in Java?

- A. Checked exceptions are less severe than unchecked exceptions
- B. Checked exceptions occur at runtime; unchecked at compile time
- C. Both checked and unchecked exceptions behave identically always
- D. Checked must be declared or caught; unchecked are not required

Answer: D

Q742. What is exception chaining and why is it useful?

- A. Catching multiple exceptions in a single catch block always
- B. Throwing the same exception multiple times in succession
- C. Creating a chain of try-catch blocks for different error types
- D. Wrapping an exception in another to preserve the original cause

Answer: D

Q743. What is the purpose of custom exception classes?

- A. They make exceptions run faster than built-in exception types
- B. They automatically fix errors without any developer intervention
- C. They replace all built-in exceptions in the language entirely
- D. They provide domain-specific error types with relevant context

Answer: D

Q744. Why should you avoid using exceptions for normal control flow?

- A. Because exceptions prevent the program from being compiled
- B. Because exceptions cannot carry any data or message with them
- C. Because exceptions are expensive and obscure program logic flow
- D. Because exceptions only work inside the main function of program

Answer: C

Q745. What is the difference between throw and throws in Java?

- A. throw is for checked exceptions only; throws for unchecked only
- B. throw creates a new exception class; throws catches exceptions
- C. throw raises an exception; throws declares exceptions a method may throw
- D. Both throw and throws perform the exact same operation always

Answer: C

Q746. What is a try-with-resources statement in Java?

- A. A try block that limits CPU resources used during execution
- B. A try block that automatically imports required library classes
- C. A try block that automatically closes resources when done
- D. A try block that retries the operation with different data

Answer: C

Q747. What is the purpose of logging exceptions?

- A. To convert exceptions into warnings that can be safely ignored
- B. To automatically send exceptions to users via email alerts
- C. To prevent exceptions from occurring in future executions
- D. To record error details for debugging and monitoring purposes

Answer: D

Q748. What happens if an exception is thrown but not caught?

- A. The exception is automatically converted to a return value
- B. The program terminates and typically displays a stack trace
- C. The exception is silently ignored and execution continues
- D. The program pauses and waits for user input to continue

Answer: B

Q749. What is a multi-catch block in Java?

- A. Nesting multiple catch blocks inside one another deeply
- B. A catch block that runs multiple times for one exception
- C. A single catch block handling multiple exception types
- D. Catching exceptions from multiple try blocks at one time

Answer: C

Q750. What is defensive programming in relation to error handling?

- A. Anticipating potential errors and handling them proactively
- B. Only handling errors after they cause the program to crash
- C. Ignoring all possible errors to keep the code base simple
- D. Delegating all error handling to the operating system only

Answer: A

Q751. What is semantic versioning (SemVer)?

- A. Dating versions with the release date as the number
- B. Numbering versions sequentially with single integers
- C. Naming versions with random words for easy recall
- D. A system using MAJOR.MINOR.PATCH format for versions

Answer: D

Q752. What is dependency injection in the context of modules?

- A. Providing module dependencies externally rather than hardcoding
- B. Injecting malicious code into a module for exploitation
- C. Removing unused dependencies from the project configuration
- D. Automatically downloading dependencies from the internet

Answer: A

Q753. What is circular dependency and why is it problematic?

- A. When a module depends on too many other external modules
- B. When two modules have functions with the same name defined
- C. When a dependency is updated too frequently breaking code
- D. When module A depends on B and B depends on A causing issues

Answer: D

Q754. What is the difference between a static and dynamic library?

- A. Static is linked at compile time; dynamic is linked at runtime
- B. Static is for C only; dynamic is for all other languages only
- C. Static is faster but dynamic uses less total memory always
- D. Static cannot be updated; dynamic updates automatically always

Answer: A

Q755. What is a lockfile in dependency management?

- A. A file that limits the number of dependencies allowed
- B. A file that encrypts the dependencies for secure storage
- C. A file that locks the source code from being modified
- D. A file recording exact versions of all installed dependencies

Answer: D

Q756. What is tree shaking in module bundling?

- A. Converting modules from one format to another at build
- B. Reorganizing module code into a tree data structure
- C. Sorting modules alphabetically for better organization
- D. Removing unused code from the final bundle to reduce size

Answer: D

Q757. What is the difference between CommonJS and ES modules?

- A. CommonJS is newer than ES modules in JavaScript evolution
- B. CommonJS uses require/exports; ES modules use import/export
- C. CommonJS uses import/export; ES modules use require/exports
- D. Both systems use identical syntax and behave the same way

Answer: B

Q758. What is a monorepo?

- A. A single repository containing multiple related projects
- B. A repository containing only one small module of code
- C. A repository that stores only documentation not code
- D. A repository that cannot have more than one contributor

Answer: A

Q759. What is the purpose of a module bundler like Webpack?

- A. It compresses images used in web application projects
- B. It automatically writes unit tests for all module code
- C. It combines multiple modules into optimized output bundles
- D. It deploys applications to production servers directly

Answer: C

Q760. What is a transitive dependency?

- A. A dependency of a dependency that is indirectly required
- B. A dependency that transitions between different versions
- C. A dependency that is optional and not always installed
- D. A dependency that is only used during testing phases

Answer: A

Q761. What is the time complexity of binary search on a sorted array?

- A. $O(n^2)$ because it uses nested loops for searching
- B. $O(n)$ because it checks every element in the array
- C. $O(1)$ because it directly accesses the target element
- D. $O(\log n)$ because it halves the search space each step

Answer: D

Q762. What is memoization and how does it improve performance?

- A. It memorizes variable names for faster variable lookup
- B. It logs all function calls to memory for later review
- C. It caches function results to avoid redundant computation
- D. It allocates extra memory to make functions run faster

Answer: C

Q763. What is the observer pattern in software design?

- A. A pattern where one object observes system memory usage
- B. A pattern where objects subscribe to and receive event updates
- C. A pattern for creating exactly one instance of a class only
- D. A pattern for sorting collections using comparison callbacks

Answer: B

Q764. What is dynamic programming?

- A. Breaking problems into overlapping subproblems and caching
- B. Creating programs that adapt to user preferences at runtime
- C. Writing code that changes its own source during execution
- D. Programming in a dynamically typed language like Python

Answer: A

Q765. What is a hash function and what makes a good one?

- A. A function that encrypts data using a secret key for security
- B. A function that compresses files to reduce their storage size
- C. A function mapping data to fixed-size values with few collisions
- D. A function that sorts data into alphabetical order always

Answer: C

Q766. What is the singleton pattern?

- A. A pattern ensuring a class has only one instance globally
- B. A pattern for connecting a single client to a server only
- C. A pattern that creates multiple instances of different classes
- D. A pattern that limits functions to one parameter at most

Answer: A

Q767. What is the factory pattern?

- A. A pattern for converting objects between different types
- B. A pattern for destroying objects when no longer needed
- C. A pattern delegating object creation to a factory method
- D. A pattern for mass producing identical objects in batches

Answer: C

Q768. What is the difference between BFS and DFS graph traversal?

- A. BFS is only for trees; DFS is only for general graph structures
- B. BFS explores depth first; DFS explores breadth first always
- C. BFS uses a stack; DFS uses a queue for storing nodes to visit
- D. BFS explores level by level; DFS explores as deep as possible

Answer: D

Q769. What is a heap data structure?

- A. A complete binary tree satisfying the heap ordering property
- B. A hash table that automatically resizes when it gets full
- C. A linked list sorted in ascending order of element values
- D. A disorganized pile of unsorted data in memory blocks

Answer: A

Q770. What is the strategy pattern?

- A. A pattern for choosing between inheritance and composition
- B. A pattern for planning the development timeline of projects
- C. A pattern encapsulating interchangeable algorithms in objects
- D. A pattern that strategically allocates memory for performance

Answer: C

Q771. What is a semaphore and how does it differ from a mutex?

- A. A semaphore allows only one thread; a mutex allows multiple
- B. A semaphore is for processes only; a mutex is for threads only
- C. A semaphore allows multiple threads up to a count; mutex allows one
- D. Both semaphores and mutexes behave identically in all situations

Answer: C

Q772. What is a thread pool and why is it used?

- A. A pool of memory shared between all threads in a program
- B. A data structure for storing thread execution results only
- C. A group of pre-created threads reused for executing tasks
- D. A debugging tool that monitors thread activity in programs

Answer: C

Q773. What are the four necessary conditions for a deadlock to occur?

- A. Speed, memory, disk, and network resource limitations
- B. Compilation, linking, loading, and execution phase errors
- C. Input, processing, output, and storage operation failures
- D. Mutual exclusion, hold-and-wait, no preemption, circular wait

Answer: D

Q774. What is the producer-consumer problem?

- A. A problem of converting data between different encoding formats
- B. A problem of coordinating threads that produce and consume data
- C. A problem of allocating memory between different program modules
- D. A problem about optimizing compilation speed of programs

Answer: B

Q775. What is an atomic operation?

- A. An operation that completes indivisibly without interruption
- B. An operation that uses atomic energy for faster processing
- C. An operation that can be divided into smaller sub-steps
- D. An operation that only works on integer data type values

Answer: A

Q776. What is the difference between preemptive and cooperative multitasking?

- A. Both use identical scheduling mechanisms in all operating systems
- B. Preemptive only works on single-core; cooperative on multi-core
- C. Preemptive OS controls scheduling; cooperative tasks yield control
- D. Preemptive is slower; cooperative is faster in all situations

Answer: C

Q777. What is a condition variable used for?

- A. Storing the condition of an if-else statement in memory
- B. Checking if a variable meets certain validation requirements
- C. Allowing threads to wait for a specific condition to become true
- D. Defining constant values that cannot be changed during runtime

Answer: C

Q778. What is the dining philosophers problem about?

- A. Managing memory allocation in object-oriented program design
- B. Optimizing database queries for better application performance
- C. Designing efficient algorithms for sorting large data arrays
- D. Illustrating deadlock and resource contention among processes

Answer: D

Q779. What is thread safety?

- A. Code that functions correctly when accessed by multiple threads
- B. A technique for making threads run faster on modern hardware
- C. A security feature that prevents threads from being terminated
- D. Writing code that runs only on a single thread at one time

Answer: A

Q780. What is the difference between blocking and non-blocking algorithms?

- A. Blocking can cause threads to wait; non-blocking guarantee progress
- B. Both types have identical performance in all concurrent scenarios
- C. Blocking algorithms are faster; non-blocking are always slower
- D. Non-blocking algorithms cannot be used with multiple threads

Answer: A

Q781. What is test-driven development (TDD)?

- A. Writing tests after all code is complete and deployed
- B. Developing tests and code simultaneously without any plan
- C. Writing tests before code then implementing to pass tests
- D. Testing only the most critical features of the application

Answer: C

Q782. What is code coverage in testing?

- A. The total lines of code in the entire project codebase
- B. The percentage of code executed during test suite runs
- C. The number of developers who have reviewed the code
- D. The amount of documentation written for the codebase

Answer: B

Q783. What is an integration test?

- A. A test verifying a single function works in isolation
- B. A test ensuring code compiles without any syntax errors
- C. A test measuring the speed of algorithm execution time
- D. A test verifying that multiple components work together

Answer: D

Q784. What is mocking in unit testing?

- A. Replacing real dependencies with controlled fake objects
- B. Ridiculing code that is poorly written by developers
- C. Copying production data to use in test environments
- D. Running tests multiple times to check for consistency

Answer: A

Q785. What is regression testing?

- A. Testing the application on different operating systems platforms
- B. Re-running tests to ensure changes have not broken existing code
- C. Testing new features before they are implemented in code
- D. Testing the program under extreme load and stress conditions

Answer: B

Q786. What is a profiler used for in optimization?

- A. Measuring performance bottlenecks in program execution
- B. Formatting source code to follow style guidelines set
- C. Writing automated test cases for the program codebase
- D. Translating code from one programming language to another

Answer: A

Q787. What is the difference between white-box and black-box testing?

- A. Both approaches test code in exactly the same way always
- B. White-box tests the UI; black-box tests the backend code
- C. White-box is manual testing; black-box is automated testing
- D. White-box tests with code knowledge; black-box without it

Answer: D

Q788. What is assertion in testing?

- A. A statement that imports external libraries into the code
- B. A statement that declares a variable with a specific type
- C. A statement that verifies a condition is true during testing
- D. A statement that prints debug information to the console

Answer: C

Q789. What is continuous integration (CI)?

- A. Integrating third-party libraries into the project once
- B. Continuously writing code without any breaks or pauses
- C. Automatically building and testing code on every commit
- D. Integrating all code at the end of the development cycle

Answer: C

Q790. What is a flaky test and why is it problematic?

- A. A test that always fails consistently every time it runs
- B. A test that only works on the original developer's machine
- C. A test that passes or fails inconsistently without code changes
- D. A test that is too slow to run during the CI pipeline build

Answer: C

Q791. What is continuous deployment (CD)?

- A. Deploying all code changes at the end of the project
- B. Continuously writing deployment scripts for the team
- C. Deploying code only during scheduled maintenance windows
- D. Automatically deploying every passing build to production

Answer: D

Q792. What is the SOLID principle in software design?

- A. A single principle about writing solid unbreakable programs
- B. A principle about making code run on solid-state hardware
- C. Five design principles for maintainable object-oriented code
- D. A principle requiring all code to be stored in solid files

Answer: C

Q793. What is technical debt?

- A. The monetary cost of purchasing development tools and licenses
- B. The amount of time spent learning new programming technologies
- C. The implied cost of future rework from choosing quick solutions
- D. The storage cost of maintaining large code repositories online

Answer: C

Q794. What is a microservice architecture?

- A. An architecture using very small functions for all operations
- B. A development practice of writing very concise source code
- C. An architecture breaking an app into small independent services
- D. A very small application with minimal functionality only

Answer: C

Q795. What is pair programming?

- A. Two developers working together on the same code actively
- B. Two programs running in parallel on the same computer
- C. Two developers working on separate tasks simultaneously
- D. Two teams competing to write the same feature faster

Answer: A

Q796. What is the purpose of a linter?

- A. To encrypt source code to protect intellectual property
- B. To convert code from one programming language to another
- C. To analyze code for style issues and potential errors found
- D. To compile source code into optimized machine code output

Answer: C

Q797. What is the difference between monolithic and microservice architectures?

- A. Monolithic is one deployable unit; microservices are many separate
- B. Both architectures deploy and scale in exactly the same way
- C. Microservices use one language; monolithic uses many languages
- D. Monolithic is faster; microservices are slower in all cases

Answer: A

Q798. What is a sprint in Scrum methodology?

- A. A fixed-length iteration for completing a set of planned work
- B. A quick fix applied to production code in emergency situations
- C. A code review session held at the end of every work day
- D. A performance test measuring application response time speed

Answer: A

Q799. What is the purpose of a retrospective in agile development?

- A. To assign tasks to team members for the next sprint period
- B. To review the final product before releasing to customers
- C. To reflect on the process and identify areas for improvement
- D. To estimate the effort required for upcoming development work

Answer: C

Q800. What is infrastructure as code (IaC)?

- A. Managing infrastructure through machine-readable definition files
- B. Converting infrastructure diagrams into programming source code
- C. Writing code that runs on physical infrastructure hardware
- D. Coding directly on production servers without version control

Answer: A

Q801. What is the difference between a compiler error and a linker error?

- A. Linker errors happen before compiler errors in the build process
- B. Compiler errors are warnings while linker errors crash the system
- C. There is no difference between them
- D. Compiler errors occur in syntax checking while linker errors occur when combining object files

Answer: D

Q802. What is the purpose of an object file in compilation?

- A. It is the final executable program
- B. It contains machine code that has not yet been linked
- C. It stores the original source code
- D. It holds only debug information

Answer: B

Q803. What is the difference between procedural and object-oriented programming?

- A. Procedural languages cannot handle large programs
- B. OOP cannot use functions at all
- C. Procedural is faster than OOP in all cases
- D. Procedural uses functions to organize code while OOP uses classes and objects

Answer: D

Q804. What is a semantic error in programming?

- A. An error where the code has incorrect punctuation
- B. An error where the code is syntactically valid but logically meaningless
- C. An error caused by using the wrong file extension
- D. An error that prevents compilation due to missing brackets

Answer: B

Q805. What is the role of a loader in program execution?

- A. It loads the executable into memory and prepares it for execution
- B. It links object files together into an executable
- C. It compiles the source code into machine code
- D. It checks the program for syntax errors

Answer: A

Q806. What is the difference between compiled and managed languages?

- A. Managed languages must be interpreted line by line
- B. Compiled languages produce native code while managed languages run on a virtual machine
- C. Compiled languages cannot use garbage collection
- D. Managed languages are always slower than compiled languages

Answer: B

Q807. What is an abstract syntax tree (AST) used for?

- A. Representing the hierarchical structure of source code during compilation
- B. Drawing flowcharts of the program
- C. Storing the final executable code
- D. Managing memory allocation at runtime

Answer: A

Q808. What is the purpose of a garbage-collected runtime environment?

- A. To optimize CPU instruction scheduling
- B. To prevent all types of programming errors
- C. To compile code faster than traditional compilers
- D. To automatically reclaim memory that is no longer in use

Answer: D

Q809. What is the difference between a scripting language and a systems language?

- A. Scripting languages cannot define functions
- B. Systems languages are always object-oriented
- C. Scripting languages prioritize ease of use while systems languages prioritize performance and control
- D. Scripting languages can only run on web browsers

Answer: C

Q810. What is the purpose of intermediate representation (IR) in compilers?

- A. To display the program output to the user
- B. To serve as a platform-independent form between source code and machine code
- C. To manage the file system during compilation
- D. To encrypt the source code for security

Answer: B

Q811. What is autoboxing in Java?

- A. Automatically compiling code without user intervention
- B. Automatically converting primitive types to their wrapper class objects
- C. Automatically resizing arrays when they are full
- D. Automatically casting all types to strings

Answer: B

Q812. What is the difference between a weakly typed and a strongly typed language?

- A. Strongly typed languages cannot perform arithmetic operations
- B. Weakly typed languages are always interpreted
- C. Weakly typed languages do not have data types at all
- D. Weakly typed languages allow implicit type conversions while strongly typed languages restrict them

Answer: D

Q813. What is a tuple data type?

- A. A type that stores exactly two values
- B. An immutable ordered collection of elements that can have different types
- C. A mutable list of elements of the same type
- D. A type used only for database records

Answer: B

Q814. What is type coercion?

- A. Explicitly casting a variable to a new type using syntax
- B. Restricting a variable to a single type permanently
- C. Automatic conversion of a value from one type to another by the language
- D. Removing type information at compile time

Answer: C

Q815. What is the difference between mutable and immutable data types?

- A. Mutable types can be changed after creation while immutable types cannot
- B. Mutable types can only store numbers
- C. Mutable types use more memory than immutable types always
- D. Immutable types are only available in functional languages

Answer: A

Q816. What is a generic type parameter?

- A. A placeholder for a type that is specified when the class or method is used
- B. A type that is automatically inferred at runtime
- C. A default type assigned when no type is specified
- D. A type that can only hold generic string values

Answer: A

Q817. What is the purpose of the void type?

- A. To represent a variable that holds zero
- B. To declare a variable with no type
- C. To create an empty array
- D. To indicate that a function returns no value

Answer: D

Q818. What is a tagged union or discriminated union?

- A. A type that stores one of several possible types along with a tag indicating which type is active
- B. A type used only in database programming
- C. A type that merges two arrays into one
- D. A type that combines multiple unrelated types randomly

Answer: A

Q819. Why does 0.1 + 0.2 not equal 0.3 in most programming languages?

- A. Because floating-point representation cannot exactly represent some decimal fractions
- B. Because addition is not supported for decimal numbers
- C. Because decimal numbers are stored as strings internally
- D. Because the compiler rounds all numbers to integers

Answer: A

Q820. What is the difference between a struct and a class in C#?

- A. Structs support inheritance but classes do not
- B. Classes cannot be instantiated while structs can
- C. Structs can have methods but classes cannot
- D. Structs are value types allocated on the stack while classes are reference types allocated on the heap

Answer: D

Q821. What is the difference between logical OR (||) and bitwise OR (|)?

- A. Logical OR works on numbers while bitwise OR works on Booleans
- B. They produce identical results for all input types
- C. Bitwise OR short-circuits but logical OR does not
- D. Logical OR works on Boolean values and short-circuits while bitwise OR operates on individual bits

Answer: D

Q822. What is the purpose of the modulus operator in determining even or odd numbers?

- A. If $n \% 2$ equals 1 the number is even, if it equals 0 the number is odd
- B. If $n \% 3$ equals 0 the number is even
- C. If $n \% 2$ equals 0 the number is even, if it equals 1 the number is odd
- D. The modulus operator cannot be used for this purpose

Answer: C

Q823. What does the bitwise NOT (~) operator do to an integer?

- A. It sets all bits to zero
- B. It converts the integer to its absolute value
- C. It flips all the bits, changing 0s to 1s and 1s to 0s
- D. It doubles the value of the integer

Answer: C

Q824. What is the optional chaining operator (?.) used for?

- A. Creating optional function parameters
- B. Safely accessing nested object properties without throwing if a reference is null
- C. Defining nullable return types
- D. Forcing a null value to throw an exception

Answer: B

Q825. What is the difference between the assignment operator (=) and the equality operator (==)?

- A. Neither operator works with numeric values
- B. The assignment operator stores a value in a variable while the equality operator compares two values
- C. They perform the same function in all contexts
- D. The equality operator stores values while the assignment operator compares them

Answer: B

Q826. How does the right shift operator (>>) work on signed integers?

- A. It shifts bits to the right and fills the leftmost bits with the sign bit
- B. It divides the number by 10
- C. It shifts bits to the right and always fills with zeros
- D. It reverses the order of all bits

Answer: A

Q827. What is the spread operator (...) used for in JavaScript?

- A. Removing duplicate values from an array
- B. Multiplying all elements in an array
- C. Expanding an iterable into individual elements
- D. Converting an object to a string

Answer: C

Q828. What is the typeof operator used for in JavaScript?

- A. Returning a string indicating the type of the operand
- B. Comparing two types for equality
- C. Creating a new type at runtime
- D. Casting a value to a specified type

Answer: A

Q829. What is the result of using the bitwise AND operator (5 & 6)?

- A. 6
- B. 4
- C. 11
- D. 5

Answer: B

Q830. What is the exponentiation operator () in Python and JavaScript?**

- A. It raises the left operand to the power of the right operand
- B. It performs multiplication of two numbers
- C. It performs integer division
- D. It calculates the square root of a number

Answer: A

Q831. What is the advantage of using a switch statement over multiple if-else chains?

- A. Switch statements automatically handle all possible cases
- B. Switch statements can evaluate complex Boolean expressions better
- C. Switch statements are clearer and more efficient when comparing a single variable against many constant values
- D. Switch statements work with all data types in all languages

Answer: C

Q832. What is the difference between a pre-test loop and a post-test loop?

- A. A pre-test loop checks the condition before each iteration while a post-test loop checks after
- B. A pre-test loop runs faster than a post-test loop
- C. A pre-test loop always runs at least twice
- D. A post-test loop cannot use counter variables

Answer: A

Q833. What is a state machine and how does it relate to control flow?

- A. A type of loop that runs indefinitely
- B. A machine that stores all program variables
- C. A hardware device that controls program execution
- D. A model where behavior is determined by the current state and transitions triggered by events

Answer: D

Q834. Why should deeply nested control structures be avoided?

- A. They reduce code readability and increase cognitive complexity for developers
- B. They are not supported by modern compilers
- C. They always cause runtime errors in production
- D. They consume significantly more memory than flat structures

Answer: A

Q835. What is short-circuit evaluation in the context of control flow?

- A. Evaluating only part of a Boolean expression when the result can be determined early
- B. Skipping the entire if block when the condition is complex
- C. Terminating the program when an error is detected
- D. Running loops with fewer iterations than specified

Answer: A

Q836. What is the purpose of the 'yield' keyword in generators?

- A. It pauses the function execution and returns a value, resuming from that point on the next call
- B. It yields control to the operating system
- C. It creates a new thread for the function
- D. It terminates the function permanently

Answer: A

Q837. What is the difference between a counted loop and a conditional loop?

- A. A conditional loop always runs at least once
- B. A counted loop can only count up while a conditional loop counts down
- C. A counted loop runs a predetermined number of times while a conditional loop runs until a condition changes
- D. A counted loop does not use any variables

Answer: C

Q838. What is the ternary operator and how does it simplify control flow?

- A. It creates three separate code branches
- B. It allows three conditions to be checked simultaneously
- C. It replaces a simple if-else with a concise single-line expression that returns one of two values
- D. It replaces for loops with a shorter syntax

Answer: C

Q839. What is the purpose of the 'pass' statement in Python?

- A. It skips to the next iteration of a loop
- B. It passes a value to a function
- C. It terminates the program gracefully
- D. It serves as a placeholder where a statement is syntactically required but no action is needed

Answer: D

Q840. What is the difference between break and return inside a nested loop?

- A. They behave identically in all situations
- B. Break exits only the innermost loop while return exits the entire function
- C. Break exits all nested loops while return exits only one
- D. Break exits the function while return exits only the loop

Answer: B

Q841. What is the difference between named and anonymous functions?

- A. Named functions are faster than anonymous functions
- B. Anonymous functions cannot accept parameters
- C. Named functions cannot be passed as arguments
- D. Named functions have an identifier while anonymous functions are defined without a name

Answer: D

Q842. What is the purpose of the *args parameter in Python?

- A. It creates a pointer to another function
- B. It converts all arguments to strings
- C. It restricts the function to accept only one argument
- D. It allows a function to accept a variable number of positional arguments as a tuple

Answer: D

Q843. What is function composition?

- A. Writing multiple functions in the same file
- B. Inheriting functions from a parent class
- C. Creating a function with many parameters
- D. Combining two or more functions so the output of one becomes the input of another

Answer: D

Q844. What is the difference between synchronous and asynchronous function calls?

- A. Synchronous calls are always faster than asynchronous calls
- B. Synchronous calls block until completion while asynchronous calls return immediately and complete later
- C. Asynchronous calls cannot return values
- D. Synchronous calls run on multiple threads simultaneously

Answer: B

Q845. What is a decorator in Python?

- A. A comment that describes a function
- B. A tool for formatting output
- C. A special syntax for defining classes
- D. A function that wraps another function to modify or extend its behavior

Answer: D

Q846. What is the **kwargs parameter in Python used for?

- A. It converts arguments to key-value pairs automatically
- B. It blocks keyword arguments from being passed
- C. It creates two copies of each argument
- D. It accepts a variable number of keyword arguments as a dictionary

Answer: D

Q847. What is the purpose of the 'static' keyword for functions in Java?

- A. It makes the function run faster
- B. It prevents the function from being modified
- C. It restricts the function to a single thread
- D. It allows the function to be called without creating an instance of the class

Answer: D

Q848. What is an immediately invoked function expression (IIFE)?

- A. A function that waits for user input before executing
- B. A function that is defined and called at the same time
- C. A function that immediately throws an exception
- D. A function that is invoked only once per program run

Answer: B

Q849. What is the difference between a procedure and a function in traditional programming terminology?

- A. A function returns a value while a procedure performs an action without returning a value
- B. A procedure is faster than a function
- C. A procedure can only contain one statement
- D. A function cannot modify global state

Answer: A

Q850. What is the difference between character streams and byte streams in Java?

- A. Character streams are faster than byte streams
- B. Byte streams can only read files while character streams can read and write
- C. Character streams handle text data using Unicode encoding while byte streams handle raw binary data
- D. Character streams cannot handle non-English text

Answer: C

Q851. What is deserialization?

- A. Deleting serialized data from disk
- B. Compressing data for faster transmission
- C. Converting a stored byte sequence back into an object in memory
- D. Converting text to binary format

Answer: C

Q852. What is the purpose of a buffer in I/O operations?

- A. To permanently store data on disk
- B. To compress files before saving
- C. To encrypt data during transfer
- D. To temporarily hold data in memory to reduce the number of actual I/O operations

Answer: D

Q853. What is the difference between text mode and binary mode when opening a file?

- A. Text mode translates line endings and handles encoding while binary mode reads raw bytes without translation
- B. Text mode only works with ASCII characters
- C. Binary mode is slower than text mode
- D. Text mode is only for reading while binary mode is only for writing

Answer: A

Q854. What is the purpose of the StringIO class in Python?

- A. It converts strings to integers
- B. It encrypts string data for file storage
- C. It provides a file-like interface for reading and writing strings in memory
- D. It counts the number of I/O operations

Answer: C

Q855. What is the purpose of I/O redirection in command-line programs?

- A. It speeds up I/O operations by using faster hardware
- B. It redirects standard input, output, or error streams to files or other programs
- C. It prevents programs from reading input
- D. It converts all output to binary format

Answer: B

Q856. What is XML parsing and what are the two main approaches?

- A. DOM is faster than SAX for all file sizes
- B. DOM parsing loads the entire document into memory while SAX parsing processes it event by event
- C. SAX loads the entire document while DOM reads it line by line
- D. XML can only be parsed using regular expressions

Answer: B

Q857. What is the difference between formatted and unformatted I/O?

- A. Formatted I/O is always faster than unformatted I/O
- B. Unformatted I/O can only handle strings
- C. Formatted I/O is only available in C
- D. Formatted I/O converts data to human-readable text while unformatted I/O reads and writes raw binary data

Answer: D

Q858. What is the purpose of the encoding parameter when opening a file in Python?

- A. It specifies how characters are converted to and from bytes when reading or writing
- B. It encrypts the file contents for security
- C. It compresses the file to save disk space
- D. It sets the maximum file size allowed

Answer: A

Q859. What is the difference between polling and event-driven I/O?

- A. Polling is always more efficient than event-driven I/O
- B. Polling cannot handle multiple I/O sources
- C. Polling repeatedly checks for I/O readiness while event-driven I/O notifies the program when I/O is ready
- D. Event-driven I/O blocks the program until data arrives

Answer: C

Q860. What is the difference between `String.valueOf()` and `toString()` in Java?

- A. `valueOf()` handles null by returning the string 'null' while `toString()` throws a `NullPointerException` on null
- B. They are identical methods with different names
- C. `valueOf()` only works with primitive types
- D. `toString()` handles null better than `valueOf()`

Answer: A

Q861. What is a raw string literal and when is it useful?

- A. A string where escape sequences are treated as literal characters, useful for regex and file paths
- B. A string that is always stored in raw binary format
- C. A string that cannot contain any special characters
- D. A string that has not been processed by the compiler

Answer: A

Q862. Why is string concatenation in a loop inefficient for immutable strings?

- A. Because each concatenation creates a new string object, copying all existing characters
- B. Because the loop runs slower with string operations
- C. Because the compiler cannot optimize string operations in loops
- D. Because strings in loops are automatically encrypted

Answer: A

Q863. What is the purpose of the join() method for strings?

- A. It connects two string variables permanently
- B. It links a string to a file on disk
- C. It concatenates an iterable of strings using a specified separator
- D. It joins a string with a number

Answer: C

Q864. What is the difference between ASCII and extended ASCII?

- A. Extended ASCII is a completely different encoding system
- B. ASCII defines 128 characters using 7 bits while extended ASCII adds 128 more characters using a full 8 bits
- C. Extended ASCII is only used for Asian languages
- D. ASCII supports more characters than extended ASCII

Answer: B

Q865. What is a template literal (template string) in JavaScript?

- A. A string enclosed in backticks that supports embedded expressions and multi-line text
- B. A string that serves as a template for HTML documents
- C. A string defined in a separate template file
- D. A string that cannot be modified after creation

Answer: A

Q866. What is the purpose of the format() method or f-strings for string formatting?

- A. They format the string as valid HTML
- B. They change the font and styling of the string
- C. They compress the string to reduce memory usage
- D. They embed variable values and expressions directly within a string template

Answer: D

Q867. What is the difference between the equals() method and == for string comparison in Java?

- A. equals() is faster than == for all cases
- B. == compares content while equals() compares references
- C. equals() compares content while == compares object references in memory
- D. They perform identical comparisons

Answer: C

Q868. What is a character class in regular expressions?

- A. A CSS class applied to text characters
- B. A function that classifies characters as letters or digits
- C. A class in OOP that represents a single character
- D. A set of characters enclosed in square brackets that matches any single character from the set

Answer: D

Q869. What is the purpose of the strip() or trim() method beyond removing spaces?

- A. It strips the string of all vowels
- B. It only removes spaces and nothing else
- C. It reduces the string to a fixed length
- D. It removes specified leading and trailing characters, not just whitespace

Answer: D

Q870. What is the difference between open addressing and separate chaining for hash table collision resolution?

- A. Separate chaining stores colliding elements in linked lists at each bucket while open addressing probes for the next empty slot
- B. Separate chaining requires a sorted hash table
- C. Open addressing uses linked lists while separate chaining probes for empty slots
- D. They always produce identical performance characteristics

Answer: A

Q871. What is the time complexity of inserting an element at the beginning of a linked list?

- A. $O(\log n)$ because it uses binary search
- B. $O(n)$ because all elements must be shifted
- C. $O(1)$ because only the head pointer needs to be updated
- D. $O(n^2)$ because of nested iterations

Answer: C

Q872. What is a circular buffer (ring buffer)?

- A. A buffer that copies data in a circular pattern
- B. A buffer shaped like a circle in memory
- C. A fixed-size buffer that wraps around when the end is reached, reusing space from the beginning
- D. A buffer that stores only circular data structures

Answer: C

Q873. What is the load factor of a hash table?

- A. The ratio of stored elements to the total number of buckets
- B. The speed at which elements are inserted
- C. The maximum number of elements it can store
- D. The amount of memory each element consumes

Answer: A

Q874. What is the difference between a min-heap and a max-heap?

- A. Max-heaps can only store positive numbers
- B. Min-heaps use less memory than max-heaps
- C. Min-heaps are smaller than max-heaps
- D. In a min-heap the smallest element is at the root while in a max-heap the largest element is at the root

Answer: D

Q875. What is the advantage of using a LinkedHashMap over a regular HashMap?

- A. LinkedHashMap is faster for all operations
- B. LinkedHashMap uses less memory than HashMap
- C. HashMap cannot handle string keys
- D. LinkedHashMap maintains insertion order or access order while HashMap does not guarantee order

Answer: D

Q876. What is an immutable collection?

- A. A collection that automatically sorts its elements
- B. A collection that cannot be modified after creation, providing thread safety and predictability
- C. A collection that is stored in read-only memory
- D. A collection that can only hold immutable objects

Answer: B

Q877. What is the difference between a stack-based and a queue-based approach to tree traversal?

- A. A queue produces depth-first while a stack produces breadth-first
- B. A stack traverses faster than a queue
- C. A stack produces depth-first traversal while a queue produces breadth-first traversal
- D. Both produce the same traversal order

Answer: C

Q878. What is the time complexity of searching for an element in a balanced binary search tree?

- A. $O(1)$ because elements are directly accessible by index
- B. $O(n)$ because every node must be checked
- C. $O(n \log n)$ because it requires sorting first
- D. $O(\log n)$ because the search space halves at each step

Answer: D

Q879. What is the `Collections.unmodifiableList()` method used for in Java?

- A. It sorts the list permanently
- B. It converts the list to an array
- C. It returns an unmodifiable view of the list that throws an exception on modification attempts
- D. It removes all elements from the list

Answer: C

Q880. What is the difference between an interface and an abstract class in Java?

- A. They are identical constructs with different keywords
- B. Abstract classes support multiple inheritance while interfaces do not
- C. Interfaces can have constructors while abstract classes cannot
- D. An interface defines method signatures without state while an abstract class can have both abstract and concrete methods with state

Answer: D

Q881. What is the Open/Closed Principle?

- A. Software entities should be open for extension but closed for modification
- B. All methods should be public (open) and fields should be private (closed)
- C. Open source code should not be modified by closed source projects
- D. Classes should be open for use and closed for testing

Answer: A

Q882. What is the purpose of the 'protected' access modifier?

- A. It makes members accessible only within the same class
- B. It allows access within the same class, subclasses, and same package
- C. It prevents any access to the member
- D. It makes members accessible only through getter methods

Answer: B

Q883. What is method hiding in Java?

- A. Making a method invisible to other classes
- B. Encrypting method implementations for security
- C. Removing a method from a class at runtime
- D. When a subclass defines a static method with the same signature as a static method in the superclass

Answer: D

Q884. What is the Single Responsibility Principle?

- A. A program should have only one main class
- B. A class should have only one reason to change, meaning it should have only one responsibility
- C. Each method should have only one parameter
- D. Each class should have only one constructor

Answer: B

Q885. What is the difference between shallow copy and deep copy of an object?

- A. Shallow copy is always slower than deep copy
- B. They produce identical results for all object types
- C. Deep copy only copies primitive fields
- D. Shallow copy copies only field values (sharing referenced objects) while deep copy recursively copies all referenced objects

Answer: D

Q886. What is a sealed class in Kotlin or Java?

- A. A class that cannot have any methods
- B. A class that restricts which classes can extend it, defining a closed set of subclasses
- C. A class that cannot be imported into other packages
- D. A class that seals its data from garbage collection

Answer: B

Q887. What is the Interface Segregation Principle?

- A. All interfaces should be in a separate file
- B. All classes must implement at least one interface
- C. Interfaces should have exactly one method each
- D. Clients should not be forced to depend on methods they do not use

Answer: D

Q888. What is the Dependency Inversion Principle?

- A. Dependencies should always flow from child to parent classes
- B. Methods should be invoked in reverse order of definition
- C. All dependencies should be inverted to remove circular references
- D. High-level modules should depend on abstractions rather than concrete implementations

Answer: D

Q889. What is a data class in Kotlin or a record in Java?

- A. A class that stores only static data
- B. A class that stores data in binary format
- C. A class that automatically generates equals, hashCode, toString, and copy methods based on constructor properties
- D. A class that can only be used for database operations

Answer: C

Q890. What is a memory-mapped file and how does it differ from traditional file I/O?

- A. It maps a file directly into the process address space, allowing file access through memory operations instead of read/write calls
- B. It stores file data in a special map data structure
- C. It maps memory addresses to file names
- D. It compresses files to fit in memory

Answer: A

Q891. What is the difference between strong references and weak references?

- A. Weak references cannot be used with any data type
- B. Strong references prevent garbage collection of the referenced object while weak references allow it
- C. Weak references are faster than strong references
- D. Strong references use more memory per reference

Answer: B

Q892. What is memory alignment and why does it matter?

- A. Aligning data at memory addresses that are multiples of their size for efficient CPU access
- B. Sorting data in memory alphabetically
- C. Aligning code comments to the same column
- D. Matching memory sizes between different programs

Answer: A

Q893. What is the purpose of the finalize() method in Java?

- A. It finalizes the class definition preventing changes
- B. It is called by the garbage collector before reclaiming an object's memory, but its use is discouraged
- C. It permanently saves an object to disk
- D. It marks an object as complete and ready for use

Answer: B

Q894. What is a double-free error?

- A. Allocating twice the required memory
- B. Attempting to free memory that has already been freed, causing undefined behavior
- C. Using two different free functions on the same object
- D. Freeing two different objects at the same time

Answer: B

Q895. What is a phantom reference in Java?

- A. A reference to an object in another program
- B. A reference that allows determining when an object has been garbage collected without preventing collection
- C. An invisible reference that cannot be accessed
- D. A reference to a deleted object

Answer: B

Q896. What is the difference between stack allocation and heap allocation in terms of performance?

- A. Stack allocation is slower because it checks for overflow
- B. They have identical performance characteristics
- C. Heap allocation is always faster than stack allocation
- D. Stack allocation is faster because it only requires moving the stack pointer, while heap allocation involves finding suitable free blocks

Answer: D

Q897. What is the purpose of soft references in Java?

- A. They allow the garbage collector to reclaim objects only when memory is low, useful for caches
- B. They create softly linked objects that cannot be modified
- C. They create references that automatically become null after a timeout
- D. They reference objects stored on soft storage like SSDs

Answer: A

Q898. What is the concept of ownership in Rust?

- A. A design pattern for managing database connections
- B. A system where each value has exactly one variable that owns it, controlling when the value is dropped
- C. A way to assign copyright to code
- D. A security feature that restricts file access

Answer: B

Q899. What is memory compaction in garbage collection?

- A. Reducing the virtual memory size of a process
- B. Moving live objects together to eliminate fragmentation and create contiguous free space
- C. Compressing data to use less memory
- D. Deleting all objects from memory

Answer: B

Q900. What is the difference between try-except-else and try-except-finally in Python?

- A. The else block runs when an exception occurs while finally runs when it does not
- B. The else block runs only when no exception occurs while the finally block runs regardless of whether an exception occurred
- C. They are identical in behavior
- D. The else block handles exceptions while finally throws them

Answer: B

Q901. Why is catching generic exceptions (like catch Exception) considered a bad practice?

- A. Because generic exceptions take more memory than specific ones
- B. Because generic exceptions are slower to catch than specific ones
- C. Because the compiler does not allow generic exception catches
- D. Because it catches all exceptions including ones the code cannot handle, hiding bugs and unexpected errors

Answer: D

Q902. What is exception filtering in C# (when clause)?

- A. Adding a condition to a catch block that must be true for the catch to handle the exception
- B. Filtering out exceptions based on their message text only
- C. Removing exceptions from the program permanently
- D. Blocking certain exception types from being thrown

Answer: A

Q903. What is the purpose of the suppressed exceptions feature in Java?

- A. It attaches secondary exceptions that were suppressed during cleanup in try-with-resources to the primary exception
- B. It reduces the number of exceptions a program can throw
- C. It silently ignores all exceptions
- D. It prevents exceptions from being logged

Answer: A

Q904. What is the difference between assertion errors and exceptions?

- A. Assertions are faster than exceptions
- B. Assertions replace the need for all exception handling
- C. Exceptions cannot be caught while assertions can
- D. Assertions check for programmer errors that should never happen while exceptions handle expected error conditions

Answer: D

Q905. What is the anti-pattern of using exceptions for control flow?

- A. Using too few exception handlers in the program
- B. Throwing exceptions from constructors
- C. Catching exceptions and logging them
- D. Using try-catch blocks for normal program flow instead of conditional checks, harming performance and readability

Answer: D

Q906. What is exception wrapping or translation?

- A. Wrapping exception handling code in a separate class
- B. Catching a low-level exception and throwing a higher-level exception that is more meaningful to the caller
- C. Converting exception messages to different languages
- D. Converting compile-time errors to runtime exceptions

Answer: B

Q907. What is the purpose of the 'finally' block in resource cleanup?

- A. It guarantees that cleanup code (like closing files or connections) executes whether or not an exception occurred
- B. It finalizes the data and sends it to the database
- C. It prevents the exception from propagating further
- D. It runs only when no exception is thrown

Answer: A

Q908. What is the difference between re-throwing and wrapping an exception?

- A. Re-throwing is faster while wrapping provides more information
- B. Re-throwing propagates the same exception unchanged while wrapping creates a new exception with the original as the cause
- C. They are identical operations with different names
- D. Re-throwing creates a new exception while wrapping keeps the same one

Answer: B

Q909. What is the role of the getMessage() method on exceptions?

- A. It returns a human-readable description of the error that caused the exception
- B. It sends an error message to the user via email
- C. It creates a new exception with a custom message
- D. It translates the exception into a log entry

Answer: A

Q910. What is the purpose of a requirements.txt file in Python?

- A. It lists the hardware requirements for the program
- B. It lists the required test cases
- C. It specifies the Python package dependencies and their versions needed for the project
- D. It contains the project documentation

Answer: C

Q911. What is the difference between a bundler and a package manager?

- A. A package manager only works with frontend code
- B. A package manager installs and manages dependencies while a bundler combines source files into optimized bundles for deployment
- C. A bundler installs packages while a package manager bundles code
- D. They are identical tools with different names

Answer: B

Q912. What is dependency resolution in package management?

- A. Removing unnecessary dependencies from the project
- B. Converting dependencies to a different programming language
- C. Resolving conflicts between developers about which packages to use
- D. The process of determining compatible versions of all direct and transitive dependencies

Answer: D

Q913. What is the purpose of namespace imports in programming?

- A. Importing only private functions from a module
- B. Creating namespace aliases for the operating system
- C. Importing all exports from a module under a single namespace to avoid naming conflicts
- D. Creating new names for existing functions

Answer: C

Q914. What is the purpose of the peerDependencies field in package.json?

- A. It lists peer-reviewed packages only
- B. It defines packages used exclusively during peer programming
- C. It specifies packages that should be provided by the consuming project rather than installed as nested dependencies
- D. It lists packages that are optional for the project

Answer: C

Q915. What is lazy loading of modules?

- A. Deferring the loading of a module until it is actually needed, reducing initial startup time
- B. Loading modules slowly to reduce CPU usage
- C. Loading only the documentation of a module
- D. Loading modules without importing them

Answer: A

Q916. What is the difference between static imports and dynamic imports?

- A. Dynamic imports can only load built-in modules
- B. Static imports are resolved at compile time while dynamic imports load modules at runtime based on conditions
- C. Static imports are faster than dynamic imports in all cases
- D. Static imports happen at runtime while dynamic imports happen at compile time

Answer: B

Q917. What is the purpose of a module's public API?

- A. The internal implementation details of the module
- B. The set of exported functions, classes, and types that external code can use
- C. The private methods used for testing
- D. The documentation comments in the module

Answer: B

Q918. What is the npm audit command used for?

- A. Scanning project dependencies for known security vulnerabilities
- B. Checking the financial cost of dependencies
- C. Auditing the code for style violations
- D. Auditing the performance of installed packages

Answer: A

Q919. What is the difference between a workspace and a standalone package?

- A. Standalone packages cannot have dependencies
- B. Workspaces can only contain two packages
- C. A workspace manages multiple related packages in a single repository with shared dependencies while standalone packages are independent
- D. Workspaces are for development only while packages are for production

Answer: C

Q920. What is the Builder pattern and when should it be used?

- A. A pattern for building database tables
- B. A pattern that compiles code faster
- C. A pattern for building user interfaces
- D. A pattern that separates object construction from representation, useful for objects with many optional parameters

Answer: D

Q921. What is the Proxy pattern?

- A. A pattern that provides a surrogate object that controls access to another object
- B. A pattern that hides all methods of an object
- C. A pattern for configuring network proxies
- D. A pattern that copies objects between servers

Answer: A

Q922. What is memoization in the context of dynamic programming?

- A. Caching results of subproblems to avoid redundant computation in recursive algorithms
- B. Storing algorithm pseudocode in comments
- C. Writing notes about algorithm design
- D. Memorizing the source code for the exam

Answer: A

Q923. What is the Chain of Responsibility pattern?

- A. A pattern where a request is passed along a chain of handlers until one handles it
- B. A pattern for linking database tables
- C. A pattern for managing responsibility among team members
- D. A pattern that chains function calls together

Answer: A

Q924. What is the difference between space complexity and time complexity?

- A. They measure the same thing using different units
- B. Space complexity is always larger than time complexity
- C. Space complexity measures memory usage growth while time complexity measures execution time growth relative to input size
- D. Time complexity only applies to recursive algorithms

Answer: C

Q925. What is the Template Method pattern?

- A. A pattern for generating boilerplate code
- B. A pattern that defines the skeleton of an algorithm in a base class, letting subclasses override specific steps
- C. A pattern that templates all method signatures
- D. A pattern for creating HTML templates

Answer: B

Q926. What is the difference between top-down and bottom-up dynamic programming?

- A. They always produce different results
- B. Top-down is always faster than bottom-up
- C. Top-down uses recursion with memoization while bottom-up builds solutions iteratively from smallest subproblems
- D. Bottom-up uses recursion while top-down uses iteration

Answer: C

Q927. What is the Mediator pattern?

- A. A pattern that converts between data formats
- B. A pattern for mediating between the CPU and memory
- C. A pattern for middle-tier server components
- D. A pattern that defines an object to encapsulate how a set of objects interact, reducing direct dependencies

Answer: D

Q928. What is amortized analysis?

- A. A technique that averages the time per operation over a sequence of operations to give a tighter bound than worst-case
- B. Analyzing the amount of memory amortized over program lifetime
- C. Financial analysis of algorithm development costs
- D. Estimating the remaining time for a running algorithm

Answer: A

Q929. What is the difference between a monitor and a semaphore?

- A. A monitor provides mutual exclusion with condition variables as a high-level construct while a semaphore is a lower-level counting mechanism
- B. Monitors are faster than semaphores
- C. They are identical synchronization mechanisms
- D. Semaphores provide mutual exclusion while monitors do not

Answer: A

Q930. What is the purpose of the volatile keyword in concurrent programming?

- A. It ensures that reads and writes to a variable are visible to all threads by preventing CPU caching of the value
- B. It makes a variable faster to access
- C. It allocates the variable on the heap instead of the stack
- D. It prevents a variable from being modified

Answer: A

Q931. What is a barrier in concurrent programming?

- A. A hardware barrier between CPU cores
- B. A security barrier preventing unauthorized thread access
- C. A synchronization point where all threads must arrive before any can proceed past it
- D. A barrier preventing memory leaks

Answer: C

Q932. What is a livelock and how does it differ from a deadlock?

- A. A livelock is a deadlock that eventually resolves itself
- B. In a livelock, threads actively respond to each other but make no progress, unlike deadlock where threads are blocked indefinitely
- C. A deadlock involves active threads while a livelock involves blocked threads
- D. A livelock uses more memory than a deadlock

Answer: B

Q933. What is the purpose of a ReadWriteLock?

- A. It locks reading while allowing concurrent writes
- B. It provides separate locks for different data types
- C. It allows multiple threads to read simultaneously but requires exclusive access for writing
- D. It separates read and write operations into different files

Answer: C

Q934. What is the difference between a Future and a Promise?

- A. A Future is a read-only handle to an asynchronous result while a Promise is the writable side that resolves the value
- B. They are identical concepts in all languages
- C. Futures are faster than Promises
- D. Promises are used only in JavaScript while Futures are used everywhere

Answer: A

Q935. What is the difference between cooperative and preemptive scheduling of threads?

- A. They produce identical scheduling outcomes
- B. Cooperative scheduling requires threads to voluntarily yield control while preemptive scheduling allows the OS to interrupt threads at any time
- C. Preemptive scheduling cannot handle more than four threads
- D. Cooperative scheduling is always more efficient

Answer: B

Q936. What is the purpose of a CountdownLatch?

- A. It counts down to thread termination
- B. It allows one or more threads to wait until a set of operations in other threads completes
- C. It limits the number of threads that can run simultaneously
- D. It counts the number of threads in the system

Answer: B

Q937. What is the actor model of concurrency?

- A. A model where one main actor directs all other threads
- B. A model where independent actors communicate only through asynchronous message passing, avoiding shared state
- C. A model where actors perform computations in a theater-like environment
- D. A model that uses actors to represent database transactions

Answer: B

Q938. What is thread-local storage?

- A. A local database connection for each thread
- B. Storage on the thread's stack only
- C. Storage that is local to the computer running the thread
- D. Memory that is private to each thread, where each thread has its own independent copy of the variable

Answer: D

Q939. What is the difference between a stub and a mock in testing?

- A. A stub provides predetermined responses to method calls while a mock also verifies that expected interactions occurred
- B. They are identical testing concepts
- C. Stubs are faster than mocks
- D. Mocks provide data while stubs verify behavior

Answer: A

Q940. What is boundary value analysis in testing?

- A. Analyzing the boundaries of the codebase
- B. Analyzing code coverage boundaries
- C. Setting boundaries on test execution time
- D. Testing at the edges of input ranges where errors are most likely to occur

Answer: D

Q941. What is the purpose of test doubles?

- A. Double-checking test results manually
- B. Creating two versions of each test
- C. Running each test twice for verification
- D. Objects that stand in for real dependencies during testing, providing isolation and control

Answer: D

Q942. What is the arrange-act-assert (AAA) pattern in testing?

- A. A pattern for asserting code style compliance
- B. A pattern for arranging test files alphabetically
- C. A pattern for testing in three environments
- D. A pattern that structures tests into setup (arrange), execution (act), and verification (assert) phases

Answer: D

Q943. What is snapshot testing?

- A. Testing at a specific snapshot in version control
- B. Capturing the output of a component and comparing it against a stored reference snapshot in future test runs
- C. Testing memory snapshots for leaks
- D. Taking screenshots of the application for visual comparison

Answer: B

Q944. What is the purpose of a code coverage tool?

- A. Measuring how much of the source code is executed by the test suite
- B. Covering up bugs so they are not visible
- C. Covering code with comments for documentation
- D. Protecting code from unauthorized access

Answer: A

Q945. What is the difference between end-to-end testing and integration testing?

- A. They test the same scope at different speeds
- B. Integration tests cover more code than end-to-end tests
- C. End-to-end tests verify the entire application workflow from user perspective while integration tests verify interactions between specific components
- D. End-to-end tests are faster than integration tests

Answer: C

Q946. What is the purpose of conditional breakpoints in a debugger?

- A. Breaking execution conditionally based on CPU temperature
- B. Breaking execution only when a programmer-specified condition is true, avoiding stopping at every iteration
- C. Creating breakpoints that expire after a certain number of hits
- D. Setting breakpoints that only work on certain days

Answer: B

Q947. What is test parameterization?

- A. Setting global parameters for all tests
- B. Adjusting test framework parameters
- C. Parameterizing the test environment
- D. Running the same test logic with different input values and expected results

Answer: D

Q948. What is the difference between a smoke test and a sanity test?

- A. Sanity tests run before smoke tests in the pipeline
- B. They are identical types of tests
- C. Smoke tests are longer than sanity tests
- D. Smoke tests verify that critical functionality works after a build while sanity tests verify that specific bug fixes or features work correctly

Answer: D

Q949. What is the difference between continuous integration and continuous delivery?

- A. They are identical processes with different names
- B. CI works with all languages while CD only works with web applications
- C. CI deploys code while CD tests it
- D. CI automatically builds and tests code changes while CD ensures code is always in a deployable state and can be released at any time

Answer: D

Q950. What is the purpose of a code linter?

- A. Compiling code faster than the standard compiler
- B. Linking code files together into an executable
- C. Removing lint from the physical computer
- D. Analyzing source code to find potential errors, style violations, and suspicious patterns before execution

Answer: D

Q951. What is the difference between a hotfix and a regular release?

- A. A hotfix is an urgent fix for a critical production issue deployed outside the normal release cycle while a regular release follows the planned schedule
- B. Hotfixes are larger than regular releases
- C. Regular releases only contain bug fixes while hotfixes contain features
- D. Hotfixes do not require testing

Answer: A

Q952. What is the purpose of semantic versioning (SemVer)?

- A. Numbering each line of source code
- B. Versioning based on release dates
- C. Communicating the nature of changes through version numbers (MAJOR.MINOR.PATCH) so users know compatibility impact
- D. Assigning version numbers based on team vote

Answer: C

Q953. What is the purpose of a CI/CD pipeline?

- A. A tool for managing database migrations
- B. A physical pipeline for connecting servers
- C. A network protocol for transferring code
- D. An automated workflow that builds, tests, and deploys code changes through defined stages

Answer: D

Q954. What is the difference between blue-green deployment and canary deployment?

- A. Blue-green is faster than canary
- B. They deploy code in the same way but at different times
- C. Blue-green switches all traffic between two identical environments while canary gradually routes a percentage of traffic to the new version
- D. Canary deployment requires two complete environments while blue-green uses one

Answer: C

Q955. What is the purpose of a code formatter like Prettier or Black?

- A. Making code execute faster
- B. Converting code between programming languages
- C. Formatting code for printing on paper
- D. Automatically formatting code to follow consistent style rules, eliminating style debates in code reviews

Answer: D

Q956. What is the purpose of a pre-commit hook?

- A. Creating a backup before each commit
- B. Sending a notification before each commit
- C. Running automated checks (like linting, tests, or formatting) before a git commit is created, preventing bad code from entering the repository
- D. Committing code before anyone else

Answer: C

Q957. What is the difference between a monolith and microservices architecture?

- A. Monoliths are always slower than microservices
- B. A monolith deploys the entire application as a single unit while microservices decompose it into independently deployable services
- C. Monoliths do not support any form of modularity
- D. Microservices cannot share databases

Answer: B

Q958. What is the difference between eager evaluation and lazy evaluation of function arguments?

- A. Lazy evaluation computes all arguments immediately
- B. Eager evaluation is always slower than lazy evaluation
- C. Eager evaluation only works with pure functions
- D. Eager evaluation computes arguments before passing them while lazy evaluation delays computation until the value is actually needed

Answer: D

Q959. What is the difference between stable and unstable sorting algorithms?

- A. Stable sorts only work on numeric data
- B. Unstable sorts produce incorrect results for duplicate values
- C. Stable sorts are always faster than unstable sorts
- D. Stable sorts preserve the relative order of equal elements while unstable sorts may rearrange them

Answer: D

Q960. What is the purpose of a definition of done (DoD) in agile development?

- A. A list of developers assigned to each task
- B. Defining when the project is completely finished
- C. A shared checklist of criteria that must be met for a work item to be considered complete
- D. A document that defines the final release date

Answer: C

Programming (C++, Java, Python) — Hard

480 questions

Q961. What is the halting problem in computer science?

- A. Optimizing a program execution speed through advanced compiler techniques
- B. Determining whether an arbitrary program will finish running or loop forever
- C. Finding and reporting memory leaks occurring during a program execution
- D. Determining whether a program contains syntax errors in its source code

Answer: B

Q962. In language theory, what is a context-free grammar (CFG)?

- A. A grammar where production rules apply regardless of context
- B. A grammar equivalent to standard regular expressions
- C. A grammar where rules depend on surrounding token context
- D. A grammar that is exclusively used in natural languages

Answer: A

Q963. What is tail call optimization (TCO)?

- A. Optimizing the first function call in a chain
- B. Parallelizing multiple function calls at once
- C. Reusing the current stack frame for a tail call
- D. Removing unused function calls from the code

Answer: C

Q964. What distinguishes a pure function in functional programming?

- A. It returns the same output for the same input always
- B. It cannot accept any parameters from the caller
- C. It modifies global state and external variables freely
- D. It must use recursion as its primary control flow

Answer: A

Q965. What is referential transparency?

- A. A function that returns void and produces no output value
- B. An expression replaceable with its value without side effects
- C. A pointer that references itself in a circular data structure
- D. A variable that can be referenced from anywhere in the code

Answer: B

Q966. What is the difference between static and dynamic dispatch?

- A. Static dispatch is always slower than dynamic dispatch in practice
- B. Static dispatch resolves methods at compile time; dynamic at runtime
- C. They are identical concepts in object-oriented programming languages
- D. Dynamic dispatch is only available in the C programming language

Answer: B

Q967. What is a closure in programming?

- A. A method used to close open file handles
- B. A function bundled with its lexical scope
- C. A technique to terminate a running loop
- D. A destructor that cleans up class state

Answer: B

Q968. What is continuation-passing style (CPS)?

- A. A structured method for handling runtime errors
- B. Passing control explicitly via continuation functions
- C. A design pattern for building user interfaces
- D. Passing all arguments by reference to functions

Answer: B

Q969. What is homoiconicity in a programming language?

- A. All variables in the program share the same type
- B. The code can run on any platform without changes
- C. Code structure uses the language own data structures
- D. Functions return the exact same type they receive

Answer: C

Q970. What is the Curry-Howard correspondence?

- A. A technique for normalizing database tables
- B. A method for performing function currying
- C. A link between mathematical proofs and programs
- D. A protocol for networking between computers

Answer: C

Q971. What is the IEEE 754 standard?

- A. A specification for database schemas
- B. A standard for security encryption
- C. A protocol for computer networking
- D. A standard for floating-point numbers

Answer: D

Q972. What is a phantom type?

- A. A type providing compile-time safety without storing data
- B. A type that has been deprecated from the language standard
- C. A type reserved for representing ghost or phantom objects
- D. A type that only exists at runtime and not at compile time

Answer: A

Q973. What is the difference between structural typing and nominal typing?

- A. They are exactly the same concept with different names
- B. Structural typing is only applicable to struct data types
- C. Nominal typing is only applicable to numeric data types
- D. Structural checks compatibility by shape; nominal by name

Answer: D

Q974. What is type erasure in Java generics?

- A. Removing type information at the compile stage
- B. Deleting all variables of a particular given type
- C. Generic type info removed at runtime as raw types
- D. Converting every type to Object at compile time

Answer: C

Q975. What is a dependent type?

- A. A type that inherits from another
- B. A type that depends on a library
- C. A type that is defined by a value
- D. A type that requires garbage collection

Answer: C

Q976. What is the bottom type in type theory?

- A. A type with no values and subtype of every type
- B. The null type representing an absence of a value
- C. The base class that all other types inherit from
- D. The standard integer type used for whole numbers

Answer: A

Q977. What is variance in type systems?

- A. Random and unpredictable behavior in type inference
- B. The size variation of different data types in memory
- C. How subtyping of complex types relates to components
- D. The total range of values a type can hold

Answer: C

Q978. What is an algebraic data type (ADT)?

- A. An array type that exclusively stores numeric values
- B. A data type designed for mathematical operations only
- C. A type used for representing algebraic math expressions
- D. A composite type combining sum and product operations

Answer: D

Q979. What is the unit type?

- A. A type with exactly one possible value
- B. A type used for measurement units
- C. The integer constant value of one
- D. A class with a single instance only

Answer: A

Q980. What is the purpose of type witnesses in advanced type systems?

- A. To provide explicit proof of type relationships
- B. To help debug complex type errors in code
- C. To convert types dynamically during execution
- D. To observe and log type changes at runtime

Answer: A

Q981. What is operator overloading?

- A. Defining custom behavior for user-defined types
- B. An error caused by incorrect operator usage
- C. Increasing the default operator precedence level
- D. Using too many operators in one expression

Answer: A

Q982. What is the result of ~5 in a system using 32-bit signed integers?

- A. 5
- B. -5
- C. 6
- D. -6

Answer: D

Q983. What is the spaceship operator (<=>)?

- A. An assignment operator for setting variables
- B. A comparison operator returning -1, 0, or 1
- C. A bitwise shift operator for binary values
- D. A logical operator for boolean expressions

Answer: B

Q984. What is the difference between arithmetic right shift (>>) and logical right shift (>>>)?

- A. There is no difference between them at all
- B. >>> preserves the sign bit; >> fills with zeros
- C. >> preserves the sign bit; >>> fills with zeros
- D. >> always fills the vacant bit positions with zeros

Answer: C

Q985. What is the Elvis operator (?:) in Kotlin?

- A. Returns left if non-null, otherwise right
- B. A type-checking operator for null safety
- C. A standard ternary conditional operator
- D. A bitwise operator for binary operations

Answer: A

Q986. How does Python's 'is' operator differ from '=='?

- A. is runs faster but produces less accurate results
- B. is checks identity; == checks value equality
- C. They are functionally identical in all cases
- D. == checks identity while is checks equality

Answer: B

Q987. What is the safe navigation operator (?.) in Kotlin/TypeScript?

- A. It forces a null pointer exception to be thrown immediately
- B. It converts any null value to the integer zero automatically
- C. It always returns a null value regardless of the receiver
- D. It accesses a member only if the receiver is non-null

Answer: D

Q988. What does the >>> operator do in Java that >> does not?

- A. It fills leftmost bits with 0, ignoring the sign
- B. It operates more slowly than the >> operator
- C. It works on floating-point numbers unlike >>
- D. Nothing different from the >> operator at all

Answer: A

Q989. What is destructuring assignment?

- A. A type of memory deallocation for objects
- B. Deleting variables from the current scope
- C. Unpacking values from arrays or objects into variables
- D. Assigning null to multiple variables at once

Answer: C

Q990. What is the pipeline operator (|>) proposed for JavaScript?

- A. A bitwise OR operation on binary values
- B. A comparison operator for ordering two values
- C. It passes left result as argument to right function
- D. A type conversion operator between data types

Answer: C

Q991. What is tail recursion and how does it relate to loops?

- A. Recursion that uses a tail pointer data structure
- B. Recursion without any base case for termination
- C. A loop structure that invokes itself recursively
- D. A tail-position recursive call optimized into a loop

Answer: D

Q992. What is the difference between structured and unstructured control flow?

- A. Structured control flow is slower than unstructured control
- B. Unstructured control flow is more readable than structured
- C. There is no meaningful difference between the two approaches
- D. Structured uses blocks like if/while; unstructured uses goto

Answer: D

Q993. What is pattern matching as a control structure?

- A. Matching strings with regular expressions
- B. A generalized switch that destructures data types
- C. A repeating pattern within a loop construct
- D. A text search algorithm for finding substrings

Answer: B

Q994. What is the difference between eager and lazy evaluation in control flow?

- A. Eager evaluates expressions immediately; lazy defers until value needed
- B. Lazy evaluation is always faster than eager evaluation in all cases
- C. Eager evaluation is only used in purely functional language paradigms
- D. There is no practical difference between the two evaluation strategies

Answer: A

Q995. What is a coroutine and how does it differ from a subroutine?

- A. A coroutine is a specialized type of loop iteration control structure
- B. A subroutine can be paused and resumed at arbitrary execution points
- C. They are the same thing with different names in different languages
- D. A coroutine can suspend and resume; a subroutine runs to completion

Answer: D

Q996. What is the purpose of guard clauses?

- A. To lock and protect shared resources from concurrent thread access
- B. To guard against invalid memory allocation during program startup
- C. To protect systems against unauthorized network access attempts
- D. To handle edge cases early by returning at the start of functions

Answer: D

Q997. What is a trampoline in the context of recursion?

- A. A technique converting recursion into iterative loops
- B. A specialized debugging tool for stack overflow
- C. A physical device used for hardware testing
- D. A tree-based data structure for fast lookup

Answer: A

Q998. What is continuation in control flow?

- A. A loop continuation condition for iterations
- B. An alternative to the break statement in loops
- C. The next sequential statement in the code
- D. An abstract representation of remaining computation

Answer: D

Q999. What is the difference between deterministic and non-deterministic control flow?

- A. There is no practical difference between them at all
- B. Deterministic follows a fixed path; non-deterministic explores many
- C. Non-deterministic control flow always causes the program to crash
- D. Deterministic control flow always executes more slowly than others

Answer: B

Q1000. What is the 'goto considered harmful' argument?

- A. Goto should be used everywhere as the primary control flow method
- B. Goto causes memory leaks and resource allocation failures in code
- C. Dijkstra argued goto makes programs hard to follow and reason about
- D. Goto is always the most efficient control flow mechanism available

Answer: C

Q1001. What is function currying?

- A. Overloading a function with additional parameter types
- B. Caching the results of previous function calls
- C. Optimizing a function for better performance
- D. Transforming multi-argument function into single-argument chain

Answer: D

Q1002. What is memoization?

- A. Writing descriptive comments inside function body code
- B. Caching expensive function results to avoid recomputation
- C. Saving a function definition to persistent disk storage
- D. Allocating additional memory for function local variables

Answer: B

Q1003. What is the difference between early binding and late binding?

- A. Early binding resolves at compile time; late binding at runtime
- B. There is no meaningful difference between the two approaches
- C. Early binding is always slower than late binding in all cases
- D. Late binding only works in the C programming language alone

Answer: A

Q1004. What is a variadic function?

- A. A function that has a variable return data type
- B. A function that returns multiple values at once
- C. A function that accepts no input parameters
- D. A function accepting variable number of arguments

Answer: D

Q1005. What is partial application?

- A. Fixing some arguments to produce a new function taking the rest
- B. Calling a function with completely wrong argument types or values
- C. A function that only partially works and fails on some inputs
- D. An incomplete function definition that is missing its body code

Answer: A

Q1006. What is a thunk?

- A. A runtime error in the application
- B. A wrapper delaying computation evaluation
- C. A pre-allocated block of heap memory
- D. A particular type of variable binding

Answer: B

Q1007. What are first-class functions?

- A. Functions assignable to variables and passable as arguments
- B. Functions that have the highest execution priority level
- C. Functions that belong to the main class of the program
- D. Functions that are defined first in source code order

Answer: A

Q1008. What is the difference between a method and a function?

- A. Methods cannot return any values while functions always can
- B. Functions cannot accept any parameters while methods always can
- C. They are always the same thing in every programming language
- D. A method belongs to an object or class; a function is standalone

Answer: D

Q1009. What is a generator function?

- A. A function that creates and returns other functions
- B. A function that yields multiple values over time lazily
- C. A function that generates random numbers on demand
- D. A constructor that initializes a new class instance

Answer: B

Q1010. What is the Y combinator?

- A. A fixed-point combinator enabling anonymous recursion
- B. A popular comparison-based sorting algorithm
- C. A well-known startup incubator company
- D. A commonly used structural design pattern

Answer: A

Q1011. What is memory-mapped I/O?

- A. Mapping files into virtual memory address space
- B. Encrypting all I/O data during transfer
- C. Storing I/O data in the processor cache
- D. I/O that uses only RAM for data storage

Answer: A

Q1012. What is the difference between blocking and non-blocking I/O?

- A. Non-blocking I/O is always slower than blocking I/O overall
- B. There is no meaningful difference between the two modes
- C. Blocking I/O never waits for the operation to finish first
- D. Blocking waits until completion; non-blocking returns immediately

Answer: D

Q1013. What is asynchronous I/O?

- A. I/O that continues without waiting, notified on completion
- B. I/O that is identical to synchronous blocking I/O
- C. I/O operations that execute in strict sequential order
- D. I/O operations that are always extremely fast

Answer: A

Q1014. What is the purpose of the select() system call?

- A. Selecting a record from a database table query
- B. Selecting files to delete from the file system
- C. Monitoring file descriptors for I/O readiness
- D. Choosing which thread the scheduler should run

Answer: C

Q1015. What is zero-copy I/O?

- A. I/O operations that have zero financial cost
- B. I/O operations that transfer no data at all
- C. Deleting data as it is being transferred out
- D. Transferring data without copying between buffers

Answer: D

Q1016. What is the reactor pattern in I/O?

- A. A nuclear-physics computing simulation pattern
- B. An event-driven dispatcher that demultiplexes I/O events
- C. A chemical reaction simulation programming model
- D. A simple loop-based sequential processing pattern

Answer: B

Q1017. What are I/O multiplexing mechanisms like epoll and kqueue?

- A. Command-line utilities designed for copying files efficiently
- B. Multi-threaded I/O libraries for parallel data processing
- C. Mechanisms that duplicate I/O devices for redundancy
- D. Scalable kernel interfaces for monitoring many file descriptors

Answer: D

Q1018. What is the difference between character devices and block devices?

- A. Character devices transfer byte by byte; block in fixed chunks
- B. Block devices are exclusively used for text file operations
- C. There is no functional difference between device types
- D. Character devices are always faster than block-based devices

Answer: A

Q1019. What is Direct Memory Access (DMA) in the context of I/O?

- A. Hardware allowing device-to-memory transfer without CPU
- B. A software programming technique for data transfer
- C. A debugging method for diagnosing hardware I/O issues
- D. A specialized type of dynamic memory allocation method

Answer: A

Q1020. What is scatter-gather I/O (vectored I/O)?

- A. Reading/writing multiple non-contiguous buffers at once
- B. Distributing I/O operations across network nodes
- C. Random I/O patterns across scattered disk sectors
- D. Fragmented file storage across multiple disk drives

Answer: A

Q1021. What is the time complexity of naive string matching?

- A. $O(n*m)$
- B. $O(\log n)$
- C. $O(n)$
- D. $O(n^2)$

Answer: A

Q1022. What is the KMP (Knuth-Morris-Pratt) string matching algorithm?

- A. A comparison-based sorting algorithm for arrays
- B. A tree traversal algorithm for binary search trees
- C. A linear-time string matcher using a failure function
- D. A hashing algorithm for generating unique hash codes

Answer: C

Q1023. What is string interning?

- A. Compressing strings to reduce their memory size
- B. Storing one copy of each distinct string value
- C. Encrypting strings for secure data transmission
- D. Storing strings in a temporary memory buffer

Answer: B

Q1024. What is the Rabin-Karp algorithm?

- A. A graph traversal algorithm for shortest paths
- B. A string matcher using hashing for efficient search
- C. A comparison-based sorting algorithm for collections
- D. A compression algorithm for reducing data file sizes

Answer: B

Q1025. What is a trie (prefix tree) and how does it relate to strings?

- A. A linked list that stores strings in sequential memory order
- B. A standard binary search tree data structure
- C. A hash table that maps string keys to their computed values
- D. A tree where each node represents a character for prefix search

Answer: D

Q1026. What is the difference between UTF-8, UTF-16, and UTF-32?

- A. UTF-32 uses the least amount of storage space per character overall
- B. They are all identical encoding formats with no differences at all
- C. UTF-8 uses 1-4 bytes, UTF-16 uses 2-4 bytes, UTF-32 uses exactly 4
- D. UTF-8 is a fixed-width encoding that always uses exactly one byte

Answer: C

Q1027. What is a rope data structure for strings?

- A. A contiguous array of individual character values
- B. A hash map that stores character frequency counts
- C. A singly linked list of character nodes
- D. A balanced tree for efficient long string operations

Answer: D

Q1028. What is the Levenshtein distance?

- A. The total length of a string in characters
- B. The count of matching characters between strings
- C. The byte size of a string in memory allocation
- D. The minimum single-character edits between strings

Answer: D

Q1029. What is the Boyer-Moore string search algorithm's key insight?

- A. Start comparing from the beginning of the pattern each time
- B. Use hashing to compare the pattern against every position
- C. Compare every single character without skipping any at all
- D. Start from pattern end and use heuristics to skip sections

Answer: D

Q1030. What are surrogate pairs in UTF-16?

- A. Duplicate characters stored in the encoding table
- B. Error codes returned by the UTF-16 encoding process
- C. Pairs of 16-bit units for characters beyond the BMP
- D. Paired delimiter characters used in string literals

Answer: C

Q1031. What is a bloom filter?

- A. A comparison-based sorting algorithm for arrays
- B. A space-efficient probabilistic set membership tester
- C. An image processing filter for visual enhancement
- D. A standard hash table with separate chaining buckets

Answer: B

Q1032. What is a skip list?

- A. A filtered array that excludes elements matching a condition
- B. A probabilistic multi-layer linked list for $O(\log n)$ search
- C. A list that skips over certain elements during iteration
- D. A type of priority queue that maintains sorted element order

Answer: B

Q1033. What is a concurrent hash map?

- A. A hash map using multiple hash functions per key
- B. A thread-safe hash map without full map locking
- C. A distributed database across network nodes only
- D. A regular hash map with no special properties

Answer: B

Q1034. What is the amortized time complexity of adding an element to a dynamic array?

- A. $O(1)$ amortized
- B. Always $O(n)$
- C. $O(\log n)$ each
- D. $O(n \log n)$ each

Answer: A

Q1035. What is a persistent data structure?

- A. An immutable structure that can never be accessed again
- B. A structure preserving previous versions on modification
- C. A relational database schema structure for queries
- D. A structure saved to persistent disk storage

Answer: B

Q1036. What is a B-tree and where is it commonly used?

- A. A self-balancing tree with multiple keys per node for databases
- B. A hash table variant that uses tree-based collision resolution
- C. A standard binary tree with two children per node
- D. A linked list type that organizes data in sequential node chains

Answer: A

Q1037. What is consistent hashing?

- A. A type of encryption used for securing sensitive data values
- B. A collision resolution method for hash table bucket overflow
- C. A hashing scheme that minimally redistributes keys on change
- D. Regular hashing with standard hash function methods

Answer: C

Q1038. What is a lock-free data structure?

- A. An unprotected array that any thread can freely modify
- B. A concurrent structure using atomic operations not locks
- C. A structure with no synchronization protection at all
- D. A read-only data structure that cannot be changed ever

Answer: B

Q1039. What is the difference between external and internal iterators?

- A. Internal iterators always consume more memory than external iterators
- B. External iterators are always faster than internal iterators in practice
- C. There is no meaningful difference between the two iterator types
- D. External iterators are client-controlled; internal are collection-controlled

Answer: D

Q1040. What is a disjoint set (union-find) data structure?

- A. A structure tracking element partitions with union and find
- B. A sorted set that maintains elements in order
- C. A hash set variant that uses separate chaining for entries
- D. A set that has no intersections with any other set at all

Answer: A

Q1041. What is the SOLID principle in OOP?

- A. A coding style guide for naming conventions in code
- B. A data structure for organizing elements efficiently
- C. Five design principles: SRP, OCP, LSP, ISP, and DIP
- D. A testing framework for validating software behavior

Answer: C

Q1042. What is a virtual method table (vtable)?

- A. A hash table for caching method results quickly
- B. A debug table for tracking method execution and logging
- C. A lookup table for resolving dynamic method calls at runtime
- D. A method list defined in the source code of the program

Answer: C

Q1043. What is the difference between association, aggregation, and composition?

- A. Association is general; aggregation weak; composition strong ownership
- B. They are all identical concepts with just different names used
- C. Composition is always the better design choice in every single case
- D. Only naming differences exist with no behavioral distinctions at all

Answer: A

Q1044. What is a mixin?

- A. A class that mixes different data types together randomly
- B. A class providing methods to others without being standalone
- C. A type of interface that defines a strict method contract
- D. A structural design pattern for building complex objects

Answer: B

Q1045. What is double dispatch and how does it differ from single dispatch?

- A. Single dispatch uses two methods while double dispatch uses only one
- B. There is no practical difference between the two dispatch mechanisms
- C. Single dispatch selects by one type; double dispatch by two types
- D. Double dispatch is twice as fast in execution as single dispatch

Answer: C

Q1046. What is metaclass programming?

- A. Creating graphical user interfaces programmatically
- B. Writing classes whose instances are themselves classes
- C. A commonly used structural design pattern for code
- D. Programming about programs in general terms

Answer: B

Q1047. What is the difference between prototype-based and class-based OOP?

- A. Prototype-based OOP is always slower than class-based OOP overall
- B. Class-based uses blueprints; prototype-based clones existing objects
- C. They are the same approach with no meaningful differences at all
- D. Class-based OOP cannot support any form of inheritance at all

Answer: B

Q1048. What is the fragile base class problem?

- A. A compilation error from missing method implementations
- B. A memory leak caused by improper object disposal code
- C. A class that is too small to be useful on its own
- D. Base class changes can break derived classes unexpectedly

Answer: D

Q1049. What is covariant return type?

- A. A generic type parameter used in collection class definitions
- B. Allowing an overriding method to return a more specific subtype
- C. Always returning the exact same type as the parent declares
- D. A type error that occurs when return types do not match up

Answer: B

Q1050. What is the object slicing problem in C++?

- A. Splitting an object into smaller component parts for storage
- B. When assigning derived to base variable, derived data is lost
- C. A memory allocation issue caused by fragmented heap storage
- D. A threading problem from concurrent access to shared objects

Answer: B

Q1051. What is the mark-and-sweep garbage collection algorithm?

- A. Marking files for deletion from the disk system
- B. A disk cleanup algorithm for removing temporary files
- C. A memory allocation method for new object creation
- D. A GC that marks reachable objects then frees unmarked

Answer: D

Q1052. What is a generational garbage collector?

- A. A manual memory manager without automatic collection
- B. A garbage collector for old legacy programs only
- C. A GC dividing objects by age, collecting young often
- D. A reference counting garbage collection implementation

Answer: C

Q1053. What is the difference between `unique_ptr`, `shared_ptr`, and `weak_ptr` in C++?

- A. `unique_ptr` is exclusive; `shared_ptr` shares; `weak_ptr` observes only
- B. They are all the same smart pointer type with no differences
- C. `shared_ptr` cannot be copied or assigned to another `shared_ptr` var
- D. `unique_ptr` is always slower than the other two smart pointer types

Answer: A

Q1054. What is memory-mapped files and how do they work?

- A. Files mapped into virtual address space for memory I/O
- B. Files that have been encrypted for secure data storage
- C. Files that are stored only in volatile RAM memory
- D. Files that have been compressed to save disk space used

Answer: A

Q1055. What is a slab allocator?

- A. A garbage collector that runs periodically in background
- B. A simple wrapper around the standard `malloc` function call
- C. A disk allocator that manages file system block assignments
- D. A mechanism pre-allocating and caching object-sized blocks

Answer: D

Q1056. What is escape analysis in JVM?

- A. Compiler analysis determining if object references escape scope
- B. Network traffic analysis for monitoring data flow patterns
- C. Analyzing program exit codes and termination conditions
- D. Exception analysis for tracking error propagation in code

Answer: A

Q1057. What is the ABA problem in lock-free memory management?

- A. When a location changes A to B to A, making CAS incorrectly succeed
- B. A problem related to the alphabet ordering of values
- C. A variant of buffer overflow caused by writing past array boundaries
- D. A type of memory leak caused by circular reference chains in objects

Answer: A

Q1058. What is a write barrier in garbage collection?

- A. A mechanism intercepting pointer writes for GC invariants
- B. A memory protection scheme against buffer overflows
- C. A disk write cache for improving I/O performance
- D. A network firewall for filtering traffic

Answer: A

Q1059. What is Copy-on-Write (COW)?

- A. Sharing memory until a write triggers actual copying
- B. Copying files manually between directories on disk
- C. A garbage collection algorithm for reclaiming memory
- D. An error handling technique for runtime exceptions

Answer: A

Q1060. What is address space layout randomization (ASLR)?

- A. Randomizing process memory addresses to prevent attacks
- B. A memory allocation strategy for dynamic array growth
- C. An addressing mode for accessing hardware registers
- D. A sorting algorithm for ordering memory addresses

Answer: A

Q1061. What is the difference between fault tolerance and exception handling?

- A. Exception handling prevents all errors from ever occurring in the program
- B. Exception handling manages code errors; fault tolerance is system resilience
- C. They are exactly the same concept with no practical differences at all
- D. Fault tolerance is always faster and easier to implement than exceptions

Answer: B

Q1062. What is an exception safety guarantee?

- A. A guarantee that no exceptions will ever occur in the program
- B. Levels of state guarantees after exception: no-throw, strong, basic
- C. A compiler feature that eliminates all runtime exceptions automatically
- D. A testing tool that verifies exception handlers work correctly always

Answer: B

Q1063. What is the 'catch-all' anti-pattern?

- A. A testing pattern for verifying exception handler correctness
- B. Catching all exceptions generically which hides bugs in code
- C. Catching only specific well-defined exceptions in the code
- D. A recommended design pattern for robust exception handling

Answer: B

Q1064. What is the difference between recoverable and unrecoverable errors?

- A. There is no meaningful difference between the two error categories
- B. All errors in every program are always recoverable with proper code
- C. All errors in every program are always unrecoverable without fixes
- D. Recoverable errors can be handled; unrecoverable require termination

Answer: D

Q1065. What is Result/Either type as an alternative to exceptions?

- A. A logging mechanism for recording exception messages
- B. A type representing either success value or error value
- C. A direct replacement for try-catch in all situations
- D. A simple boolean type with true or false values

Answer: B

Q1066. What is the problem with exceptions and concurrency?

- A. Exceptions in one thread may not be visible in another thread
- B. No problems exist between exceptions and concurrency at all
- C. Concurrency completely prevents all exceptions from occurring
- D. Exceptions are inherently thread-safe in all implementations

Answer: A

Q1067. What is a panic in Rust/Go?

- A. A regular exception like any other in the language
- B. A warning message that can be safely ignored by developers
- C. An unrecoverable error that unwinds the stack and terminates
- D. A debug log entry recorded during program execution review

Answer: C

Q1068. What is structured concurrency's approach to exception handling?

- A. Ignoring all exceptions thrown by child tasks completely
- B. Propagating child exceptions to parent and cancelling siblings
- C. Logging exceptions only without any propagation to parents
- D. Catching all exceptions silently without any further action

Answer: B

Q1069. What is the difference between resumable and non-resumable exceptions?

- A. Resumable allows continuing at throw point; non-resumable unwinds
- B. There is no meaningful difference between the two exception types
- C. Non-resumable exceptions can be resumed after the handler finishes
- D. Resumable exceptions are always slower than non-resumable exceptions

Answer: A

Q1070. What is the 'exception specification' problem in C++?

- A. Exception specifications are completely reliable and always enforced
- B. They always improve performance of exception handling in all cases
- C. They are required by the C++ standard for every function definition
- D. C++ throw() specs were problematic and replaced by noexcept keyword

Answer: D

Q1071. What is dependency injection?

- A. Installing package dependencies from the package registry
- B. A package manager for downloading and installing third-party libs
- C. Providing dependencies from outside rather than creating internally
- D. An import statement for including external module code in file

Answer: C

Q1072. What is the OSGi framework?

- A. A build tool for compiling and packaging Java applications
- B. A dynamic Java module system supporting runtime bundle mgmt
- C. An operating system for mobile devices and desktops
- D. A testing framework for writing and running automated tests

Answer: B

Q1073. What is the difference between CommonJS and ES Modules?

- A. They are identical module systems with no meaningful differences
- B. ES Modules do not support any form of async module loading
- C. CommonJS uses require synchronously; ESM uses import/export
- D. CommonJS is the newer standard developed after ES Modules

Answer: C

Q1074. What is the diamond dependency problem in package management?

- A. Two dependencies requiring different versions of the same package
- B. A class inheritance issue in object-oriented programming code
- C. A circular dependency between two modules in the same project
- D. A build error caused by missing configuration files in project

Answer: A

Q1075. What is a plugin architecture?

- A. A fixed codebase that cannot be extended or modified at all
- B. A build system for compiling and packaging application binaries
- C. A testing framework for automated integration and unit testing
- D. A design allowing runtime extension via external loaded modules

Answer: D

Q1076. What is the Inversion of Control (IoC) container?

- A. A framework component managing object creation and dependencies
- B. A data structure for storing elements in a specific order only
- C. A reversed loop that iterates in backward order through data
- D. A version control system for tracking changes to source code

Answer: A

Q1077. What is ABI (Application Binary Interface) compatibility?

- A. A network protocol for communication between distributed services
- B. The same as API compatibility with no additional constraints
- C. Binary-level interface ensuring compiled modules interact correctly
- D. A data format for serializing objects to disk or network transfer

Answer: C

Q1078. What is the Service Provider Interface (SPI) pattern?

- A. A web service for hosting API endpoints on the internet
- B. A physical network interface for connecting hardware devices
- C. A design where providers implement a service-defined interface
- D. A build tool for automating compilation and deployment tasks

Answer: C

Q1079. What is reproducible builds and why is it important for packages?

- A. Building projects as fast as possible to save developer time
- B. Building multiple modules in parallel across available cores
- C. Building the same project twice in a row for verification
- D. Ensuring same source always produces bit-identical output binary

Answer: D

Q1080. What is the difference between peer dependencies and regular dependencies?

- A. Regular dependencies are always shared between all project packages
- B. There is no meaningful difference between the two dependency types
- C. Peer dependencies require the consumer to provide a shared version
- D. Peer dependencies are completely optional and never need installing

Answer: C

Q1081. What is the Visitor pattern and when would you use it?

- A. Defining new operations on objects without changing their classes
- B. An authentication pattern for verifying user identity securely
- C. A logging pattern for recording all application events to files
- D. Tracking website visitors for analytics and usage metrics

Answer: A

Q1082. What is a monad in functional programming?

- A. A single immutable value in the program code
- B. A pattern wrapping values for chaining operations with effects
- C. A primitive data type for storing basic values in memory
- D. A loop construct for iterating through collection elements

Answer: B

Q1083. What is the difference between a microkernel and a monolithic architecture?

- A. Microkernel architecture is always larger in size than monolithic
- B. Monolithic architecture is always more modular than microkernel
- C. Microkernel has minimal core with plugins; monolithic is all-in-one
- D. There is no meaningful difference between the two architecture styles

Answer: C

Q1084. What is reactive programming?

- A. Procedural programming with sequential statement execution
- B. A paradigm focused on async data streams and change propagation
- C. Programming with user interface reactions and animations
- D. Event-driven programming using callbacks and event listeners

Answer: B

Q1085. What is the Command Query Responsibility Segregation (CQRS) pattern?

- A. An authentication method for verifying user credentials
- B. A database query for selecting records from tables
- C. A sorting algorithm for ordering elements in collections
- D. Separating read and write operations into different models

Answer: D

Q1086. What is a domain-specific language (DSL)?

- A. A general-purpose programming language for all domains
- B. A scripting language for automating repetitive tasks
- C. A specialized language designed for a particular domain
- D. An assembly language for low-level hardware programming

Answer: C

Q1087. What is the actor model?

- A. Independent units communicating via asynchronous message passing
- B. A relational database model for organizing table schemas
- C. A user interface pattern for laying out visual components
- D. A design pattern for video game character development

Answer: A

Q1088. What is aspect-oriented programming (AOP)?

- A. A subtype of object-oriented programming methodology
- B. Separating cross-cutting concerns from business logic via aspects
- C. A user interface pattern for organizing visual layout components
- D. A testing approach for validating application behavior thoroughly

Answer: B

Q1089. What is the hexagonal architecture (ports and adapters)?

- A. A network topology for connecting distributed system nodes
- B. A six-sided geometric design for user interfaces
- C. Isolating core domain from externals through ports and adapters
- D. A database design for organizing tables and relationships

Answer: C

Q1090. What is algebraic effects as a programming concept?

- A. Expressing side effects as values handled by effect handlers separately
- B. An optimization technique for improving numerical computation speed
- C. Mathematical operations performed on algebraic expressions
- D. A specialized data structure for storing algebraic computation results

Answer: A

Q1091. What is the difference between optimistic and pessimistic concurrency control?

- A. There is no meaningful difference between the two control approaches
- B. Pessimistic concurrency control never uses any locks on resources
- C. Optimistic concurrency control uses locks to prevent all conflicts
- D. Optimistic assumes no conflicts and validates; pessimistic locks first

Answer: D

Q1092. What is the Compare-and-Swap (CAS) operation?

- A. A file swap operation for exchanging file contents
- B. A string operation for replacing character subsequences
- C. A sorting operation for ordering elements in arrays
- D. An atomic operation comparing and swapping memory values

Answer: D

Q1093. What is the Java Memory Model (JMM)?

- A. A memory leak detector for finding unreleased allocations
- B. A garbage collection model for reclaiming unused object memory
- C. A specification defining thread memory interaction and visibility
- D. How Java manages memory allocation for objects on the heap

Answer: C

Q1094. What is work stealing in thread pool implementations?

- A. A security vulnerability where threads access unauthorized data
- B. Interrupting running threads to reassign their execution priority
- C. Idle threads stealing tasks from busy threads for load balancing
- D. Taking another thread source code and executing it directly

Answer: C

Q1095. What is a memory barrier/fence?

- A. A hardware instruction enforcing memory operation ordering for CPUs
- B. A memory protection boundary between user and kernel space
- C. A memory limit restricting maximum allocation size for processes
- D. A garbage collection trigger that starts when memory is low

Answer: A

Q1096. What is the happens-before relationship in concurrency?

- A. A scheduling algorithm for determining thread execution priority
- B. Simple chronological ordering of events by their timestamps
- C. A debugging concept for tracing the order of function calls
- D. A guarantee that writes by one action are visible to another

Answer: D

Q1097. What is a read-write lock (ReentrantReadWriteLock)?

- A. A file lock for preventing concurrent file access
- B. A database lock for serializing transaction execution
- C. A lock allowing multiple readers but exclusive writers
- D. A mutex variant supporting recursive lock acquisition

Answer: C

Q1098. What is lock-free programming?

- A. A type of race condition caused by missing synchronization
- B. Programming without any synchronization mechanisms at all
- C. Single-threaded programming without any concurrent access
- D. Using atomic operations instead of locks for thread progress

Answer: D

Q1099. What is the difference between green threads and OS threads?

- A. Green threads are runtime-scheduled; OS threads are kernel-managed
- B. Green threads are always faster than OS threads in every scenario
- C. They are exactly the same with no meaningful differences at all
- D. OS threads use less memory than green threads in all situations

Answer: A

Q1100. What is Software Transactional Memory (STM)?

- A. A concurrency mechanism using transactions for shared memory with retry
- B. A memory allocation strategy for managing heap memory efficiently
- C. A thread scheduler for determining execution order of all threads
- D. A relational database system for persistent data storage

Answer: A

Q1101. What is property-based testing?

- A. Specifying invariants that hold for all generated inputs
- B. Testing object properties for correct values and types
- C. Testing CSS properties for correct visual rendering
- D. Testing configuration properties for application behavior

Answer: A

Q1102. What is fuzzing?

- A. Making code intentionally unclear and hard to understand
- B. Feeding random or malformed data to find vulnerabilities
- C. A string operation for making text fuzzy and unreadable
- D. A sorting technique that introduces randomness into order

Answer: B

Q1103. What is mutation testing?

- A. Testing DNA sequences for mutations in bioinformatics apps
- B. Introducing small code changes to check if tests catch them
- C. Testing genetic algorithms for correct evolution behavior
- D. Testing code changes for backward compatibility with old code

Answer: B

Q1104. What is the difference between static and dynamic analysis?

- A. Dynamic analysis is performed before compilation of the source code
- B. Static examines code without executing; dynamic examines during execution
- C. There is no meaningful difference between the two analysis approaches
- D. Static analysis is performed at runtime while the program is executing

Answer: B

Q1105. What is a flame graph in performance analysis?

- A. An error graph displaying exception frequency over time
- B. A fire hazard visualization for safety monitoring apps
- C. A visualization of call stacks showing CPU time by code path
- D. A heat map showing temperature distribution across hardware

Answer: C

Q1106. What is symbolic execution?

- A. Executing code with special symbol characters as input data
- B. Executing code with symbolic values to explore all possible paths
- C. A debugging mode that shows variable names instead of values
- D. Running code with special characters to test encoding correctness

Answer: B

Q1107. What is the difference between correctness testing and performance testing?

- A. They are exactly the same type of testing with no differences
- B. Correctness verifies behavior; performance measures speed and resources
- C. Correctness testing measures execution speed like performance does
- D. Performance testing checks for bugs just like correctness testing

Answer: B

Q1108. What is chaos engineering?

- A. Random testing with no specific goals or expected outcomes
- B. Unstructured development without any planning or methods
- C. Creating intentionally buggy software for entertainment
- D. Intentionally introducing failures to test system resilience

Answer: D

Q1109. What is taint analysis?

- A. Analyzing memory allocation patterns for optimization purposes
- B. Tracking untrusted input flow to find security vulnerabilities
- C. Checking code for bugs using standard debugging tools
- D. Formatting code according to standard style guide conventions

Answer: B

Q1110. What is the difference between benchmarking and profiling?

- A. They are identical performance analysis methods with no differences
- B. Benchmarking measures overall metrics; profiling identifies bottlenecks
- C. Profiling gives overall performance metrics; benchmarking is specific
- D. Benchmarking provides more detailed analysis than profiling overall

Answer: B

Q1111. What is the Twelve-Factor App methodology?

- A. Twelve testing phases that every software project must go through
- B. A math-based approach to solving programming problems algorithmically
- C. Twelve best practices for building scalable cloud-native applications
- D. Twelve programming languages that every developer should learn fully

Answer: C

Q1112. What is trunk-based development?

- A. A database design pattern for normalizing relational table schemas
- B. A branching strategy with single shared branch and short-lived branches
- C. Using tree data structures for organizing hierarchical information
- D. A testing strategy for validating software behavior systematically

Answer: B

Q1113. What is the difference between monorepo and polyrepo strategies?

- A. Polyrepo is only suitable for small projects with limited source code
- B. Monorepo is always the better choice regardless of project requirements
- C. There is no meaningful difference between the two repository strategies
- D. Monorepo stores all projects in one repo; polyrepo uses separate repos

Answer: D

Q1114. What is Infrastructure as Code (IaC)?

- A. Managing infrastructure through configuration files not manual processes
- B. A programming paradigm for organizing code into functional modules
- C. Writing source code directly on production servers for deployment
- D. A coding standard for formatting and naming conventions in code

Answer: A

Q1115. What is the concept of 'shift left' in testing?

- A. Moving code directly to production without any testing or review
- B. Left-aligning all source code for consistent formatting readability
- C. Performing testing earlier in the development lifecycle for detection
- D. Moving test files to the left side of the project directory tree

Answer: C

Q1116. What is feature flagging/toggling?

- A. A branching strategy for organizing feature development in Git
- B. A testing tool for validating feature behavior before release
- C. Deleting features permanently from the codebase and all traces
- D. Enabling or disabling features at runtime without deploying code

Answer: D

Q1117. What is the strangler fig pattern for legacy modernization?

- A. Gradually replacing legacy systems by building new functionality around
- B. A refactoring tool for automatically restructuring legacy code
- C. A tree-related data structure pattern for hierarchical organization
- D. Deleting all old code at once and replacing with a new system

Answer: A

Q1118. What is observability in software systems?

- A. A debugging tool for stepping through code execution line by line
- B. Watching code run step by step in a debugger tool window
- C. Monitoring only server uptime without any additional data collection
- D. Understanding system state from external outputs like logs and metrics

Answer: D

Q1119. What is the CAP theorem and how does it affect development practices?

- A. A design pattern for organizing code into reusable modular components
- B. In distributed systems, only two of Consistency, Availability, Partition tolerance
- C. A coding principle for writing clean and maintainable source code
- D. A testing theorem for determining the minimum number of test cases

Answer: B

Q1120. What is the concept of 'toil' in Site Reliability Engineering (SRE)?

- A. Testing software to validate that it meets all specified requirements
- B. Manual repetitive automatable work scaling linearly with service size
- C. Hard work that requires significant effort and long hours from team
- D. Debugging code to find and fix errors in the application source

Answer: B

Q1121. A program compiles without errors but produces incorrect output for certain edge cases. What type of error is this?

- A. A runtime error causing immediate termination
- B. A linker error from missing library references
- C. A syntax error caught by the compiler during build
- D. A logical error in the algorithm for edge cases

Answer: D

Q1122. In a multi-pass compiler, what is the primary advantage of separating analysis into multiple passes?

- A. It allows forward references and better optimization
- B. It makes the compiler run faster by skipping checks
- C. It eliminates the need for any optimization phase
- D. It removes the requirement for a symbol table

Answer: A

Q1123. A program works on the developer machine but fails in production. Which factor is most likely the cause?

- A. The source code changed itself during the deployment
- B. Environment differences such as OS or library versions
- C. The language behaves differently on each machine
- D. The compiler randomly generates different code each run

Answer: B

Q1124. What is the key difference between ahead-of-time (AOT) and just-in-time (JIT) compilation?

- A. AOT requires an interpreter; JIT does not need one
- B. JIT always produces slower code than AOT in every case
- C. AOT compiles before execution; JIT compiles during runtime
- D. AOT is only for scripting; JIT only for systems languages

Answer: C

Q1125. Why might a developer choose a dynamically typed language for rapid prototyping?

- A. Dynamic typing makes programs run faster on all hardware
- B. Dynamic typing enables faster iteration by skipping types
- C. Dynamic typing always produces more efficient machine code
- D. Dynamic typing prevents all runtime errors from occurring

Answer: B

Q1126. What problem does the halting problem describe in computer science?

- A. The difficulty of stopping an infinite loop in a program
- B. The challenge of compiling large programs within limits
- C. The problem of optimizing code for memory constraints
- D. The impossibility of deciding if any program will halt

Answer: D

Q1127. A language uses type inference to determine variable types. What is the main benefit?

- A. It makes the language dynamically typed at runtime
- B. It prevents the compiler from performing optimization
- C. It combines static type safety with reduced verbosity
- D. It eliminates type safety entirely from the language

Answer: C

Q1128. In language design, what does homoiconicity mean?

- A. The language can only represent one data type
- B. Code and data share the same representation
- C. The language does not support metaprogramming
- D. The compiler and interpreter are identical

Answer: B

Q1129. What is the primary purpose of an abstract syntax tree (AST) in compilation?

- A. To store the final machine code for the target
- B. To handle input and output during execution
- C. To represent hierarchical syntactic structure
- D. To manage memory allocation during runtime

Answer: C

Q1130. A team debates compiled vs interpreted for a real-time system. What is the strongest argument for compiled?

- A. Compiled languages always have simpler syntax
- B. Compiled languages are easier for beginners
- C. Compiled languages need no testing before use
- D. Compiled languages produce optimized native code

Answer: D

Q1131. A financial app uses float for currency and has rounding errors. What is the recommended fix?

- A. Convert all values to strings before performing calculations
- B. Use double type since it eliminates all rounding errors
- C. Increase float precision by adding more decimal places
- D. Use integer types storing values in cents to avoid issues

Answer: D

Q1132. What is the concept of type erasure in Java generics?

- A. Type information is stored in a separate metadata file
- B. The compiler removes all type info making code untyped
- C. The JVM dynamically creates new types for each generic use
- D. Generic type parameters are replaced with bounds at compile

Answer: D

Q1133. A developer encounters NaN in JavaScript. Which operation could produce this result?

- A. Multiplying a string that cannot be parsed as a number
- B. Comparing two floating-point numbers with strict equality
- C. Concatenating two numeric strings together into one
- D. Dividing a positive integer by a negative integer value

Answer: A

Q1134. Why does comparing floating-point numbers with exact equality often fail?

- A. The equality operator does not work with numeric types
- B. Binary representation introduces tiny precision errors
- C. Floating-point numbers are stored as strings internally
- D. Floating-point comparisons are not supported by CPUs

Answer: B

Q1135. In Rust, what is the purpose of the Option<T> type?

- A. It provides thread-safe access to shared mutable data
- B. It stores multiple values of different types at once
- C. It safely represents a value that may or may not exist
- D. It optimizes memory allocation for large structures only

Answer: C

Q1136. What is a phantom type in languages with advanced type systems?

- A. A type that cannot be instantiated under any circumstances
- B. A type that stores data in an encrypted format always
- C. A type that automatically converts between other data types
- D. A type existing only at compile time using no runtime data

Answer: D

Q1137. A C program assigns a negative signed integer to an unsigned variable. What is the likely outcome?

- A. The runtime throws an overflow exception and terminates
- B. The program will refuse to compile due to a type mismatch
- C. The unsigned variable automatically stores zero as default
- D. The value becomes a large positive number via reinterpretation

Answer: D

Q1138. What problem do algebraic data types (ADTs) solve in functional programming?

- A. They eliminate the need for functions in functional code
- B. They provide automatic serialization to network formats
- C. They optimize memory by compressing data automatically
- D. They model complex data as combinations of sum and product

Answer: D

Q1139. Why does JavaScript treat typeof null as object instead of null?

- A. Because null and objects share the same memory structure
- B. Because null is implemented as an empty object internally
- C. Because of a legacy bug in the original JS implementation
- D. Because the spec intentionally groups null with objects

Answer: C

Q1140. What is the difference between structural typing and nominal typing?

- A. Structural checks type names; nominal checks type shape
- B. Nominal typing does not support inheritance between types
- C. Structural checks type shape; nominal checks type names
- D. Structural typing is slower than nominal in every case

Answer: C

Q1141. A developer uses bitwise AND to check if a number is odd. Which expression correctly identifies odd numbers?

- A. $(n \& 2) == 1$ returns true for odd numbers in binary
- B. $(n \& 1) == 0$ returns true because odd clears last bit
- C. $(n \& 0) == 1$ returns true for odd numbers correctly
- D. $(n \& 1) == 1$ returns true because odd last bit is set

Answer: D

Q1142. In C++, what potential issue arises from overloading the comma operator?

- A. It may break the guaranteed left-to-right evaluation
- B. It causes compilation errors in all template-based code
- C. The comma operator cannot be overloaded in C++ standard
- D. It automatically converts all expressions to void type

Answer: A

Q1143. What does the unsigned right shift ($\gg$) do differently from signed right shift ($\gg$)?

- A. It shifts bits right and preserves the sign bit in position
- B. It rotates bits right wrapping them around to the left side
- C. It shifts bits left instead of right despite arrow direction
- D. It shifts bits right and fills with zeros regardless of sign

Answer: D

Q1144. A Python developer writes 'a is b' vs 'a == b'. When do these give different results?

- A. When a and b reference the exact same object in memory
- B. When a and b are the same integer literal used repeatedly
- C. When a and b have equal values but are different objects
- D. When both a and b are None compared in an if statement

Answer: C

Q1145. What is the result of applying De Morgan's law to $!(A \&\& B)$?

- A. $(!A \parallel !B)$ distributes the negation and flips the operator
- B. $(!A \&\& !B)$ negation of each operand combined with AND
- C. $(!A \&\& B)$ negates only the first operand in expression
- D. $(A \parallel B)$ simply removes the negation from the expression

Answer: A

Q1146. Why should developers avoid relying on operator precedence in complex expressions?

- A. Because parentheses are not allowed with operator precedence
- B. Because complex precedence reduces readability and causes bugs
- C. Because operator precedence varies between every language
- D. Because operator precedence only applies to arithmetic ops

Answer: B

Q1147. What is the purpose of the pipe operator ($|>$) in functional languages like Elixir?

- A. It passes the result of one expression as input to next
- B. It redirects standard output to a file on the system
- C. It creates a new parallel process for each function call
- D. It performs bitwise OR operations on integer operands

Answer: A

Q1148. In JavaScript, what is the result of {} + [] when evaluated as a statement?

- A. A type error because objects and arrays cannot be added
- B. The string '{}[]' because both operands convert to string
- C. An empty object combined with empty array resulting in NaN
- D. The number 0 because {} is empty block and +[] coerces to 0

Answer: D

Q1149. What advantage does the walrus operator (:=) provide in Python 3.8+?

- A. It performs parallel assignment of multiple variables at once
- B. It creates constants that cannot be reassigned after init
- C. It defines default parameter values for function declarations
- D. It assigns a value and returns it within a single expression

Answer: D

Q1150. How does operator precedence interact with associativity when precedence is equal?

- A. Associativity determines evaluation order for same precedence
- B. The compiler randomly selects which operator to evaluate first
- C. Operators with equal precedence always evaluate left to right
- D. Operators with equal precedence cannot appear in one expression

Answer: A

Q1151. A function uses recursion with no base case. What will happen during execution?

- A. The compiler automatically adds a base case to stop
- B. The call stack overflows causing a stack overflow error
- C. The function returns null after one recursive call
- D. The function loops exactly once then exits normally

Answer: B

Q1152. In a switch with intentional fall-through, what is the best practice for maintainability?

- A. Convert the entire switch to unrelated if statements
- B. Remove all break statements for code consistency
- C. Add explicit comments documenting intentional fall-through
- D. Use goto statements instead of fall-through for clarity

Answer: C

Q1153. A developer replaces nested if-else with a strategy pattern. What is the main benefit?

- A. Memory usage is reduced because fewer objects are created
- B. The code executes faster because polymorphism is quicker
- C. The number of classes in the codebase decreases greatly
- D. Each strategy is encapsulated making code more extensible

Answer: D

Q1154. What is the time complexity of a triply nested loop where each iterates n times?

- A. $O(n^3)$ because each nesting multiplies by factor n
- B. $O(n)$ because only the innermost loop affects complexity
- C. $O(n^2)$ because the outer loop does not count
- D. $O(\log n)$ because nested loops divide problem recursively

Answer: A

Q1155. How does continuation-passing style (CPS) change program control flow?

- A. It uses goto statements to jump between code sections
- B. It eliminates all function calls using only loops instead
- C. It converts recursive functions into iterative ones always
- D. It passes a callback representing the next step to each

Answer: D

Q1156. A loop uses floating-point comparison as its termination condition. Why is this risky?

- A. The loop will always terminate early before full completion
- B. Floating-point loops always execute one extra iteration
- C. Floating-point numbers cannot be used in loop conditions
- D. Precision errors may prevent termination condition triggering

Answer: D

Q1157. What is the advantage of iterators over index-based loops for data structures?

- A. Iterators use less memory than index-based loops in all cases
- B. Iterators always process elements faster than index access
- C. Iterators prevent any modifications to underlying collection
- D. Iterators provide a uniform interface across data structures

Answer: D

Q1158. What problem does the loop-and-a-half pattern solve?

- A. It creates loops executing exactly one and a half times
- B. It converts while loops into for loops for performance
- C. It handles loops needing to check condition in middle
- D. It optimizes loops to run at half the time of regular

Answer: C

Q1159. How do higher-order functions like map and filter replace loops in functional programming?

- A. They convert loops into goto statements for faster speed
- B. They apply transformations declaratively without iteration
- C. They eliminate the need for conditional logic in programs
- D. They only work with numeric data not other types at all

Answer: B

Q1160. Why might a compiler transform a recursive function into an iterative loop?

- A. Because iteration uses more memory improving caching
- B. Because recursive functions always produce wrong output
- C. Because iteration avoids stack overhead and overflow
- D. Because recursion is not supported by modern CPUs

Answer: C

Q1161. A recursive function for Fibonacci has exponential time complexity. How can this be optimized?

- A. By converting the function parameters to global variables
- B. By using memoization to cache previously computed values
- C. By adding more recursive calls to speed up computation
- D. By removing the base case to reduce function calls

Answer: B

Q1162. What is the Y combinator and what problem does it solve in lambda calculus?

- A. It validates the types of arguments passed to functions
- B. It combines two functions into a single optimized one
- C. It enables recursion in languages without named functions
- D. It converts iterative loops into recursive function calls

Answer: C

Q1163. A function modifies a global variable as a side effect. Why is this considered poor practice?

- A. Because side effects always cause the program to crash immediately
- B. Because side effects make functions unpredictable and harder to test
- C. Because global variables use too much memory in programs
- D. Because programming languages prohibit modifying global variables

Answer: B

Q1164. What is the difference between currying and partial application?

- A. Currying combines multiple functions; partial splits them into pieces
- B. Currying splits a function into single-arg chain; partial fixes some args
- C. Currying and partial application are exactly the same technique always
- D. Currying only works with two arguments; partial works with any count

Answer: B

Q1165. How does tail call optimization (TCO) prevent stack overflow in recursive functions?

- A. It moves recursive calls to a separate memory heap location
- B. It converts all recursive calls into asynchronous operations
- C. It limits the depth of recursion to a fixed maximum count
- D. It reuses the current stack frame for the next recursive call

Answer: D

Q1166. What is a thunk in the context of lazy evaluation?

- A. A deferred computation wrapped in a zero-argument function
- B. A data structure that stores function parameter information
- C. A compiled machine code fragment stored in memory cache
- D. An error handler that catches exceptions in lazy code blocks

Answer: A

Q1167. A developer uses function composition $f(g(x))$. What is the main benefit of this approach?

- A. It builds complex operations from simple reusable functions
- B. It eliminates the need for any error handling in code
- C. It guarantees the composed functions execute in parallel
- D. It ensures both functions always return the same data type

Answer: A

Q1168. What is a functor in functional programming?

- A. An object that stores multiple functions in an internal array
- B. A function that can only be called once during execution
- C. A special variable that holds a reference to a function pointer
- D. A type that can be mapped over with a structure-preserving function

Answer: D

Q1169. Why are first-class functions important for functional programming paradigms?

- A. They enable functions to be passed, returned, and assigned like data
- B. They prevent functions from having any parameters or return values
- C. They allow functions to be stored in databases for persistence
- D. They make all functions execute faster than in other paradigms

Answer: A

Q1170. What problem does the trampolining technique solve for recursive functions?

- A. It automatically parallelizes recursive calls across CPU cores
- B. It prevents stack overflow by converting recursion to loop steps
- C. It speeds up recursive calculations by using hardware acceleration
- D. It eliminates the need for base cases in recursive algorithms

Answer: B

Q1171. A program reads a large file line by line instead of loading it all at once. What is the main advantage?

- A. It prevents the file from being modified by other processes
- B. Line-by-line reading is always faster than loading at once
- C. It automatically handles character encoding conversion issues
- D. It reduces memory usage by processing one line at a time

Answer: D

Q1172. What is memory-mapped I/O and when is it beneficial?

- A. Allocating extra RAM to speed up sequential file reading
- B. Mapping files into virtual memory for faster random access
- C. Storing all program variables in external files on disk
- D. Converting memory addresses to file paths automatically

Answer: B

Q1173. A multi-threaded program writes to the same file simultaneously. What problem can occur?

- A. Race conditions cause corrupted or interleaved file data
- B. Each thread creates its own separate copy of the file
- C. The file automatically becomes read-only preventing writes
- D. The operating system duplicates the file for each thread

Answer: A

Q1174. What is the difference between blocking and non-blocking I/O in server design?

- A. Blocking handles more connections; non-blocking handles fewer
- B. Non-blocking cannot read from network sockets at all ever
- C. Blocking waits for completion; non-blocking returns immediately
- D. Blocking is only for files; non-blocking only for networks

Answer: C

Q1175. What is the purpose of a write-ahead log (WAL) in database I/O?

- A. It logs all read operations for performance monitoring only
- B. It encrypts data before writing to prevent unauthorized access
- C. It compresses write operations to reduce disk space usage
- D. It records changes before applying them for crash recovery

Answer: D

Q1176. A developer needs to parse a 10GB XML file. Why would a SAX parser be preferred over DOM?

- A. DOM parsers cannot handle files containing nested elements
- B. SAX processes XML as a stream without loading it all in memory
- C. SAX is more accurate at parsing XML structure than DOM
- D. SAX automatically validates XML against its schema definition

Answer: B

Q1177. What is the reactor pattern in asynchronous I/O programming?

- A. A technique for converting synchronous code to asynchronous
- B. A single-threaded event loop dispatching I/O events to handlers
- C. A method of buffering all output until program termination
- D. A pattern that creates a new thread for every I/O request

Answer: B

Q1178. How does zero-copy I/O improve performance compared to traditional I/O?

- A. It avoids copying data between kernel and user space buffers
- B. It compresses data to zero bytes before transferring it out
- C. It uses zero system calls by directly accessing hardware
- D. It eliminates the need for any disk storage by using RAM only

Answer: A

Q1179. What is the purpose of the fsync system call?

- A. It synchronizes file access between multiple running threads
- B. It converts file system format between different OS types
- C. It synchronizes file timestamps with the system clock value
- D. It ensures file data is physically written to storage device

Answer: D

Q1180. A program handles both file I/O and network I/O. Why might it use an event-driven architecture?

- A. Event-driven guarantees faster disk I/O than any alternative
- B. Event-driven efficiently handles many concurrent I/O sources
- C. Event-driven converts network data to file format automatically
- D. Event-driven eliminates all errors in I/O operations always

Answer: B

Q1181. A program compares user passwords using == instead of a constant-time comparison. What is the security risk?

- A. The program will crash when processing long password strings
- B. The comparison may return incorrect results for all input
- C. The passwords will be stored in plain text in memory always
- D. Timing attacks can reveal password characters one by one

Answer: D

Q1182. Why is the KMP algorithm preferred over naive string matching for large texts?

- A. KMP requires the pattern to be sorted before searching starts
- B. KMP uses more memory but is simpler to implement correctly
- C. KMP avoids rechecking characters using a failure function table
- D. KMP only works with numeric strings not alphabetic text data

Answer: C

Q1183. What is the purpose of string interning and when can it cause problems?

- A. It translates strings between languages causing encoding issues
- B. It encrypts strings for security but slows down all comparisons
- C. It caches unique strings saving memory but can cause leaks if overused
- D. It compresses strings for storage but prevents substring extraction

Answer: C

Q1184. How does the Rabin-Karp algorithm use hashing for string matching?

- A. It creates a hash table of all possible patterns before searching
- B. It hashes the entire text once and compares with pattern hash
- C. It uses rolling hash to efficiently compare pattern with substrings
- D. It hashes each character individually and sorts them for matching

Answer: C

Q1185. What is a rope data structure and when is it better than a standard string?

- A. A rope compresses strings to use minimal memory allocation
- B. A rope converts strings to binary for faster CPU processing
- C. A rope stores strings in encrypted form for better security
- D. A rope is a tree of small strings efficient for large text edits

Answer: D

Q1186. What is the difference between UTF-8 and UTF-16 encoding?

- A. UTF-8 uses 8 bits per character always; UTF-16 uses 16 always
- B. UTF-8 only supports English; UTF-16 supports all world languages
- C. UTF-8 uses variable 1-4 bytes; UTF-16 uses 2 or 4 byte units
- D. UTF-8 is binary format; UTF-16 is a text-based format only

Answer: C

Q1187. A developer finds that reversing a string containing emoji produces corrupted output. What is the cause?

- A. The reversal algorithm has a bug unrelated to character type
- B. Emoji use multi-code-unit sequences broken by naive reversal
- C. Emoji characters are not valid in any programming language
- D. The string library does not support any non-ASCII characters

Answer: B

Q1188. What is the purpose of normalization forms (NFC, NFD) in Unicode string processing?

- A. They ensure equivalent character sequences have canonical form
- B. They convert strings from Unicode to ASCII encoding format
- C. They validate that strings contain only printable characters
- D. They compress Unicode strings to reduce their storage size

Answer: A

Q1189. Why might a trie data structure be preferred for autocomplete over a hash map of strings?

- A. Tries automatically sort strings in reverse alphabetical order
- B. Tries use less memory than hash maps for all data sizes
- C. Tries are faster for exact match lookups than hash maps
- D. Tries support efficient prefix-based searching and traversal

Answer: D

Q1190. What is a homoglyph attack in the context of string security?

- A. An attack that overflows string buffers with excess data
- B. An attack that deletes all string data from the database
- C. An attack that encrypts strings making them unreadable
- D. An attack using visually similar characters to deceive users

Answer: D

Q1191. A developer uses a HashMap with mutable objects as keys. What problem can occur?

- A. The map automatically prevents mutation of all key objects
- B. The map will throw an error when any key is first inserted
- C. Mutable keys cause the map to sort entries in reverse order
- D. Mutating keys changes their hash, making entries unretrievable

Answer: D

Q1192. What is the amortized time complexity of adding an element to a dynamic array (ArrayList)?

- A. $O(\log n)$ because dynamic arrays use binary search for insert
- B. $O(n^2)$ because each insertion triggers a full array sort
- C. $O(n)$ because every insertion requires copying all elements
- D. $O(1)$ amortized because resizing cost is spread over operations

Answer: D

Q1193. What is a bloom filter and what tradeoff does it make?

- A. A compression technique that trades accuracy for file size
- B. A probabilistic structure allowing false positives not negatives
- C. A sorting algorithm that trades stability for speed
- D. A search algorithm trading memory for faster query execution

Answer: B

Q1194. Why is a red-black tree preferred over a simple BST for implementing TreeMap?

- A. Red-black trees guarantee $O(\log n)$ by maintaining balance always
- B. Red-black trees use less memory than simple BSTs in all cases
- C. Red-black trees do not require any comparison between elements
- D. Simple BSTs cannot store key-value pairs only single values

Answer: A

Q1195. What is the difference between a B-tree and a binary search tree?

- A. B-trees cannot be balanced while binary search trees can be
- B. B-trees have multiple children per node optimized for disk I/O
- C. B-trees store only two children per node like binary trees
- D. B-trees only work with string keys not with numeric values

Answer: B

Q1196. A concurrent program uses a regular HashMap and encounters data corruption. What is the fix?

- A. Use ConcurrentHashMap which provides thread-safe operations
- B. Use a larger initial capacity for the HashMap allocation
- C. Add more entries to the HashMap to prevent hash collisions
- D. Convert the HashMap to a TreeMap for sorted data access

Answer: A

Q1197. What is a skip list and what advantage does it offer?

- A. A list that skips every other element for faster traversal
- B. A list that automatically removes duplicate elements on insert
- C. A probabilistic structure providing $O(\log n)$ search in linked list
- D. A list optimized for storing very large string values only

Answer: C

Q1198. What is the difference between a weak reference map and a regular map?

- A. Weak maps allow garbage collection of keys not strongly referenced
- B. Weak maps are faster for lookup operations than regular maps
- C. Weak maps cannot store null values while regular maps can
- D. Weak maps store values with less precision than regular maps

Answer: A

Q1199. What is cache-oblivious algorithm design and how does it relate to array processing?

- A. It disables caching to ensure consistent performance results
- B. It ignores CPU cache entirely for simpler algorithm design
- C. It fills the CPU cache completely before processing any data
- D. It optimizes memory access patterns without knowing cache size

Answer: D

Q1200. Why does a hash table with a poor hash function degrade to $O(n)$ lookup time?

- A. Because many collisions cluster entries into the same bucket
- B. Because poor hash functions produce negative index values
- C. Because poor hash functions prevent the table from resizing
- D. Because the hash table automatically switches to linear search

Answer: A

Q1201. A developer violates the Single Responsibility Principle by creating a class that handles both data validation and database storage. What is the consequence?

- A. The class automatically becomes thread-safe for concurrent use
- B. The compiler will refuse to compile the class with both concerns
- C. The class will execute faster due to fewer method calls needed
- D. Changes to validation logic may unintentionally break storage code

Answer: D

Q1202. What is the visitor pattern and when should it be used?

- A. A pattern for creating multiple instances of a single class
- B. A pattern to add operations to objects without modifying classes
- C. A pattern for sorting objects in a collection by their type
- D. A pattern for managing user sessions in web applications

Answer: B

Q1203. How does the SOLID Open/Closed Principle guide class design?

- A. Classes should be open for both extension and modification always
- B. Classes should be closed for both extension and modification
- C. Classes should be open for extension and closed for modification
- D. Classes should be open for modification and closed for extension

Answer: C

Q1204. What is the difference between covariance and contravariance in generic types?

- A. Both covariance and contravariance preserve subtype ordering
- B. Neither covariance nor contravariance affects subtype ordering
- C. Covariance preserves subtype ordering; contravariance reverses it
- D. Covariance reverses subtype ordering; contravariance preserves it

Answer: C

Q1205. What is the template method pattern and how does it use inheritance?

- A. It generates code templates for creating new class files quickly
- B. It replaces all inheritance with composition for better flexibility
- C. It creates multiple templates of the same class for different uses
- D. It defines algorithm skeleton in base class deferring steps to sub

Answer: D

Q1206. Why should developers prefer composition over inheritance in most cases?

- A. Because inheritance is not supported in modern languages anymore
- B. Because composition provides more flexibility and looser coupling
- C. Because composition always uses less memory than inheritance does
- D. Because inheritance cannot be used with interfaces in any language

Answer: B

Q1207. What is double dispatch and how does the visitor pattern implement it?

- A. Resolving a method call based on two object types at runtime
- B. Dispatching a method call to two different servers simultaneously
- C. Creating two instances of the same class for parallel processing
- D. Calling the same method twice with different arguments each time

Answer: A

Q1208. What problem does the fragile base class issue describe?

- A. Changes to base class may unexpectedly break derived classes
- B. Base classes consume too much memory in the inheritance chain
- C. Base classes cannot have more than one level of inheritance
- D. Derived classes always override every method in the base class

Answer: A

Q1209. How does the CQRS pattern separate read and write operations in OOP design?

- A. It combines all reads and writes into a single unified model
- B. It uses separate models for reading queries and writing commands
- C. It prevents any reading operations during write transactions
- D. It uses different programming languages for reads and writes

Answer: B

Q1210. What is the expression problem and how do OOP and FP approach it differently?

- A. Both OOP and FP solve the expression problem in identical ways
- B. OOP easily adds new operations; FP easily adds new types always
- C. Neither OOP nor FP can add new types or operations at all ever
- D. OOP easily adds new types; FP easily adds new operations on types

Answer: D

Q1211. A program has a memory leak that only manifests after running for several days. What approach best identifies the leak?

- A. Adding more RAM to the server to delay the out-of-memory crash
- B. Using a memory profiler to track allocation growth over time
- C. Reviewing source code manually for every allocation and free
- D. Restarting the program more frequently to prevent the leak

Answer: B

Q1212. What is the generational hypothesis in garbage collection?

- A. Older objects are more likely to be collected than newer ones
- B. Most objects die young so collecting young generation is efficient
- C. Garbage collection should run only once per program execution
- D. All objects have the same lifetime regardless of creation time

Answer: B

Q1213. What is a use-after-free vulnerability and how can it be exploited?

- A. Freeing memory twice causing the allocator to add extra memory
- B. Accessing freed memory that may contain attacker-controlled data
- C. Using free software after a trial period has expired completely
- D. Using a variable after its scope ends causing a compiler error

Answer: B

Q1214. How does Rust's ownership system prevent memory safety bugs at compile time?

- A. Each value has one owner and memory is freed when owner goes out of scope
- B. It prevents all heap allocation forcing everything onto the stack
- C. It uses garbage collection that runs at compile time instead
- D. It requires programmers to manually free every allocation made

Answer: A

Q1215. What is the difference between mark-and-sweep and mark-and-compact garbage collection?

- A. Mark-sweep leaves objects in place; mark-compact moves them together
- B. Mark-sweep moves live objects; mark-compact leaves them in place
- C. Mark-compact cannot handle circular references unlike mark-sweep
- D. Mark-sweep is faster than mark-compact in all possible scenarios

Answer: A

Q1216. What is a slab allocator and when is it used in operating systems?

- A. An allocator that only works with very large memory blocks
- B. An allocator that stores data on disk instead of in memory
- C. An allocator caching fixed-size objects for frequent kernel allocations
- D. An allocator that compresses memory to fit more data in RAM

Answer: C

Q1217. How does address space layout randomization (ASLR) improve memory security?

- A. It encrypts all memory contents to prevent unauthorized reading
- B. It randomizes memory layout making exploit addresses unpredictable
- C. It prevents any program from allocating memory on the heap
- D. It locks memory pages so they cannot be written to by code

Answer: B

Q1218. What is false sharing in multi-threaded programs and how does it affect performance?

- A. Threads falsely reporting memory as free when it is still in use
- B. Sharing read-only data between threads causing unnecessary copying
- C. Multiple threads sharing the same variable causing data corruption
- D. Threads on different cores invalidate cache lines they do not truly share

Answer: D

Q1219. What is the purpose of memory-mapped files and how do they differ from traditional file I/O?

- A. They compress files in memory to reduce the total memory used
- B. They encrypt files in memory to prevent unauthorized data reading
- C. They store files entirely in CPU registers for faster access
- D. They map file contents to virtual memory allowing direct pointer access

Answer: D

Q1220. What is a garbage collection pause and why is it problematic for real-time applications?

- A. A delay caused by the GC writing collection logs to disk
- B. A delay when the GC starts running for the very first time
- C. A stop-the-world pause halting application threads during collection
- D. A pause occurring only when the program first allocates memory

Answer: C

Q1221. A developer wraps every method in try-catch blocks that silently swallow exceptions. What is the consequence?

- A. Memory usage decreases because exceptions are not stored
- B. The application runs faster due to fewer exception propagations
- C. The application becomes completely immune to all runtime errors
- D. Errors are hidden making debugging extremely difficult and slow

Answer: D

Q1222. What is the difference between fault tolerance and error recovery in system design?

- A. Fault tolerance continues despite faults; recovery restores state
- B. Fault tolerance prevents all errors; recovery fixes them after
- C. Fault tolerance is for hardware; recovery is for software only
- D. Both terms mean exactly the same thing in all system contexts

Answer: A

Q1223. How does the circuit breaker pattern prevent cascading failures in distributed systems?

- A. It duplicates every request to ensure at least one succeeds
- B. It physically disconnects failing servers from the network
- C. It stops calling a failing service after threshold allowing recovery
- D. It encrypts all communication to prevent transmission errors

Answer: C

Q1224. What is the difference between fail-fast and fail-safe error handling strategies?

- A. Fail-fast is for testing; fail-safe is for production environments
- B. Both strategies handle errors in exactly the same manner always
- C. Fail-fast ignores errors; fail-safe reports them immediately
- D. Fail-fast reports errors immediately; fail-safe continues operating

Answer: D

Q1225. A microservice architecture experiences a cascading failure. Which pattern best prevents this?

- A. Increasing timeout values to give failing services more time
- B. Using bulkhead isolation to contain failures within boundaries
- C. Adding more servers to handle the increased error load
- D. Removing all error handling to let services fail naturally

Answer: B

Q1226. What is the purpose of structured error handling using Result types instead of exceptions?

- A. Result types are faster to execute than exception handling code
- B. Result types prevent any errors from occurring during runtime
- C. Result types make error handling explicit in function signatures
- D. Result types automatically fix errors without developer action

Answer: C

Q1227. How does the retry pattern with exponential backoff improve error handling?

- A. It retries immediately without any delay between each attempt
- B. It retries a fixed number of times with constant delay between
- C. It decreases delay between retries to resolve errors faster
- D. It increases delay between retries reducing load on failing service

Answer: D

Q1228. What is the difference between recoverable and unrecoverable errors in Rust?

- A. Recoverable use Result type; unrecoverable use panic macro
- B. Rust does not distinguish between recoverable and unrecoverable
- C. Recoverable use panic macro; unrecoverable use Result type
- D. Both types are handled using the same mechanism in Rust

Answer: A

Q1229. What is the supervisor pattern in Erlang/OTP for fault tolerance?

- A. A design that prevents any process from ever failing at all
- B. A thread that monitors CPU usage and kills slow processes
- C. A process that monitors child processes and restarts on failure
- D. A function that validates all input before processing it

Answer: C

Q1230. Why is it important to include contextual information when throwing custom exceptions?

- A. Context helps developers quickly diagnose the root cause of errors
- B. Context information makes exceptions run faster in production
- C. Context reduces the memory footprint of exception objects created
- D. Context automatically fixes the error without human intervention

Answer: A

Q1231. A project has conflicting versions of the same dependency required by two libraries. How is this typically resolved?

- A. The project must choose one library and remove the other
- B. Dependency resolution algorithms select a compatible version
- C. Both versions are deleted and the dependency is removed
- D. All dependencies are updated to the latest version available

Answer: B

Q1232. What is the difference between a plugin architecture and a monolithic architecture?

- A. Plugins require more memory than monolithic architecture always
- B. Monolithic supports more features than plugin architecture can
- C. Plugins allow extending functionality without modifying core code
- D. Plugins are slower; monolithic is always faster in execution

Answer: C

Q1233. What is dependency hell and what strategies mitigate it?

- A. Dependencies that are too large consuming excessive disk space
- B. Conflicting version requirements across the dependency graph
- C. Having too few dependencies causing limited functionality
- D. Dependencies that require payment to use in production apps

Answer: B

Q1234. How does lazy loading of modules improve application performance?

- A. It defers loading modules until they are actually needed
- B. It compresses modules to reduce their memory footprint
- C. It loads all modules at startup to prevent later delays
- D. It caches all modules permanently in browser local storage

Answer: A

Q1235. What is the difference between a micro-frontend architecture and a traditional monolithic frontend?

- A. Micro-frontends split UI into independently deployable modules
- B. Monolithic frontends load faster than micro-frontend solutions
- C. Micro-frontends cannot share any state between components
- D. Micro-frontends use only one framework; monolithic uses many

Answer: A

Q1236. What is a supply chain attack in the context of package management?

- A. Compromising a dependency to inject malicious code downstream
- B. An attack that corrupts the package manager configuration file
- C. An attack that prevents packages from being downloaded online
- D. An attack that targets the physical delivery of hardware

Answer: A

Q1237. What is the purpose of code splitting in modern web applications?

- A. Separating code comments from executable statements in files
- B. Splitting code across multiple programming languages for speed
- C. Dividing code equally among team members for parallel coding
- D. Breaking application code into chunks loaded on demand for speed

Answer: D

Q1238. How does the facade pattern simplify complex module interactions?

- A. It removes dependencies between modules by duplicating all code
- B. It hides all modules and prevents any external access to them
- C. It converts all module functions into a single large function
- D. It provides a simplified interface to a complex subsystem of modules

Answer: D

Q1239. What is the hexagonal architecture and how does it relate to module organization?

- A. An architecture using exactly six modules in the application
- B. An architecture isolating core logic from external adapters and ports
- C. An architecture requiring six layers of abstraction in modules
- D. An architecture that arranges modules in a hexagonal grid pattern

Answer: B

Q1240. What is the difference between hard and soft dependencies in module design?

- A. Hard dependencies are faster; soft dependencies are always slower
- B. Hard dependencies are for production; soft are for development only
- C. Hard dependencies are required; soft are optional or gracefully absent
- D. Both hard and soft dependencies behave identically in all cases

Answer: C

Q1241. What is the difference between greedy algorithms and dynamic programming approaches?

- A. Both approaches always produce identical results for all problems
- B. Greedy makes locally optimal choices; DP considers all subproblems
- C. Greedy always finds the optimal solution; DP finds approximate
- D. DP makes locally optimal choices; greedy considers all subproblems

Answer: B

Q1242. What is a monad in functional programming and what problem does it solve?

- A. A pattern for chaining operations while managing side effects
- B. A single unit of memory allocation in functional programs
- C. A function that converts between different data type formats
- D. A data type that can only hold one value at any given time

Answer: A

Q1243. How does the A* algorithm combine features of Dijkstra and greedy best-first search?

- A. A* runs both algorithms separately and picks the better result
- B. A* ignores heuristics and relies only on actual path cost found
- C. A* uses actual cost plus heuristic estimate to guide its search
- D. A* uses Dijkstra for graphs and greedy for trees exclusively

Answer: C

Q1244. What is the CAP theorem and how does it affect distributed system design?

- A. It states only two of consistency, availability, partition tolerance
- B. It applies only to single-machine systems not distributed systems
- C. It guarantees all three properties simultaneously in all systems
- D. It states systems can achieve consistency, availability, and partitions

Answer: A

Q1245. What is the difference between a trie and a balanced BST for string operations?

- A. Tries cannot store strings with common prefixes unlike BSTs
- B. Tries are slower for prefix queries than balanced BSTs always
- C. Tries provide $O(m)$ lookup by key length; BSTs provide $O(m \log n)$
- D. BSTs provide faster prefix-based searching than tries in all cases

Answer: C

Q1246. What is the actor model for concurrent programming?

- A. A model that uses only one thread for all program operations
- B. A model that prevents all concurrency to avoid race conditions
- C. A model where threads share mutable state through locks
- D. A model where actors communicate through asynchronous messages

Answer: D

Q1247. What problem does the flyweight pattern solve in software design?

- A. It prevents objects from growing beyond a certain memory limit
- B. It shares common state among many objects to reduce memory use
- C. It makes heavy objects lighter by removing unnecessary methods
- D. It converts heavyweight processes into lightweight threads only

Answer: B

Q1248. How does consistent hashing improve distributed hash table scalability?

- A. It minimizes key redistribution when nodes are added or removed
- B. It sorts all keys alphabetically across all nodes in the system
- C. It eliminates all hash collisions in the distributed hash table
- D. It ensures all hash values are consecutive integers in sequence

Answer: A

Q1249. What is the difference between optimistic and pessimistic concurrency control?

- A. Optimistic locks before access; pessimistic assumes no conflict
- B. Optimistic is for reads only; pessimistic is for writes exclusively
- C. Both approaches use identical locking mechanisms for all access
- D. Optimistic assumes no conflict and validates; pessimistic locks first

Answer: D

Q1250. What is event sourcing and how does it differ from traditional CRUD?

- A. Event sourcing uses events only for logging not for state recovery
- B. Event sourcing stores only the current state like CRUD does
- C. Event sourcing stores all state changes as immutable event sequence
- D. Event sourcing deletes old data automatically unlike CRUD approach

Answer: C

Q1251. What is lock-free programming and why is it challenging to implement correctly?

- A. A technique that prevents any thread from ever being interrupted
- B. A style that avoids all shared memory between concurrent threads
- C. Using atomic operations for thread safety without traditional locks
- D. Programming without any synchronization for maximum speed always

Answer: C

Q1252. What is the ABA problem in lock-free programming?

- A. When a thread reads value A but another writes B before commit
- B. When two threads both try to set the same value simultaneously
- C. When atomic operations fail due to hardware memory limitations
- D. When a value changes from A to B to A, making change undetectable

Answer: D

Q1253. How does the fork-join framework improve parallel computation?

- A. It recursively splits tasks and merges results for parallelism
- B. It joins multiple programs into one before executing them all
- C. It creates a new thread for every single instruction to execute
- D. It forks the program into separate processes that never rejoin

Answer: A

Q1254. What is the happens-before relationship in the Java memory model?

- A. A rule preventing any reordering of instructions by the compiler
- B. A requirement that all code must execute in written source order
- C. A guarantee that one action's effects are visible to another action
- D. A constraint requiring all threads to start in a specific order

Answer: C

Q1255. What is a memory barrier or fence instruction?

- A. A mechanism for allocating memory in a thread-safe manner
- B. A technique for protecting memory from unauthorized access
- C. A hardware limit on the total amount of memory a program uses
- D. An instruction ensuring memory operations complete in order

Answer: D

Q1256. How does software transactional memory (STM) simplify concurrent programming?

- A. It allows composable atomic memory transactions without locks
- B. It automatically detects and removes all race conditions found
- C. It stores all thread data in a transactional database system
- D. It converts multi-threaded code into single-threaded execution

Answer: A

Q1257. What is the thundering herd problem in concurrent systems?

- A. When threads are distributed unevenly across available CPU cores
- B. When a single thread consumes all available CPU resources alone
- C. When many waiting threads wake simultaneously overwhelming resources
- D. When many threads are created causing the system to slow down

Answer: C

Q1258. What is work stealing in parallel task scheduling?

- A. Idle threads take tasks from busy threads' queues for balance
- B. Threads duplicating work done by other threads for redundancy
- C. Threads stealing CPU time from other threads for faster execution
- D. The scheduler removing low-priority tasks to free up resources

Answer: A

Q1259. How does the read-copy-update (RCU) mechanism work in concurrent systems?

- A. Readers lock data before reading; writers wait for all readers
- B. All operations require exclusive locks for both reads and writes
- C. Readers access data without locks; writers create updated copies
- D. Data is copied to each thread's private memory before any access

Answer: C

Q1260. What is the difference between data parallelism and task parallelism?

- A. Both data and task parallelism are identical concepts in practice
- B. Task parallelism applies different operations to same data always
- C. Data parallelism is for CPUs only; task parallelism for GPUs only
- D. Data parallelism applies same operation to different data portions

Answer: D

Q1261. A test suite has 95% code coverage but critical bugs still reach production. What does this indicate?

- A. That code coverage is a useless metric that should be abandoned
- B. That coverage alone does not guarantee test quality or edge cases
- C. That 95% coverage is too low and 100% is always required
- D. That the remaining 5% of uncovered code contains all the bugs

Answer: B

Q1262. What is mutation testing and how does it evaluate test suite quality?

- A. It tests the program with mutated or corrupted input data files
- B. It modifies the test execution order to find order dependencies
- C. It introduces small code changes to check if tests detect them
- D. It randomly changes test inputs to find flaky test failures

Answer: C

Q1263. What is the difference between performance testing and load testing?

- A. Performance tests measure speed; load tests measure correctness
- B. Performance measures response under conditions; load tests capacity
- C. Performance testing is manual; load testing is always automated
- D. Both terms describe identical testing activities in all contexts

Answer: B

Q1264. A developer uses a memory profiler and finds objects that are referenced but never used. What is this problem called?

- A. A segmentation fault caused by accessing invalid memory address
- B. A buffer overflow caused by writing past array boundaries
- C. A logical memory leak where references prevent garbage collection
- D. A stack overflow caused by deep recursive function calls

Answer: C

Q1265. What is property-based testing and how does it differ from example-based testing?

- A. Property-based verifies invariants with generated inputs randomly
- B. Both approaches use identical methodologies for test generation
- C. Property-based is slower and less effective than example-based
- D. Property-based uses fixed examples; example-based uses random ones

Answer: A

Q1266. What is chaos engineering and how does it improve system reliability?

- A. It intentionally injects failures to discover system weaknesses
- B. It randomly deletes source code files to test backup systems
- C. It introduces random code changes to test developer adaptability
- D. It creates chaotic user interfaces to test user error handling

Answer: A

Q1267. What is the difference between static analysis and dynamic analysis?

- A. Static runs code to find bugs; dynamic examines code without running
- B. Static examines code without running; dynamic analyzes during execution
- C. Static is for compiled languages only; dynamic for interpreted only
- D. Both static and dynamic analysis require running the program code

Answer: B

Q1268. How does delta debugging systematically minimize a failing test case?

- A. It deletes random lines until the test no longer demonstrates failure
- B. It adds more test cases until the failure becomes easier to understand
- C. It compares failing test with passing test to find code differences
- D. It systematically removes portions to find minimal failure-inducing input

Answer: D

Q1269. What is the purpose of fuzzing as a testing technique?

- A. Generating clear and readable test cases for documentation
- B. Measuring the fuzziness or uncertainty of test result accuracy
- C. Feeding random or malformed input to find crashes and vulnerabilities
- D. Testing the user interface with blurred or fuzzy screen rendering

Answer: C

Q1270. What is the testing pyramid and what does it recommend for test distribution?

- A. More unit tests at base, fewer integration, fewest UI tests at top
- B. More UI tests at top, fewer unit tests at bottom of pyramid
- C. Only unit tests with no integration or end-to-end tests at all
- D. Equal numbers of each test type across all levels of the pyramid

Answer: A

Q1271. A team has a large codebase with no tests. What is the recommended approach to introduce testing?

- A. Write tests for all code at once before making any changes
- B. Hire a separate testing team to write all tests independently
- C. Start by adding tests around code that changes most frequently
- D. Delete all existing code and rewrite it with tests from scratch

Answer: C

Q1272. What is the strangler fig pattern for legacy system modernization?

- A. Gradually replacing legacy components while keeping system running
- B. Wrapping the legacy system in a new interface without changes
- C. Immediately replacing the entire legacy system with new code
- D. Documenting the legacy system before shutting it down entirely

Answer: A

Q1273. How does trunk-based development differ from GitFlow?

- A. Both approaches use identical branching and merging strategies
- B. Trunk-based uses short-lived branches merging to main frequently
- C. Trunk-based does not use version control; GitFlow requires it
- D. Trunk-based uses long-lived branches; GitFlow uses short-lived ones

Answer: B

Q1274. What is the twelve-factor app methodology?

- A. A methodology requiring exactly twelve source code files per app
- B. A set of best practices for building scalable cloud-native apps
- C. A testing methodology requiring twelve types of tests minimum
- D. A security framework with twelve levels of access permissions

Answer: B

Q1275. What is observability in software systems and how does it differ from monitoring?

- A. Observability infers internal state from outputs; monitoring tracks metrics
- B. Observability and monitoring are identical concepts in practice
- C. Monitoring provides deeper insight than observability in all cases
- D. Observability only tracks errors; monitoring tracks all system events

Answer: A

Q1276. What is the purpose of feature flags in software deployment?

- A. They indicate which programming language features are being used
- B. They enable toggling features on or off without code deployment
- C. They flag code that contains bugs for developers to fix later
- D. They mark features as complete in the project management tool

Answer: B

Q1277. How does domain-driven design (DDD) improve complex software projects?

- A. It restricts development to a single programming language only
- B. It aligns code structure with business domain using ubiquitous language
- C. It eliminates the need for any testing in the development process
- D. It focuses on optimizing database domain and query performance

Answer: B

Q1278. What is the difference between horizontal and vertical scaling strategies?

- A. Horizontal adds more machines; vertical adds resources to existing
- B. Horizontal adds resources to one machine; vertical adds more machines
- C. Horizontal is for databases only; vertical is for applications only
- D. Both strategies add resources in exactly the same way to systems

Answer: A

Q1279. What is the concept of eventual consistency in distributed systems?

- A. All nodes always have the exact same data at every point in time
- B. Only the primary node has consistent data; replicas do not
- C. Data consistency is never achieved in any distributed system
- D. All nodes will converge to the same state given enough time

Answer: D

Q1280. What is the purpose of a service mesh in microservice architectures?

- A. It handles service-to-service communication, security, and observability
- B. It combines all microservices into a single monolithic application
- C. It stores all microservice data in a single shared database
- D. It generates source code for new microservices automatically

Answer: A

Q1281. How does the LLVM compiler infrastructure enable language-independent optimization?

- A. By using a common intermediate representation that multiple language frontends can target
- B. By interpreting all programs at runtime without compilation
- C. By compiling all languages directly to x86 machine code
- D. By requiring all languages to use the same syntax

Answer: A

Q1282. What is the difference between structural and nominal type equivalence?

- A. Structural equivalence is used only in compiled languages
- B. Nominal equivalence ignores all type information
- C. They produce identical results in every programming language
- D. Structural equivalence compares type structure while nominal equivalence compares type names

Answer: D

Q1283. What is the Church-Turing thesis and its significance for programming?

- A. It proves that all programs will eventually halt
- B. It states that any computable function can be computed by a Turing machine
- C. It defines the maximum speed of any computer program
- D. It establishes that compiled code is always faster than interpreted code

Answer: B

Q1284. What is the expression problem in programming language design?

- A. The challenge of parsing mathematical expressions efficiently
- B. The problem of evaluating boolean expressions with short-circuit logic
- C. The difficulty of adding both new data types and new operations without modifying existing code
- D. The difficulty of handling operator precedence in complex expressions

Answer: C

Q1285. How does dependent typing differ from traditional static typing?

- A. Dependent typing eliminates the need for any type annotations
- B. Dependent types only work with integer values
- C. Dependent typing is another name for dynamic typing
- D. Dependent types allow types to depend on runtime values, enabling more precise type specifications

Answer: D

Q1286. What is the difference between call-by-name and call-by-need evaluation strategies?

- A. Call-by-name re-evaluates arguments each time they are used while call-by-need caches the result after first evaluation
- B. Call-by-need always evaluates arguments before passing them
- C. Call-by-name prevents all side effects in function calls
- D. Call-by-name is the same as call-by-value

Answer: A

Q1287. What is effect typing and how does it extend traditional type systems?

- A. It measures the performance impact of type checking
- B. It tracks computational side effects in the type system alongside value types
- C. It replaces all runtime errors with compile-time type errors
- D. It only applies to visual effects in graphical programming

Answer: B

Q1288. What is the significance of the lambda calculus as a foundation for programming languages?

- A. It establishes the rules for object-oriented inheritance hierarchies
- B. It provides a minimal formal system that can express any computable function using only function abstraction and application
- C. It is a specialized calculus for mathematical computation only
- D. It defines how compilers optimize machine code instructions

Answer: B

Q1289. What is gradual typing and what problem does it solve?

- A. It allows mixing statically typed and dynamically typed code in the same program
- B. It gradually removes type annotations as the program matures
- C. It automatically converts all types to strings at runtime
- D. It types variables based on the order they are declared

Answer: A

Q1290. What is the difference between total and partial functions in programming language theory?

- A. Total functions are only found in imperative languages
- B. Total functions handle all data types while partial functions handle only one
- C. Partial functions are faster because they skip some computations
- D. Total functions return a result for every possible input while partial functions may not terminate or may be undefined for some inputs

Answer: D

Q1291. What is a higher-kinded type and which languages support it?

- A. A type reserved for advanced mathematical computations
- B. A type that is always at the top of the inheritance hierarchy
- C. A type constructor that takes other type constructors as parameters, supported by Haskell and Scala
- D. A type that can store higher-precision decimal numbers

Answer: C

Q1292. What is the difference between covariance and contravariance in generic type parameters?

- A. Covariance applies to input parameters and contravariance to output parameters
- B. They both produce identical behavior in all type systems
- C. Covariance preserves the type ordering while contravariance reverses it
- D. Covariance is used only in dynamically typed languages

Answer: C

Q1293. What is a refinement type in programming languages?

- A. A type that restricts a base type with a predicate specifying allowed values
- B. A type that refines a class through inheritance only
- C. A type that can only be used in pattern matching
- D. A type that has been optimized for faster runtime performance

Answer: A

Q1294. How does Rust's ownership model interact with data types to ensure memory safety?

- A. Rust uses garbage collection for all heap-allocated types
- B. Rust copies all values automatically without restriction
- C. All data types in Rust are allocated on the stack
- D. Each value has exactly one owner, and the value is dropped when the owner goes out of scope

Answer: D

Q1295. What is the difference between existential types and universal types?

- A. Existential types are used only for numbers while universal types handle strings
- B. Universal types must be defined at compile time but existential at runtime
- C. They are different names for the same concept
- D. Universal types work for all types (forall) while existential types hide the concrete type (there exists)

Answer: D

Q1296. What is type-level programming and how is it used in practice?

- A. Using types only for documentation purposes
- B. Automatically generating types from database schemas
- C. Programming without any type annotations
- D. Writing programs that execute entirely at the type level during compilation to enforce complex constraints

Answer: D

Q1297. What is the difference between nominal subtyping and structural subtyping for interfaces?

- A. Structural subtyping requires classes to declare they implement an interface
- B. Nominal subtyping is only used in interpreted languages
- C. Nominal subtyping requires explicit declaration of subtype relationships while structural subtyping checks compatibility by shape
- D. They produce identical behavior in all object-oriented languages

Answer: C

Q1298. What is a linear type system and what guarantee does it provide?

- A. A type system that arranges all types in a linear hierarchy
- B. A type system that evaluates type expressions in linear time
- C. A type system that only supports linear data structures
- D. A type system that ensures every value is used exactly once, preventing resource leaks

Answer: D

Q1299. What is a GADT (Generalized Algebraic Data Type)?

- A. A data type that can only hold algebraic expressions
- B. An algebraic data type where constructors can specify different type parameters for the resulting type
- C. A type alias for complex generic signatures
- D. A standard generic type available in all languages

Answer: B

Q1300. What is the bottom type (Nothing/Never) and how is it used in type systems?

- A. A type that is the subtype of all types and has no possible values, used for functions that never return
- B. The default type assigned to uninitialized variables
- C. A type used exclusively for error codes
- D. A type representing the smallest possible integer value

Answer: A

Q1301. Why can operator overloading lead to maintenance problems in large codebases?

- A. Because operator overloading always causes runtime performance degradation
- B. Because most languages prohibit any form of operator overloading
- C. Because overloaded operators cannot be used with primitive types
- D. Because overloaded operators may behave in unexpected ways that differ from their conventional meaning

Answer: D

Q1302. How can bitwise operators be used to swap two variables without a temporary variable?

- A. Bitwise operators cannot be used to swap variables
- B. By using AND: `a &= b; b &= a; a &= b` swaps the values
- C. By using XOR: `a ^= b; b ^= a; a ^= b` swaps the values
- D. By using OR: `a |= b; b |= a; a |= b` swaps the values

Answer: C

Q1303. What is the purpose of the three-way comparison (spaceship) operator in modern languages?

- A. It returns negative, zero, or positive to indicate less than, equal, or greater than in a single operation
- B. It performs three sequential comparisons automatically
- C. It compares three variables simultaneously
- D. It creates a three-element tuple from two values

Answer: A

Q1304. What is the significance of De Morgan's laws for simplifying Boolean expressions?

- A. They state that AND and OR operators are always interchangeable
- B. They define the order of evaluation for nested Boolean expressions
- C. They prove that all Boolean expressions can be reduced to a single operator
- D. They state that $\text{NOT}(A \text{ AND } B)$ equals $(\text{NOT } A) \text{ OR } (\text{NOT } B)$ and $\text{NOT}(A \text{ OR } B)$ equals $(\text{NOT } A) \text{ AND } (\text{NOT } B)$

Answer: D

Q1305. How does the nullish coalescing operator (??) differ from the logical OR (||) in JavaScript?

- A. ?? only considers null and undefined as falsy while || considers 0, empty string, and false as falsy too
- B. ?? works only with strings while || works with all types
- C. They behave identically for all input values
- D. || only considers null as falsy while ?? considers all falsy values

Answer: A

Q1306. What is the role of the comma operator in C and JavaScript?

- A. It evaluates both operands from left to right and returns the value of the rightmost operand
- B. It concatenates two strings with a comma separator
- C. It splits a string into multiple parts
- D. It creates an array containing both operands

Answer: A

Q1307. Why is using bit shifting for multiplication and division by powers of two considered an optimization?

- A. Because bit shifting produces more accurate results than multiplication
- B. Because multiplication operators are not supported for integers
- C. Because bit shifting operations are executed in a single CPU cycle while multiplication may take multiple cycles
- D. Because bit shifting works on all data types including strings

Answer: C

Q1308. What problem does the double-checked locking pattern solve using comparison operators?

- A. It compares two objects twice to ensure deep equality
- B. It validates input parameters using two different comparison methods
- C. It reduces the overhead of synchronization by checking a condition before and after acquiring a lock
- D. It performs two sequential comparisons for sorting stability

Answer: C

Q1309. What is the difference between strict equality and abstract equality in type-coercing languages?

- A. Abstract equality never performs type conversion
- B. Strict equality is only available in statically typed languages
- C. Strict equality is slower than abstract equality in all cases
- D. Strict equality compares value and type without coercion while abstract equality converts types before comparing

Answer: D

Q1310. How does the Elvis operator (?:) in Kotlin simplify null handling compared to traditional null checks?

- A. It automatically throws an exception when null is encountered
- B. It provides a default value when the left expression is null in a single concise expression
- C. It converts null values to zero for arithmetic operations
- D. It eliminates all null values from the program

Answer: B

Q1311. How does the visitor pattern implement double dispatch to avoid complex conditional logic?

- A. By using reflection to dynamically invoke methods
- B. By storing all behavior in a single switch statement
- C. By using two virtual method calls to determine behavior based on both the visitor and element types at runtime
- D. By using nested if-else chains with type checking

Answer: C

Q1312. What is the difference between structured exception handling and setjmp/longjmp in C?

- A. Structured exceptions cannot cross function boundaries
- B. setjmp/longjmp automatically free allocated memory during the jump
- C. Structured exceptions use language constructs with proper stack unwinding while setjmp/longjmp perform non-local jumps without cleanup
- D. setjmp/longjmp provide better type safety than structured exceptions

Answer: C

Q1313. How does Rust's match expression enforce exhaustive pattern matching?

- A. It uses a default catch-all case automatically
- B. The compiler requires all possible patterns to be covered, refusing to compile if any case is missing
- C. It only works with integer values
- D. It throws a runtime exception for unmatched patterns

Answer: B

Q1314. What is the difference between cooperative and preemptive multitasking in terms of control flow?

- A. Cooperative multitasking automatically prevents deadlocks
- B. Preemptive multitasking cannot handle more than two tasks
- C. Cooperative multitasking requires tasks to explicitly yield control while preemptive multitasking allows the scheduler to interrupt tasks
- D. Cooperative multitasking is always faster than preemptive

Answer: C

Q1315. How do algebraic effects provide a more composable alternative to traditional control flow mechanisms?

- A. They eliminate all side effects from programs
- B. They replace all loops with recursive functions
- C. They separate the description of an effect from its handler, allowing different interpretations of the same effectful code
- D. They automatically parallelize sequential code

Answer: C

Q1316. What is the purpose of the computed goto (or indirect branch) in low-level optimization?

- A. It computes the result of a mathematical expression before branching
- B. It creates multiple goto statements in a single line
- C. It jumps to a target address computed at runtime, enabling efficient dispatch tables for interpreters
- D. It optimizes recursive function calls by removing the call stack

Answer: C

Q1317. What is the relationship between monadic composition and control flow in functional programming?

- A. Monads eliminate the need for conditional statements
- B. Monads chain computations where each step depends on the previous result, encoding control flow as data transformations
- C. Monads replace all loops with simple function calls
- D. Monads are only used for mathematical computations

Answer: B

Q1318. How does speculative execution in modern CPUs interact with branch prediction for control structures?

- A. The CPU always executes both branches simultaneously
- B. The CPU predicts which branch will be taken and speculatively executes instructions along that path before the condition is resolved
- C. Branch prediction only works with simple if-else structures
- D. Speculative execution skips all conditional checks for performance

Answer: B

Q1319. What is the difference between call/cc (call with current continuation) and standard exception handling?

- A. They are identical mechanisms with different names
- B. call/cc is simpler to implement than exception handling
- C. call/cc captures the entire remaining computation as a first-class value that can be invoked multiple times, while exceptions only unwind the stack
- D. Exception handling can capture continuations but call/cc cannot

Answer: C

Q1320. Why is the halting problem relevant to static analysis of loops and control flow?

- A. It proves that recursive functions always terminate
- B. It shows that all loops will eventually halt if written correctly
- C. It proves that no algorithm can determine in general whether an arbitrary loop will terminate, limiting what static analyzers can guarantee
- D. It only applies to infinite loops in interpreted languages

Answer: C

Q1321. How does memoization improve the performance of recursive functions?

- A. By executing recursive calls in parallel threads
- B. By reducing the number of parameters a function accepts
- C. By converting all recursive calls into iterative loops automatically
- D. By caching previously computed results and returning them for repeated inputs instead of recomputing

Answer: D

Q1322. What is the difference between currying and partial application in functional programming?

- A. Currying fixes arguments while partial application chains functions
- B. They are identical techniques with different names
- C. Currying transforms a function with multiple arguments into a chain of single-argument functions, while partial application fixes some arguments of a function
- D. Currying only works with two arguments

Answer: C

Q1323. What is a continuation and how is it used in functional programming?

- A. A representation of the rest of the computation at a given point, passed as an argument to control execution flow
- B. A function that continues executing after an error
- C. A loop that never terminates
- D. A variable that persists between function calls

Answer: A

Q1324. What is the difference between a free variable and a bound variable in a closure?

- A. Free variables are faster to access than bound variables
- B. Free variables must be global variables
- C. A bound variable is defined within the function while a free variable is defined in an enclosing scope and captured by the closure
- D. Bound variables cannot be modified inside the function

Answer: C

Q1325. What is a functor in functional programming and how does it relate to function application?

- A. A type that implements a map operation, allowing a function to be applied to wrapped values without unwrapping
- B. A debugging tool for tracking function calls
- C. A special type of constructor in OOP
- D. A function that creates other functions

Answer: A

Q1326. How does the trampolining technique prevent stack overflow in languages without tail call optimization?

- A. It increases the stack size dynamically at runtime
- B. It wraps recursive calls in thunks and uses a loop to evaluate them iteratively, keeping the stack flat
- C. It converts all functions to asynchronous operations
- D. It splits recursive calls across multiple threads

Answer: B

Q1327. What is referential transparency and why is it important for reasoning about functions?

- A. It ensures functions can reference variables in any scope
- B. It requires all functions to be visible from any part of the program
- C. It means all function names must clearly describe their purpose
- D. An expression is referentially transparent if it can be replaced with its value without changing program behavior

Answer: D

Q1328. What is the difference between applicative functors and monads in terms of function application?

- A. Applicative functors apply independent wrapped functions to wrapped values while monads allow dependent sequential computations
- B. Applicative functors are always faster than monads
- C. Applicative functors require mutable state
- D. Monads cannot be used for error handling

Answer: A

Q1329. What is a fixed-point combinator and what role does it play in lambda calculus?

- A. It finds the optimal parameters for a function
- B. It converts any function into a constant function
- C. It enables recursion in lambda calculus where functions cannot refer to themselves by name
- D. It fixes errors in function definitions automatically

Answer: C

Q1330. How does defunctionalization transform higher-order functions into first-order programs?

- A. By converting functions into global variables
- B. By removing all function definitions from the program
- C. By replacing function values with data constructors and using an apply function to dispatch on them
- D. By inlining all function calls at compile time

Answer: C

Q1331. How does the proactor pattern differ from the reactor pattern in asynchronous I/O?

- A. The reactor only works with network I/O while the proactor works with file I/O
- B. The proactor blocks on I/O while the reactor does not
- C. They are identical patterns with different names
- D. The reactor demultiplexes events and dispatches to handlers while the proactor initiates operations and handles completion notifications

Answer: D

Q1332. What is the purpose of the sendfile() system call and how does it optimize file transfer?

- A. It sends files as email attachments
- B. It transfers data between file descriptors directly in kernel space without copying to user space
- C. It encrypts files before sending them over the network
- D. It splits large files into smaller chunks for parallel transfer

Answer: B

Q1333. What is the io_uring interface in Linux and why is it significant?

- A. It encrypts I/O data using ring signatures
- B. It is a circular buffer for logging I/O errors
- C. It manages USB device connections in a ring topology
- D. It provides a high-performance asynchronous I/O interface using shared ring buffers between user and kernel space

Answer: D

Q1334. What is the write amplification problem in storage I/O?

- A. When the actual amount of data written to storage is significantly larger than the logical data being written
- B. When multiple programs write to the same file simultaneously
- C. When write operations are slower than read operations
- D. When data is written in the wrong encoding format

Answer: A

Q1335. How does memory-mapped I/O differ from port-mapped I/O at the hardware level?

- A. Port-mapped I/O uses the same address space as main memory
- B. Memory-mapped I/O uses the same address space for memory and I/O devices while port-mapped I/O uses a separate address space
- C. Memory-mapped I/O is slower in all cases
- D. Memory-mapped I/O is only available on ARM processors

Answer: B

Q1336. What is the purpose of the O_DIRECT flag when opening files on Linux?

- A. It opens the file in a special directory
- B. It encrypts all data written to the file
- C. It bypasses the operating system page cache and performs direct I/O between user buffers and the storage device
- D. It directs output to the console instead of the file

Answer: C

Q1337. What is the difference between level-triggered and edge-triggered I/O notification?

- A. Level-triggered reports readiness as long as the condition holds while edge-triggered reports only when the state changes
- B. Edge-triggered is always more efficient than level-triggered
- C. Level-triggered only works with read operations
- D. Edge-triggered requires blocking I/O

Answer: A

Q1338. What is the purpose of vectored I/O (scatter-gather) and when is it beneficial?

- A. It reads or writes data from multiple non-contiguous buffers in a single system call, reducing overhead
- B. It gathers I/O statistics from multiple processes
- C. It splits a single buffer into multiple files
- D. It encrypts data by scattering bits across storage

Answer: A

Q1339. What is the durability guarantee provided by fsync versus fdatasync?

- A. fsync flushes both file data and metadata to disk while fdatasync flushes only the file data and essential metadata
- B. They provide identical durability guarantees in all cases
- C. fdatasync provides stronger guarantees than fsync
- D. fsync only flushes metadata while fdatasync flushes data

Answer: A

Q1340. What is the thundering herd problem in the context of I/O multiplexing with accept()?

- A. When multiple threads waiting on the same listening socket all wake up for a single connection, but only one succeeds
- B. When too many connections arrive simultaneously and crash the server
- C. When I/O buffers overflow due to rapid data arrival
- D. When multiple files are opened simultaneously causing disk contention

Answer: A

Q1341. What is the Aho-Corasick algorithm and when is it used for string matching?

- A. It searches for multiple patterns simultaneously in a text by building an automaton from all patterns
- B. It sorts strings in alphabetical order efficiently
- C. It matches a single pattern against multiple texts
- D. It compresses strings using pattern-based encoding

Answer: A

Q1342. What is the suffix array data structure and how does it improve string operations?

- A. A linked list of all string prefixes
- B. An array that stores only the last character of each string
- C. A hash table mapping string suffixes to their positions
- D. A sorted array of all suffixes of a string that enables efficient substring search, longest common prefix computation, and other operations

Answer: D

Q1343. How does the Burrows-Wheeler Transform improve text compression?

- A. It rearranges characters so that similar characters are grouped together, making the result more compressible
- B. It removes duplicate characters from the text
- C. It encrypts the text to reduce its size
- D. It replaces all characters with shorter binary codes

Answer: A

Q1344. What is the difference between greedy and lazy quantifiers in regular expressions?

- A. Greedy quantifiers match as much text as possible while lazy quantifiers match as little as possible
- B. Lazy quantifiers match more text than greedy quantifiers
- C. They produce identical matches in all cases
- D. Greedy quantifiers are faster than lazy quantifiers

Answer: A

Q1345. What is catastrophic backtracking in regular expressions?

- A. When a regex engine takes exponential time due to ambiguous patterns that cause excessive backtracking
- B. When backtracking causes the program to run out of memory
- C. When a regex accidentally matches the wrong string
- D. When a regex fails to match any string

Answer: A

Q1346. What is the Z-algorithm for string matching?

- A. It constructs a Z-array where $Z[i]$ is the length of the longest substring starting at position i that matches a prefix of the string
- B. It finds the last occurrence of a pattern in a string
- C. It sorts characters in a string alphabetically
- D. It compresses strings using a zigzag encoding

Answer: A

Q1347. How do persistent string data structures maintain previous versions efficiently?

- A. By compressing all previous versions into a single archive
- B. By keeping complete copies of every version in memory
- C. By storing only the most recent version and discarding older ones
- D. By sharing unchanged portions between versions and only allocating new memory for changed parts

Answer: D

Q1348. What is Unicode normalization and why can two visually identical strings have different byte representations?

- A. Unicode only affects non-English text
- B. Different byte representations indicate file corruption
- C. Normalization removes all special characters from a string
- D. Characters can be represented as single code points or as combining character sequences, and normalization converts between these forms

Answer: D

Q1349. What is the difference between DFA and NFA-based regex engines?

- A. DFA engines are always slower than NFA engines
- B. DFA engines guarantee linear time matching but cannot support backreferences while NFA engines support backreferences but may have exponential worst case
- C. NFA engines cannot match basic patterns
- D. DFA engines support all regex features including backreferences

Answer: B

Q1350. What is the longest common subsequence (LCS) problem and its string application?

- A. Determining which string is longer
- B. Finding the longest palindrome within a string
- C. Finding the longest sequence of characters common to two strings in the same order but not necessarily contiguous, used in diff algorithms
- D. Finding the longest substring that appears in both strings at the same position

Answer: C

Q1351. What is a cuckoo hash table and how does it achieve $O(1)$ worst-case lookups?

- A. It stores elements in sorted order for binary search
- B. It uses a single hash function with no collisions
- C. It uses two hash functions and two tables, displacing existing elements on collision, guaranteeing each key is in one of two possible locations
- D. It uses three hash functions and three tables

Answer: C

Q1352. What is the difference between an AVL tree and a red-black tree in terms of balancing?

- A. They maintain identical balance guarantees
- B. AVL trees maintain stricter balance with at most 1 height difference, while red-black trees allow up to $2x$ height difference, making AVL faster for lookups but slower for insertions
- C. AVL trees do not support deletion operations
- D. Red-black trees are always perfectly balanced

Answer: B

Q1353. What is a Fibonacci heap and why is it used in graph algorithms?

- A. A heap with exactly 5 children per node
- B. A heap used only for sorting Fibonacci sequences
- C. A heap that stores Fibonacci numbers
- D. A heap with amortized $O(1)$ decrease-key operation, making it optimal for algorithms like Dijkstra's shortest path

Answer: D

Q1354. What is the difference between an R-tree and a B-tree for indexing?

- A. B-trees can index spatial data better than R-trees
- B. R-trees index multi-dimensional spatial data using bounding rectangles while B-trees index one-dimensional ordered data
- C. R-trees are used for text search while B-trees are for spatial data
- D. They provide identical functionality for all data types

Answer: B

Q1355. How does a count-min sketch provide approximate frequency counts with bounded error?

- A. It uses multiple hash functions mapping to a matrix of counters, providing an upper bound estimate with configurable error probability
- B. It stores only the most frequent elements exactly
- C. It counts exactly by using unlimited memory
- D. It compresses all data into a single counter

Answer: A

Q1356. What is a treap and how does it combine properties of trees and heaps?

- A. A special array that combines tree and heap indices
- B. A heap that supports tree traversal operations
- C. A tree that stores elements in heap order only
- D. A BST where each node has a random priority, maintaining BST property on keys and heap property on priorities for probabilistic balance

Answer: D

Q1357. What is the difference between LSM trees and B-trees for write-heavy workloads?

- A. LSM trees do not support read operations
- B. They perform identically for all workload types
- C. B-trees are always faster than LSM trees for writes
- D. LSM trees buffer writes in memory and merge sorted runs sequentially, optimizing write throughput, while B-trees update pages in place with random I/O

Answer: D

Q1358. What is a HAT-trie and why is it more cache-friendly than a standard trie?

- A. It stores multiple characters per node in contiguous arrays, reducing pointer chasing and improving CPU cache utilization
- B. It stores the entire trie in a single array
- C. It hashes all keys before insertion into the trie
- D. It uses hardware acceleration for trie operations

Answer: A

Q1359. What is the difference between wait-free and lock-free concurrent data structures?

- A. They provide identical progress guarantees
- B. Wait-free structures are slower than lock-free structures always
- C. Lock-free guarantees system-wide progress while wait-free guarantees per-thread progress in bounded steps
- D. Lock-free structures use locks internally while wait-free does not

Answer: C

Q1360. What is a CRDT (Conflict-free Replicated Data Type) and why is it important for distributed collections?

- A. A compressed data type for reducing network bandwidth
- B. A data type that prevents all conflicts by using locks
- C. A data type designed so that concurrent updates on different replicas always converge to the same state without coordination
- D. A data type that replicates data by copying entire collections

Answer: C

Q1361. How does the bridge pattern separate abstraction from implementation?

- A. By using two separate inheritance hierarchies connected by composition, allowing them to vary independently
- B. By merging all implementations into a single abstract class
- C. By using global variables to connect unrelated classes
- D. By creating a bridge between two interfaces using static methods

Answer: A

Q1362. What is the Liskov Substitution Principle's impact on exception handling in subclasses?

- A. Subclasses must throw the exact same exceptions as the superclass
- B. Subclasses should never throw any exceptions
- C. Subclass methods should not throw broader exceptions than the superclass method specifies to maintain substitutability
- D. The LSP does not apply to exception handling

Answer: C

Q1363. What is the difference between the state pattern and the strategy pattern?

- A. The state pattern is for concurrency while the strategy pattern is for sorting
- B. They are identical patterns with different names
- C. The state pattern encapsulates state-dependent behavior with transitions between states while the strategy pattern selects an algorithm without transitions
- D. The strategy pattern manages state transitions while the state pattern selects algorithms

Answer: C

Q1364. What is the object pool pattern and when is it beneficial?

- A. A pattern that distributes objects across multiple servers
- B. A pattern that creates objects in a swimming pool metaphor
- C. A pattern that pools all objects into a single collection for easy access
- D. A pattern that maintains a pool of reusable objects to avoid expensive creation and destruction costs

Answer: D

Q1365. What is the problem with inheritance-based reuse known as the 'yo-yo problem'?

- A. When multiple inheritance creates an infinite loop
- B. When understanding a class requires tracing through many levels of inheritance, bouncing up and down the hierarchy
- C. When objects repeatedly create and destroy themselves
- D. When a class inherits from both a parent and grandchild class

Answer: B

Q1366. How does the specification pattern encapsulate business rules as recombining objects?

- A. By storing rules as string expressions evaluated at runtime
- B. By hardcoding rules into if-else chains
- C. By writing all business rules as database queries
- D. By representing each business rule as an object with an isSatisfiedBy method that can be combined using AND, OR, and NOT operations

Answer: D

Q1367. What is the difference between type classes in Haskell and interfaces in Java?

- A. Type classes are slower than interfaces
- B. They are identical concepts in different languages
- C. Java interfaces support more features than type classes
- D. Type classes can add behavior to existing types without modifying them while Java interfaces require implementation at class definition time

Answer: D

Q1368. What is the memento pattern and how does it support undo functionality?

- A. It stores all objects in persistent memory
- B. It creates memory snapshots of the entire application
- C. It memorizes method call history in a static log
- D. It captures an object's internal state in a separate memento object that can be used to restore the state later

Answer: D

Q1369. How does the entity-component-system (ECS) architecture differ from traditional OOP?

- A. ECS requires all entities to share the same base class
- B. ECS is just another name for the MVC pattern
- C. ECS uses deep inheritance hierarchies for entity behavior
- D. ECS separates data (components) from behavior (systems) and uses composition over inheritance, favoring cache-friendly memory layouts

Answer: D

Q1370. What is the expression problem and how do OOP and functional programming approach it differently?

- A. OOP cannot handle any expressions while FP handles all expressions
- B. Both paradigms solve the expression problem identically
- C. FP cannot define new types while OOP cannot define new functions
- D. OOP easily adds new types but struggles to add new operations, while FP easily adds new operations but struggles to add new types

Answer: D

Q1371. What is the tri-color marking algorithm in garbage collection?

- A. It colors memory pages for security purposes
- B. It marks objects based on their data type using three categories
- C. It uses three passes to determine which objects to collect
- D. It classifies objects as white (unreachable), gray (to be scanned), and black (scanned) to safely identify garbage during concurrent collection

Answer: D

Q1372. How does the Boehm-Demers-Weiser conservative garbage collector work for C programs?

- A. It requires the programmer to annotate all pointers
- B. It only collects objects on the stack
- C. It converts C programs to Java for garbage collection
- D. It treats any value in memory that could be a pointer as a potential reference, preventing collection of possibly-referenced objects

Answer: D

Q1373. What is the difference between stop-the-world and concurrent garbage collection?

- A. Concurrent GC always takes longer than stop-the-world
- B. Stop-the-world only pauses one thread at a time
- C. They produce identical pause times in all scenarios
- D. Stop-the-world pauses all application threads during collection while concurrent GC runs alongside the application with minimal pauses

Answer: D

Q1374. What is the ZGC garbage collector in Java designed to achieve?

- A. Maximum throughput at the expense of longer pauses
- B. Sub-millisecond pause times regardless of heap size by using colored pointers and load barriers
- C. Smaller memory footprint than serial GC
- D. Compatibility with C code through JNI

Answer: B

Q1375. What is the purpose of Rust's borrow checker?

- A. It enforces rules at compile time ensuring references do not outlive the data they point to and preventing data races
- B. It monitors runtime memory usage for debugging
- C. It automatically frees memory when borrowing ends
- D. It checks if the program borrows external libraries correctly

Answer: A

Q1376. What is the difference between arena allocation and general-purpose heap allocation?

- A. Arena allocation is slower than heap allocation for all use cases
- B. Arena allocation requires garbage collection to function
- C. Arena allocation groups objects with similar lifetimes and frees them all at once, avoiding individual deallocation overhead
- D. General-purpose allocation always uses less memory than arena allocation

Answer: C

Q1377. What is the Cheney's algorithm for copying garbage collection?

- A. It copies garbage to a separate memory region for analysis
- B. It creates backup copies of all objects for safety
- C. It uses three memory spaces in a rotating pattern
- D. It uses two equal-sized semi-spaces, copying live objects from one to the other, compacting memory and eliminating fragmentation

Answer: D

Q1378. What is the problem of false sharing in concurrent programs and how is it mitigated?

- A. When multiple threads allocate memory from the same pool
- B. When different threads modify variables on the same cache line, causing unnecessary cache invalidation and performance degradation
- C. When shared memory is falsely reported as freed
- D. When threads share incorrect data due to synchronization bugs

Answer: B

Q1379. What is the Shenandoah GC's approach to concurrent compaction?

- A. It compacts memory only during application idle time
- B. It uses stop-the-world pauses for all compaction operations
- C. It relies on the operating system for compaction
- D. It uses Brooks forwarding pointers to redirect references to relocated objects, enabling compaction concurrent with the application

Answer: D

Q1380. What is the concept of escape analysis and how does it optimize memory allocation?

- A. It analyzes memory escape routes during buffer overflow attacks
- B. It determines when variables escape from their declared scope
- C. It detects when objects escape the program's memory space
- D. It analyzes whether an object escapes the method or thread that created it, enabling stack allocation or elimination for non-escaping objects

Answer: D

Q1381. What is the difference between the Result type in Rust and traditional exception handling?

- A. Result encodes success or failure in the return type, forcing callers to handle errors explicitly at compile time, while exceptions propagate implicitly
- B. Result types can only represent two types of errors
- C. Exceptions provide better type safety than Result
- D. Result types are slower than exception handling

Answer: A

Q1382. What is the concept of exception safety in C++ and its three guarantee levels?

- A. Low, medium, and high exception severity levels
- B. Compile-time safety, runtime safety, and thread safety
- C. Type safety, memory safety, and null safety
- D. Basic guarantee (no leaks), strong guarantee (transaction-like rollback), and no-throw guarantee (never throws exceptions)

Answer: D

Q1383. How does Kotlin's approach to checked exceptions differ from Java's?

- A. Kotlin uses a different keyword for checked exceptions
- B. Kotlin requires all exceptions to be checked
- C. Kotlin only allows checked exceptions in coroutines
- D. Kotlin does not have checked exceptions, treating all exceptions as unchecked, because studies showed checked exceptions often lead to empty catch blocks

Answer: D

Q1384. What is the relationship between exceptions and monadic error handling using Either/Result?

- A. Monadic error handling replaces all try-catch blocks with recursion
- B. Exceptions are a subset of monadic error handling
- C. Monadic error handling cannot represent multiple error types
- D. Both represent computations that may fail, but monadic error handling makes failure explicit in types and composable through flatMap/bind

Answer: D

Q1385. What is the zero-cost exception handling model used in modern C++ compilers?

- A. It eliminates all exception-related code at compile time
- B. Exceptions are completely free even when thrown
- C. It replaces exceptions with return codes automatically
- D. It adds no performance overhead when exceptions are not thrown, using table-based stack unwinding only when exceptions occur

Answer: D

Q1386. What is the problem with exceptions in asynchronous or concurrent code?

- A. Exceptions cannot cross thread or async boundaries naturally and may be lost if not properly captured and propagated
- B. Concurrent code cannot throw any exceptions
- C. Exceptions are always thread-safe and work identically in concurrent code
- D. Exceptions in threads automatically propagate to the main thread

Answer: A

Q1387. How does structured concurrency improve exception handling compared to unstructured async code?

- A. It handles exceptions by restarting failed tasks indefinitely
- B. It ensures that all concurrent tasks are scoped to a parent, and exceptions from child tasks automatically cancel siblings and propagate to the parent
- C. It defers all exceptions until the program exits
- D. It eliminates all exceptions in concurrent code

Answer: B

Q1388. What is the effect handler approach to exception handling and how does it generalize try-catch?

- A. Effect handlers define how effects (including exceptions) are interpreted, allowing the same code to have different error handling strategies depending on the handler
- B. They add hardware-level exception handling support
- C. Effect handlers only work in statically typed languages
- D. Effect handlers replace all exception types with a single error type

Answer: A

Q1389. What is the supervisor pattern in actor systems for handling failures?

- A. A thread monitor that prevents all exceptions from occurring
- B. A design pattern that replaces all try-catch with logging
- C. A parent actor monitors child actors and applies a restart, stop, or escalate strategy when a child fails with an exception
- D. A senior developer supervises code reviews for exception handling

Answer: C

Q1390. What is the module resolution algorithm in Node.js and how does it find modules?

- A. It uses a random search through all directories
- B. It only searches the current working directory
- C. It always looks in the global installation directory first
- D. It searches node_modules directories ascending from the requiring file's location, then global locations, checking for files, directories with index.js, or package.json main fields

Answer: D

Q1391. What is the purpose of import maps in modern JavaScript?

- A. They provide a map visualization of module dependencies
- B. They create documentation maps for modules
- C. They allow controlling how module specifiers are resolved in the browser without a build step
- D. They map imports to different programming languages

Answer: C

Q1392. What is the difference between side-effect-free and side-effectful module imports?

- A. Side-effect-free imports are slower
- B. They produce identical behavior in all bundlers
- C. Side-effectful imports cannot export any values
- D. Side-effect-free imports only use the exported values while side-effectful imports execute code that modifies global state on import

Answer: D

Q1393. What is the challenge of diamond dependency conflicts in package management?

- A. When a package depends on itself recursively
- B. When two packages have identical names in different registries
- C. When a dependency graph forms a perfect diamond shape
- D. When two dependencies require different incompatible versions of the same transitive dependency

Answer: D

Q1394. How does the Java Platform Module System (JPMS) differ from traditional classpath-based loading?

- A. Classpath provides stronger encapsulation than JPMS
- B. JPMS uses module declarations to explicitly define dependencies and encapsulated packages, unlike the flat classpath that exposes everything
- C. JPMS is slower than classpath loading
- D. JPMS only works with Java 7 and earlier

Answer: B

Q1395. What is supply chain security in the context of package management?

- A. Protecting against malicious code injected into dependencies through compromised packages or maintainer accounts
- B. Managing the physical supply chain of hardware components
- C. Ensuring packages are delivered in the correct order
- D. Securing the network connection used to download packages

Answer: A

Q1396. What is the purpose of a Software Bill of Materials (SBOM) for packages?

- A. It describes the hardware requirements for the software
- B. It lists the financial cost of each dependency
- C. It provides a comprehensive inventory of all components, libraries, and dependencies in a software project for security and compliance
- D. It documents the development timeline of each package

Answer: C

Q1397. How does Deno's module system differ from Node.js in terms of dependency management?

- A. Deno uses the same npm and node_modules system as Node.js
- B. Deno requires all dependencies to be bundled before execution
- C. Deno does not support external modules
- D. Deno imports modules directly from URLs without a centralized package manager or node_modules directory

Answer: D

Q1398. What is the purpose of module federation in webpack 5?

- A. Allowing multiple independently deployed applications to share modules at runtime without rebuilding
- B. Federating modules across different government agencies
- C. Merging all modules into a single federated package
- D. Creating federal backups of all modules

Answer: A

Q1399. What is the difference between vendoring and using a package manager for dependencies?

- A. Vendoring copies dependency source code directly into the project repository while package managers download them on demand
- B. Package managers copy all code into the repository
- C. Vendoring only works with interpreted languages
- D. Vendoring is always faster than using a package manager

Answer: A

Q1400. What is the difference between the Abstract Factory and Factory Method patterns?

- A. Abstract Factory is a simpler version of Factory Method
- B. Abstract Factory creates families of related objects through an interface while Factory Method delegates creation of a single object to subclasses
- C. They are identical patterns used in different programming languages
- D. Factory Method creates multiple objects while Abstract Factory creates one

Answer: B

Q1401. How does the CQRS pattern improve scalability in distributed systems?

- A. By combining all queries and commands into a single optimized model
- B. By caching all queries on the client side
- C. By separating read and write models, allowing each to be scaled, optimized, and evolved independently
- D. By reducing the number of database tables

Answer: C

Q1402. What is the difference between compile-time and runtime metaprogramming?

- A. Compile-time metaprogramming is always slower
- B. They both execute at the same phase
- C. Compile-time metaprogramming generates or transforms code during compilation while runtime metaprogramming modifies program structure during execution
- D. Runtime metaprogramming cannot create new functions

Answer: C

Q1403. What is the difference between event sourcing and traditional state storage?

- A. Event sourcing stores every state change as an immutable event, deriving current state by replaying events, while traditional storage saves only current state
- B. Traditional storage keeps complete history while event sourcing does not
- C. They store data in identical formats
- D. Event sourcing only stores the final state

Answer: A

Q1404. How does the saga pattern manage distributed transactions across microservices?

- A. It retries all failed transactions indefinitely
- B. It uses a single distributed lock across all services
- C. It orchestrates a sequence of local transactions with compensating actions to undo completed steps if a later step fails
- D. It rolls back all services using two-phase commit

Answer: C

Q1405. What is the difference between the interpreter pattern and an actual language interpreter?

- A. The interpreter pattern compiles code while a language interpreter interprets it
- B. The interpreter pattern defines a grammar representation and an interpreter for sentences in that grammar, used for simple DSLs, while a language interpreter is a full program execution engine
- C. A language interpreter is simpler than the interpreter pattern
- D. They are identical in scope and implementation

Answer: B

Q1406. What is the purpose of the circuit breaker pattern in distributed systems?

- A. It manages electrical circuits in server hardware
- B. It circuits all requests through a central router
- C. It breaks circular dependencies between services
- D. It prevents cascading failures by stopping requests to a failing service and allowing it time to recover

Answer: D

Q1407. What is the difference between a red-black tree and a 2-3-4 tree?

- A. A red-black tree is a binary representation of a 2-3-4 tree where red links represent nodes merged within 2-3-4 tree nodes
- B. Red-black trees have more nodes than 2-3-4 trees for the same data
- C. They are completely unrelated data structures
- D. 2-3-4 trees are always faster for all operations

Answer: A

Q1408. What is backpressure in reactive programming and why is it important?

- A. Pressure applied to undo the last operation
- B. Network congestion that slows down all requests
- C. Memory pressure from running out of heap space
- D. A mechanism that allows consumers to signal producers to slow down when they cannot keep up with the data flow rate

Answer: D

Q1409. What is the difference between optimistic and pessimistic concurrency control in application design?

- A. They are identical strategies with different names
- B. Pessimistic control never uses any locks
- C. Optimistic control assumes conflicts are rare and detects them at commit time, while pessimistic control prevents conflicts by acquiring locks before access
- D. Optimistic control is always faster than pessimistic control

Answer: C

Q1410. What is the difference between lock-free and wait-free algorithms?

- A. Lock-free algorithms are always faster than wait-free
- B. Lock-free guarantees system-wide progress while wait-free guarantees bounded completion for every thread
- C. They provide identical progress guarantees
- D. Wait-free algorithms use locks while lock-free algorithms do not

Answer: B

Q1411. What is the epoch-based reclamation (EBR) technique for concurrent memory management?

- A. It reclaims memory at fixed time intervals
- B. It creates backup copies of memory at each epoch
- C. It uses epoch timestamps to sort memory allocations
- D. It divides time into epochs and defers memory reclamation until all threads have advanced past the epoch when the memory was freed

Answer: D

Q1412. How does the Disruptor pattern achieve high-performance inter-thread communication?

- A. It disrupts thread scheduling to gain priority
- B. It uses a lock-free ring buffer with sequencing to enable producer-consumer communication with mechanical sympathy for CPU caches
- C. It uses multiple locks for each data element
- D. It stores all data in a shared database

Answer: B

Q1413. What is the difference between linearizability and serializability?

- A. They are identical consistency models
- B. Linearizability requires operations to appear atomic at a single point in real time while serializability requires the result to be equivalent to some serial execution
- C. Serializability is stronger than linearizability
- D. Linearizability only applies to single-threaded programs

Answer: B

Q1414. What is the purpose of hazard pointers in lock-free data structures?

- A. They mark memory locations that are unsafe to write to
- B. They protect nodes from being reclaimed while other threads still reference them, solving the safe memory reclamation problem
- C. They point to hazardous code sections that should not be executed
- D. They point to threads that may cause deadlocks

Answer: B

Q1415. What is the difference between strong and weak memory ordering on modern CPUs?

- A. Strong ordering prevents all concurrency bugs automatically
- B. Weak ordering is always faster than strong ordering
- C. Strong ordering guarantees operations appear in program order to all threads while weak ordering allows reordering for performance, requiring explicit barriers
- D. They only affect floating-point operations

Answer: C

Q1416. What is structured concurrency and why is it considered safer than unstructured thread creation?

- A. It requires all threads to run the same structured code
- B. It only allows creating threads in specific code structures
- C. It ensures concurrent tasks are scoped to a defined lifetime, guaranteeing all tasks complete or are cancelled before the scope exits
- D. It structures threads in a hierarchical tree permanently

Answer: C

Q1417. What is the consensus problem in distributed concurrent systems?

- A. Agreeing on the number of threads to allocate
- B. Getting multiple processes to agree on a single value despite failures, foundational for distributed coordination
- C. Getting all stakeholders to agree on the software architecture
- D. Reaching consensus on which programming language to use

Answer: B

Q1418. What is the FLP impossibility result and its significance for concurrent systems?

- A. It demonstrates that threads cannot share memory safely
- B. It shows that lock-free algorithms are impossible
- C. It proves that deterministic consensus is impossible in an asynchronous system where even one process can crash
- D. It proves that all concurrent programs eventually deadlock

Answer: C

Q1419. What is the difference between virtual threads (Project Loom) and platform threads in Java?

- A. Virtual threads cannot perform I/O operations
- B. Platform threads use less memory than virtual threads
- C. Virtual threads are lightweight user-mode threads managed by the JVM that can be created in millions, while platform threads are OS threads with higher overhead
- D. Virtual threads are faster for all workloads

Answer: C

Q1420. What is concolic testing and how does it combine concrete and symbolic execution?

- A. It tests concurrent code with symbolic debuggers
- B. It uses concrete test data and symbolic names for variables
- C. It runs concrete execution alongside symbolic execution, using concrete values to guide path exploration and symbolic constraints for input generation
- D. It combines console output testing with logic testing

Answer: C

Q1421. What is the difference between deterministic and non-deterministic testing of concurrent programs?

- A. They produce identical results for concurrent code
- B. Deterministic testing controls thread scheduling to explore specific interleavings while non-deterministic relies on random scheduling
- C. Non-deterministic testing is more thorough than deterministic
- D. Deterministic tests always pass while non-deterministic may fail

Answer: B

Q1422. What is metamorphic testing and when is it useful?

- A. Testing where outputs of related test inputs are compared using known relationships when expected outputs are hard to determine
- B. Testing the metadata of test files
- C. Testing morphological transformations of data structures
- D. Testing that code changes (morphs) correctly across versions

Answer: A

Q1423. How does reverse debugging (time-travel debugging) work?

- A. It debugs code before it is compiled
- B. It reverses the order of test execution
- C. It runs the program in reverse to undo changes
- D. It records the entire program execution state, allowing developers to step backward through execution to find the root cause of bugs

Answer: D

Q1424. What is the difference between generation-based and mutation-based fuzzing?

- A. Generation-based fuzzing creates inputs from grammar specifications while mutation-based fuzzing modifies existing valid inputs
- B. They produce identical test coverage
- C. Mutation-based fuzzing requires manual test case writing
- D. Generation-based fuzzing is always faster

Answer: A

Q1425. What is the purpose of formal verification compared to testing?

- A. Formal verification is faster than testing
- B. Testing provides stronger guarantees than formal verification
- C. Formal verification replaces the need for all testing
- D. Formal verification mathematically proves correctness for all possible inputs while testing can only check a finite number of cases

Answer: D

Q1426. What is coverage-guided fuzzing and how does it improve upon random fuzzing?

- A. It focuses only on covering previously tested paths
- B. It uses code coverage feedback to guide input mutations toward unexplored code paths, achieving deeper code coverage
- C. It guides tests toward code that already has high coverage
- D. It reduces coverage to speed up fuzzing

Answer: B

Q1427. What is the oracle problem in software testing?

- A. The difficulty of choosing the right testing framework
- B. A problem with Oracle database testing
- C. The challenge of determining the correct expected output for a given test input, especially for complex systems
- D. A problem with predicting which tests will fail

Answer: C

Q1428. What is the difference between white-box fuzzing and black-box fuzzing?

- A. White-box fuzzing only tests visible code while black-box tests hidden code
- B. They explore the same paths using different input formats
- C. White-box fuzzing uses program analysis to systematically generate inputs that explore specific paths while black-box fuzzing treats the program as opaque
- D. Black-box fuzzing requires source code access while white-box does not

Answer: C

Q1429. What is the purpose of differential testing (differential fuzzing)?

- A. Testing the difference between old and new code versions only
- B. Testing the performance difference between two algorithms
- C. Testing equivalent implementations of the same specification against each other to find discrepancies indicating bugs
- D. Testing differences between development and production environments

Answer: C

Q1430. What is the purpose of chaos engineering and how does Netflix's Chaos Monkey work?

- A. Chaos Monkey introduces syntax errors into code
- B. It randomly assigns developers to different teams
- C. Chaos Monkey randomly terminates production instances to verify that the system can tolerate failures, building confidence in resilience
- D. It chaotically reorders database records for testing

Answer: C

Q1431. What is the difference between observability and monitoring?

- A. They are identical practices with different names
- B. Monitoring tracks predefined metrics and alerts while observability enables exploring and understanding system behavior from its outputs without predefined queries
- C. Monitoring provides more insight than observability
- D. Observability is simpler than monitoring

Answer: B

Q1432. How does GitOps extend DevOps practices using Git as the single source of truth?

- A. GitOps uses Git repositories as the declarative source of truth for infrastructure and application configuration, with automated reconciliation
- B. GitOps eliminates the need for CI/CD pipelines
- C. GitOps only applies to frontend development
- D. GitOps replaces Git with a new version control system

Answer: A

Q1433. What is the purpose of a service mesh and how does it relate to microservice development?

- A. It creates a mesh network between physical servers
- B. It provides a dedicated infrastructure layer handling service-to-service communication including load balancing, encryption, and observability without changing application code
- C. It provides networking hardware for services
- D. It meshes all services into a single application

Answer: B

Q1434. What is the concept of developer experience (DX) and why does it matter?

- A. DX focuses on making tools, APIs, documentation, and workflows efficient and pleasant for developers, directly impacting productivity and code quality
- B. DX only applies to designing developer-facing UI
- C. DX measures how many years of experience a developer has
- D. DX is about decorating the developer workspace

Answer: A

Q1435. What is the difference between Platform Engineering and traditional DevOps?

- A. Platform Engineering builds internal developer platforms as a product to enable self-service capabilities, while DevOps focuses on culture and practices for collaboration
- B. DevOps provides self-service while Platform Engineering requires manual ops
- C. They are identical approaches to infrastructure management
- D. Platform Engineering replaces DevOps entirely

Answer: A

Q1436. What is the DORA metrics framework for measuring software delivery performance?

- A. It measures developer satisfaction and code quality only
- B. It measures server uptime and network latency
- C. It measures the number of lines of code written per day
- D. It measures deployment frequency, lead time for changes, change failure rate, and time to restore service

Answer: D

Q1437. What is the purpose of dependency injection containers and how do they implement inversion of control?

- A. They invert the direction of data flow in the application
- B. They contain all project dependencies in a single file
- C. They inject security vulnerabilities for penetration testing
- D. They manage object creation and lifecycle, automatically resolving and injecting dependencies based on configuration, inverting control of dependency creation from application code to the container

Answer: D

Q1438. What is the concept of immutable infrastructure?

- A. Infrastructure that cannot be deleted once created
- B. Infrastructure that uses immutable data structures only
- C. Infrastructure where servers are never modified after deployment; changes require building and deploying entirely new instances
- D. Infrastructure that never changes its IP addresses

Answer: C

Q1439. What is the difference between vertical and horizontal scaling and when is each appropriate?

- A. Horizontal scaling requires a single powerful server
- B. Vertical scaling is always better than horizontal
- C. Vertical scaling adds more resources to a single machine while horizontal scaling adds more machines, with horizontal being better for distributed workloads and vertical for simpler architectures
- D. They achieve identical results in all scenarios

Answer: C

Q1440. How does the error kernel pattern improve fault tolerance in distributed systems?

- A. It creates a kernel-level exception handler for all processes
- B. It prevents all errors by running code in a secure kernel
- C. It logs all errors to a central kernel database
- D. It isolates error-prone code in separate processes supervised by a minimal, highly reliable kernel that handles restarts

Answer: D

SUBJECT

Software Engineering

1,440 Multiple Choice Questions

Software Engineering — Easy

480 questions

Q1. What does software engineering primarily deal with?

- A. Systematic software development
- B. Configuring computer network infrastructure
- C. Managing enterprise database systems
- D. Building and maintaining hardware systems

Answer: A

Q2. Who is considered the first person to use the term 'software engineering'?

- A. Alan Turing
- B. Edsger Dijkstra
- C. Dennis Ritchie
- D. Margaret Hamilton

Answer: D

Q3. Which of the following is NOT a characteristic of good software?

- A. Dependability
- B. Complexity
- C. Maintainability
- D. Efficiency

Answer: B

Q4. What is the main goal of software engineering?

- A. To incorporate the latest emerging technology
- B. To minimize the overall development team size
- C. To produce quality software within budget and schedule
- D. To write source code as quickly as possible

Answer: C

Q5. Software engineering is a branch of which field?

- A. Computer science and engineering
- B. Electrical engineering and circuits
- C. Mechanical engineering and design
- D. Civil engineering and structures

Answer: A

Q6. What does the term 'software crisis' refer to?

- A. Rapidly increasing hardware component costs
- B. Insufficient programming language features
- C. Problems in developing large-scale software
- D. A critical shortage of trained programmers

Answer: C

Q7. Which of the following is a software product?

- A. Operating System
- B. Mechanical keyboard
- C. Monitor display unit
- D. Laser printer device

Answer: A

Q8. What is a software process?

- A. A standalone computer program in the development
- B. A relational database system within the project scope
- C. A set of activities for software development
- D. A physical hardware component during the software

Answer: C

Q9. Which attribute is essential for software quality?

- A. Large file size
- B. Complex algorithms
- C. Reliability
- D. Maximum features

Answer: C

Q10. What is the purpose of software documentation?

- A. To increase overall project file size
- B. To help understand and maintain the software
- C. To make onboarding harder for new hires
- D. To intentionally slow down the development

Answer: B

Q11. What is the Waterfall model?

- A. A circular iterative development process
- B. A random unstructured coding process
- C. An agile methodology with sprint cycles
- D. A linear sequential software development approach

Answer: D

Q12. Which phase comes first in the Waterfall model?

- A. Implementation phase
- B. Testing the system
- C. Requirements analysis
- D. Deployment to users

Answer: C

Q13. What is an iterative process model?

- A. A model that omits all testing phases
- B. A model that executes only a single run
- C. A hardware-focused design model only
- D. A model that repeats development cycles

Answer: D

Q14. What does the 'V-Model' emphasize?

- A. Only source code implementation tasks
- B. Hardware testing procedures and checks
- C. Only requirements documentation writing
- D. Verification and validation at each development stage

Answer: D

Q15. In which model are requirements expected to be fully known at the start?

- A. Spiral
- B. Prototype
- C. Waterfall
- D. Agile

Answer: C

Q16. What is a software development life cycle (SDLC)?

- A. The process of planning, creating, testing, and deploying software
- B. A physical computing hardware component unit in all development efforts
- C. A specific type of programming language syntax for all project stakeholders
- D. A specialized category of database system type within the given constraints

Answer: A

Q17. Which model is also known as the 'Classic Life Cycle'?

- A. Waterfall
- B. RAD to achi
- C. Spiral by
- D. Agile and

Answer: A

Q18. What is prototyping in software development?

- A. Building a preliminary version to explore requirements
- B. Delivering the final production product in the engineering
- C. Performing hardware acceptance testing as part of the methodology
- D. Removing outdated legacy source code for effective project outcomes

Answer: A

Q19. Which process model involves risk analysis in each iteration?

- A. RAD
- B. Spiral
- C. V-Model
- D. Waterfall

Answer: B

Q20. What is the main disadvantage of the Waterfall model?

- A. Spending too much time on testing activities
- B. Being excessively flexible and unstructured
- C. Moving too fast through development phases
- D. Difficult to accommodate changes after a phase is completed

Answer: D

Q21. What is Agile software development?

- A. A hardware design and manufacturing method as defined by standards
- B. An iterative approach emphasizing flexibility and customer collaboration
- C. A documentation-only waterfall approach for all project stakeholders
- D. A strict sequential step-by-step approach over the entire lifecycle

Answer: B

Q22. Which of the following is an Agile methodology?

- A. Scrum
- B. CMMI
- C. V-Model
- D. Waterfall

Answer: A

Q23. What is a 'Sprint' in Scrum?

- A. A formal documentation writing phase only
- B. A time-boxed iteration typically lasting 2-4 weeks
- C. A physical running exercise for team health
- D. A specific category of software defect type

Answer: B

Q24. Who is the Product Owner in Scrum?

- A. The quality assurance tester on the project in all development efforts
- B. The person responsible for maximizing the value of the product
- C. The database administrator managing storage by the project team members
- D. The chief executive officer of the company within the given constraints

Answer: B

Q25. What is the Agile Manifesto?

- A. A declaration of four core values and twelve principles for Agile development
- B. A programming language specification document to achieve project objectives
- C. A legally binding contractual document only and related components
- D. A hardware installation and setup manual in the engineering discipline

Answer: A

Q26. How many values does the Agile Manifesto contain?

- A. 2
- B. 8
- C. 4
- D. 6

Answer: C

Q27. What does Scrum emphasize?

- A. Hardware design and manufacturing work
- B. Individual solo development practices only
- C. Self-organizing teams and iterative progress
- D. Producing extensive formal documentation

Answer: C

Q28. What is a 'User Story' in Agile?

- A. A fictional novel written about end users for effective project outcomes
- B. A comprehensive automated test case only in the development process
- C. A formal detailed software bug report only as part of the methodology
- D. A short description of a feature from the user's perspective

Answer: D

Q29. What is a Daily Standup meeting?

- A. A brief daily meeting where team members share progress and blockers
- B. An extensive annual strategic planning event for the development team
- C. A comprehensive weekly team meeting session during the software lifecycle
- D. A detailed monthly project review session within the project scope

Answer: A

Q30. Which Agile value prioritizes 'working software' over what?

- A. Individual interactions
- B. Responding to change
- C. Customer collaboration
- D. Comprehensive documentation

Answer: D

Q31. What is requirements engineering?

- A. Performing software testing and validation in the development process
- B. Designing hardware system architectures during the software lifecycle
- C. Writing application source code for features as part of the methodology
- D. The process of defining, documenting, and maintaining software requirements

Answer: D

Q32. What is a functional requirement?

- A. A constraint on the system design choices
- B. A measurable performance goal for the system
- C. A requirement that describes what the system should do
- D. A detailed hardware specification document

Answer: C

Q33. What is a non-functional requirement?

- A. A requirement that specifies quality attributes like performance or security
- B. A requirement that is impossible to implement for the development team
- C. A broken or flawed software requirement only within the project scope
- D. A specification for hardware system components throughout the project

Answer: A

Q34. What is requirements elicitation?

- A. The process of gathering requirements from stakeholders
- B. Performing systematic software testing runs
- C. Writing application source code for modules
- D. Deploying software to production servers

Answer: A

Q35. Which of the following is a requirements elicitation technique?

- A. Interview sessions
- B. Debugging errors
- C. Coding source logic
- D. Code compilation

Answer: A

Q36. What is a Software Requirements Specification (SRS)?

- A. The application source code implementation
- B. A comprehensive software test plan document
- C. A complete end-user operational and help manual
- D. A document describing all requirements of the system

Answer: D

Q37. Who are stakeholders in requirements engineering?

- A. Only the programmers writing the source code
- B. Only the paying customers of the software product
- C. Only the project managers overseeing the work
- D. Anyone who has an interest in or is affected by the system

Answer: D

Q38. What is the difference between user requirements and system requirements?

- A. They are fundamentally the same concept entirely in the system context
- B. User requirements are high-level; system requirements are detailed technical specifications
- C. System requirements are always much simpler overall during implementation
- D. User requirements are far more technically detailed for quality purposes

Answer: B

Q39. What is requirements validation?

- A. Implementing requirements directly into code by the organization
- B. Initially writing new requirements documents in practice typically
- C. Checking that requirements are correct, complete, and consistent
- D. Permanently deleting unnecessary requirements as a standard approach

Answer: C

Q40. What is a use case?

- A. A description of how a user interacts with the system to achieve a goal
- B. A category of reported software defect type across all phases
- C. A specific automated test case scenario only at every stage
- D. A step in the deployment process workflow in a systematic way

Answer: A

Q41. What is software project management?

- A. Planning, organizing, and controlling software development activities
- B. Designing hardware system architectures throughout the project
- C. Writing application source code and modules within the project scope
- D. Performing software testing and debugging for the development team

Answer: A

Q42. What is a Gantt chart?

- A. A bar chart showing project schedule and task dependencies
- B. A relational database schema design diagram
- C. A specific type of source code format only
- D. A specialized software testing automation tool

Answer: A

Q43. What does WBS stand for?

- A. Work Based Standards
- B. Web Browser Software
- C. Work Breakdown Structure
- D. Wireless Broadcasting System

Answer: C

Q44. What is a project milestone?

- A. An established coding standard and guideline
- B. A significant point or event in a project timeline
- C. A physical unit of measured distance only
- D. A specific category of software defect type

Answer: B

Q45. What is the critical path in project management?

- A. The longest sequence of dependent tasks determining the minimum project duration
- B. The path with the fewest total listed tasks for quality purposes
- C. The shortest possible project path for tasks during implementation
- D. The set of easiest tasks in the project plan in the system context

Answer: A

Q46. What is project scope?

- A. The total number of developers on the team as a standard approach
- B. The total size of the source code base only in practice typically
- C. The defined boundaries of what the project will and will not deliver
- D. Only the financial budget for the project by the organization

Answer: C

Q47. What is a deliverable in project management?

- A. A tangible or intangible output produced as a result of project work
- B. A scheduled meeting agenda and minutes record in a systematic way
- C. A food delivery order to the project team at every stage
- D. A formal software defect bug report document across all phases

Answer: A

Q48. What is project estimation?

- A. Measuring physical hardware performance speeds
- B. Predicting the effort, cost, and duration needed for a project
- C. Simply counting total lines of application code
- D. Making random guesses without any real basis

Answer: B

Q49. What is resource allocation?

- A. Deleting unnecessary project data and files
- B. Assigning available resources to project tasks
- C. Writing comprehensive project documentation
- D. Purchasing new computing hardware equipment

Answer: B

Q50. What are the three constraints of project management?

- A. Design, Build, Ship within the system boundary
- B. Code, Test, Deploy for the project goals
- C. Scope, Time, Cost (the triple constraint)
- D. Plan, Do, Check according to best practices

Answer: C

Q51. What is software design?

- A. Writing detailed end-user operation manual docs as a standard approach
- B. Deploying software to production server systems in practice typically
- C. Writing comprehensive test case scenario suites for quality purposes
- D. The process of defining the architecture, components, and interfaces of a system

Answer: D

Q52. What is modularity in software design?

- A. Using only one file for all the source code
- B. Dividing a system into separate, independent modules
- C. Writing one single large monolithic program
- D. Avoiding the use of functions in the code

Answer: B

Q53. What is cohesion in software design?

- A. Unnecessary duplication of source code and logic
- B. Multiple modules depending on each other heavily
- C. A specific category of software defect and bug
- D. The degree to which elements within a module belong together

Answer: D

Q54. What is coupling in software design?

- A. A high degree of internal module cohesion only
- B. A specific type of data structure or element
- C. The degree of interdependence between modules
- D. A programming language built-in feature type

Answer: C

Q55. Which is preferred: high cohesion or low cohesion?

- A. No cohesion at all
- B. Both are equally good
- C. Low cohesion is preferred
- D. High cohesion is preferred

Answer: D

Q56. Which is preferred: tight coupling or loose coupling?

- A. Loose coupling is preferred
- B. No coupling at all works
- C. Tight coupling is preferred
- D. Both are equally good

Answer: A

Q57. What is a design pattern?

- A. A specific coding style and convention set
- B. A user interface visual color scheme choice
- C. A reusable solution to a commonly occurring design problem
- D. A comprehensive software testing method type

Answer: C

Q58. What is abstraction in design?

- A. A specific type of runtime software error type
- B. Adding increasingly more implementation detail
- C. Removing all application source code and files
- D. Hiding complex details and showing only essential features

Answer: D

Q59. What is information hiding?

- A. Concealing internal design decisions within a module from other modules
- B. Permanently deleting stored information across all phases
- C. Hiding files on the operating system disk at every stage
- D. Encrypting sensitive data during transit by the organization

Answer: A

Q60. What is the difference between high-level and low-level design?

- A. They are fundamentally the same thing exactly in a systematic way
- B. High-level defines architecture; low-level defines detailed component design
- C. Low-level design is always more important overall for the project goals
- D. High-level design is done after coding is finished according to best practices

Answer: B

Q61. What is software architecture?

- A. The high-level structure of a software system defining its components and their relationships
- B. A specific type of programming language syntax in the engineering discipline
- C. A comprehensive software testing framework type for effective project outcomes
- D. The actual application source code itself only and related components to achieve project objectives

Answer: A

Q62. What is a client-server architecture?

- A. A standalone desktop application design only during the software lifecycle
- B. An architecture where clients request services from a central server
- C. A decentralized peer-to-peer network system in the development process
- D. A single standalone computer system by itself as part of the methodology

Answer: B

Q63. What is a layered architecture?

- A. An architecture organized into horizontal layers, each providing services to the layer above
- B. A completely flat system structure layout design within the project scope
- C. A random unorganized structure for the code throughout the project in the system context
- D. A circular system structure design layout type for the development team

Answer: A

Q64. What is a monolithic architecture?

- A. A client-server communication system setup only in practice typically
- B. A single-tiered application where all components are tightly integrated
- C. A fully distributed computing system setup design during implementation
- D. A microservices-based architecture approach type for quality purposes

Answer: B

Q65. What is the purpose of software architecture?

- A. To write application source code much faster as a standard approach
- B. To provide a blueprint for the system and enable communication among stakeholders
- C. To completely eliminate the software testing phase by the organization
- D. To reduce the overall project development team at every stage

Answer: B

Q66. What is a three-tier architecture?

- A. A three-story office building design layout across all phases
- B. An architecture with presentation, business logic, and data tiers
- C. Three different programming languages used for the project goals
- D. A development team of three programmers in a systematic way

Answer: B

Q67. What is a repository architecture?

- A. A network infrastructure architecture type only by the development process
- B. A Git source code version control repository according to best practices
- C. An architecture where components share data through a central data store
- D. An operating system file storage system only within the system boundary

Answer: C

Q68. What is a pipe-and-filter architecture?

- A. An access security management system type only as defined by standards
- B. A physical plumbing infrastructure system design and its related activities
- C. A network traffic filtering system setup only over the entire lifecycle
- D. An architecture where data flows through a sequence of processing components

Answer: D

Q69. What is an event-driven architecture?

- A. An architecture where components communicate through events
- B. A sequential processing architecture style only
- C. A scheduling calendar management system only
- D. A batch data processing system design only

Answer: A

Q70. What is a distributed system architecture?

- A. An architecture where components run on multiple networked computers
- B. A standalone desktop application design approach within the given constraints
- C. A single standalone computer hardware unit only for all project stakeholders
- D. A local desktop application installation setup in all development efforts

Answer: A

Q71. What does UI stand for?

- A. User Information
- B. User Interface
- C. Universal Integration
- D. Unified Input

Answer: B

Q72. What does UX stand for?

- A. User Experience
- B. Universal Exchange
- C. Unified Expression
- D. User Extension

Answer: A

Q73. What is usability?

- A. The ease with which users can learn and use a software system
- B. Application processing execution speed rating
- C. System memory consumption and usage metrics
- D. Internal source code logic complexity metrics

Answer: A

Q74. What is a wireframe?

- A. A basic visual guide showing the layout and structure of a web page or app
- B. A type of network cable connector wire type during the software lifecycle
- C. A general-purpose programming language tool set within the project scope
- D. A specialized software testing tool suite only for the development team

Answer: A

Q75. What is responsive design?

- A. Websites that load extremely fast for users throughout the project
- B. A design that responds to voice command input in the system context
- C. A type of programming language paradigm type during implementation
- D. Design that adapts the layout to different screen sizes and devices

Answer: D

Q76. What is a mockup in UI design?

- A. A relational database schema design model type in practice typically
- B. A specific automated test case scenario setup for quality purposes
- C. A source code implementation template file set as a standard approach
- D. A static visual representation of how the final product will look

Answer: D

Q77. What is navigation in UI design?

- A. The system that allows users to move between different sections of an application
- B. A physical GPS location navigation system tool by the organization
- C. Network routing and traffic configuration setup at every stage
- D. Operating system file directory browsing tool across all phases

Answer: A

Q78. What is consistency in UI design?

- A. Using random and varying page layouts designs in a systematic way
- B. Making every page look distinctly different each for the project goals
- C. Changing interface colors very frequently always according to best practices
- D. Maintaining uniform design elements and behaviors throughout the interface

Answer: D

Q79. What is a call-to-action (CTA) button?

- A. A phone call button for customer support line
- B. A window close or dismiss button on the page
- C. A button designed to prompt users to take a specific action
- D. A help documentation access button on the page

Answer: C

Q80. What is accessibility in UI design?

- A. Database access control and permissions setup
- B. Access control to application source code files
- C. Designing software so it can be used by people with disabilities
- D. Internet access and connectivity service plans

Answer: C

Q81. What is coding in software development?

- A. Performing software testing and debugging work
- B. Writing comprehensive project documentation work
- C. Assembling physical hardware components together
- D. The process of writing source code to implement software design

Answer: D

Q82. What is a coding standard?

- A. A comprehensive software testing standard only
- B. A hardware specification and standard set only
- C. A specific type of source code module and file
- D. A set of guidelines for writing consistent and readable code

Answer: D

Q83. What is code review?

- A. Compiling source code into executable files by the project team members
- B. Writing new feature source code logic only in all development efforts
- C. Permanently deleting source code files only within the given constraints
- D. Systematic examination of source code by peers to find defects

Answer: D

Q84. What is refactoring?

- A. Permanently deleting old source code and files
- B. Completely rewriting everything from scratch again
- C. Adding entirely new features and functionality
- D. Restructuring existing code without changing its external behavior

Answer: D

Q85. What is version control?

- A. Managing changes to source code over time
- B. Controlling hardware component versions only
- C. A type of software testing methodology only
- D. Controlling software versions in a retail store

Answer: A

Q86. What is debugging?

- A. Deploying software to production server systems
- B. Intentionally adding bugs to the source code
- C. Writing brand new application source code only
- D. The process of finding and fixing defects in code

Answer: D

Q87. What is an IDE?

- A. Integrated Development Environment — a software tool for writing and debugging code
- B. An internet web hosting service provider only and related components
- C. A standard network communication protocol type in the engineering discipline
- D. A relational database management system only to achieve project objectives

Answer: A

Q88. What is source code?

- A. Human-readable instructions written in a programming language
- B. Binary data stored on disk drives and memory
- C. Low-level machine instruction code and binary
- D. Hardware specifications and schematics diagrams

Answer: A

Q89. What is a build in software development?

- A. A physical construction building project plan
- B. The process of converting source code into executable software
- C. A comprehensive design document and artifact
- D. A specific type of software test execution run

Answer: B

Q90. What is technical documentation?

- A. End user reviews and feedback comment forms as part of the methodology
- B. Sales reports and revenue projections summary in the development process
- C. Written descriptions of code structure, APIs, and system behavior
- D. Marketing and promotional campaign materials for effective project outcomes

Answer: C

Q91. What is software testing?

- A. Designing software system architecture diagrams as defined by standards
- B. The process of evaluating software to find defects and verify it meets requirements
- C. Writing application source code for new features and its related activities
- D. Deploying software to production server systems over the entire lifecycle

Answer: B

Q92. What is a test case?

- A. A set of conditions and expected results for testing a specific aspect of software
- B. A specific type of source code module and library in all development efforts
- C. A suitcase for carrying test documents and files within the given constraints
- D. A legal court case and proceedings in a courtroom for all project stakeholders

Answer: A

Q93. What is unit testing?

- A. Testing graphical user interface screens and flows
- B. Testing physical hardware units and device parts
- C. Testing the entire integrated system all at once
- D. Testing individual components or functions in isolation

Answer: D

Q94. What is integration testing?

- A. Testing documentation to achieve project objectives
- B. Testing in isolation and related components
- C. Testing the interaction between integrated components
- D. Unit testing by the project team members

Answer: C

Q95. What is a bug in software?

- A. A recognized software design pattern and approach
- B. A new feature request from project stakeholders
- C. A defect or error in software that causes incorrect behavior
- D. An insect found inside a computer case housing

Answer: C

Q96. What is black-box testing?

- A. Testing with a completely black display screen view
- B. Testing physical hardware system components only
- C. Testing software without knowledge of its internal code structure
- D. Testing in a completely dark room with no light

Answer: C

Q97. What is white-box testing?

- A. Testing with a completely white display screen view
- B. Testing software with knowledge of its internal code structure
- C. Testing comprehensive project documentation items
- D. Testing written on white paper documents and sheets

Answer: B

Q98. What is regression testing?

- A. Going backwards in the development lifecycle flow in the engineering discipline
- B. Testing the software for the very first time ever for effective project outcomes
- C. Testing the user registration feature flow only as part of the methodology
- D. Re-testing software after changes to ensure existing functionality still works

Answer: D

Q99. What is the purpose of a test plan?

- A. To outline the testing strategy, objectives, and schedule
- B. To design the user interface screen layout pages
- C. To deploy software to production server systems
- D. To write application source code for new modules

Answer: A

Q100. What is acceptance testing?

- A. Testing to determine if the software meets user requirements and is ready for delivery
- B. Accepting all reported software bugs as features in the development process
- C. Testing performed by internal developers only ever within the project scope
- D. Testing document formatting and acceptance rules during the software lifecycle

Answer: A

Q101. What is software maintenance?

- A. Writing brand new software from scratch entirely in the development process
- B. Maintaining physical hardware equipment and parts during the software lifecycle
- C. Installing software on new workstations and PCs within the project scope
- D. Modifying software after delivery to correct faults, improve performance, or adapt to changes

Answer: D

Q102. What is corrective maintenance?

- A. Improving overall system performance speed rates
- B. Fixing defects and bugs discovered after deployment
- C. Adapting software to new operating environments
- D. Adding new features and functionality to the app

Answer: B

Q103. What is adaptive maintenance?

- A. Adding brand new features to the product
- B. Modifying software to work in a changed environment
- C. Fixing bugs found after release for the development team
- D. Improving performance and speed throughout the project

Answer: B

Q104. What is perfective maintenance?

- A. Adapting software to new operating platforms only for quality purposes
- B. Making software absolutely perfect and flawless in the system context
- C. Fixing only critical security bugs and flaws only during implementation
- D. Enhancing software functionality and performance based on user feedback

Answer: D

Q105. What is preventive maintenance?

- A. Preventing all use of the software product system in practice typically
- B. Making changes to prevent future problems and improve maintainability
- C. Preventing all user access to the entire system by the organization
- D. Preventing only software bugs from occurring only as a standard approach

Answer: B

Q106. What percentage of software cost is typically spent on maintenance?

- A. 60-80%
- B. 25% acr
- C. 10% at
- D. 5% in

Answer: A

Q107. What is a software patch?

- A. A physical hardware fix and repair procedure doc
- B. A completely new software release version number
- C. A piece of code designed to fix or update existing software
- D. A fabric patch for clothing and textile repairs

Answer: C

Q108. What is legacy software?

- A. Brand new recently developed software application
- B. Open source freely available software components
- C. Older software that is still in use but may be difficult to maintain
- D. Mobile application software for phones and tablets

Answer: C

Q109. What is software evolution?

- A. Darwin's biological theory of species evolution for the project goals
- B. Permanently deleting outdated software and files according to best practices
- C. Physical hardware evolution over many years time within the system boundary
- D. The process of continuously changing and adapting software over its lifecycle

Answer: D

Q110. What is a software release?

- A. Permanently deleting the software product files
- B. A press release about the software product launch
- C. A version of software made available to users
- D. A specific type of software defect or bug found

Answer: C

Q111. What does QA stand for in software?

- A. Quantitative Assessment
- B. Query Analysis
- C. Quick Application
- D. Quality Assurance

Answer: D

Q112. What is software quality?

- A. The overall speed of software development activities
- B. The total number of features available in the app
- C. The degree to which software meets requirements and user expectations
- D. The total length of all source code lines and files

Answer: C

Q113. What is the difference between QA and testing?

- A. The two concepts are fundamentally the same exact thing across all phases
- B. QA focuses on preventing defects through processes; testing focuses on finding defects
- C. Testing is always much broader than QA practices are in a systematic way
- D. QA only involves software testing activities and nothing else for the project goals

Answer: B

Q114. What is a software standard?

- A. A physical flag or banner displayed for events according to best practices
- B. A physical computing hardware component and part by the development process
- C. A specific type of application source code module within the system boundary
- D. A documented agreement containing rules and guidelines for software development

Answer: D

Q115. What is a code inspection?

- A. A formal review of source code by peers to find defects
- B. Running the compiled application directly for users
- C. Compiling the source code into binary executables
- D. Looking at a computer display screen only visually

Answer: A

Q116. What is a software defect?

- A. An intended software feature or product capability and its related activities
- B. A scheduled software update or patch file release over the entire lifecycle
- C. A recognized software design pattern and approach as defined by standards
- D. A flaw in software that causes it to produce incorrect or unexpected results

Answer: D

Q117. What is the purpose of quality control?

- A. Checking that deliverables meet quality standards
- B. Controlling the overall project scope boundary
- C. Controlling physical hardware components and parts
- D. Controlling the overall project team size numbers

Answer: A

Q118. What is a review in software QA?

- A. A systematic examination of software artifacts to find issues
- B. A movie review and rating evaluation by critics
- C. An employee performance review and appraisal form
- D. A consumer product rating and evaluation review

Answer: A

Q119. What is ISO 9001?

- A. An international standard for quality management systems
- B. A specialized software testing tool suite and program
- C. A general-purpose programming language tool and set
- D. A specific type of application source code file only

Answer: A

Q120. What is the cost of quality?

- A. The total cost of ensuring quality plus the cost of poor quality
- B. Only the prevention costs of quality effort spent
- C. Only the rework and bug fixing cost totals only
- D. Only the software testing costs and effort totals

Answer: A

Q121. What are software metrics?

- A. A comprehensive software testing method type only for effective project outcomes
- B. Quantitative measures used to assess software quality, process, and productivity
- C. A specific type of application source code file type in the engineering discipline
- D. Physical measurements of hardware component parts to achieve project objectives

Answer: B

Q122. What is Lines of Code (LOC)?

- A. A poem or literary piece of creative written work as part of the methodology
- B. A metric that measures software size by counting the number of lines in source code
- C. A specialized debugging and tracing tool and suite during the software lifecycle
- D. A specific type of software coding error and defect in the development process

Answer: B

Q123. What is a product metric?

- A. A process metric for development workflow tracking
- B. A metric used for measuring sales revenue totals
- C. A measure of some attribute of the software product itself
- D. A hardware performance metric for device benchmarks

Answer: C

Q124. What is a process metric?

- A. A network performance metric for bandwidth speed
- B. A physical hardware component metric measurement
- C. A measure of the software product itself directly
- D. A measure of some attribute of the software development process

Answer: D

Q125. What is productivity in software metrics?

- A. The total number of meetings held per week time
- B. Working overtime hours every single week constantly
- C. The ratio of output (software produced) to input (resources used)
- D. The typing speed of individual developer teammates

Answer: C

Q126. What is defect rate?

- A. A metric for software testing execution speed rate
- B. Speed of creating new defects per work hour only
- C. The number of defects found per unit of software or time
- D. A specific type of reported software bug and defect

Answer: C

Q127. What does KLOC stand for?

- A. King Lines of Code within the
- B. Known Lines of Code throughout
- C. Thousands of Lines of Code
- D. Key Lines of Code for the

Answer: C

Q128. What is code complexity?

- A. The total number of declared variable names only for quality purposes
- B. The total length of all source code lines and files during implementation
- C. How hard source code is to initially write out in the system context
- D. A metric measuring how difficult code is to understand, test, and maintain

Answer: D

Q129. What is mean time between failures (MTBF)?

- A. The average time between system failures, indicating reliability
- B. The maximum allowed software testing time duration
- C. The average time to repair a system from failure
- D. The minimum testing frequency and interval schedule

Answer: A

Q130. What is mean time to repair (MTTR)?

- A. The average time between system failures and outages
- B. The maximum time allowed for running test suites
- C. The average time it takes to fix a failure and restore the system
- D. The minimum time to respond to a user web request

Answer: C

Q131. What is Software Configuration Management (SCM)?

- A. A creational object-oriented design pattern type
- B. The process of tracking and controlling changes to software
- C. A source code implementation coding technique only
- D. A comprehensive software testing method type only

Answer: B

Q132. What is version control?

- A. A software architecture and design method approach in the engineering discipline
- B. Controlling video game software versions and builds and related components
- C. A type of software testing methodology and approach to achieve project objectives
- D. A system that records changes to files over time so specific versions can be recalled

Answer: D

Q133. What is a baseline in configuration management?

- A. A specific line of application source code and logic as part of the methodology
- B. A starting position in a sports competition event for effective project outcomes
- C. A software testing benchmark result point and score in the development process
- D. A formally reviewed and agreed-upon specification serving as the basis for further development

Answer: D

Q134. What is a configuration item (CI)?

- A. Any artifact placed under configuration management control
- B. A user story in the product backlog items listing
- C. A piece of physical hardware equipment and component
- D. A specific automated test case scenario and script

Answer: A

Q135. What is Git?

- A. A programming language within the project
- B. A testing tool for the development team
- C. A distributed version control system
- D. A command during the software lifecycle

Answer: C

Q136. What is a commit in version control?

- A. A personal promise or commitment to other people
- B. A production deployment to end users and clients
- C. A recorded change or set of changes to the repository
- D. A test execution run and its result log report

Answer: C

Q137. What is a branch in version control?

- A. A physical tree branch in nature and the outdoors
- B. A separate line of development diverging from the main codebase
- C. A specific type of reported software bug and defect
- D. A deployment stage in the release pipeline process

Answer: B

Q138. What is merging in version control?

- A. Deleting branches in the system context
- B. Creating new files during implementation
- C. Combining changes throughout the project
- D. Integrating changes from one branch into another

Answer: D

Q139. What is a repository?

- A. A specific software testing environment and setup only in practice typically
- B. A central location where version-controlled files and their history are stored
- C. A production deployment server location and address as a standard approach
- D. A physical library building for books and resources for quality purposes

Answer: B

Q140. What is a change request?

- A. A formal proposal to modify a software configuration item
- B. A new feature request from end users and clients
- C. A request to change current employment job position
- D. A software defect bug report document filed only

Answer: A

Q141. What is risk in software project management?

- A. An uncertain event that could positively or negatively affect the project
- B. A task that has already been fully completed and done in a systematic way
- C. A documented project functional requirement and spec for the project goals
- D. An absolute certainty in the project plan document across all phases

Answer: A

Q142. What is risk identification?

- A. The process of determining which risks might affect the project
- B. Intentionally creating new project risks and issues
- C. Completely eliminating all identified project risks
- D. Deliberately ignoring all project risks completely

Answer: A

Q143. What is risk analysis?

- A. A source code implementation coding technique only
- B. Analyzing application source code logic and syntax
- C. A comprehensive software testing method approach
- D. Evaluating identified risks to determine their probability and impact

Answer: D

Q144. What is risk mitigation?

- A. Accepting all identified risks without any action
- B. Taking actions to reduce the probability or impact of a risk
- C. Deliberately ignoring all identified project risks
- D. Avoiding all projects and work activities entirely

Answer: B

Q145. What is a risk register?

- A. A physical cash register in a retail store space according to best practices
- B. A hardware register for storing CPU data and values by the development process
- C. A document listing identified risks, their analysis, and planned responses
- D. A specific type of application source code file only within the system boundary

Answer: C

Q146. What are the two components used to assess risk?

- A. Code and tests over the
- B. Scope and quality
- C. Probability and impact
- D. Time and money and its

Answer: C

Q147. What is a risk response strategy?

- A. Deleting the risk register and all documents too
- B. Running away from all identified project risks now
- C. A planned approach to address an identified risk
- D. Deliberately ignoring the identified risk entirely

Answer: C

Q148. What is risk avoidance?

- A. Accepting the risk without any planned action taken
- B. Changing the project plan to eliminate a risk or its impact
- C. Deliberately ignoring the identified risk in full
- D. Avoiding the project entirely from the very start

Answer: B

Q149. What is risk transfer?

- A. Transferring application source code ownership
- B. Moving data to a different storage location setup
- C. Shifting the impact of a risk to a third party
- D. Moving to a new office building location downtown

Answer: C

Q150. What is risk acceptance?

- A. Accepting all reported bugs as intended features only
- B. Accepting all proposed requirement changes submitted
- C. Accepting that the project will certainly fail overall
- D. Acknowledging a risk and deciding to take no preemptive action

Answer: D

Q151. What is software security?

- A. Locking the physical computer server room and door across all phases
- B. Using strong passwords for user accounts only ever for the project goals
- C. Protecting software from unauthorized access, use, modification, and destruction
- D. Installing antivirus software applications only ever in a systematic way

Answer: C

Q152. What is authentication?

- A. Verifying the identity of a user or system
- B. Encryption within the system boundary
- C. Firewall configuration by the development process
- D. Authorization according to best practices

Answer: A

Q153. What is authorization?

- A. Encryption over the entire lifecycle as defined by standards
- B. Authentication and its related activities
- C. Determining what permissions an authenticated user has
- D. Decryption for all project stakeholders within the given

Answer: C

Q154. What is encryption?

- A. Authentication in the engineering discipline for effective project outcomes
- B. Converting data into a coded form to prevent unauthorized access
- C. Compression in all development efforts by the project team members
- D. Decryption and related components to achieve project objectives

Answer: B

Q155. What is a vulnerability in software?

- A. A software defect fix and patch release and update
- B. A specific automated test case and scenario and script
- C. An intended software feature and product capability
- D. A weakness that can be exploited to compromise security

Answer: D

Q156. What is a security threat?

- A. A scheduled software update release event and launch as part of the methodology
- B. A formal source code peer review session and meeting during the software lifecycle
- C. A potential cause of an unwanted incident that may result in harm to a system
- D. A new feature request from stakeholders and clients in the development process

Answer: C

Q157. What is SQL injection?

- A. A database installation and setup procedure and guide
- B. An attack that inserts malicious SQL code into application queries
- C. A scheduled database backup and restore job and task
- D. A standard SQL database query feature and capability

Answer: B

Q158. What is cross-site scripting (XSS)?

- A. An attack where malicious scripts are injected into web pages viewed by other users
- B. A JavaScript library for DOM manipulation and events throughout the project
- C. A CSS framework for styling web content and layouts for the development team
- D. A source code implementation coding technique and tool within the project scope

Answer: A

Q159. What is a firewall?

- A. A type of computer virus or malware threat and attack during implementation
- B. A general-purpose programming language tool and system for quality purposes
- C. A physical wall constructed of fire bricks and mortar in the system context
- D. A network security device that monitors and controls incoming and outgoing traffic

Answer: D

Q160. What is HTTPS?

- A. A general-purpose programming language type and tool in practice typically
- B. A specific type of database storage system and service as a standard approach
- C. A comprehensive software testing framework and tool set by the organization
- D. HTTP Secure — a protocol for secure communication over a network using encryption

Answer: D

Q161. What is the primary purpose of software engineering practices?

- A. To eliminate the need for any project documentation
- B. To maximize the number of programming languages used
- C. To focus exclusively on writing source code quickly
- D. To apply systematic methods to software development

Answer: D

Q162. Which term describes software that is free from defects?

- A. Reliable software
- B. Recursive software
- C. Redundant software
- D. Reflective software

Answer: A

Q163. What is a software development life cycle (SDLC)?

- A. A programming language used for building systems
- B. A network protocol for transferring project data
- C. A framework defining steps in software development
- D. A hardware specification for running applications

Answer: C

Q164. Which of these is a key challenge in software engineering?

- A. Increasing the number of operating systems used
- B. Managing changing project requirements over time
- C. Ensuring hardware components are always upgraded
- D. Reducing the physical size of server equipment

Answer: B

Q165. What does the term 'stakeholder' mean in software projects?

- A. A testing tool used for automated regression tests
- B. A developer who writes the majority of source code
- C. A server that hosts the application in production
- D. Anyone with an interest in the software project outcome

Answer: D

Q166. What is meant by software maintainability?

- A. Number of features included in the first release
- B. Amount of memory the software consumes at runtime
- C. Ease of modifying software after initial deployment
- D. Speed at which software executes user operations

Answer: C

Q167. Which activity comes first in most software development processes?

- A. Performance testing and tuning
- B. Source code implementation phase
- C. Requirements gathering and analysis
- D. System deployment and release

Answer: C

Q168. What is the role of a software engineer?

- A. To design, develop, and maintain software systems
- B. To manage financial accounting for the company
- C. To install electrical wiring in office buildings
- D. To manufacture physical computer hardware units

Answer: A

Q169. What does 'scalability' mean for a software system?

- A. Requirement to rewrite code for every new feature
- B. Tendency to decrease in performance under any load
- C. Ability to handle growing amounts of work efficiently
- D. Capability to run on a single device without issues

Answer: C

Q170. Which of the following best describes a software prototype?

- A. The final production release of the application
- B. An early working model of the proposed system
- C. A hardware device used during system testing
- D. A legal contract between client and developers

Answer: B

Q171. What is the waterfall model in software development?

- A. A random unstructured method without any planning steps
- B. A circular iterative approach with no defined ending point
- C. A sequential linear development approach with distinct phases
- D. A parallel development technique using multiple teams only

Answer: C

Q172. Which phase typically follows design in the waterfall model?

- A. Implementation and coding
- B. System deployment step
- C. Requirements gathering
- D. Project initiation phase

Answer: A

Q173. What is an iterative development process?

- A. Repeating development cycles to refine the software product
- B. Skipping testing phases to accelerate delivery timelines
- C. Completing all development in a single pass without changes
- D. Eliminating design activities to focus on coding tasks

Answer: A

Q174. What does the V-model emphasize in software development?

- A. Eliminating all forms of documentation from the process
- B. Combining all phases into a single continuous activity
- C. Removing verification activities to save project time
- D. Corresponding testing phase for each development stage

Answer: D

Q175. Which model allows returning to previous phases if needed?

- A. Incremental model
- B. Pure waterfall model
- C. Big bang model only
- D. Code and fix model

Answer: A

Q176. What is the main advantage of using a process model?

- A. It eliminates the need for skilled software developers
- B. It guarantees the software will have zero defect count
- C. It removes all project risks from the development work
- D. It provides structure and predictability to development

Answer: D

Q177. In which model is software delivered in small functional pieces?

- A. Ad hoc coding method
- B. Big bang development way
- C. Incremental delivery model
- D. Pure waterfall approach

Answer: C

Q178. What is a spiral model primarily concerned with?

- A. Eliminating documentation
- B. Avoiding testing phases
- C. Reducing team size needs
- D. Risk analysis and management

Answer: D

Q179. What is the prototyping model used for?

- A. Replacing all testing with user observation sessions
- B. Building early versions to clarify system requirements
- C. Eliminating the need for any design documentation
- D. Skipping requirements and jumping to final deployment

Answer: B

Q180. Which model is best suited for well-understood requirements?

- A. Spiral model technique
- B. Prototyping model plan
- C. Waterfall model approach
- D. Agile methodology way

Answer: C

Q181. What is the core principle of agile software development?

- A. Following a strict sequential development process
- B. Completing comprehensive documentation before coding
- C. Avoiding customer collaboration during development
- D. Responding to change over following a rigid plan

Answer: D

Q182. What is a sprint in Scrum methodology?

- A. A method for deploying applications to production servers
- B. A tool for measuring source code execution performance
- C. A document describing all system requirements upfront
- D. A fixed time period for completing a set of work items

Answer: D

Q183. Who is the Product Owner in Scrum?

- A. The developer who writes the most lines of source code
- B. The tester who executes all automated regression test suites
- C. The person who manages the server infrastructure setup
- D. The person responsible for maximizing product value delivered

Answer: D

Q184. What is a user story in agile development?

- A. A formal contract between the client and development team
- B. A detailed technical specification for database schema design
- C. A short description of a feature from the user perspective
- D. A comprehensive test plan covering all system components

Answer: C

Q185. What does the Agile Manifesto value most?

- A. Contract negotiation over customer collaboration efforts
- B. Following a plan over responding to change requests
- C. Individuals and interactions over processes and tools
- D. Comprehensive documentation over working software delivery

Answer: C

Q186. What is a daily standup meeting in agile?

- A. A monthly review of the project budget and resources
- B. A brief daily meeting to share progress and obstacles
- C. A quarterly assessment of organizational strategy goals
- D. A weekly planning session for the entire project scope

Answer: B

Q187. What is the purpose of a product backlog?

- A. To list all desired features and work items for the product
- B. To track employee attendance and working hours precisely
- C. To document the organizational chart of the company fully
- D. To store source code versions in a repository management tool

Answer: A

Q188. How long is a typical Scrum sprint?

- A. One to two days only
- B. One to four weeks
- C. Eight to ten months
- D. Six to twelve months

Answer: B

Q189. What is the Scrum Master's primary role?

- A. To approve all design choices
- B. To manage project budget
- C. To facilitate the Scrum process
- D. To write most of the code

Answer: C

Q190. What happens during a sprint review?

- A. All unfinished tasks are permanently removed from scope
- B. The team demonstrates completed work to stakeholders
- C. The development environment is reset to default state
- D. The project budget is reduced by senior management

Answer: B

Q191. What is requirements engineering in software development?

- A. The process of discovering and documenting system needs
- B. A strategy for deploying applications to cloud servers
- C. A technique for writing efficient source code algorithms
- D. A method for testing software performance under load

Answer: A

Q192. What is a functional requirement?

- A. A specification of what the system should do for users
- B. A plan for marketing the software product to customers
- C. A schedule for training end users on system operations
- D. A description of the hardware needed to run the system

Answer: A

Q193. What is a non-functional requirement?

- A. A project schedule milestone for the delivery deadline
- B. A feature that the system must implement for the user
- C. A programming language choice for the development team
- D. A constraint on how the system performs its functions

Answer: D

Q194. Who are the primary sources of software requirements?

- A. Hardware manufacturers and network equipment vendors
- B. Operating system kernel developers and maintainers
- C. Compiler designers and programming language creators
- D. Stakeholders and end users of the system being built

Answer: D

Q195. What is a Software Requirements Specification (SRS)?

- A. A network diagram showing system topology and connections
- B. A database schema defining all tables and relationships
- C. A source code file containing the main application logic
- D. A document that describes all system requirements in detail

Answer: D

Q196. What does 'requirements elicitation' mean?

- A. Testing the software against its documented requirements
- B. Gathering requirements from stakeholders and other sources
- C. Deploying the finished software to the production server
- D. Deleting unnecessary requirements from the project scope

Answer: B

Q197. Which of these is an example of a non-functional requirement?

- A. The system shall generate monthly financial reports
- B. The system shall send email notifications to users
- C. The system shall respond to queries within two seconds
- D. The system shall allow users to create new accounts

Answer: C

Q198. What is the purpose of requirements validation?

- A. To ensure requirements accurately reflect stakeholder needs
- B. To train end users on how to operate the new system
- C. To deploy the software to the staging environment first
- D. To write source code that implements the system features

Answer: A

Q199. What is a use case in requirements engineering?

- A. A report on the system's performance test results
- B. A list of programming languages used in the project
- C. A diagram of the database schema and relationships
- D. A description of how users interact with the system

Answer: D

Q200. Why is requirements traceability important?

- A. It links requirements to design, code, and test artifacts
- B. It eliminates the need for user acceptance testing work
- C. It reduces the amount of memory the application uses
- D. It speeds up the compilation of the source code files

Answer: A

Q201. What is the primary goal of software project management?

- A. Installing and configuring the production server hardware equipment
- B. Writing all the source code for the project without any team members
- C. Designing the user interface layout and choosing the color schemes
- D. Delivering software on time, within budget, and meeting quality standards

Answer: D

Q202. What is a project milestone?

- A. A hardware component required for the server infrastructure
- B. A significant checkpoint marking completion of a project phase
- C. A type of programming error found during code review sessions
- D. A daily task assigned to individual software developers on team

Answer: B

Q203. What is the purpose of a Gantt chart in project management?

- A. To document all the source code changes made by developers
- B. To track the number of defects found during system testing
- C. To visually represent the project schedule and task timelines
- D. To calculate the total budget required for hardware purchases

Answer: C

Q204. What does WBS stand for in project management?

- A. Workflow Business Strategy
- B. Wireless Bandwidth System
- C. Web Based Software tool
- D. Work Breakdown Structure

Answer: D

Q205. Who is responsible for overall project success?

- A. The quality assurance tester only
- B. The external client exclusively
- C. The project manager leading the team
- D. Each individual developer equally

Answer: C

Q206. What is project scope in software development?

- A. The defined boundaries of what the project will deliver
- B. The programming language selected for implementation
- C. The total number of developers assigned to the team
- D. The number of servers allocated for production use

Answer: A

Q207. What is resource allocation in project management?

- A. Assigning people, tools, and budget to project activities
- B. Designing the database schema for the system storage
- C. Deleting unnecessary files from the project repository
- D. Running automated tests on the application source code

Answer: A

Q208. What is the critical path in a project schedule?

- A. A shortcut method for skipping non-essential project activities
- B. The easiest set of tasks that any team member can complete alone
- C. A backup plan executed only when the main project plan fails
- D. The longest sequence of dependent tasks determining project duration

Answer: D

Q209. What is a project stakeholder?

- A. Only the software developers who write the application source code
- B. Just the project manager who oversees the entire development team
- C. Exclusively the client who pays for the software development work
- D. Anyone who has an interest in or is affected by the project outcome

Answer: D

Q210. What does project monitoring involve?

- A. Tracking progress and comparing it against the original plan
- B. Writing source code for the application features and modules
- C. Conducting final user acceptance testing before deployment
- D. Designing the user interface screens and navigation flows

Answer: A

Q211. What is software design in the development process?

- A. A technique for deploying software to production servers
- B. A strategy for gathering requirements from stakeholders
- C. A method for testing software under heavy load conditions
- D. The process of defining system structure and components

Answer: D

Q212. What is cohesion in software design?

- A. The degree to which elements within a module belong together
- B. The total number of lines of code in a software application
- C. The speed at which the application responds to user actions
- D. The amount of dependencies between different software modules

Answer: A

Q213. What is coupling in software design?

- A. The speed at which the application compiles from source
- B. The degree of interdependence between software modules
- C. The total number of classes in the application codebase
- D. The amount of documentation written for each module

Answer: B

Q214. What is the goal of good software design?

- A. Neither cohesion nor coupling matter in quality design work
- B. High coupling between modules and low cohesion within them
- C. High cohesion within modules and low coupling between them
- D. Equal cohesion and coupling across all system components

Answer: C

Q215. What is an interface in software design?

- A. A network protocol for transmitting data between servers
- B. A contract defining how components interact with each other
- C. A database table storing application configuration settings
- D. A visual screen layout designed for end user interactions only

Answer: B

Q216. What is modularity in software design?

- A. Writing all code in a single file without any organization
- B. Maximizing dependencies between all parts of the system
- C. Dividing software into separate independent functional units
- D. Avoiding any separation of concerns in the code structure

Answer: C

Q217. What is abstraction in software design?

- A. Avoiding any simplification of the system's internal logic
- B. Making all internal data structures visible to every module
- C. Including every implementation detail in the public interface
- D. Hiding complex details and showing only essential features

Answer: D

Q218. What is the purpose of a design pattern?

- A. To increase the complexity of the software architecture
- B. To provide reusable solutions to common design problems
- C. To replace the need for requirements analysis activities
- D. To create unique solutions that cannot be reused anywhere

Answer: B

Q219. What is encapsulation in object-oriented design?

- A. Exposing all internal data to every part of the program
- B. Allowing unrestricted access to all object internal states
- C. Bundling data and methods together and restricting access
- D. Separating data from the methods that operate on the data

Answer: C

Q220. What is a class diagram used for in software design?

- A. To display the project timeline and scheduling milestones
- B. To show the structure of classes and their relationships
- C. To illustrate network topology and server configurations
- D. To represent database queries and their execution plans

Answer: B

Q221. What is software architecture?

- A. A detailed line-by-line description of the source code logic
- B. A specific programming language used for building applications
- C. A testing framework for validating user interface interactions
- D. The high-level structure of a software system and its components

Answer: D

Q222. What is a client-server architecture?

- A. A design pattern for sorting data in ascending order quickly
- B. A system where all processing occurs on the user device only
- C. A system where clients request services from a central server
- D. A testing methodology for validating database query results

Answer: C

Q223. What is a layered architecture?

- A. An architecture using a single flat structure without layers
- B. An architecture organizing software into hierarchical layers
- C. A deployment strategy for distributing code across servers
- D. A testing technique for validating multi-step user workflows

Answer: B

Q224. What is a monolithic architecture?

- A. A single unified application where all components are interconnected
- B. A database architecture using multiple separate data storage nodes
- C. A network topology with redundant connections between all nodes
- D. A distributed system with many independent microservice components

Answer: A

Q225. What is the purpose of an architectural pattern?

- A. To define specific algorithms for data processing and sorting
- B. To provide proven solutions for common structural design problems
- C. To specify the exact programming language for implementation
- D. To determine the number of developers needed for a project

Answer: B

Q226. What is a microservices architecture?

- A. Small independent services that communicate over network protocols
- B. A single large application built as one deployable unit together
- C. A programming style using only one function for all logic needs
- D. A database design using only a single table for all data items

Answer: A

Q227. What is a pipe-and-filter architecture?

- A. Data flows through processing components connected in sequence
- B. A user interface pattern for displaying forms and input fields
- C. All components execute simultaneously without any data flow order
- D. A database technique for filtering query results by conditions

Answer: A

Q228. What is a repository architecture?

- A. Data is distributed randomly across components without any order
- B. Each component maintains its own isolated copy of all system data
- C. Components never share any data and operate fully independently
- D. Components share data through a central data store or repository

Answer: D

Q229. What does 'architectural style' refer to?

- A. A family of systems sharing structural and behavioral properties
- B. A method for writing unit tests for individual code functions
- C. A specific color scheme used for the application user interface
- D. A technique for compressing files to reduce storage space used

Answer: A

Q230. What is the role of an architect in software development?

- A. Writing all the unit tests for the application source code now
- B. Making high-level design decisions about system structure choices
- C. Managing the project budget and timeline schedules for delivery
- D. Conducting user interviews to gather system requirements only

Answer: B

Q231. What is user interface (UI) design?

- A. Managing version control repositories for source code
- B. Writing backend database queries for data retrieval tasks
- C. Designing visual elements that users interact with on screen
- D. Configuring server infrastructure for application hosting

Answer: C

Q232. What is usability in UI design?

- A. The total number of colors used in the application theme
- B. The programming language used to implement the interface
- C. How easy and efficient it is for users to use the interface
- D. The amount of data stored in the application database

Answer: C

Q233. What is a wireframe in UI design?

- A. A basic visual layout showing the structure of a page
- B. A database diagram showing tables and their relationships
- C. A network map displaying server connections and topology
- D. A fully functional prototype with all features implemented

Answer: A

Q234. What does 'responsive design' mean?

- A. Design that adapts to different screen sizes and devices
- B. Design that responds quickly to user input and actions
- C. Design that avoids any visual elements on the interface
- D. Design that uses only one fixed layout for all screens

Answer: A

Q235. What is the purpose of user feedback in UI design?

- A. To reduce the number of features available to end users
- B. To understand user needs and improve the interface design
- C. To increase the amount of code in the application files
- D. To eliminate the need for any future design improvements

Answer: B

Q236. What is a call-to-action (CTA) in UI design?

- A. A prompt encouraging users to take a specific action step
- B. A backend process running on the application server side
- C. A database constraint enforcing data integrity rules now
- D. A technical error message displayed during system failure

Answer: A

Q237. What is navigation in user interface design?

- A. The server-side logic that processes user form inputs
- B. The encryption algorithm protecting sensitive user data
- C. The means by which users move through an application
- D. The database indexing strategy for faster data queries

Answer: C

Q238. What is visual hierarchy in UI design?

- A. Hiding important elements behind multiple menu layers
- B. Removing all visual distinctions between page elements
- C. Arranging elements to guide user attention by importance
- D. Placing all elements at the same size and visual weight

Answer: C

Q239. What is the goal of accessibility in UI design?

- A. Making interfaces only usable by experienced tech users
- B. Reducing the number of interactive elements on a page
- C. Restricting access to the interface based on user roles
- D. Making interfaces usable by people with diverse abilities

Answer: D

Q240. What is a tooltip in user interface design?

- A. A small popup providing information when hovering over an element
- B. A large modal window requiring user input before proceeding further
- C. A background service processing data without any user interaction
- D. A database record storing user preferences and account settings

Answer: A

Q241. What is source code in software development?

- A. A network configuration file for setting up server connections
- B. The physical hardware on which software applications run today
- C. Human-readable instructions written in a programming language
- D. A project management document describing the timeline and budget

Answer: C

Q242. What is a compiler used for in software development?

- A. Translating source code into machine-executable binary code
- B. Creating database schemas and defining table relationships
- C. Designing the user interface layout for web applications only
- D. Managing project team members and assigning work to them

Answer: A

Q243. What is code refactoring?

- A. Removing all comments and documentation from the source code
- B. Deploying the application to a new production server environment
- C. Restructuring existing code without changing its external behavior
- D. Adding new features and functionality to the software product

Answer: C

Q244. What is version control in software development?

- A. A system for tracking and managing changes to source code files
- B. A tool for managing hardware inventory in the server data room
- C. A method for testing the application performance under load now
- D. A technique for designing the user interface with color themes

Answer: A

Q245. What is the purpose of code comments?

- A. To explain the intent and logic of code for future readers
- B. To encrypt the source code to prevent unauthorized access now
- C. To increase the execution speed of the compiled application
- D. To add new features without modifying any existing source code

Answer: A

Q246. What is an Integrated Development Environment (IDE)?

- A. A database management system for storing application data now
- B. A software tool combining editor, compiler, and debugger features
- C. A physical room where developers work together on a project
- D. A hardware device used for testing software on mobile devices

Answer: B

Q247. What is a syntax error in programming?

- A. A violation of the programming language grammar and structure rules
- B. A performance issue that causes the application to run very slowly
- C. A logical mistake that produces incorrect output from the program
- D. A design flaw in the system architecture causing scalability issues

Answer: A

Q248. What is debugging in software development?

- A. Finding and fixing defects and errors in the source code logic
- B. Deploying the application to the production server environment
- C. Writing new source code to add features to the application
- D. Designing the database schema for the application data store

Answer: A

Q249. What is a code review?

- A. Running automated performance tests on the production environment
- B. Deleting all source code and starting the project from scratch now
- C. Installing software updates on all developer workstation machines
- D. Examining source code by peers to find defects and improve quality

Answer: D

Q250. What is the purpose of coding standards?

- A. To prevent any developer from modifying code written by other people
- B. To ensure consistent code style and quality across the development team
- C. To eliminate the need for testing by ensuring code is always correct
- D. To restrict developers from using any external libraries or frameworks

Answer: B

Q251. What is the main purpose of software testing?

- A. To design the user interface layout and visual style elements
- B. To deploy the application to the production server environment
- C. To find defects and verify the software meets its requirements
- D. To write the source code for the application being developed

Answer: C

Q252. What is a test case in software testing?

- A. A set of conditions and expected results for verifying functionality
- B. A hardware component required for running the test environment
- C. A programming language used for writing application source code
- D. A project management document listing all team member assignments

Answer: A

Q253. What is unit testing?

- A. Testing network connectivity between all production server nodes
- B. Testing the entire integrated system from end to end completely
- C. Testing the user interface for visual design and layout accuracy
- D. Testing individual components or functions in isolation from others

Answer: D

Q254. What is integration testing?

- A. Testing the project schedule against the planned delivery dates
- B. Testing how multiple components work together as a combined unit
- C. Testing individual functions in complete isolation from all others
- D. Testing the application user interface for accessibility standards

Answer: B

Q255. What is a bug or defect in software?

- A. A new feature requested by the client during the development phase
- B. A hardware failure in the production server causing system downtime
- C. A scheduled maintenance window for updating the operating system
- D. A flaw causing the software to produce incorrect or unexpected results

Answer: D

Q256. What is regression testing?

- A. Testing new features for the first time before release to production
- B. Removing features from the software to simplify the application now
- C. Retesting software after changes to ensure existing features still work
- D. Writing new source code to implement additional system requirements

Answer: C

Q257. What is system testing?

- A. Testing a single function or method in isolation from all other code
- B. Testing only the network configuration of the production servers
- C. Testing only the database queries for correctness and performance
- D. Testing the complete integrated system against specified requirements

Answer: D

Q258. What is the difference between manual and automated testing?

- A. Manual testing is done by humans; automated testing uses scripts and tools
- B. Automated testing cannot find any bugs that manual testing would miss
- C. Manual and automated testing produce identical results in all scenarios
- D. Manual testing is always faster than automated testing in every situation

Answer: A

Q259. What is acceptance testing?

- A. Testing the network bandwidth between the client and server nodes
- B. Testing to determine if the system satisfies the acceptance criteria
- C. Testing individual database tables for data integrity constraints
- D. Testing the internal code structure for compliance with style guides

Answer: B

Q260. What is a test plan?

- A. A database query for retrieving user account information data
- B. A document describing the testing scope, approach, and schedule
- C. A network diagram showing the production server architecture
- D. A source code file containing the main application entry point

Answer: B

Q261. What is software maintenance?

- A. Modifying software after delivery to correct faults or improve performance
- B. Writing the initial source code for a brand new software application
- C. Designing the system architecture before any development work begins
- D. Gathering user requirements during the early project planning phases

Answer: A

Q262. What is corrective maintenance?

- A. Fixing bugs and defects discovered after software has been deployed
- B. Adapting software to work in a new operating system environment
- C. Improving performance without changing any existing functionality
- D. Adding brand new features requested by the users after deployment

Answer: A

Q263. What is adaptive maintenance?

- A. Adding new features that were not in the original requirements set
- B. Fixing defects found during production use of the deployed system
- C. Modifying software to work in a changed or new environment context
- D. Restructuring code to improve readability without behavior change

Answer: C

Q264. What is perfective maintenance?

- A. Fixing critical bugs that prevent the system from functioning at all
- B. Removing the software from production and replacing it with a new one
- C. Modifying the software to run on a completely different platform now
- D. Enhancing software by adding new features or improving existing ones

Answer: D

Q265. What percentage of total software costs does maintenance typically represent?

- A. About twenty percent of total lifecycle costs
- B. Over sixty percent of total lifecycle costs
- C. Less than five percent of total lifecycle costs
- D. Exactly ten percent of total lifecycle costs

Answer: B

Q266. What is preventive maintenance?

- A. Making changes to prevent future problems and improve maintainability
- B. Adding features only when explicitly requested by the project sponsor
- C. Fixing bugs only after users report them in the production environment
- D. Removing all documentation from the codebase to simplify the system

Answer: A

Q267. What is a software patch?

- A. A full version upgrade adding many new features and capabilities now
- B. A small update released to fix specific bugs or security issues quickly
- C. A complete rewrite of the entire application from scratch using new code
- D. A hardware replacement required to run updated software on the server

Answer: B

Q268. Why is documentation important for software maintenance?

- A. It increases the file size of the project without adding any value
- B. It is only useful during initial development and never after that
- C. It helps maintainers understand the system to make changes correctly
- D. It slows down maintenance by requiring reading before any changes

Answer: C

Q269. What is software evolution?

- A. The continuous process of changing software throughout its lifetime
- B. A one-time event that happens only during initial development phase
- C. The process of permanently discontinuing a software product fully
- D. A testing technique applied only before the first release of code

Answer: A

Q270. What is legacy software?

- A. Brand new software developed using the latest programming frameworks
- B. A type of testing framework used exclusively for mobile applications
- C. Older software that remains in use despite outdated technology choices
- D. A project management methodology for small startup development teams

Answer: C

Q271. What is quality assurance (QA) in software engineering?

- A. A programming language designed for writing high-quality source code automatically
- B. A specific testing technique used only during the final phase of development work
- C. A hardware specification defining minimum requirements for running applications
- D. Planned activities ensuring the development process produces quality software products

Answer: D

Q272. What is the difference between quality assurance and quality control?

- A. QA focuses on processes; quality control focuses on product inspection and testing
- B. QA and quality control are identical concepts with no meaningful differences at all
- C. QA is done after release; quality control is done before development starts now
- D. QA tests the product; quality control improves the development process activities

Answer: A

Q273. What is a software quality standard?

- A. A hardware requirement for running the quality assurance testing tools now
- B. A documented set of criteria that software must meet for acceptable quality
- C. A project management template for scheduling development team meetings
- D. A specific programming language recommended for building quality software

Answer: B

Q274. What is the purpose of a quality audit?

- A. To deploy the application to the production environment for user access
- B. To evaluate whether quality processes are being followed correctly by teams
- C. To design the database schema and define table relationships for storage
- D. To write source code for the application features and user interface now

Answer: B

Q275. What does ISO 9001 relate to in software engineering?

- A. A quality management system standard applicable to software development
- B. A specific programming language standard for embedded system software
- C. A hardware interface specification for connecting peripheral devices now
- D. A network protocol standard for transmitting data between computers

Answer: A

Q276. What is a defect in the context of quality assurance?

- A. A positive user review about the software application performance
- B. Any deviation from the expected behavior defined in the requirements
- C. A new feature added to the software after the initial release date
- D. A planned system downtime for performing scheduled maintenance work

Answer: B

Q277. What is the role of a QA engineer?

- A. Managing the project budget and negotiating contracts with vendors now
- B. Ensuring software quality through testing, reviews, and process improvement
- C. Designing the network infrastructure for the production environment
- D. Writing all the source code for the application without any supervision

Answer: B

Q278. What is peer review in quality assurance?

- A. A management review of project budget and timeline progress reports
- B. A single developer reviewing their own code without any external feedback
- C. Team members examining each other's work to find defects and improvements
- D. An automated tool scanning source code for syntax errors and warnings

Answer: C

Q279. What is a quality metric?

- A. A hardware sensor measuring server temperature in the data center
- B. A type of source code comment used for documenting class interfaces
- C. A network monitoring tool tracking bandwidth usage and latency now
- D. A measurable indicator of software quality attributes and characteristics

Answer: D

Q280. Why is quality important in software engineering?

- A. It increases the number of defects found during production operations
- B. It makes the source code longer and more difficult to understand now
- C. It ensures user satisfaction, reduces costs, and builds customer trust
- D. It slows down development without providing any measurable benefits

Answer: C

Q281. What are software metrics?

- A. Quantitative measures used to assess software quality and processes
- B. A type of programming language for building web applications today
- C. A hardware specification for measuring processor clock speed rates
- D. A project management tool for scheduling team meetings and events

Answer: A

Q282. What does Lines of Code (LOC) measure?

- A. The amount of memory the software uses during runtime execution
- B. The speed at which the application executes user requests today
- C. The number of team members working on the software project now
- D. The size of a software program by counting its source code lines

Answer: D

Q283. What is cyclomatic complexity?

- A. A method for designing user interfaces for mobile application screens
- B. A strategy for managing project budgets and resource allocations today
- C. A technique for deploying software to cloud computing environments now
- D. A measure of the number of independent paths through a program's code

Answer: D

Q284. Why are software metrics important?

- A. They replace the need for experienced developers on the team now
- B. They eliminate the need for any testing or quality assurance work
- C. They provide objective data for making informed project decisions
- D. They guarantee that all software products will be defect-free always

Answer: C

Q285. What is a product metric?

- A. A measure of some attribute of the software product being developed
- B. A financial metric tracking company revenue from software licenses
- C. A hardware metric measuring server uptime and availability rates
- D. A measure of the development process performance and efficiency now

Answer: A

Q286. What is a process metric?

- A. A type of user interface element used in dashboard design layouts
- B. A measure of the final product features and characteristics directly
- C. A network measurement tool for tracking data transfer speed rates
- D. A measure of the software development process effectiveness and quality

Answer: D

Q287. What does defect density measure?

- A. The physical density of storage media used by the application today
- B. The number of defects per unit size of the software being measured
- C. The visual density of elements on the user interface screen layout
- D. The network traffic density between the client and server systems

Answer: B

Q288. What is the purpose of measuring code coverage?

- A. To determine what percentage of code is exercised by test suites
- B. To count the total number of source code files in the project now
- C. To track the number of developers contributing to the codebase
- D. To measure the physical disk space occupied by the application

Answer: A

Q289. What is mean time between failures (MTBF)?

- A. The time it takes to write one line of source code on average
- B. The average time a system operates between consecutive failures
- C. The average salary of developers working on the software project
- D. The number of features added to each software release on average

Answer: B

Q290. What is a baseline in software measurement?

- A. A default user interface theme applied when the application starts
- B. A specific line of code that forms the foundation of the application
- C. A reference point for comparing future measurements and tracking change
- D. A minimum hardware requirement for running the software application

Answer: C

Q291. What is software configuration management (SCM)?

- A. Managing and tracking changes to software throughout its lifecycle
- B. A hardware setup process for servers in the data center facility
- C. A programming technique for writing configuration file parsers now
- D. A testing method for verifying system configuration settings only

Answer: A

Q292. What is version control in configuration management?

- A. A technique for designing user interfaces with multiple theme options
- B. A strategy for managing team member schedules and work assignments
- C. A method for testing different versions of the software simultaneously
- D. A system that records changes to files over time for later retrieval

Answer: D

Q293. What is a software baseline in configuration management?

- A. A preliminary draft of the source code before any reviews take place
- B. An informal agreement between developers about coding style preferences
- C. A formally reviewed and agreed-upon specification serving as a reference
- D. A temporary test environment used during development and testing only

Answer: C

Q294. What is a configuration item (CI)?

- A. An artifact that is placed under configuration management and tracked
- B. A physical server component that needs regular hardware maintenance
- C. A project management tool for scheduling team meetings and events
- D. A financial metric tracking development costs against the project budget

Answer: A

Q295. What is the purpose of a change control board (CCB)?

- A. To conduct performance testing on the application before release to users
- B. To design the user interface screens and navigation flows for the system
- C. To evaluate and approve or reject proposed changes to the project baseline
- D. To write all the source code for the application without any team input

Answer: C

Q296. What is a code repository?

- A. A customer support center handling user complaints and feature requests
- B. A central location where source code is stored and version controlled
- C. A training facility where developers learn new programming languages
- D. A physical library containing printed copies of programming textbooks

Answer: B

Q297. What is branching in version control?

- A. Deploying the application to multiple production servers simultaneously
- B. Creating a separate line of development from the main codebase path
- C. Merging all team members' work into a single file without any review
- D. Deleting old versions of the source code to save disk storage space

Answer: B

Q298. What is merging in version control?

- A. Combining changes from different branches into a single unified branch
- B. Splitting a single file into multiple smaller files for better organization
- C. Creating a new repository from scratch without any existing source code
- D. Deleting all branches except the main branch to simplify the repository

Answer: A

Q299. What is a commit in version control?

- A. A snapshot of changes saved to the repository at a specific point in time
- B. A financial transaction related to purchasing software licenses for tools
- C. A verbal agreement between team members about the project scope and goals
- D. A hardware upgrade performed on the development workstation machines now

Answer: A

Q300. What is a release in software configuration management?

- A. A meeting where the team discusses upcoming features and priorities now
- B. A training session for new developers joining the project team currently
- C. A testing phase where all automated test suites are executed and reviewed
- D. A version of the software distributed to users for installation and use

Answer: D

Q301. What is risk in software project management?

- A. A completed task that has already been delivered successfully now
- B. An uncertain event that could negatively impact the project outcome
- C. A team meeting scheduled to discuss upcoming project milestones
- D. A guaranteed problem that will definitely occur during development

Answer: B

Q302. What is risk identification?

- A. The process of eliminating all possible risks from the project scope
- B. The process of finding and documenting potential project risks early
- C. A technique for writing source code without any potential defects
- D. A method for deploying software to production without any testing

Answer: B

Q303. What is risk mitigation?

- A. Taking actions to reduce the probability or impact of identified risks
- B. Ignoring identified risks and hoping they will not occur during work
- C. Removing all project constraints to eliminate potential risk factors
- D. Increasing the likelihood of risks to test the team's response ability

Answer: A

Q304. What is a risk register used for?

- A. Tracking employee attendance and work hours for payroll processing
- B. Storing source code backups on external hard drives for recovery
- C. Managing customer support tickets and feature request submissions
- D. Recording identified risks with their assessments and response plans

Answer: D

Q305. What is risk probability?

- A. The likelihood that a specific risk event will actually occur today
- B. The guaranteed certainty that a risk will happen in every project
- C. The cost of fixing a problem after it occurs in production later
- D. The time required to complete the project from start to finish

Answer: A

Q306. What is risk impact?

- A. The total budget allocated for the software development project
- B. The number of team members assigned to work on the project now
- C. The severity of consequences if a risk event actually materializes
- D. The programming language selected for implementing the features

Answer: C

Q307. What is risk avoidance?

- A. Increasing the risk probability to learn from potential failure events
- B. Changing the project plan to eliminate the risk or its impact entirely
- C. Accepting the risk and taking no action to address it at all currently
- D. Transferring the risk to another team without any formal agreement

Answer: B

Q308. What is risk acceptance?

- A. Actively working to increase the probability of the risk event occurring
- B. Denying that any risks exist in the project scope and schedule plan
- C. Transferring all project risks to external contractors without oversight
- D. Acknowledging a risk exists and deciding not to take action against it

Answer: D

Q309. What is a contingency plan?

- A. A training program for teaching developers new programming skills
- B. A budget report showing how much money has been spent so far now
- C. A predefined plan to execute if a specific risk event actually occurs
- D. A daily standup meeting where team members share progress updates

Answer: C

Q310. Who is responsible for managing risks in a software project?

- A. The project manager with support from the entire project team
- B. The hardware vendor who provides the production server equipment
- C. The external client who commissioned the software development
- D. Only the most junior developer on the project team exclusively

Answer: A

Q311. What is software security engineering?

- A. A strategy for managing team schedules and project milestone dates
- B. Building software that continues to function correctly under attack
- C. A technique for optimizing database query execution and performance
- D. A method for testing user interface responsiveness and visual design

Answer: B

Q312. What is authentication in software security?

- A. Designing the user interface for accessibility compliance only
- B. Verifying the identity of a user or system requesting access
- C. Testing the application for performance under heavy load now
- D. Encrypting all data stored in the application database tables

Answer: B

Q313. What is authorization in software security?

- A. Encrypting data during transmission between client and the server
- B. Verifying the identity of a user attempting to log into the system
- C. Testing the application for compatibility across web browsers now
- D. Determining what actions an authenticated user is permitted to do

Answer: D

Q314. What is encryption used for in software security?

- A. Compressing files to reduce storage space on the server drives
- B. Protecting data by converting it into an unreadable format safely
- C. Speeding up database query execution for better performance now
- D. Formatting user interface elements for better visual appearance

Answer: B

Q315. What is a firewall in security?

- A. A programming language designed for building secure applications
- B. A system that monitors and controls network traffic based on rules
- C. A project management tool for tracking team progress and tasks
- D. A testing framework for validating software against requirements

Answer: B

Q316. What is a software vulnerability?

- A. A documentation gap that makes the codebase harder to read now
- B. A hardware component that needs replacement in the server rack
- C. A weakness in software that can be exploited by an attacker now
- D. A planned feature that has not been implemented yet in the code

Answer: C

Q317. What is the principle of least privilege?

- A. Giving all users full administrative access to the entire system now
- B. Granting maximum privileges to reduce the number of access errors
- C. Removing all access controls to simplify the system architecture
- D. Giving users only the minimum access rights needed for their tasks

Answer: D

Q318. What is a security patch?

- A. A software update that fixes known security vulnerabilities quickly
- B. A project management document describing the sprint backlog items
- C. A decorative design element added to the user interface layout now
- D. A hardware component that improves server processing speed rates

Answer: A

Q319. What is malware?

- A. A project management metric tracking team productivity and output
- B. A testing tool for measuring application performance under load
- C. Software designed to damage, disrupt, or gain unauthorized access
- D. A feature enhancement requested by users during product feedback

Answer: C

Q320. What is a security audit?

- A. A performance test measuring application response time and speed
- B. A code review focused only on visual design and layout of pages
- C. A financial review of the project budget expenditure and savings
- D. A systematic evaluation of security measures and vulnerabilities now

Answer: D

Q321. What is the primary goal of software engineering?

- A. To write code as fast as possible
- B. To develop reliable software systematically
- C. To eliminate all bugs
- D. To use the latest technology

Answer: B

Q322. Which of the following is a key attribute of good software?

- A. Complexity
- B. Maintainability
- C. Obscurity
- D. Rigidity

Answer: B

Q323. What is the software crisis?

- A. A shortage of programmers
- B. The difficulty of writing correct, understandable software on time and within budget
- C. A virus that destroys software
- D. The rapid decline in hardware prices

Answer: B

Q324. What does the term 'software process' refer to?

- A. The act of compiling code
- B. A set of activities that leads to the production of a software product
- C. The process of installing software
- D. A method for deleting software

Answer: B

Q325. Which activity is NOT part of a generic software process?

- A. Specification
- B. Development
- C. Validation
- D. Hardware manufacturing

Answer: D

Q326. What is a software product?

- A. Only the executable code
- B. Code, documentation, and associated data
- C. Only the user manual
- D. Only the source code files

Answer: B

Q327. What is meant by software dependability?

- A. Software that depends on hardware
- B. Software that is reliable, secure, and safe
- C. Software that requires constant updates
- D. Software written by a single developer

Answer: B

Q328. Who are the stakeholders in a software project?

- A. Only the programmers
- B. Anyone affected by or involved in the project
- C. Only the customers
- D. Only the project manager

Answer: B

Q329. What does the acronym CASE stand for in software engineering?

- A. Computer-Aided Software Engineering
- B. Central Application System Environment
- C. Code Analysis and Structured Editing
- D. Comprehensive Automated System Evaluation

Answer: A

Q330. What is the difference between generic and customized software products?

- A. There is no difference
- B. Generic products are sold to any customer; customized products are built for a specific client
- C. Generic products are always free
- D. Customized products cannot be modified

Answer: B

Q331. What is a software process model?

- A. A physical model of a computer
- B. A simplified representation of a software development process
- C. A type of programming language
- D. A hardware design specification

Answer: B

Q332. In the Waterfall model, can you return to a previous phase easily?

- A. Yes, it encourages going back
- B. No, it follows a strict linear sequence
- C. Only during testing
- D. Only on weekends

Answer: B

Q333. What is a prototype in software development?

- A. The final product
- B. An early working model used to explore ideas and gather feedback
- C. A type of bug report
- D. A deployment tool

Answer: B

Q334. What is the main advantage of the iterative development model?

- A. It never requires documentation
- B. It allows for early delivery of partial systems and gradual improvement
- C. It eliminates the need for testing
- D. It requires no planning

Answer: B

Q335. Which phase of the Waterfall model involves defining what the system should do?

- A. Implementation
- B. Requirements analysis and definition
- C. System design
- D. Maintenance

Answer: B

Q336. What is incremental development?

- A. Developing the entire system at once
- B. Developing and delivering the system in small portions called increments
- C. Developing only the user interface
- D. Delaying all development until requirements are final

Answer: B

Q337. Which process model is best suited when requirements are well understood and unlikely to change?

- A. Agile model
- B. Waterfall model
- C. Spiral model
- D. Prototype model

Answer: B

Q338. What does the term 'process activity' mean in software engineering?

- A. A recreational team event
- B. A fundamental step in a software development process such as specification or testing
- C. A marketing campaign
- D. A hardware installation step

Answer: B

Q339. What distinguishes throwaway prototyping from evolutionary prototyping?

- A. They are the same approach
- B. Throwaway prototypes are discarded after gathering feedback; evolutionary prototypes are refined into the final product
- C. Throwaway prototyping takes longer
- D. Evolutionary prototyping is always discarded

Answer: B

Q340. What is the purpose of the design phase in a process model?

- A. To write marketing materials
- B. To define the system architecture and component structure
- C. To recruit developers
- D. To delete old code

Answer: B

Q341. How many values are stated in the Agile Manifesto?

- A. 2
- B. 4
- C. 6
- D. 10

Answer: B

Q342. What does the Agile Manifesto value over comprehensive documentation?

- A. Detailed contracts
- B. Working software
- C. Extensive planning
- D. Formal processes

Answer: B

Q343. What is a Sprint Review in Scrum?

- A. A meeting to fire underperforming team members
- B. A meeting at the end of a Sprint where the team demonstrates completed work to stakeholders
- C. A daily coding session
- D. A meeting to write documentation

Answer: B

Q344. What is a retrospective in Agile?

- A. A look at competitor products
- B. A meeting where the team reflects on the process and identifies improvements
- C. A code compilation step
- D. A customer presentation

Answer: B

Q345. Which Agile principle emphasizes face-to-face communication?

- A. Documentation first
- B. The most efficient method of conveying information is face-to-face conversation
- C. Email is preferred
- D. All communication should be written

Answer: B

Q346. What is the role of the Product Owner in Scrum?

- A. Writing all the code
- B. Representing the customer and managing the Product Backlog
- C. Managing servers
- D. Designing the user interface

Answer: B

Q347. What is the purpose of a Software Requirements Specification (SRS)?

- A. To write code
- B. To document what the software should do and its constraints
- C. To deploy software
- D. To test the final product

Answer: B

Q348. What is a stakeholder in requirements engineering?

- A. Only the CEO
- B. Any person or organization that has an interest in or is affected by the system
- C. Only the end user
- D. Only the developer

Answer: B

Q349. What is a feasibility study in requirements engineering?

- A. A study on team morale
- B. An assessment of whether a proposed system is technically, economically, and operationally viable
- C. A review of competitor products
- D. A performance test

Answer: B

Q350. What does 'requirements validation' ensure?

- A. That the code compiles
- B. That the documented requirements accurately reflect stakeholder needs
- C. That the hardware is working
- D. That the team is on schedule

Answer: B

Q351. What is an example of a non-functional requirement?

- A. The system shall allow users to log in
- B. The system shall respond to queries within 2 seconds
- C. The system shall display a list of products
- D. The system shall allow users to add items to a cart

Answer: B

Q352. What is a requirements interview?

- A. A job interview for developers
- B. A technique where analysts ask stakeholders questions to understand their needs
- C. A code review session
- D. A testing technique

Answer: B

Q353. Why is requirements documentation important?

- A. It is not important
- B. It provides a shared reference that reduces misunderstandings and supports verification
- C. It replaces the need for testing
- D. It is only needed for legal reasons

Answer: B

Q354. What is the main purpose of project planning?

- A. To write code
- B. To define the project scope, schedule, and resources needed
- C. To test the software
- D. To deploy the final product

Answer: B

Q355. What is a milestone in project management?

- A. A unit of distance
- B. A significant point or event in the project schedule marking the completion of a major deliverable
- C. A type of programming error
- D. A software tool

Answer: B

Q356. What does the critical path represent in project scheduling?

- A. The easiest path through the project
- B. The longest sequence of dependent tasks that determines the minimum project duration
- C. The shortest task in the project
- D. A backup plan

Answer: B

Q357. What is effort estimation in software projects?

- A. Estimating the weight of hardware
- B. Predicting the amount of work (person-hours or person-months) needed to complete the project
- C. Counting the number of employees
- D. Estimating electricity costs

Answer: B

Q358. What is a project schedule?

- A. A list of holidays
- B. A timeline showing when project activities will be performed and deliverables completed
- C. A team roster
- D. A list of programming languages

Answer: B

Q359. What is the purpose of a project status report?

- A. To advertise the product
- B. To communicate the current state of the project, including progress, issues, and risks to stakeholders
- C. To write code documentation
- D. To create test cases

Answer: B

Q360. What does the term 'scope' mean in project management?

- A. The range of a telescope
- B. The boundaries of the project defining what will and will not be delivered
- C. The number of developers on the team
- D. The programming language used

Answer: B

Q361. What is a PERT chart?

- A. A bar chart showing resource usage
- B. A network diagram showing task dependencies and estimated durations for project scheduling
- C. A pie chart showing budget allocation
- D. A flow chart for code logic

Answer: B

Q362. What is the role of a project manager in software development?

- A. Writing all the code
- B. Planning, organizing, and controlling project activities to meet objectives within constraints
- C. Only attending meetings
- D. Designing the database

Answer: B

Q363. What is information hiding in software design?

- A. Encrypting all data
- B. Restricting access to module internals so that other modules interact only through well-defined interfaces
- C. Hiding code from developers
- D. Deleting documentation

Answer: B

Q364. What is a class diagram in object-oriented design?

- A. A classroom schedule
- B. A UML diagram showing classes, their attributes, methods, and relationships
- C. A network topology diagram
- D. A flowchart of code execution

Answer: B

Q365. What is the purpose of architectural design?

- A. To choose office furniture
- B. To define the overall structure and organization of the software system
- C. To write unit tests
- D. To deploy the software

Answer: B

Q366. What is stepwise refinement?

- A. Refining sugar for cooking
- B. A top-down design strategy that progressively decomposes a high-level design into more detailed components
- C. Optimizing database queries
- D. A testing technique

Answer: B

Q367. What is an interface in object-oriented design?

- A. A monitor screen
- B. A contract that specifies a set of methods that a class must implement
- C. A type of database
- D. A network protocol

Answer: B

Q368. What does 'high cohesion' mean in software design?

- A. Modules that do many unrelated things
- B. A module whose elements are closely related and focused on a single well-defined purpose
- C. High coupling between modules
- D. A module with many dependencies

Answer: B

Q369. What is a sequence diagram used for?

- A. Showing the order of DNA sequences
- B. Showing the order of interactions between objects over time
- C. Showing the project schedule
- D. Showing the database schema

Answer: B

Q370. What is a peer-to-peer architecture?

- A. An architecture with a central server
- B. An architecture where each node acts as both client and server, sharing resources directly
- C. A hierarchical architecture
- D. A single-user system

Answer: B

Q371. What is the role of middleware in software architecture?

- A. To design user interfaces
- B. To provide services that enable communication and data exchange between different applications or components
- C. To write documentation
- D. To test software

Answer: B

Q372. What is the pipe-and-filter architectural style?

- A. A plumbing design pattern
- B. An architecture where data flows through a series of processing components (filters) connected by channels (pipes)
- C. A testing methodology
- D. A security technique

Answer: B

Q373. What is scalability in software architecture?

- A. The size of the codebase
- B. The ability of a system to handle increased workload by adding resources without fundamental changes
- C. The number of files in a project
- D. The font size in the UI

Answer: B

Q374. What is a repository architectural style?

- A. A version control system
- B. An architecture where components share data through a central data store
- C. A type of database
- D. A code hosting platform

Answer: B

Q375. What is the purpose of an architectural view?

- A. A scenic view from the office
- B. A representation of the system architecture from a particular perspective or concern
- C. A code review
- D. A test report

Answer: B

Q376. What is a component in software architecture?

- A. A hardware chip
- B. A modular, deployable, and replaceable unit of software that encapsulates a set of related functions
- C. A test case
- D. A project document

Answer: B

Q377. What is the purpose of user-centered design?

- A. Designing for developers only
- B. Placing the needs, abilities, and limitations of end users at the center of the design process
- C. Designing only the database
- D. Focusing exclusively on aesthetics

Answer: B

Q378. What is the purpose of visual hierarchy in UI design?

- A. Making all elements the same size
- B. Organizing and prioritizing elements so users naturally see the most important information first
- C. Hiding important information
- D. Making the design more complex

Answer: B

Q379. What is a tooltip in UI design?

- A. A physical tool tip
- B. A small informational pop-up that appears when a user hovers over or focuses on an element
- C. A debugging message
- D. A server error message

Answer: B

Q380. What is white space (negative space) in UI design?

- A. A bug in the interface
- B. The empty areas between UI elements that improve readability and reduce visual clutter
- C. Unused storage space
- D. Deleted content areas

Answer: B

Q381. What is a variable naming convention?

- A. A meeting about variable names
- B. A set of rules for how variables should be named in code to improve readability
- C. A type of data structure
- D. A debugging technique

Answer: B

Q382. What is an IDE (Integrated Development Environment)?

- A. A type of programming language
- B. A software application that provides comprehensive tools for coding, debugging, testing, and building software in one interface
- C. A database management system
- D. A project management tool

Answer: B

Q383. What is the difference between a compiler and an interpreter?

- A. They are the same thing
- B. A compiler translates the entire source code at once into machine code; an interpreter translates and executes code line by line
- C. Compilers are slower
- D. Interpreters produce machine code files

Answer: B

Q384. What is version control used for in implementation?

- A. Controlling the version of the operating system
- B. Tracking and managing changes to source code over time, allowing collaboration and rollback
- C. Controlling access to the internet
- D. Managing database versions

Answer: B

Q385. What is a build process?

- A. Constructing a building
- B. The automated process of compiling source code, linking dependencies, and producing executable software
- C. A meeting to plan features
- D. A manual testing procedure

Answer: B

Q386. What is a test suite?

- A. A hotel room for testers
- B. A collection of test cases organized together to test a particular feature or aspect of the system
- C. A single test case
- D. A testing tool

Answer: B

Q387. What is a test oracle?

- A. A fortune-telling tool
- B. A mechanism or reference used to determine whether a test has passed or failed by comparing actual results to expected results
- C. A database management tool
- D. A type of programming language

Answer: B

Q388. What is the difference between a bug and a defect?

- A. They are completely different concepts
- B. They are generally used interchangeably to refer to a flaw in the software that causes incorrect behavior
- C. Bugs are in hardware; defects are in software
- D. Defects are intentional; bugs are not

Answer: B

Q389. What is preventive maintenance in software?

- A. Preventing users from using the software
- B. Making changes to software to prevent potential future problems before they occur
- C. Preventing new features from being added
- D. Uninstalling the software

Answer: B

Q390. What percentage of a software system's total lifecycle cost is typically spent on maintenance?

- A. Less than 10%
- B. 60% to 80% or more
- C. Exactly 50%
- D. Less than 5%

Answer: B

Q391. What is a patch in software maintenance?

- A. A fabric patch for clothing
- B. A small update released to fix specific bugs or security vulnerabilities without a full new version
- C. A complete system rewrite
- D. A new feature release

Answer: B

Q392. What triggers corrective maintenance?

- A. A request for new features
- B. The discovery of bugs or defects that need to be fixed
- C. A change in management
- D. A system upgrade

Answer: B

Q393. What triggers adaptive maintenance?

- A. User complaints about the color scheme
- B. Changes in the operating environment such as new hardware, operating systems, or regulations
- C. Developer preferences
- D. Marketing campaigns

Answer: B

Q394. What is a legacy system?

- A. A system that was recently built
- B. An older software system that remains in use because it fulfills a critical business need despite using outdated technology
- C. A system with no users
- D. A brand-new prototype

Answer: B

Q395. What is the purpose of a maintenance request?

- A. Requesting a day off
- B. A formal document describing a needed change, fix, or enhancement to an existing software system
- C. Requesting new hardware
- D. Requesting a team meeting

Answer: B

Q396. What is the purpose of quality assurance in software?

- A. To find all bugs
- B. To establish processes and standards that prevent defects and ensure quality throughout development
- C. To replace testing
- D. To increase project costs

Answer: B

Q397. What is a quality standard?

- A. A high-quality flag
- B. A documented set of criteria, guidelines, and best practices that products or processes must meet
- C. A type of programming language
- D. A testing tool

Answer: B

Q398. What is a software audit?

- A. A financial audit of the company
- B. An independent examination of software products or processes to assess compliance with standards and specifications
- C. A code compilation step
- D. A deployment procedure

Answer: B

Q399. What is a checklist in QA?

- A. A grocery list
- B. A predefined list of items or criteria to be verified during quality reviews to ensure completeness
- C. A list of team members
- D. A deployment schedule

Answer: B

Q400. What is process improvement in QA?

- A. Buying better hardware
- B. Systematically analyzing and enhancing development processes to increase product quality and efficiency
- C. Hiring more developers
- D. Increasing project budgets

Answer: B

Q401. What is a defect report?

- A. A report card for developers
- B. A document that describes a found defect, including steps to reproduce, expected behavior, and actual behavior
- C. A financial report
- D. A project plan

Answer: B

Q402. What is peer review in the context of QA?

- A. A review of peer-to-peer networks
- B. A process where colleagues examine each other's work products to identify defects and share knowledge
- C. A performance review by managers
- D. A customer review

Answer: B

Q403. What is the purpose of software metrics?

- A. To count developers
- B. To quantitatively measure attributes of software products and processes for decision-making and improvement
- C. To measure hardware performance
- D. To count meetings

Answer: B

Q404. What is the difference between size metrics and complexity metrics?

- A. They are the same
- B. Size metrics measure the quantity of code (like LOC); complexity metrics measure how difficult the code is to understand and maintain
- C. Size is always more important
- D. Complexity cannot be measured

Answer: B

Q405. What is defect removal efficiency (DRE)?

- A. The speed of removing code
- B. The percentage of defects found before software release compared to the total defects found before and after release
- C. The efficiency of the development team
- D. The speed of code compilation

Answer: B

Q406. What is the purpose of measuring lines of code (LOC)?

- A. To determine developer pay
- B. To estimate the size of a software system, which can be used for effort estimation and productivity measurement
- C. To count the number of files
- D. To measure code quality directly

Answer: B

Q407. What is a project metric?

- A. A metric about the project manager
- B. A measurement that tracks project progress and performance, such as schedule variance, cost variance, or team velocity
- C. A metric about the building
- D. A metric about company revenue

Answer: B

Q408. Why can LOC be misleading as a productivity metric?

- A. LOC is always accurate
- B. It penalizes efficient code that achieves more with fewer lines and varies greatly by programming language
- C. LOC is never misleading
- D. All lines of code are equally valuable

Answer: B

Q409. What is a baseline in software metrics?

- A. A fishing line
- B. A set of reference measurements collected from past projects used for comparison and estimation in future projects
- C. A line of code
- D. A base class

Answer: B

Q410. What is a direct metric versus an indirect metric?

- A. They are the same
- B. A direct metric is measured directly (like LOC); an indirect metric is derived from other measurements (like defect density)
- C. Direct metrics are more complex
- D. Indirect metrics are always inaccurate

Answer: B

Q411. What is the purpose of a version control system?

- A. To control the version of the operating system
- B. To track and manage changes to files over time, enabling collaboration and history tracking
- C. To control file access permissions
- D. To manage database versions

Answer: B

Q412. What is the purpose of a release in configuration management?

- A. Releasing employees
- B. A formally distributed version of the software that has been tested and approved for deployment to users
- C. Releasing disk space
- D. Releasing a news article

Answer: B

Q413. What is a repository in version control?

- A. A physical storage room
- B. A central storage location that contains all files, their version history, and metadata managed by the version control system
- C. A library building
- D. A data center

Answer: B

Q414. What is a change request in configuration management?

- A. A request to change office furniture
- B. A formal proposal to modify a configuration item, documenting the reason, impact, and details of the requested change
- C. A request for a new computer
- D. A request for vacation

Answer: B

Q415. What is a merge in version control?

- A. Merging two companies
- B. Combining changes from different branches or versions into a single unified version
- C. Merging database tables
- D. Merging network connections

Answer: B

Q416. What is the difference between risk avoidance and risk acceptance?

- A. They are the same strategy
- B. Risk avoidance changes plans to eliminate the risk; risk acceptance acknowledges the risk and chooses not to take action unless it occurs
- C. Risk avoidance is always possible
- D. Risk acceptance means ignoring risks

Answer: B

Q417. What is the difference between a risk and an issue?

- A. They are the same thing
- B. A risk is a potential future problem; an issue is a problem that has already occurred
- C. Risks are always worse than issues
- D. Issues cannot be managed

Answer: B

Q418. What is the probability of a risk?

- A. Always 100%
- B. The likelihood that the risk event will actually occur, typically expressed as a percentage or qualitative rating
- C. Always 0%
- D. The cost of the risk

Answer: B

Q419. Why is early risk identification important?

- A. It is not important
- B. Identifying risks early allows more time and options for developing effective responses before they become costly problems
- C. Early identification wastes time
- D. Risks should only be identified at the end

Answer: B

Q420. What is encryption in software security?

- A. Compressing files
- B. The process of converting data into a coded format that can only be read by someone with the proper decryption key
- C. Deleting files permanently
- D. Backing up data

Answer: B

Q421. What is a firewall in network security?

- A. A wall that is on fire
- B. A security system that monitors and controls incoming and outgoing network traffic based on predetermined security rules
- C. A physical wall in a building
- D. A type of antivirus software

Answer: B

Q422. What is a vulnerability in software security?

- A. A strong point in the software
- B. A weakness or flaw in the software that could be exploited by an attacker to compromise the system
- C. A feature request
- D. A type of software update

Answer: B

Q423. What is multi-factor authentication (MFA)?

- A. Using the same password multiple times
- B. A security mechanism requiring users to provide two or more different types of verification to prove their identity
- C. Multiple usernames
- D. Multiple login pages

Answer: B

Q424. What is the purpose of access control in software?

- A. Controlling physical building access
- B. Restricting who can access what resources and what actions they can perform within the system
- C. Controlling internet speed
- D. Controlling screen brightness

Answer: B

Q425. What is phishing?

- A. Going fishing
- B. A social engineering attack that tricks users into revealing sensitive information by impersonating a trustworthy entity
- C. A type of programming
- D. A network protocol

Answer: B

Q426. What is data integrity in security?

- A. The amount of data stored
- B. Ensuring that data has not been altered, tampered with, or corrupted in an unauthorized manner
- C. The speed of data transfer
- D. The format of data storage

Answer: B

Q427. What is the purpose of security logging?

- A. Logging into secure systems
- B. Recording security-relevant events and actions to enable monitoring, auditing, and forensic analysis
- C. Logging trees for wood
- D. Creating log files for debugging

Answer: B

Q428. What is a Sprint Backlog in Scrum?

- A. A list of all features ever requested for the product
- B. The set of product backlog items selected for the current sprint plus a plan for delivering them
- C. A document that tracks bugs found during testing
- D. A report generated after each sprint is completed

Answer: B

Q429. What does the term 'iteration' mean in Agile development?

- A. A single line of code written by a developer
- B. A fixed-length time period during which a team works on a set of features and produces a working increment
- C. The process of hiring new team members
- D. A method for deploying software to production servers

Answer: B

Q430. What is the purpose of a burndown chart in Agile?

- A. To track the budget spent on a project
- B. To show the remaining work in a sprint or project over time
- C. To list all the bugs found during development
- D. To assign tasks to individual team members

Answer: B

Q431. What is the recommended size of a Scrum team?

- A. Exactly 3 members
- B. Between 20 and 50 members
- C. Typically 5 to 9 members plus Scrum Master and Product Owner
- D. There is no recommended size for Scrum teams

Answer: C

Q432. What is the difference between functional and non-functional requirements?

- A. Functional requirements describe what the system should do while non-functional requirements describe how well it should perform
- B. Functional requirements are optional while non-functional requirements are mandatory
- C. Non-functional requirements describe features while functional requirements describe constraints
- D. There is no difference; they are interchangeable terms

Answer: A

Q433. What is a requirements review meeting?

- A. A meeting where developers write code based on requirements
- B. A meeting where stakeholders and the team examine requirements documents to identify errors, ambiguities, and omissions
- C. A meeting held only after the software is deployed to users
- D. A meeting to decide the project budget and timeline

Answer: B

Q434. What is a work package in project management?

- A. The entire project plan document
- B. The smallest unit of work in a Work Breakdown Structure that can be estimated, scheduled, and monitored
- C. A software package used for managing projects
- D. The final deliverable shipped to the customer

Answer: B

Q435. What is the principle of separation of concerns in software design?

- A. Keeping the development team separated into isolated groups
- B. Dividing a program into distinct sections where each section addresses a separate concern or functionality
- C. Ensuring that all code is written in a single file for simplicity
- D. Separating the testing phase from the development phase

Answer: B

Q436. What is an activity diagram in software design?

- A. A diagram showing the physical hardware components of a system
- B. A flowchart-like diagram that shows the flow of activities and actions in a process or system
- C. A chart showing project team member availability
- D. A diagram that displays database table relationships

Answer: B

Q437. What does 'low coupling' mean in software design?

- A. Modules have many dependencies on each other and share data extensively
- B. Modules are relatively independent with minimal dependencies between them
- C. The software has very few modules in total
- D. All modules are combined into a single large component

Answer: B

Q438. What is a distributed system in software architecture?

- A. A system where all code runs on a single computer
- B. A system whose components are located on different networked computers that communicate and coordinate actions by passing messages
- C. A system that is distributed to users via USB drives
- D. A system that only works without an internet connection

Answer: B

Q439. What is a prototype in UI design?

- A. The final production-ready version of the software
- B. An early model or simulation of the user interface used to explore and test design ideas before development
- C. A type of programming language used for building interfaces
- D. A document listing all the bugs found in the interface

Answer: B

Q440. What is a modal dialog in UI design?

- A. A full-screen window that replaces the entire application
- B. A window or box that appears on top of the main content and requires user interaction before they can return to the main interface
- C. A type of navigation menu that appears at the bottom of the screen
- D. A background process that runs without any visual indication

Answer: B

Q441. What is a breadcrumb in UI design?

- A. A type of animation effect applied to buttons
- B. A navigation aid that shows the user's current location within the site hierarchy and allows navigation back to parent pages
- C. A small cookie stored in the user's browser
- D. A debugging tool for tracking user errors

Answer: B

Q442. What does 'affordance' mean in UI design?

- A. The cost of developing the user interface
- B. A design property that suggests how an element should be used, such as a button looking clickable
- C. The number of colors used in the interface
- D. The amount of text displayed on a single screen

Answer: B

Q443. What is user-centered design (UCD)?

- A. A design approach where developers make all design decisions without user input
- B. A design philosophy that places the end user's needs, preferences, and limitations at the center of each stage of the design process
- C. A technique for designing database schemas based on user data
- D. A project management method for scheduling design tasks

Answer: B

Q444. What is a runtime error in software development?

- A. An error that occurs only during the design phase
- B. An error that occurs while the program is executing, causing it to crash or produce incorrect results
- C. An error in the project management timeline
- D. An error that can only be detected by reading the source code manually

Answer: B

Q445. What is a logic error in programming?

- A. An error that prevents the program from compiling
- B. An error where the program runs without crashing but produces incorrect results due to a flaw in the algorithm or logic
- C. An error caused by hardware malfunction
- D. An error that only occurs when the program is deployed to production

Answer: B

Q446. What is a code repository used for?

- A. Storing only compiled binary files for deployment
- B. Storing, managing, and tracking changes to source code files in a centralized or distributed location
- C. Providing a graphical interface for writing code
- D. Running automated tests on the production server

Answer: B

Q447. What is pseudocode?

- A. A real programming language used for production systems
- B. An informal, high-level description of an algorithm using natural language and basic programming constructs
- C. Encrypted source code that cannot be read by humans
- D. A tool for automatically generating documentation from code

Answer: B

Q448. What is the purpose of indentation in source code?

- A. It makes the program run faster by optimizing memory usage
- B. It improves code readability by visually representing the structure and nesting of code blocks
- C. It is required for all programming languages to compile successfully
- D. It adds comments to the code automatically

Answer: B

Q449. What is a test environment?

- A. The office where testers sit and work
- B. A configured setup of hardware, software, and network used to execute tests that simulates the production environment
- C. A programming language used exclusively for writing tests
- D. The final version of the software deployed to end users

Answer: B

Q450. What is the purpose of smoke testing?

- A. To test the software under extreme stress conditions
- B. To perform a quick preliminary test to check whether the most critical functions of the software work before deeper testing begins
- C. To test the software's behavior when the hardware overheats
- D. To measure the performance of the software under normal load

Answer: B

Q451. What is alpha testing?

- A. Testing performed by end users in their own environment
- B. Testing performed by internal staff at the developer's site before release to external users
- C. Testing performed automatically without any human involvement
- D. The first test case written for any software project

Answer: B

Q452. What is beta testing?

- A. Testing performed only by the development team using automated tools
- B. Testing performed by a limited number of external users in their real environment before the final release
- C. Testing that occurs after the software has been fully deployed and is in production
- D. A type of unit testing focused on mathematical calculations

Answer: B

Q453. What is a test fixture in software testing?

- A. A physical device used to test hardware components
- B. A fixed state or set of preconditions used as a baseline for running tests to ensure consistent and repeatable results
- C. A metric that measures the percentage of code tested
- D. A report generated after all tests have been executed

Answer: B

Q454. What does 'test coverage' measure?

- A. The number of testers assigned to a project
- B. The degree to which the source code or requirements are exercised by a set of tests
- C. The amount of time spent on testing activities
- D. The cost of testing relative to the total project budget

Answer: B

Q455. What is software re-engineering?

- A. Building completely new software from scratch without referencing existing systems
- B. The process of examining and restructuring existing software to reconstitute it in a new form while preserving its functionality
- C. A marketing strategy for promoting outdated software products
- D. The process of removing all documentation from a legacy system

Answer: B

Q456. Why is regression testing important during software maintenance?

- A. It only tests new features and ignores existing functionality
- B. It ensures that changes or fixes to the software have not introduced new defects in previously working functionality
- C. It measures the financial cost of maintaining the software
- D. It is only performed when the entire software is being rewritten

Answer: B

Q457. What is the difference between software maintenance and software support?

- A. They are identical terms with no difference
- B. Maintenance involves modifying the software code and system while support involves helping users resolve issues and answering questions
- C. Support involves changing the code while maintenance involves only documentation updates
- D. Maintenance is performed by users while support is performed by developers

Answer: B

Q458. What is a code walkthrough in quality assurance?

- A. The physical act of walking around the office while reviewing code
- B. An informal review where a developer leads team members through code to find errors and share knowledge
- C. An automated tool that scans code for security vulnerabilities
- D. A formal audit conducted by external quality assessors

Answer: B

Q459. What is the difference between a defect and a failure in quality assurance?

- A. They are exactly the same thing with no difference
- B. A defect is a flaw in the software code or design, while a failure is the observable incorrect behavior when a defect is executed
- C. A failure occurs during development while a defect occurs during testing
- D. Defects are found by users while failures are found by developers

Answer: B

Q460. What is the purpose of a quality management plan?

- A. To list all the features the software should have
- B. To define the quality standards, processes, metrics, and responsibilities for ensuring that the project meets quality requirements
- C. To schedule marketing activities for the software release
- D. To calculate the total budget for the development project

Answer: B

Q461. What is defect density as a software metric?

- A. The total number of developers on a project
- B. The number of defects found per unit size of software, typically per thousand lines of code (KLOC)
- C. The time it takes to fix a single defect
- D. The percentage of the project budget spent on defect prevention

Answer: B

Q462. What is a pull request in version control?

- A. A request to download the latest version of the software
- B. A proposal to merge changes from one branch into another, allowing team members to review and discuss the changes before integration
- C. A request to remove files from the repository
- D. A notification that the server is pulling updates automatically

Answer: B

Q463. What is the purpose of a .gitignore file?

- A. To list all the files that should be included in every commit
- B. To specify patterns of files and directories that Git should not track or include in the repository
- C. To store Git configuration settings for the repository
- D. To record the history of all commits made to the repository

Answer: B

Q464. What is a merge conflict in version control?

- A. When two repositories have the same name and cannot coexist
- B. When two developers have made changes to the same part of a file on different branches and the system cannot automatically combine them
- C. When a developer accidentally deletes the entire repository
- D. When the version control server runs out of storage space

Answer: B

Q465. What is a software build in configuration management?

- A. The physical construction of hardware to run the software
- B. The process of compiling source code, linking dependencies, and packaging the result into an executable or deployable artifact
- C. A meeting where the team discusses the next sprint
- D. A document describing the software architecture

Answer: B

Q466. What is a contingency plan in risk management?

- A. The original project plan that never changes
- B. A predefined set of actions to be taken if a specific risk event occurs
- C. A plan for hiring additional developers
- D. A marketing strategy for the software product

Answer: B

Q467. What is risk transfer in software project management?

- A. Moving the project to a different department
- B. Shifting the responsibility or financial impact of a risk to a third party, such as through insurance or outsourcing
- C. Ignoring the risk and hoping it does not occur
- D. Transferring all project files to a new repository

Answer: B

Q468. What is the purpose of risk monitoring in a software project?

- A. To eliminate all risks from the project permanently
- B. To continuously track identified risks, watch for new risks, and evaluate the effectiveness of risk responses throughout the project
- C. To create a final report after the project is complete
- D. To assign blame when things go wrong

Answer: B

Q469. What is risk impact in software risk management?

- A. The likelihood that a risk will occur during the project
- B. The severity of the consequences if the risk event actually occurs
- C. The number of team members affected by project risks
- D. The total number of risks identified in the risk register

Answer: B

Q470. What is a risk assessment matrix?

- A. A spreadsheet listing all team members and their salaries
- B. A grid that maps risks by their probability of occurrence and severity of impact to prioritize which risks need attention
- C. A mathematical formula for calculating project costs
- D. A chart showing the project timeline and milestones

Answer: B

Q471. What is a security threat in software engineering?

- A. A feature that improves the software's performance
- B. Any potential danger that could exploit a vulnerability to cause harm to a system, organization, or users
- C. A type of software license agreement
- D. A coding standard that developers must follow

Answer: B

Q472. What is input sanitization in software security?

- A. Physically cleaning the keyboard and mouse used for development
- B. The process of cleaning and validating user input to remove or neutralize potentially harmful characters or code before processing
- C. Deleting all user data from the database periodically
- D. A network configuration technique for blocking external connections

Answer: B

Q473. What is the purpose of a requirements baseline in software engineering?

- A. To set a maximum number of requirements allowed in a project
- B. To establish a formally approved set of requirements that serves as the agreed-upon foundation for development
- C. To create a backup copy of the source code repository
- D. To determine the programming language for the project

Answer: B

Q474. What is a service in service-oriented architecture (SOA)?

- A. A physical server that hosts websites
- B. A self-contained unit of functionality that can be accessed remotely and used independently
- C. A customer support team that helps users
- D. A database table that stores user information

Answer: B

Q475. What is fault tolerance in software architecture?

- A. The ability of a system to blame developers for errors
- B. The ability of a system to continue operating properly even when some of its components fail
- C. A technique for writing code without any bugs
- D. The process of testing software until no faults remain

Answer: B

Q476. What is a dropdown menu in user interface design?

- A. A menu that permanently displays all options on screen at all times
- B. A UI element that reveals a list of options when activated, allowing the user to select one from the list
- C. A type of error message that drops down from the top of the screen
- D. A navigation bar fixed at the bottom of the page

Answer: B

Q477. What is a defect report (bug report) in software testing?

- A. A marketing document describing software features
- B. A formal document that records the details of a discovered defect including steps to reproduce, expected behavior, and actual behavior
- C. A project schedule showing testing deadlines
- D. A contract between the client and the development company

Answer: B

Q478. What is code churn as a software metric?

- A. The total number of developers who have ever worked on the project
- B. A measure of the amount of code that is added, modified, or deleted over a period of time
- C. The speed at which the compiler processes source code files
- D. The number of comments in the source code divided by total lines

Answer: B

Q479. What is a diff in version control?

- A. A type of programming error that is difficult to find
- B. A representation showing the differences between two versions of a file, highlighting added, removed, and changed lines
- C. A command for deleting files from the repository permanently
- D. A way to create a new branch in the repository

Answer: B

Q480. What is risk mitigation in software project management?

- A. Completely eliminating all risks from the project
- B. Taking proactive steps to reduce the probability or impact of an identified risk
- C. Ignoring risks until they become actual problems
- D. Transferring the entire project to another team

Answer: B

Software Engineering — Medium

480 questions

Q481. Which IEEE standard defines software engineering?

- A. IEEE 754
- B. IEEE 610.12
- C. IEEE 802.11
- D. IEEE 1394

Answer: B

Q482. What is the difference between software engineering and computer science?

- A. CS primarily focuses on project management
- B. They are fundamentally the same discipline
- C. CS focuses on theory; SE focuses on practical development
- D. SE primarily focuses on hardware architecture

Answer: C

Q483. Which of the following is a generic software process activity?

- A. Software specification
- B. Network setup
- C. Hardware installation
- D. Printing reports

Answer: A

Q484. What is 'software reuse'?

- A. Using existing software components in new systems
- B. Writing all source code from scratch every time
- C. Permanently deleting outdated legacy software
- D. Copying proprietary code without any permission

Answer: A

Q485. Which factor contributed most to the software crisis?

- A. Overly simple user requirements specifications
- B. Rapidly advancing hardware computing capabilities
- C. An oversaturation of qualified trained programmers
- D. Increasing software complexity outpacing development methods

Answer: D

Q486. What is a software engineering method?

- A. A specific type of programming language tool
- B. A classification of computing hardware types
- C. A structured approach for developing software
- D. A specialized device for testing code logic

Answer: C

Q487. What is the significance of the NATO Conference 1968?

- A. The initial public launch of the internet era for the development team
- B. The initial development and creation of the UNIX system
- C. The original invention of the C programming language throughout the project
- D. First formal discussion of the software crisis and software engineering

Answer: D

Q488. Which is NOT a type of software application?

- A. Hardware software
- B. System software
- C. Web applications
- D. Embedded software

Answer: A

Q489. What does CASE stand for in software engineering?

- A. Computer-Aided Software Engineering
- B. Common Architecture System Engineering
- C. Coded Application Standard Environment
- D. Central Application Software Environment

Answer: A

Q490. What is the role of a software engineer?

- A. Writing source code exclusively and nothing else
- B. Administering and managing databases exclusively
- C. Designing, developing, testing, and maintaining software systems
- D. Performing software testing activities exclusively

Answer: C

Q491. What are the four quadrants of the Spiral model?

- A. Planning, Risk Analysis, Engineering, Evaluation
- B. Gather, Analyze, Build, Ship within the project scope
- C. Design, Code, Test, Review during the software lifecycle
- D. Code, Test, Deploy, Maintain in the development process

Answer: A

Q492. What is RAD (Rapid Application Development)?

- A. An incremental model emphasizing short development cycles of 60-90 days
- B. A standardized software testing framework tool for the development team
- C. A formal documentation standard for projects throughout the project
- D. A general-purpose programming language format in the system context

Answer: A

Q493. Which model combines elements of the Waterfall and Iterative models?

- A. Spiral model
- B. Big Bang model
- C. RAD
- D. Incremental model

Answer: A

Q494. What is the incremental model?

- A. A model that skips the planning phase for quality purposes
- B. Building the complete system all at once during implementation
- C. A model used only for testing purposes in practice typically
- D. Delivering the system in increments, each adding functionality

Answer: D

Q495. In the V-Model, what corresponds to the coding phase?

- A. Integration testing
- B. System testing
- C. Unit testing
- D. Acceptance testing

Answer: C

Q496. What is evolutionary development?

- A. Developing software that remains unchanged permanently as a standard approach
- B. Only creating project plans without writing code by the organization
- C. Developing an initial version, exposing it to users, and refining through versions
- D. Focusing solely on hardware evolution processes at every stage

Answer: C

Q497. What is the main advantage of the incremental model?

- A. The model requires absolutely no testing phase in a systematic way
- B. Working software is produced early and feedback can be incorporated
- C. No formal planning phase is ever required across all phases
- D. It is always the fastest model available today for the project goals

Answer: B

Q498. What is a phase gate in process models?

- A. An established coding standard for the team
- B. A specific category of reported software bug
- C. A review point at the end of each phase before proceeding
- D. A physical entry gate for developer access

Answer: C

Q499. Which model is best suited for large, mission-critical projects?

- A. Code and Fix
- B. Big Bang accord
- C. Spiral model
- D. RAD within the

Answer: C

Q500. What is throwaway prototyping?

- A. A category of finalized production product only and its related activities
- B. Retaining and maintaining all prototypes built by the development process
- C. A formal software testing methodology approach over the entire lifecycle
- D. Building a prototype to understand requirements, then discarding it

Answer: D

Q501. What is the role of the Scrum Master?

- A. Writing all of the source code for the team
- B. Designing the relational database architecture
- C. Managing the overall project budget allocations
- D. Facilitating the Scrum process and removing impediments

Answer: D

Q502. What is Extreme Programming (XP)?

- A. A video game development methodology only throughout the project
- B. An Agile methodology emphasizing engineering practices like pair programming and TDD
- C. A project management scheduling tool only in the system context
- D. A unit testing automation framework only during implementation for quality purposes

Answer: B

Q503. What is a Product Backlog?

- A. A finalized list of all completed work tasks
- B. An ordered list of everything needed in the product
- C. A comprehensive list of known bugs and issues
- D. A physical hardware inventory listing sheet

Answer: B

Q504. What is the purpose of a Sprint Retrospective?

- A. To reflect on the sprint and identify improvements
- B. To plan all the tasks for the upcoming sprint
- C. To demonstrate completed features to stakeholders
- D. To write all required project documentation

Answer: A

Q505. What is Kanban in software development?

- A. A visual workflow management method using boards and cards
- B. A general-purpose programming language tool
- C. A comprehensive testing automation framework
- D. A relational database management system

Answer: A

Q506. What is 'velocity' in Agile?

- A. The amount of work a team can complete in a sprint
- B. The physical typing speed of the developers
- C. The network communication transfer bandwidth
- D. The code compilation and build execution speed

Answer: A

Q507. What is pair programming?

- A. Two programmers working together at one workstation
- B. Two people using the same computer for different tasks
- C. Two separate teams coding on different features
- D. Programming in pairs of related source files

Answer: A

Q508. What is a Sprint Review?

- A. A comprehensive team evaluation process event
- B. A meeting to demonstrate completed work to stakeholders
- C. A formal source code review session for the team
- D. An individual performance review meeting session

Answer: B

Q509. What is the 'Definition of Done' in Scrum?

- A. A shared understanding of what it means for work to be complete
- B. When the source code compiles without errors
- C. When the project manager gives final approval
- D. When all documentation has been written and reviewed

Answer: A

Q510. What is Continuous Integration in Agile?

- A. Integrating code once at the very end of the project in practice typically
- B. Frequently merging code changes into a shared repository with automated testing
- C. Performing only manual source code reviews periodically as a standard approach
- D. Conducting annual system integration events for release by the organization

Answer: B

Q511. What is the difference between verification and validation in requirements?

- A. Verification checks if we're building the product right; validation checks if we're building the right product
- B. They are fundamentally identical concepts entirely for the project goals according to best practices
- C. Verification is testing code; validation is writing code within the system boundary
- D. Validation always occurs before verification ever does by the development process and its related activities

Answer: A

Q512. What is requirements traceability?

- A. Tracking individual developer working hours over the entire lifecycle
- B. The ability to trace each requirement through design, implementation, and testing
- C. An established coding standard for the project for all project stakeholders
- D. A specific debugging and troubleshooting technique as defined by standards

Answer: B

Q513. What is a use case diagram?

- A. A process execution flowchart diagram only within the given constraints
- B. A relational database schema diagram layout in all development efforts
- C. A UML diagram showing actors and their interactions with the system
- D. A network topology and routing diagram map by the project team members

Answer: C

Q514. What is requirements prioritization?

- A. Permanently deleting all unimportant requirements
- B. Ranking requirements by importance, risk, or business value
- C. Grouping requirements by assigned developer team
- D. Alphabetical ordering of all requirements listed

Answer: B

Q515. What is prototyping used for in requirements engineering?

- A. Delivering the final production product to users
- B. Performing application security penetration tests
- C. Conducting system performance load testing runs
- D. Helping stakeholders visualize and validate requirements

Answer: D

Q516. What is a scenario in requirements engineering?

- A. A real-life example describing how a system might be used
- B. An established coding guideline for the teams
- C. A comprehensive deployment plan document only
- D. A specific automated test execution script only

Answer: A

Q517. What is requirements negotiation?

- A. Rejecting all stakeholder proposed requirements
- B. Resolving conflicts between stakeholders regarding requirements
- C. Demanding that all requirements absolutely be met
- D. Deliberately ignoring all stakeholder input

Answer: B

Q518. What is domain analysis in requirements engineering?

- A. Studying specific programming language syntax in the engineering discipline
- B. Studying the application domain to understand its concepts and constraints
- C. Analyzing the database schema domain structure and related components
- D. Analyzing network domain name configurations to achieve project objectives

Answer: B

Q519. What is the IEEE 830 standard?

- A. A networking communication protocol standard
- B. A standard for writing software requirements specifications
- C. A software testing methodology and standard
- D. A hardware design and manufacturing standard

Answer: B

Q520. What is a constraint in requirements engineering?

- A. A defect discovered in the source code base for effective project outcomes
- B. A restriction or limitation on the system or development process
- C. A programming language built-in feature set in the development process
- D. A specific type of software test case only as part of the methodology

Answer: B

Q521. What is the COCOMO model?

- A. A comprehensive software testing methodology and its related activities
- B. A creational object-oriented design pattern over the entire lifecycle
- C. A popular chocolate confectionery brand name by the development process
- D. A model for estimating software development effort based on lines of code

Answer: D

Q522. What is function point analysis?

- A. A method for measuring software size based on functionality delivered to users
- B. A specific software testing automation technique for all project stakeholders
- C. A systematic debugging and tracing methodology within the given constraints
- D. Counting total functions in application code as defined by standards

Answer: A

Q523. What is Earned Value Management (EVM)?

- A. A corporate employee salary payment system in all development efforts
- B. A technique for measuring project performance by comparing planned vs actual progress
- C. An established coding standard and guideline set and related components
- D. A specific software testing coverage metric only by the project team members

Answer: B

Q524. What is the difference between PERT and CPM?

- A. PERT does not produce any chart output at all for effective project outcomes
- B. The two techniques are completely identical overall to achieve project objectives
- C. CPM is exclusively used for software projects in the engineering discipline
- D. PERT uses probabilistic time estimates; CPM uses deterministic time estimates

Answer: D

Q525. What is slack (float) in project scheduling?

- A. A specific category of software defect type during the software lifecycle
- B. Additional unnecessary source code that was written in the development process
- C. Being unproductive and lazy at work always as part of the methodology
- D. The amount of time a task can be delayed without delaying the project

Answer: D

Q526. What is a project baseline?

- A. The very first line of application source code for the development team
- B. An approved version of the project plan used to measure performance
- C. A performance testing benchmark result value throughout the project
- D. The initial starting point for all coding work within the project scope

Answer: B

Q527. What is the purpose of a project charter?

- A. To write application source code for features in the system context
- B. To perform software testing and validation runs during implementation
- C. To formally authorize the project and define its objectives and stakeholders
- D. To deploy software to production server systems for quality purposes

Answer: C

Q528. What is the difference between effort and duration in project estimation?

- A. Effort is total person-hours; duration is calendar time
- B. The two terms are exactly the same thing always
- C. Effort is always measured in monetary currency
- D. Duration is always longer than effort estimates

Answer: A

Q529. What is a risk in project management?

- A. An uncertain event that could positively or negatively affect the project
- B. An absolute certainty in the project plan only in practice typically
- C. A documented project functional requirement by the organization
- D. A task that has already been fully completed as a standard approach

Answer: A

Q530. What is the purpose of a status report?

- A. To design relational database architectures in a systematic way
- B. To write application source code for features at every stage
- C. To compile and build software executable files across all phases
- D. To communicate project progress, issues, and plans to stakeholders

Answer: D

Q531. What is the Singleton design pattern?

- A. A pattern that ensures a class has only one instance with a global access point
- B. A comprehensive software testing design pattern by the development process
- C. A relational database schema design pattern type and its related activities
- D. A pattern designed for creating multiple instances within the system boundary

Answer: A

Q532. What is the Observer design pattern?

- A. A pattern where objects are notified of state changes in another object
- B. A creational object instantiation design pattern as defined by standards
- C. A structural composition and wrapping pattern for all project stakeholders
- D. A system security and access control pattern over the entire lifecycle

Answer: A

Q533. What is the Factory Method pattern?

- A. A pattern that defines an interface for creating objects, letting subclasses decide which class to instantiate
- B. A specialized debugging and tracing tool suite to achieve project objectives in the engineering discipline
- C. A comprehensive software testing method and tool by the project team members and related components
- D. A physical manufacturing production process flow within the given constraints in all development efforts

Answer: A

Q534. What is the MVC design pattern?

- A. A comprehensive software testing pattern type as part of the methodology
- B. A specific type of program variable declaration for effective project outcomes
- C. Model-View-Controller, separating data, presentation, and control logic
- D. A software deployment automation pattern type in the development process

Answer: C

Q535. What is the Strategy design pattern?

- A. A specific software testing strategy and plan within the project scope
- B. A project management planning approach method during the software lifecycle
- C. A pattern that defines a family of algorithms and makes them interchangeable
- D. A software deployment automation strategy type for the development team

Answer: C

Q536. What are the three categories of GoF design patterns?

- A. Input, Process, Output in the system
- B. Creational, Structural, Behavioral
- C. Design, Code, Test during implementation
- D. Easy, Medium, Hard throughout the project

Answer: B

Q537. What is the Open/Closed Principle?

- A. Software entities should be open for extension but closed for modification
- B. Never closing running desktop applications ever as a standard approach
- C. Always keeping file handles open in the code in practice typically
- D. Using only open source software for projects for quality purposes

Answer: A

Q538. What is the Liskov Substitution Principle?

- A. A specific naming convention for all variables by the organization at every stage
- B. A comprehensive software testing principle type across all phases in a systematic way
- C. A relational database normalization principle for the project goals according to best practices
- D. Objects of a superclass should be replaceable with objects of a subclass without affecting correctness

Answer: D

Q539. What is the DRY principle?

- A. Keep code dry and clean within the system boundary
- B. Delete Redundant Yields by the development process
- C. Don't Repeat Yourself — avoid code duplication
- D. Develop Robust Yields and its related activities

Answer: C

Q540. What is a class diagram?

- A. A project timeline for all project stakeholders within the given constraints
- B. A network diagram in all development efforts by the project team members
- C. A classroom layout over the entire lifecycle as defined by standards
- D. A UML diagram showing classes, attributes, methods, and relationships

Answer: D

Q541. What is microservices architecture?

- A. An architecture where the application is composed of small, independent services
- B. A traditional monolithic design approach to apps and related components
- C. A single centralized database system approach to achieve project objectives
- D. One single large monolithic service application by the project team members

Answer: A

Q542. What is the difference between SOA and microservices?

- A. SOA is a much newer architectural approach style as part of the methodology
- B. SOA uses enterprise service bus for communication; microservices use lightweight protocols
- C. Microservices are generally larger in scope overall for effective project outcomes
- D. The two approaches are completely identical overall in the engineering discipline

Answer: B

Q543. What is an architectural style?

- A. A named collection of architectural design decisions applicable to recurring design problems
- B. A user interface visual style and theme design during the software lifecycle
- C. A specific coding style and convention guideline in the development process
- D. A project documentation format and standard set within the project scope

Answer: A

Q544. What is the Model-View-ViewModel (MVVM) pattern?

- A. An architectural pattern separating UI, presentation logic, and business logic/data
- B. A relational database normalization pattern type for the development team
- C. A comprehensive software testing pattern design throughout the project
- D. A production deployment automation pattern type in the system context

Answer: A

Q545. What are architectural views?

- A. Database views defined with SQL query language for quality purposes
- B. User interface screens and layout page designs in practice typically
- C. Different perspectives of a system architecture for different stakeholders
- D. Screen views displayed to the end users only during implementation

Answer: C

Q546. What is the 4+1 View Model?

- A. A comprehensive software testing model and plan by the organization at every stage
- B. A relational database normalization model type across all phases in a systematic way
- C. An architecture description using Logical, Process, Development, Physical views plus Scenarios
- D. Five physical computer display monitors together as a standard approach

Answer: C

Q547. What is a message broker in architecture?

- A. A general-purpose programming language and tool within the system boundary
- B. A specific type of database storage system tool according to best practices
- C. An intermediary program that translates messages between different systems
- D. A person who physically delivers messages by hand for the project goals

Answer: C

Q548. What is the publish-subscribe pattern?

- A. A messaging pattern where publishers send messages to subscribers through a topic
- B. An established coding standard and guideline set over the entire lifecycle
- C. A relational database design pattern type only and its related activities
- D. A physical magazine subscription service plan by the development process

Answer: A

Q549. What is a service-oriented architecture (SOA)?

- A. An architecture where services communicate over a network to provide functionality
- B. A hardware system architecture blueprint only within the given constraints
- C. A single isolated standalone service application as defined by standards
- D. A traditional monolithic application design style for all project stakeholders

Answer: A

Q550. What is an API Gateway?

- A. A physical security entry gate device system in all development efforts
- B. A hardware network switching device and unit and related components
- C. A server that acts as a single entry point for a set of microservices
- D. A relational database gateway and adapter tool by the project team members

Answer: C

Q551. What are Nielsen's 10 usability heuristics?

- A. Ten specific software testing methodology types
- B. Ten established coding standards and rule sets
- C. Guidelines for evaluating user interface design quality
- D. Ten strict programming coding rules for teams

Answer: C

Q552. What is Fitts' Law?

- A. An established coding law and standard rule only and its related activities over the entire lifecycle
- B. A fundamental law of classical physics and motion within the system boundary by the development process
- C. A relational database normalization law and rule as defined by standards for all project stakeholders
- D. The time to reach a target is proportional to the distance and inversely proportional to the target size

Answer: D

Q553. What is a heuristic evaluation?

- A. A mathematical algorithm analysis method only within the given constraints
- B. A usability inspection method where experts evaluate UI against established principles
- C. A formal source code peer review session only by the project team members
- D. A system performance load testing method only in all development efforts

Answer: B

Q554. What is the difference between UI and UX?

- A. UX is only about writing source code logic only as part of the methodology in the development process
- B. The two concepts are fundamentally the same thing and related components to achieve project objectives
- C. UI focuses on visual design; UX encompasses the entire user experience including emotions and satisfaction
- D. UI is always more important than UX design work in the engineering discipline for effective project outcomes

Answer: C

Q555. What is A/B testing in UI design?

- A. Testing against two different database systems
- B. Testing on two different mobile device models
- C. Comparing two versions of a design to see which performs better
- D. Testing two different API endpoints separately

Answer: C

Q556. What is the concept of 'affordance' in UI design?

- A. The financial cost of design work and effort
- B. A particular application color scheme and theme
- C. A specific type of user interface animation style
- D. A visual clue that suggests how an element can be used

Answer: D

Q557. What is a design system?

- A. A comprehensive software testing system and suite for the development team
- B. A computer operating system installation setup during the software lifecycle
- C. A collection of reusable components, guidelines, and standards for consistent design
- D. A specific application coding framework tool type within the project scope

Answer: C

Q558. What is the principle of 'progressive disclosure'?

- A. Showing every option and detail at once to users throughout the project
- B. Hiding absolutely all system information entirely in the system context
- C. Showing only essential information initially, revealing more as needed
- D. Disclosing security information to all end users during implementation

Answer: C

Q559. What is a user persona?

- A. A specific type of user profile avatar image in practice typically
- B. A user login name and credentials identifier for quality purposes
- C. A user uploaded profile picture image file as a standard approach
- D. A fictional character representing a key user group based on research data

Answer: D

Q560. What is the Z-pattern layout?

- A. A relational database query plan pattern and type across all phases in a systematic way
- B. A comprehensive software testing pattern approach for the project goals
- C. A layout following the natural eye scanning pattern from top-left to bottom-right in a Z shape
- D. A zigzag code structure in the code base layout by the organization at every stage

Answer: C

Q561. What is continuous integration (CI)?

- A. Manually integrating code without any automation for the development team
- B. Occasional and infrequent code merging events during the software lifecycle
- C. The practice of frequently merging code changes with automated building and testing
- D. Annual scheduled system integration events only within the project scope

Answer: C

Q562. What is continuous deployment (CD)?

- A. Annual scheduled deployment release events only in the system context
- B. Deployment without any automated testing at all during implementation
- C. Manual deployment to production server systems throughout the project
- D. Automatically deploying every code change that passes automated tests to production

Answer: D

Q563. What is the difference between compilation and interpretation?

- A. Compilation is always done only at runtime always at every stage across all phases in a systematic way
- B. Compilation translates entire code to machine code before execution; interpretation translates line by line during execution
- C. Interpretation is always faster than compilation is as a standard approach by the organization
- D. The two processes are fundamentally the same thing for quality purposes in practice typically

Answer: B

Q564. What is pair programming?

- A. Programming for exactly two hours at a time only
- B. Programming on two separate display monitors only
- C. Two developers working together at one workstation
- D. Two separate teams coding independently on tasks

Answer: C

Q565. What is code smell?

- A. An actual physical smell from hardware component
- B. A simple code formatting and indentation error
- C. A symptom in the source code that may indicate a deeper problem
- D. A type of computer virus or malware threat only

Answer: C

Q566. What is static code analysis?

- A. Executing and running the application code live
- B. Analyzing code without executing it to find potential errors
- C. Dynamic runtime testing of the application code
- D. Performance load testing of the application code

Answer: B

Q567. What is Git branching?

- A. A comprehensive software testing method type
- B. Creating separate lines of development within a repository
- C. A type of sorting algorithm technique approach
- D. Physical branches of a tree in nature settings

Answer: B

Q568. What is code coverage?

- A. A project documentation completeness metric type
- B. Physical coverage of printed code documents only
- C. A metric measuring the percentage of code executed during testing
- D. A security protection and control measure type

Answer: C

Q569. What is the DRY principle in coding?

- A. Delete Redundant Yields according to best practices
- B. Develop Robust Yields within the system boundary
- C. Keep code moisture-free for the project goals
- D. Don't Repeat Yourself — avoid duplicating code or logic

Answer: D

Q570. What is a linter?

- A. A tool that analyzes code for potential errors, style violations, and best practices
- B. A type of clothing fabric material only exactly by the development process
- C. A source code compilation tool program and build over the entire lifecycle
- D. A specialized debugging and tracing tool for code and its related activities

Answer: A

Q571. What is boundary value analysis?

- A. A testing technique focusing on values at the edges of input ranges
- B. A software architecture design technique approach throughout the project
- C. Analyzing overall project scope boundary limits for the development team
- D. A source code implementation technique approach in the system context

Answer: A

Q572. What is equivalence partitioning?

- A. A creational object-oriented design pattern type for quality purposes
- B. Dividing input data into groups that should be treated the same by the software
- C. A source code implementation technique approach in practice typically
- D. Equal distribution of work across team members during implementation

Answer: B

Q573. What is the difference between verification and validation in testing?

- A. Verification is entirely optional in all practice in a systematic way for the project goals
- B. The two concepts are fundamentally the same thing as a standard approach by the organization
- C. Verification checks if the product is built right; validation checks if the right product is built
- D. Validation always occurs before verification process at every stage across all phases

Answer: C

Q574. What is system testing?

- A. Testing just the operating system installation only within the system boundary
- B. Testing only individual isolated code units first according to best practices
- C. Testing the complete integrated system to verify it meets requirements
- D. Testing physical hardware system components only by the development process

Answer: C

Q575. What is performance testing?

- A. Testing how the system performs under various conditions of load and stress
- B. Testing the overall user experience quality level over the entire lifecycle
- C. Testing overall source code quality metrics only and its related activities
- D. Testing application security vulnerabilities only as defined by standards

Answer: A

Q576. What is stress testing?

- A. Testing comprehensive project documentation files in all development efforts
- B. Testing the stress levels of developer team members for all project stakeholders
- C. Testing the system beyond normal operational capacity to find its breaking point
- D. Testing with easy and simple valid input data only within the given constraints

Answer: C

Q577. What is test-driven development (TDD)?

- A. Writing tests before writing the code to make them pass
- B. Performing testing without any source code at all
- C. Performing all testing after coding is fully done
- D. Developing test plans and nothing else at all

Answer: A

Q578. What is a test stub?

- A. A test case that has previously failed and broke and related components
- B. A very short and brief test case scenario only by the project team members
- C. A comprehensive test report and document summary to achieve project objectives
- D. A simplified implementation that provides predetermined responses during testing

Answer: D

Q579. What is a mock object?

- A. A specific type of database storage system type as part of the methodology
- B. A decorative object for the office workspace area for effective project outcomes
- C. A simulated object that mimics the behavior of real objects in a controlled way for testing
- D. A counterfeit fake commercial product item only in the engineering discipline

Answer: C

Q580. What is alpha testing vs beta testing?

- A. Alpha is internal testing by developers; beta is external testing by end users
- B. Alpha testing is always done by external end users within the project scope
- C. The two types of testing are the same thing exactly in the development process
- D. Beta testing always comes before alpha testing does during the software lifecycle

Answer: A

Q581. What is Lehman's Law of Increasing Complexity?

- A. As a system evolves, its complexity increases unless work is done to reduce it
- B. Complexity is completely irrelevant to all quality over the entire lifecycle
- C. Complexity always stays perfectly constant overall and its related activities
- D. Complexity always decreases over time by itself by the development process

Answer: A

Q582. What is reverse engineering in software maintenance?

- A. Analyzing existing software to understand its design and implementation
- B. Completely uninstalling the software product system within the given constraints
- C. Writing application source code backwards manually for all project stakeholders
- D. Building and manufacturing physical cars and parts as defined by standards

Answer: A

Q583. What is re-engineering in software?

- A. Deploying software to production server systems now and related components
- B. Restructuring or rewriting existing software to improve its quality without changing functionality
- C. Permanently deleting the software product entirely by the project team members
- D. Starting completely over from scratch and beginning in all development efforts

Answer: B

Q584. What is the impact analysis in maintenance?

- A. Conducting performance benchmarking analysis tests
- B. Analyzing user behavior and usage pattern metrics
- C. Performing financial analysis of the project budget
- D. Assessing the effects of proposed changes on the rest of the system

Answer: D

Q585. What is software aging?

- A. Developer aging and gaining experience over career for effective project outcomes
- B. Physical hardware aging over many years of running in the engineering discipline
- C. The degradation of software quality over time due to accumulated changes and environmental evolution
- D. Software physically getting old and worn down over time to achieve project objectives

Answer: C

Q586. What is the difference between maintenance and evolution?

- A. Maintenance is always more comprehensive overall during the software lifecycle
- B. Maintenance fixes and adapts; evolution is the broader process of continuous change
- C. The two concepts are completely identical overall as part of the methodology
- D. Evolution is just a type of maintenance activity in the development process

Answer: B

Q587. What is program comprehension?

- A. Reading a technical programming textbook chapter
- B. Understanding existing code to perform maintenance or modifications
- C. Compiling source code into executable file formats
- D. Writing brand new application source code logic

Answer: B

Q588. What is a maintenance request?

- A. A physical hardware equipment purchase request doc
- B. A purchase order for new office supply materials
- C. A formal request to change, fix, or enhance existing software
- D. A personal vacation and time off request document

Answer: C

Q589. What is the concept of 'software entropy'?

- A. The energy consumption of running software systems for the development team
- B. The tendency of software to become more disordered and difficult to maintain over time
- C. A concept from physics thermodynamics discipline within the project scope
- D. An established coding standard and guideline type throughout the project

Answer: B

Q590. What is migration in software maintenance?

- A. Moving to a new country or city location abroad in the system context
- B. Moving software from one environment, platform, or technology to another
- C. Installing new physical hardware equipment devices for quality purposes
- D. Permanently deleting old legacy software and files during implementation

Answer: B

Q591. What is the difference between a walkthrough and an inspection?

- A. A walkthrough always has defined formal roles setup in all development efforts
- B. A walkthrough is informal and author-led; an inspection is formal with defined roles
- C. An inspection is always informal and casual in tone within the given constraints
- D. The two approaches are completely identical overall for all project stakeholders

Answer: B

Q592. What is the ISO/IEC 25010 quality model?

- A. A standard defining software quality characteristics like functionality, reliability, and usability
- B. A relational database normalization model and form in the engineering discipline
- C. A specific automobile car model type and category by the project team members
- D. A software architecture design model pattern type and related components to achieve project objectives

Answer: A

Q593. What is formal technical review (FTR)?

- A. A regular source code coding session for features as part of the methodology
- B. A structured review technique with specific roles and procedures to find defects
- C. An informal casual team meeting with no structure for effective project outcomes
- D. A production deployment meeting and schedule plan in the development process

Answer: B

Q594. What is the defect density metric?

- A. Number of defects per individual developer person during the software lifecycle
- B. The number of defects found per unit size of software (e.g., per KLOC)
- C. The severity rating of individual defects and bugs for the development team
- D. The total number of all defects combined overall within the project scope

Answer: B

Q595. What is the concept of 'shift left' in quality?

- A. Moving the computer screen display to the left side throughout the project
- B. Moving deployment earlier in the schedule timeline during implementation
- C. Moving testing and quality activities earlier in the development lifecycle
- D. Shifting team members to different project assignments in the system context

Answer: C

Q596. What is a quality audit?

- A. An independent evaluation of software processes and products against standards
- B. A financial accounting audit process only exclusively for quality purposes
- C. A system performance load test execution and report as a standard approach
- D. A formal source code peer review session and meeting in practice typically

Answer: A

Q597. What is the Pareto principle in QA?

- A. All modules have an equal number of bugs and issues
- B. Equal distribution of bugs across all code modules
- C. Approximately 80% of defects come from 20% of the modules
- D. Testing all modules equally and consistently overall

Answer: C

Q598. What is defect prevention?

- A. Activities aimed at identifying and addressing the root causes of defects before they occur
- B. Testing only at the very end of the whole project in a systematic way
- C. Finding defects only after production release launch at every stage across all phases
- D. Deliberately ignoring all known defects and issues by the organization

Answer: A

Q599. What is the role of a QA engineer?

- A. Ensuring software quality through planning, testing, and process improvement
- B. Only writing project documentation materials only within the system boundary
- C. Only writing application source code exclusively for the project goals
- D. Only managing projects and schedules exclusively according to best practices

Answer: A

Q600. What are the internal and external failure costs?

- A. Internal failure costs are always much higher overall for all project stakeholders within the given constraints
- B. External failure costs are always much cheaper overall over the entire lifecycle as defined by standards
- C. The two cost types are fundamentally the same thing by the development process and its related activities
- D. Internal failures are found before release (rework); external failures are found after release (support, reputation)

Answer: D

Q601. What is cyclomatic complexity?

- A. A complexity metric related to bicycles and riding in practice typically
- B. The total count of all declared variable name types by the organization
- C. The total length of all source code files and lines as a standard approach
- D. A metric measuring the number of linearly independent paths through a program's source code

Answer: D

Q602. What is function point analysis?

- A. Counting total functions in source code and files at every stage
- B. A method of measuring software size based on the functionality delivered to users
- C. A specific software testing technique approach only across all phases
- D. An established coding standard and guideline set only in a systematic way

Answer: B

Q603. What is the Halstead complexity metric?

- A. A set of metrics based on counting operators and operands in source code
- B. A comprehensive software testing metric set and type according to best practices
- C. A subjective difficulty rating for the code base for the project goals
- D. A software architecture design metric set and type within the system boundary

Answer: A

Q604. What is the difference between direct and indirect metrics?

- A. Indirect metrics are always more accurate and precise over the entire lifecycle as defined by standards
- B. Direct metrics are always calculated from other data for all project stakeholders
- C. Direct metrics are measured directly (LOC); indirect metrics are derived from other measures (productivity)
- D. The two metric types are completely the same thing by the development process and its related activities

Answer: C

Q605. What is the maintainability index?

- A. A comprehensive software testing score and rating
- B. An index at the back of a textbook publication
- C. A production deployment readiness rating and score
- D. A composite metric indicating how easy it is to maintain software

Answer: D

Q606. What is cohesion metric?

- A. A coupling metric between different system modules in all development efforts
- B. A comprehensive software testing metric and score by the project team members
- C. A measure of how strongly related the responsibilities within a module are
- D. A measure of team togetherness and morale level within the given constraints

Answer: C

Q607. What is coupling metric?

- A. A production deployment readiness metric and score
- B. A measure of how strongly one module depends on other modules
- C. A comprehensive software testing metric and score
- D. A cohesion metric within a single code module only

Answer: B

Q608. What is the fan-in and fan-out metric?

- A. A cooling system performance metric for hardware and related components
- B. A relational database performance metric and score in the engineering discipline
- C. A network bandwidth performance metric measurement to achieve project objectives
- D. Fan-in counts modules calling a module; fan-out counts modules called by a module

Answer: D

Q609. What is the depth of inheritance tree (DIT) metric?

- A. The depth of a relational database schema and tables as part of the methodology
- B. The height of a tree in a garden setting outdoors for effective project outcomes
- C. The depth of software testing coverage and breadth in the development process
- D. The maximum length of the path from a class to the root class in an inheritance hierarchy

Answer: D

Q610. What is the weighted methods per class (WMC) metric?

- A. The total number of classes in the project source
- B. The total number of source code files in project
- C. The average physical weight of all methods defined
- D. The sum of complexities of all methods defined in a class

Answer: D

Q611. What is the purpose of a Change Control Board (CCB)?

- A. A physical whiteboard in the office space and room by the organization
- B. A dedicated software testing team or group of people at every stage
- C. A development source coding team or group of people across all phases
- D. A group responsible for evaluating and approving or rejecting change requests

Answer: D

Q612. What is a merge conflict?

- A. A situation where two branches have competing changes that cannot be automatically merged
- B. A software testing failure or assertion error result according to best practices
- C. A disagreement or argument between the team members in a systematic way
- D. A production deployment failure or error issue event for the project goals

Answer: A

Q613. What is continuous integration in the context of SCM?

- A. Frequently merging code changes to a shared repository with automated builds and tests
- B. Manually merging code without any automation tools by the development process
- C. Performing one-time integration at the very end and its related activities
- D. Performing annual scheduled integration events only within the system boundary

Answer: A

Q614. What is a tag in version control?

- A. A price tag attached to a retail product for sale over the entire lifecycle
- B. A commit message attached to a changeset in version control for all project stakeholders
- C. A marker that identifies a specific point in the repository's history, typically a release
- D. A specific type of development branch in the project as defined by standards

Answer: C

Q615. What is the difference between centralized and distributed version control?

- A. Centralized version control is the newer approach today by the project team members
- B. Distributed version control has only one server overall in all development efforts
- C. The two systems are completely identical in every way within the given constraints
- D. Centralized has one server; distributed gives each developer a full copy of the repository

Answer: D

Q616. What is a build system in configuration management?

- A. A physical construction building system and process and related components
- B. A project documentation management system and process in the engineering discipline
- C. A tool that automates the process of compiling source code into executable software
- D. A comprehensive software testing system and tool suite to achieve project objectives

Answer: C

Q617. What is release management?

- A. A comprehensive software testing method and approach type in the development process
- B. Releasing employees from their positions and roles for effective project outcomes
- C. A source code implementation coding technique and method as part of the methodology
- D. The process of planning, scheduling, and controlling software builds through different stages

Answer: D

Q618. What is a software configuration audit?

- A. A financial accounting audit process and procedure only during the software lifecycle
- B. A system performance load test execution and run report for the development team
- C. A formal source code peer review session and meeting only within the project scope
- D. A review to verify that configuration items conform to specifications and standards

Answer: D

Q619. What is GitFlow?

- A. A software testing flow and process model and approach in the system context
- B. A water flow through pipes and channels and drains throughout the project
- C. A production deployment pipeline and flow and process during implementation
- D. A branching model for Git that defines a strict branching structure for projects

Answer: D

Q620. What is trunk-based development?

- A. A branching strategy where developers merge small, frequent updates to a core trunk or main branch
- B. Developing software in physical tree trunks outdoors for quality purposes in practice typically
- C. A comprehensive software testing approach and method only as a standard approach
- D. A software architecture and design method approach only by the organization

Answer: A

Q621. What is a risk exposure (risk magnitude)?

- A. A production deployment issue and concern problem
- B. A type of security vulnerability and exposure found
- C. The product of risk probability and risk impact
- D. Physical exposure to sunlight and weather outdoors

Answer: C

Q622. What is the difference between known risks and unknown risks?

- A. Known risks cannot be managed effectively at all within the given constraints
- B. Unknown risks are far more common than known risks for all project stakeholders
- C. Known risks can be identified and planned for; unknown risks cannot be anticipated
- D. The two categories are fundamentally the same thing as defined by standards

Answer: C

Q623. What is Boehm's risk management approach?

- A. A production deployment automation approach and plan
- B. A framework with two main components: risk assessment and risk control
- C. A source code implementation coding approach and plan
- D. A comprehensive software testing approach and process

Answer: B

Q624. What is the RMMM plan?

- A. A comprehensive software testing plan and document only and related components
- B. Risk Mitigation, Monitoring, and Management Plan — a document for managing project risks
- C. A source code implementation coding plan and document by the project team members
- D. A plan based on Roman numeral notation and numbering in all development efforts

Answer: B

Q625. What is a contingency plan?

- A. The main and primary project plan document overall
- B. A production deployment plan and schedule doc
- C. A comprehensive software testing plan document
- D. A backup plan activated when a risk materializes

Answer: D

Q626. What are the top 10 software risks according to Boehm?

- A. Ten common source code coding errors and mistakes to achieve project objectives
- B. Ten common software testing risk types and categories in the engineering discipline
- C. Ten common deployment issues and concerns and problems for effective project outcomes
- D. Risks including personnel shortfalls, unrealistic schedules, wrong requirements, and gold plating

Answer: D

Q627. What is risk monitoring?

- A. Continuously tracking identified risks and identifying new risks throughout the project
- B. Monitoring application source code changes and edits as part of the methodology
- C. Monitoring production server health status and uptime in the development process
- D. Monitoring individual team member performance metrics during the software lifecycle

Answer: A

Q628. What is a risk trigger?

- A. A physical gun trigger mechanism device component only
- B. A software testing trigger or activation event and signal
- C. An indicator or symptom that a risk event is about to occur or has occurred
- D. A production deployment trigger or activation signal only

Answer: C

Q629. What is the difference between qualitative and quantitative risk analysis?

- A. Quantitative analysis is always purely subjective in nature for the development team
- B. The two approaches are completely identical in all ways within the project scope
- C. Qualitative analysis always uses only numeric data values throughout the project
- D. Qualitative uses subjective assessment (high/medium/low); quantitative uses numerical probabilities

Answer: D

Q630. What is a risk breakdown structure?

- A. A comprehensive software testing structure and model for quality purposes
- B. A source code structure and organization model layout during implementation
- C. A broken system that needs repair work and fixing in the system context
- D. A hierarchical organization of identified risks categorized by risk source

Answer: D

Q631. What is the OWASP Top 10?

- A. A popular music chart and ranking system and listing
- B. A list of the ten most critical web application security risks
- C. An established coding standard and guideline and rule
- D. A comprehensive software testing framework tool and set

Answer: B

Q632. What is the principle of least privilege?

- A. Maximum access permissions for all system users always across all phases
- B. No access permissions for any user in the system ever in a systematic way
- C. Users should only have the minimum access rights necessary to perform their tasks
- D. Giving everyone full administrator access to everything at every stage

Answer: C

Q633. What is a buffer overflow?

- A. Having too much data stored in a buffer data area for the project goals
- B. A database table overflow and storage issue and error within the system boundary
- C. A vulnerability where data written beyond a buffer's boundary can overwrite adjacent memory
- D. A network traffic overflow and congestion issue event according to best practices

Answer: C

Q634. What is input validation?

- A. Accepting all user input without any validation checks by the development process
- B. Validating only system output and response data values and its related activities
- C. Testing user interface elements for usability and flow over the entire lifecycle
- D. Checking and sanitizing user input to prevent malicious data from entering the system

Answer: D

Q635. What is defense in depth?

- A. A security approach using multiple layers of security controls throughout a system
- B. Defending only the database from attacks and threats within the given constraints
- C. Having one single strong defense mechanism and barrier for all project stakeholders
- D. A military strategy used only in warfare and battles as defined by standards

Answer: A

Q636. What is a security audit?

- A. A financial accounting audit process and procedure type in all development efforts
- B. A systematic evaluation of a system's security by examining compliance with security policies
- C. A system performance load test execution and run process and related components
- D. A formal source code peer review session and process only by the project team members

Answer: B

Q637. What is penetration testing?

- A. Physical building break-in testing attempt and effort
- B. Database query performance testing and benchmarking
- C. Network communication speed testing and measurement
- D. Simulating cyberattacks to identify security vulnerabilities in a system

Answer: D

Q638. What is a denial-of-service (DoS) attack?

- A. A specific type of computer virus or malware and threat in the engineering discipline
- B. Denying customer service requests politely and kindly to achieve project objectives
- C. An attack that overwhelms a system with traffic to make it unavailable to users
- D. A network configuration and setup procedure and process for effective project outcomes

Answer: C

Q639. What is the CIA triad in security?

- A. A government intelligence agency abbreviation and name as part of the methodology
- B. An established coding standard and guideline and rule set in the development process
- C. Confidentiality, Integrity, and Availability — the three core principles of information security
- D. A comprehensive software testing framework and tool suite during the software lifecycle

Answer: C

Q640. What is secure coding?

- A. Writing code in a physically secure room and building
- B. Encrypting all application source code files and folders
- C. Development practices that protect software from security vulnerabilities
- D. Using complex passwords directly in application code

Answer: C

Q641. What is the difference between verification and validation?

- A. Verification writes documentation; validation writes test code
- B. Verification deploys software; validation monitors production
- C. Verification checks correctness; validation checks usefulness
- D. Verification tests performance; validation tests user interface

Answer: C

Q642. Which factor most contributes to the software crisis?

- A. Decreasing cost of modern hardware and peripherals
- B. Growing availability of trained software professionals
- C. Increasing system complexity and changing requirements
- D. Widespread adoption of standardized coding practices

Answer: C

Q643. What is the significance of the NATO 1968 conference?

- A. It launched the initial version of the world wide web
- B. It formally introduced software engineering as a discipline
- C. It created the first commercial operating system product
- D. It established the first programming language standard

Answer: B

Q644. How does software engineering differ from programming?

- A. It focuses only on writing source code in languages
- B. It deals exclusively with database design and queries
- C. It encompasses the entire development lifecycle process
- D. It concerns itself primarily with hardware maintenance

Answer: C

Q645. What is the purpose of a feasibility study in software projects?

- A. To write the final source code for the application
- B. To conduct user acceptance testing after development
- C. To deploy the software to the production environment
- D. To determine if the project is viable and worthwhile

Answer: D

Q646. Which software characteristic ensures it works across platforms?

- A. Obscurity of the source code implementation
- B. Dependency on a single hardware configuration
- C. Portability across different operating environments
- D. Complexity of the underlying algorithm design

Answer: C

Q647. What is meant by 'technical debt' in software engineering?

- A. The monetary budget allocated for purchasing hardware
- B. Revenue generated from software licensing agreements
- C. Total salary expenses for the development team members
- D. Accumulated cost of shortcuts taken during development

Answer: D

Q648. Why are software engineering ethics considered important?

- A. They guide responsible behavior in professional practice
- B. They reduce the physical size of deployed software
- C. They increase the execution speed of the application
- D. They eliminate the need for comprehensive testing

Answer: A

Q649. What is the relationship between software and hardware engineering?

- A. Both disciplines use identical processes and methodologies
- B. Software engineering is a direct subset of hardware engineering
- C. Hardware engineering completely depends on software practices
- D. Software engineering builds on but differs from hardware methods

Answer: D

Q650. What is the primary benefit of using CASE tools?

- A. They automate and support software development activities
- B. They convert hardware specifications into software code
- C. They eliminate all possible defects from the source code
- D. They replace the need for any human software developers

Answer: A

Q651. What is the key difference between incremental and iterative models?

- A. Incremental refines features; iterative adds new functionality
- B. Both approaches are completely identical in every possible way
- C. Neither approach involves any form of planning or documentation
- D. Incremental adds functionality; iterative refines existing features

Answer: D

Q652. In the spiral model, what happens at the end of each spiral loop?

- A. Testing is permanently skipped for all remaining spiral loops
- B. The entire project is immediately deployed to production servers
- C. A review determines whether to continue to the next iteration
- D. All source code is discarded and rewritten from scratch again

Answer: C

Q653. Why might a team choose the RAD model for development?

- A. When no user involvement is needed throughout development
- B. When requirements are completely fixed and will never change
- C. When the project has unlimited budget and timeline available
- D. When rapid delivery of a working system is the top priority

Answer: D

Q654. What is a disadvantage of the waterfall model?

- A. Excessive flexibility leading to uncontrolled scope expansion
- B. Difficulty accommodating requirement changes after phase completion
- C. Lack of any structured phases making progress hard to track
- D. Too many iterative cycles causing confusion among developers

Answer: B

Q655. How does the unified process (UP) organize development activities?

- A. Into four phases: inception, elaboration, construction, transition
- B. Into two phases: planning and coding without any other steps
- C. Into six phases that must always be executed sequentially only
- D. Into a single continuous phase without any distinct milestones

Answer: A

Q656. What role does risk assessment play in the spiral model?

- A. It drives decisions about which activities to perform each cycle
- B. It has no significant impact on how the project is conducted
- C. It replaces all other engineering activities in the entire model
- D. It is performed only once at the very beginning of the project

Answer: A

Q657. When is the evolutionary development model most appropriate?

- A. When initial requirements are unclear and need exploration work
- B. When all requirements are perfectly defined before starting
- C. When the project has a very short and fixed delivery deadline
- D. When the team prefers no stakeholder interaction during work

Answer: A

Q658. What is the main criticism of the code-and-fix development approach?

- A. It requires too much upfront planning and design documentation
- B. It mandates excessive testing before any code is ever written
- C. It leads to poorly structured software that is hard to maintain
- D. It forces teams to follow overly rigid process model guidelines

Answer: C

Q659. What is the purpose of the elaboration phase in the Unified Process?

- A. To deploy the completed software to end users in production now
- B. To analyze the problem domain and establish architectural foundation
- C. To gather initial business case and scope project requirements
- D. To perform final acceptance testing before system goes live now

Answer: B

Q660. How does the component-based development model reduce development effort?

- A. By eliminating all testing and quality assurance activities fully
- B. By avoiding any form of architectural design or documentation
- C. By reusing pre-built and tested software components extensively
- D. By skipping the requirements gathering phase of development work

Answer: C

Q661. What is the difference between Scrum and Kanban?

- A. Scrum avoids visualization; Kanban prohibits all meetings
- B. Scrum has no backlog; Kanban requires sprint planning
- C. Scrum has no roles; Kanban requires a Scrum Master role
- D. Scrum uses time-boxed sprints; Kanban uses continuous flow

Answer: D

Q662. What is velocity in agile project management?

- A. The frequency of deploying code to the production servers
- B. The amount of work a team completes in one iteration period
- C. The speed at which the application executes user requests
- D. The rate at which new developers join the project team

Answer: B

Q663. What is the purpose of a sprint retrospective?

- A. To create the initial project charter and scope document
- B. To identify improvements for the team process going forward
- C. To demonstrate completed features to external stakeholders
- D. To assign blame for any defects found during the sprint

Answer: B

Q664. How does pair programming benefit agile teams?

- A. It allows one developer to do all work while the other rests
- B. It doubles the number of features developed per time unit
- C. It increases code quality through continuous peer review work
- D. It eliminates the need for any automated testing frameworks

Answer: C

Q665. What is the definition of 'done' in Scrum?

- A. A legal document signed by the client at project conclusion
- B. A shared understanding of when a work item is complete enough
- C. A report generated automatically by the build server system
- D. A metric measuring the total number of lines of code written

Answer: B

Q666. What is a burndown chart used for in agile?

- A. Documenting all architectural decisions made during the project
- B. Tracking remaining work against time in a sprint or release
- C. Measuring the total revenue generated by the software product
- D. Recording the number of team members joining or leaving work

Answer: B

Q667. What is the purpose of story points in agile estimation?

- A. To measure the physical size of the application in megabytes
- B. To estimate the relative effort required for completing tasks
- C. To count the exact number of hours each task will take to do
- D. To track the number of bugs reported by end users each week

Answer: B

Q668. How does continuous integration support agile development?

- A. By eliminating the need for version control systems entirely
- B. By delaying all testing until the final sprint of the project
- C. By preventing developers from committing code to the repository
- D. By frequently merging and testing code to detect issues early

Answer: D

Q669. What is the role of acceptance criteria in user stories?

- A. They list the hardware requirements for running the software
- B. They specify the programming language to be used for coding
- C. They define conditions that must be met for story completion
- D. They document the salary expectations of the development team

Answer: C

Q670. What is a spike in agile development?

- A. A formal review process conducted by external quality auditors
- B. A sudden increase in the number of production bugs per release
- C. A time-boxed research task to reduce uncertainty or risk factors
- D. A mandatory overtime period for all team members on a project

Answer: C

Q671. What is the difference between user requirements and system requirements?

- A. User requirements are high-level needs; system requirements are detailed specifications
- B. User requirements are optional; system requirements are always mandatory for the project
- C. User requirements come from developers; system requirements come from end user groups
- D. User requirements are written in code; system requirements are written in plain language

Answer: A

Q672. What is requirements prioritization and why is it necessary?

- A. It ranks requirements by importance to guide development order and resource allocation
- B. It randomly assigns requirements to different sprints without any logical ordering method
- C. It permanently removes low-priority requirements from the project scope and backlog
- D. It ensures all requirements are implemented simultaneously in a single development phase

Answer: A

Q673. What challenges arise during requirements elicitation with multiple stakeholders?

- A. All stakeholders always agree perfectly on every requirement without any discussion
- B. Multiple stakeholders simplify the process by reducing the total number of requirements
- C. Having more stakeholders eliminates the need for formal requirements documentation
- D. Conflicting viewpoints and priorities that must be negotiated and resolved carefully

Answer: D

Q674. How do prototypes help in requirements engineering?

- A. They replace the need for any written requirements documentation in the project
- B. They help stakeholders visualize the system and clarify unclear requirements early
- C. They guarantee that all requirements are correctly implemented in the final system
- D. They eliminate the possibility of any requirements changes during the development

Answer: B

Q675. What is the MoSCoW method in requirements prioritization?

- A. A geographic approach assigning requirements to teams based on physical location
- B. A technique categorizing requirements as Must, Should, Could, and Won't have items
- C. A testing strategy for validating requirements against acceptance criteria only
- D. A programming paradigm for implementing requirements in object-oriented languages

Answer: B

Q676. What is the difference between requirements verification and validation?

- A. Verification writes test cases; validation executes marketing campaigns for launch
- B. Verification checks document consistency; validation checks stakeholder satisfaction
- C. Verification trains users; validation calculates the return on investment figures
- D. Verification deploys the system; validation monitors production performance metrics

Answer: B

Q677. Why is natural language problematic for writing requirements?

- A. It is too precise and does not allow any flexibility in requirement statements
- B. It is only understood by technical staff and not by any business stakeholders
- C. It is inherently ambiguous and can lead to multiple conflicting interpretations
- D. It cannot express complex requirements due to limited vocabulary available

Answer: C

Q678. What is a requirements review and who should participate?

- A. A formal inspection of requirements involving stakeholders and technical staff
- B. A marketing meeting to discuss product positioning and pricing strategies
- C. An informal chat between two developers about their coding preferences only
- D. A hardware procurement session to select servers for the production use

Answer: A

Q679. How does requirements management handle change requests?

- A. By delegating all change decisions to a single junior developer on the team
- B. By rejecting all change requests to maintain the original project scope fully
- C. By immediately implementing all requested changes without any analysis first
- D. Through a formal change control process that evaluates impact before approval

Answer: D

Q680. What is the role of domain analysis in requirements engineering?

- A. It focuses exclusively on selecting the programming language for development
- B. It calculates the exact project budget and timeline before development starts
- C. It determines the number of servers needed for production deployment only
- D. It helps understand the application domain to identify relevant requirements

Answer: D

Q681. What is the purpose of earned value management (EVM)?

- A. To measure project performance by comparing planned versus actual progress and costs
- B. To determine which programming language is most suitable for the project requirements
- C. To calculate the total number of lines of code written by each developer on the team
- D. To design the system architecture and select appropriate technology stack components

Answer: A

Q682. How does the critical path method (CPM) help project managers?

- A. It guarantees that no project will ever experience delays or budget overruns
- B. It identifies tasks that directly impact project completion date if delayed at all
- C. It randomly assigns priorities to tasks without considering their dependencies now
- D. It eliminates the need for any project scheduling or timeline management work

Answer: B

Q683. What is scope creep and how can it be managed?

- A. An issue that resolves itself automatically without any management intervention
- B. A beneficial practice that should be encouraged for maximum project flexibility
- C. A problem that only affects projects using agile development methodologies now
- D. Uncontrolled scope expansion managed through formal change control processes

Answer: D

Q684. What is the difference between effort estimation and duration estimation?

- A. Effort measures code quality; duration measures the number of features implemented
- B. Effort is only relevant for testing; duration is only relevant for coding activities
- C. Effort is total work hours needed; duration is calendar time to complete the task
- D. Effort and duration are identical concepts with no meaningful difference between them

Answer: C

Q685. What is the COCOMO model used for in software project management?

- A. Managing source code versions and branches in the project code repository
- B. Testing software applications for security vulnerabilities and performance issues
- C. Estimating software development effort and cost based on project size parameters
- D. Designing user interfaces for web-based applications and mobile platforms only

Answer: C

Q686. How do project managers handle team conflicts effectively?

- A. By immediately removing the conflicting team members from the project team
- B. By assigning all disputed tasks to a single person to avoid disagreements
- C. By facilitating open communication and finding mutually acceptable solutions
- D. By ignoring conflicts and hoping they resolve themselves without intervention

Answer: C

Q687. What is a risk register in project management?

- A. A source code repository containing all versions of the application being built
- B. A database of all employee records including salary and performance evaluations
- C. A financial ledger tracking all project expenses and revenue over the timeline
- D. A documented list of identified risks with their probability and impact assessments

Answer: D

Q688. What is the purpose of a project post-mortem or lessons learned session?

- A. To calculate the final salary bonuses for the project team members after launch
- B. To assign blame to team members who made mistakes during the project execution
- C. To identify what went well and what can be improved for future project efforts
- D. To permanently archive all project documents and prevent any future access now

Answer: C

Q689. How does function point analysis help in project estimation?

- A. It measures the physical size of the deployed application in storage megabytes
- B. It tracks the number of team members assigned to the project at any given time
- C. It counts the number of programming languages used in the project source code
- D. It measures software size based on functional requirements independent of technology

Answer: D

Q690. What is the triple constraint in project management?

- A. The balance between scope, time, and cost that defines project boundaries
- B. The three phases of testing: unit, integration, and system level testing
- C. The minimum number of developers needed to start any software project
- D. The three programming languages required for every software project build

Answer: A

Q691. What is the Single Responsibility Principle (SRP)?

- A. A class should handle all possible system responsibilities now
- B. A function should always return multiple values simultaneously
- C. A class should have only one reason to change its behavior
- D. A module should depend on as many other modules as possible

Answer: C

Q692. What is the Open-Closed Principle in software design?

- A. Software entities should never be extended or modified after initial release
- B. Software entities should be open for extension but closed for modification
- C. Software entities should be both open and closed for all types of changes
- D. Software entities should be closed for extension and open for modification

Answer: B

Q693. What is the difference between the Strategy and State design patterns?

- A. Strategy and State are identical patterns with no differences between them at all
- B. Strategy selects algorithms; State changes behavior based on internal object state
- C. Strategy manages object creation; State handles object destruction and cleanup work
- D. Strategy is for UI design only; State is exclusively for database operations work

Answer: B

Q694. How does the Observer pattern support loose coupling?

- A. Observers must inherit from the subject class to receive any notification updates
- B. Observers must know the complete internal state of the subject at all times now
- C. Subjects notify observers without knowing their concrete implementation details
- D. The pattern requires tight coupling between all observers and the subject object

Answer: C

Q695. What is the Dependency Inversion Principle?

- A. High-level modules should depend on abstractions not on low-level module details
- B. Dependencies should always flow from abstractions down to concrete implementations
- C. High-level modules should directly depend on the implementation of low-level modules
- D. Low-level modules should control the behavior of all high-level modules in system

Answer: A

Q696. What is the Factory Method pattern used for?

- A. Connecting to external databases and executing queries against stored data records
- B. Sorting collections of objects based on their attributes in ascending or descending order
- C. Destroying objects and freeing memory after they are no longer needed by the program
- D. Creating objects without specifying the exact class to instantiate at compile time

Answer: D

Q697. How does the MVC pattern separate concerns in application design?

- A. It eliminates the need for any user interface by focusing solely on data storage
- B. It only separates the database from the rest of the application into two layers
- C. It separates data (Model), presentation (View), and logic (Controller) into layers
- D. It combines all application logic into a single layer without any separation at all

Answer: C

Q698. What is the Liskov Substitution Principle?

- A. Base types should be substitutable for their subtypes in all program situations
- B. Subtypes should completely change the behavior defined by their base type classes
- C. Subtypes must be substitutable for their base types without altering correctness
- D. Inheritance should be avoided entirely in favor of composition in every case

Answer: C

Q699. What is the purpose of the Adapter design pattern?

- A. To make incompatible interfaces work together through a wrapper component layer
- B. To add new functionality to objects dynamically at runtime without inheritance use
- C. To ensure a class has only one instance throughout the entire application system
- D. To create new objects by cloning existing ones without using constructors directly

Answer: A

Q700. What is the difference between composition and inheritance in design?

- A. Composition is always inferior to inheritance for all software design scenarios today
- B. Composition uses has-a relationships; inheritance uses is-a relationships between types
- C. Composition and inheritance are identical concepts with no meaningful differences at all
- D. Inheritance provides more flexibility than composition in every possible design context

Answer: B

Q701. What is the difference between software architecture and software design?

- A. Architecture and design are identical concepts with no meaningful differences between them
- B. Architecture focuses on code details; design focuses on system-level structural decisions
- C. Architecture defines high-level structure; design details component internals and algorithms
- D. Architecture is only relevant for small systems; design is only for large-scale systems

Answer: C

Q702. How does the Model-View-Controller architecture support separation of concerns?

- A. It eliminates the need for a user interface by focusing exclusively on data processing
- B. It only separates the database from the application without any further decomposition
- C. It separates data management, user interface, and control logic into distinct components
- D. It combines all concerns into a single component for simplicity and ease of maintenance

Answer: C

Q703. What are the trade-offs of microservices versus monolithic architecture?

- A. Microservices offer scalability and independence but add complexity in communication and deployment
- B. There are no meaningful trade-offs because both approaches produce identical system qualities
- C. Microservices are always superior to monolithic architectures in every possible project scenario
- D. Monolithic architectures are always more scalable than microservices in all system contexts today

Answer: A

Q704. What is an event-driven architecture?

- A. Components communicate by producing and consuming events asynchronously via message brokers
- B. All system events are processed in a strictly sequential order without any parallel execution
- C. Components communicate exclusively through synchronous direct function calls without any delay
- D. Event-driven systems cannot handle more than one event at any given point in time today

Answer: A

Q705. How does the service-oriented architecture (SOA) promote reusability?

- A. Services are designed as reusable, self-contained units with well-defined standard interfaces
- B. SOA discourages reusability by requiring unique implementations for every system integration
- C. Services in SOA are always built from scratch and never leverage any existing functionality
- D. Services are tightly coupled and cannot be reused outside their original application context

Answer: A

Q706. What is an architectural decision record (ADR)?

- A. A project management artifact tracking team member assignments and work allocations
- B. A source code file containing the main application entry point and initialization logic
- C. A database schema definition file used for creating tables and relationships in storage
- D. A document capturing the context, decision, and consequences of an architectural choice

Answer: D

Q707. What is the strangler fig pattern in architecture migration?

- A. Running both old and new systems permanently without ever completing the migration work
- B. Immediately shutting down the old system and deploying the new one all at once together
- C. Abandoning the migration and keeping the legacy system unchanged for all future operation
- D. Gradually replacing parts of a legacy system with new components until fully migrated

Answer: D

Q708. How does a message queue improve system architecture?

- A. It tightly couples all components requiring synchronous communication between every service
- B. It decouples producers and consumers allowing asynchronous communication and load leveling
- C. It reduces system reliability by introducing additional points of failure without any benefit
- D. It eliminates the need for any communication between system components or services at all

Answer: B

Q709. What is the CQRS (Command Query Responsibility Segregation) pattern?

- A. Separating read and write operations into different models for optimization flexibility
- B. Using the same data model and access path for both queries and commands in all cases
- C. Combining all read and write operations into a single model for maximum simplicity only
- D. Eliminating all read operations and only allowing write operations in the system design

Answer: A

Q710. What is the role of middleware in a distributed architecture?

- A. It only stores data and has no role in communication or transformation between services
- B. It provides communication, data transformation, and integration between distributed components
- C. It eliminates the need for any networking between distributed system components entirely
- D. It replaces all business logic components and handles every processing requirement directly

Answer: B

Q711. What are Nielsen's ten usability heuristics used for?

- A. Evaluating user interface designs for common usability problems and issues
- B. Configuring network firewalls for protecting against security threats
- C. Designing database schemas for storing application data efficiently now
- D. Writing backend server code for processing user authentication requests

Answer: A

Q712. What is the difference between UI design and UX design?

- A. UI and UX are identical disciplines with no differences in scope or focus areas
- B. UI covers project management; UX handles server infrastructure configuration
- C. UI focuses on visual elements; UX encompasses the entire user experience journey
- D. UI deals with backend logic; UX focuses exclusively on database design work now

Answer: C

Q713. What is Fitts' law and how does it apply to UI design?

- A. Smaller and distant targets are easier to click, so buttons should be minimized always
- B. Fitts' law only applies to physical devices and has no relevance to digital interfaces
- C. Larger and closer targets are faster to select, influencing button size and placement
- D. Target size has no impact on user interaction speed or accuracy in any interface now

Answer: C

Q714. How does the principle of consistency improve user interfaces?

- A. It makes the interface unpredictable so users must think carefully about every action
- B. It allows users to apply learned behaviors across the interface reducing cognitive load
- C. It forces every page to look completely different from all other pages in the system
- D. It increases the time users need to learn how to navigate and use the application

Answer: B

Q715. What is a design system in the context of UI development?

- A. A single page design template that cannot be customized or reused across the application
- B. A collection of reusable components and guidelines ensuring visual and behavioral consistency
- C. A database management system for storing user interface configuration data and settings
- D. A server monitoring tool for tracking application performance metrics and error rates

Answer: B

Q716. What is the purpose of A/B testing in UI design?

- A. Simultaneously releasing both design variants as permanent features for all users
- B. Testing the application's backend server performance under heavy load conditions
- C. Validating the database schema against the documented requirements specification
- D. Comparing two design variants to determine which performs better with real users

Answer: D

Q717. How does progressive disclosure improve complex user interfaces?

- A. It reveals information gradually, showing only what is needed at each step of interaction
- B. It hides all information permanently and never reveals anything to the user at any time
- C. It displays all possible options and information simultaneously on a single crowded screen
- D. It forces users to read complete documentation before they can interact with the system

Answer: A

Q718. What is the importance of color contrast in UI design?

- A. It ensures text and elements are readable for users including those with visual impairments
- B. High contrast should be avoided because it makes interfaces harder for everyone to read
- C. Color contrast only matters for printed materials and is irrelevant for digital interfaces
- D. It is purely decorative and has no impact on readability or accessibility of the interface

Answer: A

Q719. What is a persona in user-centered design?

- A. A software bot that automatically generates user interface layouts without human input
- B. A legal document defining the terms of service between the company and its end users
- C. A real person hired to test the application before it is released to the general public
- D. A fictional character representing a target user group based on research data analysis

Answer: D

Q720. What is the purpose of user journey mapping in UI design?

- A. To create the database schema for storing all application data and configurations
- B. To write the server-side code for processing API requests from the client side
- C. To configure the deployment pipeline for releasing the application to production
- D. To visualize the complete user experience across all touchpoints and interactions

Answer: D

Q721. What is the difference between compiled and interpreted languages?

- A. Compiled languages run slower than interpreted languages in all performance benchmarks and tests
- B. Compiled languages translate code before execution; interpreted languages translate during execution
- C. Compiled and interpreted languages are identical with no differences in execution approach at all
- D. Interpreted languages cannot be used for building any production-grade software applications

Answer: B

Q722. What is continuous integration (CI) in software development?

- A. Manually combining all code at the end of the project without any automated testing or builds
- B. Writing all source code in a single file to avoid any merge conflicts between team members
- C. Deploying code directly to production without any testing or review processes being performed
- D. Frequently merging code changes into a shared repository with automated build and test validation

Answer: D

Q723. How does pair programming improve code quality?

- A. It reduces code quality because two developers disagree and produce inconsistent solutions
- B. It doubles development speed by having two people type simultaneously on the same keyboard
- C. It eliminates the need for any code reviews since the code is already reviewed during writing
- D. Real-time review catches errors immediately and promotes knowledge sharing between developers

Answer: D

Q724. What is technical debt and how does it affect implementation?

- A. Accumulated shortcuts requiring future rework that slows down development and increases costs
- B. A financial investment made by the company to purchase better development tools and hardware
- C. A type of software license that requires payment for each line of code written in the project
- D. A project management metric measuring the total budget remaining for the development effort

Answer: A

Q725. What is the purpose of unit testing during implementation?

- A. To check the project schedule against the planned milestones and delivery date targets
- B. To test the entire application system from end to end including all integrated components
- C. To validate the user interface design against the approved wireframes and mockup designs
- D. To verify individual code components work correctly in isolation from other system parts

Answer: D

Q726. What is dependency injection and why is it useful?

- A. Providing dependencies externally rather than creating them internally for better testability
- B. Creating all dependencies inside the class constructor to simplify initialization logic now
- C. Injecting database records directly into the application without using any query language
- D. Removing all dependencies from a class so it operates completely independently of others

Answer: A

Q727. What is the purpose of a build automation tool?

- A. To create project management reports showing team productivity and budget utilization data
- B. To manually compile each source code file individually without any automated processes at all
- C. To design user interfaces by automatically generating HTML and CSS from wireframe images
- D. To automate compilation, testing, packaging, and deployment of software applications reliably

Answer: D

Q728. How does static code analysis help during implementation?

- A. It analyzes the application user interface for usability problems and accessibility violations
- B. It replaces the need for any manual code review by automatically fixing all detected issues
- C. It identifies potential bugs, security vulnerabilities, and style violations without running code
- D. It only detects errors that occur during runtime execution of the application in production

Answer: C

Q729. What is the difference between a framework and a library?

- A. Frameworks are always superior to libraries for all software development projects and tasks
- B. Frameworks and libraries are identical concepts with no meaningful differences between them
- C. A framework controls the flow and calls your code; a library is called by your code directly
- D. A library controls application flow; a framework provides utility functions called by code

Answer: C

Q730. What is containerization and how does it benefit implementation?

- A. Packaging applications with dependencies into isolated containers ensuring consistent environments
- B. Running applications directly on the host operating system without any isolation or packaging
- C. Manually configuring each server individually to match the development environment settings
- D. Removing all application dependencies to reduce the size of the deployed software package

Answer: A

Q731. What is the difference between black-box and white-box testing?

- A. Black-box testing is always superior to white-box testing for finding all types of software bugs
- B. Black-box tests functionality without code knowledge; white-box tests with code structure knowledge
- C. Black-box and white-box testing are identical approaches with no meaningful differences at all
- D. Black-box testing examines code structure; white-box testing ignores internal implementation fully

Answer: B

Q732. What is boundary value analysis in software testing?

- A. Avoiding all edge cases and testing only the most common inputs that users are expected to use
- B. Testing only the maximum possible input value and ignoring all other boundary conditions now
- C. Testing at the edges of input ranges where defects are most likely to occur in the software
- D. Testing random values from the middle of input ranges without focusing on any specific values

Answer: C

Q733. What is equivalence partitioning?

- A. Testing every possible input value individually to ensure complete coverage of the system
- B. Randomly selecting test inputs without any systematic approach to input space partitioning
- C. Removing all invalid inputs from the test suite to focus only on valid expected scenarios
- D. Dividing inputs into groups where all values in a group are expected to behave identically

Answer: D

Q734. What is code coverage and why does it matter?

- A. It measures the percentage of code exercised by tests indicating thoroughness of test suite
- B. It counts the total number of lines of code in the project regardless of testing activities
- C. It measures the speed at which the application executes under normal operating conditions
- D. It tracks the number of developers who have reviewed each source code file in the project

Answer: A

Q735. What is mutation testing?

- A. Changing the programming language of the application to improve its overall test coverage
- B. Introducing small changes to code to check if tests can detect the introduced modifications
- C. Modifying the project requirements to make them easier to test with existing test suites
- D. Deleting all test cases and relying on manual testing for all quality assurance activities

Answer: B

Q736. What is performance testing?

- A. Checking the visual design of the user interface against approved wireframes and mockups
- B. Validating that the project schedule aligns with the planned milestones and target dates
- C. Verifying that individual functions return correct results for all valid input combinations
- D. Evaluating system behavior under various load conditions to identify bottlenecks and limits

Answer: D

Q737. What is the purpose of test-driven development (TDD)?

- A. Writing tests before code to drive design and ensure comprehensive test coverage from start
- B. Using test results to determine developer salary bonuses at the end of each project phase
- C. Eliminating all testing activities to focus developer time on feature implementation only
- D. Writing code first and then deciding which tests to write based on the bugs found later

Answer: A

Q738. What is exploratory testing?

- A. Simultaneous test design, execution, and learning without predefined scripts or test cases
- B. Following a strictly predefined test script without any deviation or creativity in approach
- C. Running only automated tests without any human involvement in the testing process at all
- D. Exploring the source code without executing the application to find potential defect areas

Answer: A

Q739. What is the testing pyramid concept?

- A. Only unit tests are needed and integration and end-to-end tests should be completely avoided
- B. More end-to-end tests at the base and fewer unit tests at the top of the testing pyramid
- C. More unit tests at the base, fewer integration tests, and fewest end-to-end tests at the top
- D. Equal numbers of unit, integration, and end-to-end tests at every level of the testing pyramid

Answer: C

Q740. What is the difference between verification and validation in testing?

- A. Verification and validation are identical activities with no differences in their purposes
- B. Verification focuses on user satisfaction; validation focuses on technical correctness only
- C. Verification checks if built correctly; validation checks if the right product was built
- D. Verification is done after deployment; validation is done before development starts only

Answer: C

Q741. What is Lehman's law of increasing complexity?

- A. Software complexity naturally decreases over time without any deliberate effort being made
- B. Software complexity increases with each change unless work is done to reduce it actively
- C. Software complexity remains constant throughout the entire lifetime of the application now
- D. Software complexity is only affected by the initial design and never by later modifications

Answer: B

Q742. How does software aging affect system reliability over time?

- A. Aging has no measurable effect on software reliability because code does not physically degrade
- B. Software reliability improves automatically as hardware becomes faster and more powerful now
- C. Software always becomes more reliable over time as more bugs are found and fixed in system
- D. Accumulated changes and environmental drift gradually degrade system reliability and performance

Answer: D

Q743. What is the difference between re-engineering and reverse engineering?

- A. Re-engineering recovers design knowledge; reverse engineering restructures for improvements now
- B. Reverse engineering and re-engineering are identical processes with no differences between them
- C. Reverse engineering adds new features; re-engineering only fixes bugs in the existing codebase
- D. Reverse engineering recovers design; re-engineering restructures the system for improvement work

Answer: D

Q744. What is the impact of code smells on software maintenance effort?

- A. Code smells always improve code quality by making the codebase more interesting to read now
- B. Code smells indicate design problems that increase maintenance difficulty and defect likelihood
- C. Code smells have no impact on maintenance because they do not represent actual defects today
- D. Code smells are only cosmetic issues that never affect functionality or maintenance costs

Answer: B

Q745. How does modular design facilitate software maintenance?

- A. Modular design makes maintenance harder because changes must be applied to every module now
- B. Modular design has no relationship to maintenance effort or quality of the maintenance work
- C. Modules prevent any changes from being made to the software after initial deployment to users
- D. Changes can be isolated to specific modules without affecting the rest of the system at all

Answer: D

Q746. What is the ripple effect in software maintenance?

- A. A technique for efficiently propagating beneficial changes across all system modules quickly
- B. A design pattern that ensures changes in one module are automatically applied everywhere now
- C. Changes in one part of the system causing unintended effects in seemingly unrelated parts
- D. A method for testing software by sending signals through all connected components sequentially

Answer: C

Q747. What is program comprehension and why is it critical for maintenance?

- A. Understanding existing code is essential for making correct modifications without introducing bugs
- B. Program comprehension is unnecessary because maintainers can change code without understanding it
- C. Program comprehension only matters for documentation purposes and not for actual code changes
- D. Understanding code is only important during initial development and never during maintenance work

Answer: A

Q748. What role does impact analysis play in software maintenance?

- A. It identifies all system areas affected by a proposed change before implementation begins now
- B. It only identifies positive impacts and ignores any negative consequences of proposed changes
- C. It is only performed after changes are deployed to assess the damage caused by the change
- D. It has no role in maintenance because changes can be made without considering their effects

Answer: A

Q749. How does automated testing support software maintenance activities?

- A. It eliminates the need for any human involvement in the maintenance decision-making process
- B. It replaces all maintenance activities including code modification and deployment processes
- C. It quickly validates that changes have not broken existing functionality through regression tests
- D. It automatically fixes all bugs found during testing without any developer intervention needed

Answer: C

Q750. What is software rejuvenation?

- A. Adding cosmetic visual updates to the user interface without any functional changes at all
- B. Permanently retiring software and replacing it with an entirely new system from scratch now
- C. Periodically restarting or refreshing software to counteract performance degradation over time
- D. Increasing the hardware resources allocated to the software without changing any code at all

Answer: C

Q751. What is the CMMI (Capability Maturity Model Integration)?

- A. A hardware testing tool for measuring processor speed and memory performance now
- B. A database management system used exclusively for storing quality audit records
- C. A specific programming language designed for building quality management applications
- D. A framework for assessing and improving organizational process maturity and capability

Answer: D

Q752. What are the five maturity levels in CMMI?

- A. Requirements, Design, Implementation, Testing, and Deployment of system products
- B. Planning, Coding, Testing, Deploying, and Maintaining software applications now
- C. Beginning, Intermediate, Advanced, Expert, and Master skill certification levels
- D. Initial, Managed, Defined, Quantitatively Managed, and Optimizing across all processes

Answer: D

Q753. How does statistical process control apply to software quality assurance?

- A. It has no application to software because software processes cannot be measured now
- B. It only applies to manufacturing processes and is irrelevant to software development
- C. It uses statistical methods to monitor and control process performance and variation
- D. It replaces all testing activities with purely statistical models of defect prediction

Answer: C

Q754. What is the cost of quality (CoQ) model?

- A. A pricing strategy for selling software products to customers at premium prices now
- B. A budget allocation method assigning equal funding to all project activities always
- C. A salary structure for paying quality assurance engineers based on defects found
- D. A model categorizing quality costs into prevention, appraisal, and failure categories

Answer: D

Q755. How do formal technical reviews improve software quality?

- A. They replace all automated testing by manually verifying every line of source code now
- B. They systematically examine work products to find defects before they propagate further
- C. They only focus on coding style and have no impact on finding functional defect issues
- D. They are informal conversations without structure that rarely produce useful findings

Answer: B

Q756. What is the relationship between process quality and product quality?

- A. Poor processes always produce high-quality products because developers compensate fully
- B. Good processes increase the likelihood of producing good products consistently over time
- C. Process quality has no correlation with product quality in software engineering at all
- D. Product quality depends solely on individual developer skill and not on process at all

Answer: B

Q757. What is a quality gate in the development process?

- A. A software tool that automatically generates test cases from requirements documents
- B. A checkpoint where quality criteria must be met before proceeding to the next phase
- C. A physical entrance to the quality assurance department office in the building now
- D. A coding standard that restricts the number of lines per function in source code

Answer: B

Q758. How does root cause analysis improve quality assurance?

- A. It identifies underlying causes of defects to prevent similar issues from recurring again
- B. It has no practical application in software quality assurance or defect prevention work
- C. It replaces all testing activities by predicting where defects will occur in source code
- D. It only identifies surface-level symptoms without investigating the deeper reasons for them

Answer: A

Q759. What is the difference between verification and validation in QA?

- A. Verification only applies to code; validation only applies to documentation artifacts
- B. Verification ensures correct construction; validation ensures the right product is built
- C. Verification and validation are identical activities performed at the same project stage
- D. Verification is done by users; validation is done by developers during the development

Answer: B

Q760. What is a software quality assurance plan (SQAP)?

- A. A document defining quality activities, standards, and procedures for a project effort
- B. A hardware procurement list for purchasing testing equipment and tools for the lab
- C. A project schedule showing only the testing phase timeline without any other phases
- D. A source code file containing automated tests for the application being developed now

Answer: A

Q761. How is function point analysis different from LOC as a size metric?

- A. Function points measure functionality delivered; LOC measures physical source code volume size
- B. LOC is technology-independent while function points depend on the programming language used
- C. Function points are less accurate than LOC because they ignore the actual code implementation
- D. Function points and LOC always produce identical size measurements for any given software system

Answer: A

Q762. What is the Halstead complexity metric based on?

- A. The number of operators and operands in the source code of the program being measured today
- B. The number of developers who have contributed to the source code repository over project life
- C. The amount of time required to execute the software on a standard hardware configuration now
- D. The physical size of the compiled executable file measured in megabytes of storage space used

Answer: A

Q763. What is the difference between leading and lagging metrics?

- A. Lagging metrics predict future outcomes; leading metrics only describe what happened previously
- B. Leading and lagging metrics are identical concepts measuring the same things at the same time
- C. Leading metrics predict future outcomes; lagging metrics report on past results already achieved
- D. Leading metrics are less useful than lagging metrics because predictions are always inaccurate

Answer: C

Q764. How does the Goal-Question-Metric (GQM) approach structure measurement programs?

- A. It starts by collecting all possible metrics and then determines if any goals can be derived
- B. It defines goals first, then questions to assess goals, then metrics to answer those questions
- C. It requires implementing all metrics simultaneously before defining any organizational goals
- D. It randomly selects metrics without any connection to organizational goals or questions asked

Answer: B

Q765. What is coupling metrics and how does it relate to maintainability?

- A. Coupling can only be measured for object-oriented systems and not for any other paradigms
- B. Higher coupling always improves maintainability because modules share more common knowledge
- C. Coupling metrics have no relationship with software maintainability or quality in any context
- D. Coupling metrics measure interdependence between modules; higher coupling reduces maintainability

Answer: D

Q766. What is the purpose of software reliability metrics?

- A. To track the number of meetings held by the development team during the project timeline
- B. To measure the visual attractiveness of the user interface design and layout of the system
- C. To count the number of programming languages used in the software project source code now
- D. To quantify the probability of failure-free operation over a specified time period and context

Answer: D

Q767. How does maintainability index combine multiple metrics?

- A. It combines cyclomatic complexity, LOC, and Halstead volume into a single composite score value
- B. It measures only the number of comments in code as the sole indicator of maintainability level
- C. It only uses lines of code without considering any other metrics for maintainability assessment
- D. It counts the number of developers who can understand the code without any other measurements

Answer: A

Q768. What is the difference between objective and subjective software metrics?

- A. Subjective metrics can be fully automated while objective metrics always require human judgment
- B. Objective and subjective metrics are identical and always produce the same measurement results
- C. Objective metrics are less accurate than subjective metrics in all measurement scenarios today
- D. Objective metrics are directly measurable; subjective metrics require human judgment and opinion

Answer: D

Q769. What is technical debt measurement and how can it be quantified?

- A. Technical debt cannot be measured or quantified in any meaningful way using any known methods
- B. Technical debt measurement counts the number of developers who have left the project team
- C. It only measures the financial debt of the company and has no relation to software quality now
- D. It estimates the cost of fixing shortcuts and workarounds accumulated in the codebase over time

Answer: D

Q770. How do dashboard metrics help project managers make decisions?

- A. They automatically make all project decisions without requiring any human judgment or input
- B. They only display historical data and cannot help with current or future decision making work
- C. They replace the need for any communication between project managers and development teams
- D. They provide real-time visibility into project health enabling data-driven decision making today

Answer: D

Q771. What is the difference between centralized and distributed version control?

- A. Centralized systems are always superior to distributed systems for all development scenarios now
- B. Distributed systems require a constant network connection while centralized systems do not need
- C. Centralized uses one server repository; distributed gives each developer a full local repository
- D. Centralized and distributed version control systems are identical in architecture and capability

Answer: C

Q772. What is a merge conflict and how is it resolved?

- A. Merge conflicts only happen in centralized systems and never in distributed version control
- B. It occurs when changes overlap and requires manual resolution by choosing the correct changes
- C. Merge conflicts never occur in modern version control systems due to advanced automation today
- D. Conflicts are always resolved automatically without any human intervention being required now

Answer: B

Q773. What is continuous delivery (CD) in configuration management?

- A. Manually deploying software to production once every year during the scheduled maintenance window
- B. Automating the release process so software can be deployed to production at any time reliably
- C. Delivering hardware components continuously to the data center for server infrastructure setup
- D. Preventing any changes from reaching production until all features are completely implemented

Answer: B

Q774. What is the purpose of tagging in version control?

- A. Labeling hardware components in the server rack for identification during maintenance periods
- B. Assigning developers to specific tasks using labels in the project management tracking system
- C. Adding metadata to database records for improved search and retrieval performance in queries
- D. Marking specific points in the repository history as important like releases or milestones now

Answer: D

Q775. How does GitFlow branching strategy organize development work?

- A. It uses feature, develop, release, and hotfix branches with defined merge workflows for teams
- B. It creates a new repository for each feature instead of using branches within one repository
- C. It uses a single branch for all development without any branching or merging strategies now
- D. It prohibits branching entirely and requires all developers to commit directly to the main line

Answer: A

Q776. What is the role of build automation in configuration management?

- A. It eliminates the need for any testing by guaranteeing that automated builds are always correct
- B. It only manages the visual design of the user interface without any involvement in code builds
- C. It ensures consistent, repeatable builds by automating compilation, testing, and packaging steps
- D. It replaces version control by automatically generating source code from requirements documents

Answer: C

Q777. What is a configuration audit?

- A. A performance test measuring the application response time under peak load conditions only now
- B. A code review focused exclusively on the visual appearance of the user interface elements now
- C. A financial audit of the project budget to ensure money is being spent according to the plan
- D. A review verifying that configuration items conform to their documented specifications and records

Answer: D

Q778. What is trunk-based development?

- A. All development occurs in release branches without any central trunk or main branch at all now
- B. Developers never commit code and instead share changes through email attachments exclusively
- C. All developers commit to a single main branch with short-lived feature branches if needed now
- D. Each developer maintains their own permanent long-lived branch that is never merged with others

Answer: C

Q779. How does semantic versioning (SemVer) communicate change significance?

- A. All version numbers are randomly assigned without any meaning related to the type of changes
- B. Only the major version number is meaningful while minor and patch numbers are decorative only
- C. Semantic versioning uses alphabetical labels instead of numbers to indicate version differences
- D. Major version for breaking changes, minor for features, and patch for bug fixes are indicated

Answer: D

Q780. What is the purpose of a deployment pipeline?

- A. To design the user interface mockups that will be implemented by the development team later
- B. To generate project management reports for stakeholder review at monthly status update meetings
- C. To manually review every line of code before it is committed to the version control repository
- D. To automate and validate the stages between code commit and production deployment reliably

Answer: D

Q781. How is a risk exposure value calculated?

- A. By adding the total project budget to the number of team members assigned
- B. By multiplying the probability of occurrence by the estimated impact value
- C. By subtracting the actual cost from the planned budget at each milestone
- D. By dividing the number of defects found by the total lines of code written

Answer: B

Q782. What is the difference between risk transfer and risk avoidance?

- A. Transfer shifts risk to another party; avoidance eliminates the risk from the project entirely
- B. Transfer and avoidance are identical strategies with no differences in their approaches at all
- C. Neither transfer nor avoidance has any effect on the actual risk profile of the project work
- D. Transfer eliminates risks completely; avoidance shifts them to another responsible party now

Answer: A

Q783. What is a risk-driven development approach?

- A. Prioritizing development activities based on the highest identified risk areas in the project
- B. Using risk management only during the final testing phase and not during development work
- C. Ignoring all risks and developing features in random order without any priority criteria now
- D. Developing only low-risk features and permanently deferring all high-risk items from scope

Answer: A

Q784. How does the Boehm risk management model categorize project risks?

- A. Into top ten risk items that are continuously tracked and mitigated throughout the project
- B. Into risks that only affect the schedule and risks that only affect the budget separately
- C. Into a single list of risks that is created once and never updated during the project life
- D. Into exactly two categories of risks that are either accepted or completely eliminated now

Answer: A

Q785. What is sensitivity analysis in risk management?

- A. Measuring the emotional response of team members to project deadlines and pressure levels
- B. Determining which risks have the greatest potential impact on project outcomes and results
- C. Testing the application's response time under varying network latency and bandwidth levels
- D. Evaluating the sensitivity of hardware components to environmental conditions in data center

Answer: B

Q786. What is the purpose of risk monitoring and control?

- A. Creating the initial risk register at the beginning of the project and never updating it again
- B. Focusing exclusively on positive risks and ignoring all negative risk events in the project
- C. Tracking identified risks, detecting new risks, and evaluating risk response effectiveness now
- D. Eliminating all monitoring activities to reduce overhead and speed up project delivery today

Answer: C

Q787. How do qualitative and quantitative risk analysis differ?

- A. Quantitative analysis does not use any numerical values and relies solely on expert opinion
- B. Qualitative and quantitative analysis produce identical results using the same methods always
- C. Qualitative assesses risks subjectively by category; quantitative assigns numerical probabilities
- D. Qualitative analysis is always more accurate than quantitative analysis in every situation now

Answer: C

Q788. What is a risk response strategy?

- A. A technique for ignoring risks to maintain team morale and project momentum during development
- B. A financial strategy for investing project savings into high-risk stock market opportunities
- C. A planned approach for addressing identified risks through avoidance, mitigation, or transfer
- D. A reactive action taken only after a risk has already materialized and caused project damage

Answer: C

Q789. What is the role of risk communication in project management?

- A. Ensuring all stakeholders are informed about risks and their management throughout the project
- B. Delegating all risk communication to the most junior team member without any oversight given
- C. Communicating only positive project news and avoiding any mention of potential risk events now
- D. Hiding risk information from stakeholders to prevent unnecessary worry about project outcomes

Answer: A

Q790. What is a risk breakdown structure (RBS)?

- A. A technique for splitting the project budget into smaller allocations for each team member
- B. A method for breaking down source code into smaller functions for better risk assessment work
- C. A physical diagram showing the location of server hardware in the data center for risk review
- D. A hierarchical categorization of project risks organized by source and type for clarity now

Answer: D

Q791. What is SQL injection and how can it be prevented?

- A. An approved method for administrators to directly modify production database records quickly
- B. A database optimization technique that improves query performance through code injection method
- C. A testing technique where SQL queries are injected into the test suite for validation work
- D. Inserting malicious SQL through user inputs; prevented by parameterized queries and validation

Answer: D

Q792. What is cross-site scripting (XSS) vulnerability?

- A. A method for testing website compatibility across different web browsers and platforms now
- B. A programming pattern for organizing JavaScript code across multiple source code files today
- C. Injecting malicious scripts into web pages that execute in other users' browsers unexpectedly
- D. A legitimate technique for sharing scripts between different websites for code reuse purposes

Answer: C

Q793. What is the OWASP Top Ten?

- A. A ranking of the ten best programming languages for building secure web applications now
- B. A certification program for security engineers with ten levels of professional expertise
- C. A list of the most critical web application security risks updated periodically by experts
- D. A collection of ten design patterns used exclusively for building user interface components

Answer: C

Q794. What is defense in depth as a security strategy?

- A. Removing all security controls to simplify the application architecture and reduce costs
- B. Implementing security only at the network perimeter without any application-level controls
- C. Relying on a single strong security control to protect all application assets and data now
- D. Using multiple layers of security controls so that if one layer fails others still protect

Answer: D

Q795. What is the purpose of threat modeling?

- A. Creating mathematical models of the project schedule to predict the delivery date precisely
- B. Designing the database schema using entity-relationship diagrams for data storage planning
- C. Modeling the visual appearance of the application user interface in three-dimensional space
- D. Systematically identifying potential threats and vulnerabilities in the system architecture now

Answer: D

Q796. What is secure coding practice?

- A. Writing code as fast as possible without considering any security implications or standards
- B. Writing code that follows security guidelines to prevent common vulnerabilities and exploits
- C. Avoiding all input validation to improve the user experience and reduce form friction now
- D. Using only deprecated libraries because they have been tested longer than newer versions

Answer: B

Q797. What is the role of penetration testing in security engineering?

- A. Replacing all other security measures because penetration testing finds every possible issue
- B. Evaluating the visual design of the user interface for compliance with branding guidelines
- C. Simulating real-world attacks to identify exploitable vulnerabilities before attackers find them
- D. Testing the application performance under heavy network traffic load conditions and stress

Answer: C

Q798. What is a security token used for?

- A. Storing application source code in an encrypted format within the version control system
- B. Designing the database schema with proper normalization and referential integrity rules
- C. Providing a secure method for authenticating and authorizing user sessions and API access
- D. Measuring the performance of the application under concurrent user load conditions now

Answer: C

Q799. What is the difference between symmetric and asymmetric encryption?

- A. There is no difference between symmetric and asymmetric encryption in practice or theory now
- B. Symmetric encryption is always stronger than asymmetric encryption in every possible scenario
- C. Asymmetric encryption uses one shared key; symmetric encryption uses a public-private key pair
- D. Symmetric uses one shared key for both operations; asymmetric uses a public and private key pair

Answer: D

Q800. What is input validation and why is it critical for security?

- A. Input validation is a visual design concern that has no relationship to application security
- B. Input validation slows down the application and should be avoided for performance reasons now
- C. Checking user inputs against expected formats to prevent injection attacks and data corruption
- D. It only applies to database inputs and is not needed for any other type of user interaction

Answer: C

Q801. What is the key difference between software engineering and system engineering?

- A. System engineering only deals with hardware
- B. Software engineering focuses on software; system engineering covers the entire socio-technical system
- C. They are the same discipline
- D. System engineering is a subset of software engineering

Answer: B

Q802. What is the purpose of a software engineering code of ethics?

- A. To enforce programming standards
- B. To guide professional behavior and responsibility toward society
- C. To define coding syntax rules
- D. To set salary standards for engineers

Answer: B

Q803. Which type of software application involves processing large volumes of business data?

- A. Embedded software
- B. System software
- C. Business application software
- D. Real-time software

Answer: C

Q804. What does software evolution refer to?

- A. Rewriting software from scratch
- B. The process of changing software after delivery to correct faults or improve performance
- C. Installing new software
- D. The initial development phase

Answer: B

Q805. Which challenge is specific to web-based software engineering?

- A. Lack of programming languages
- B. Need for scalability, security, and rapid deployment
- C. Inability to use databases
- D. No need for user interfaces

Answer: B

Q806. What is the significance of the NATO Software Engineering Conference of 1968?

- A. It introduced the Java programming language
- B. It formally recognized software engineering as a discipline
- C. It created the internet
- D. It established the first software company

Answer: B

Q807. What is meant by software reuse?

- A. Recycling old hardware
- B. Using existing software components to build new systems
- C. Copying code without permission
- D. Reinstalling software on a different machine

Answer: B

Q808. What is emergent system property?

- A. A property visible in source code
- B. A property that arises from the interaction of system components and cannot be attributed to any single component
- C. A bug that appears randomly
- D. A feature added after release

Answer: B

Q809. Which factor has the MOST influence on software project failure?

- A. Using the wrong programming language
- B. Poor requirements management
- C. Using too many monitors
- D. Team members being in the same office

Answer: B

Q810. What is the relationship between software engineering and computer science?

- A. They are identical fields
- B. Computer science provides theoretical foundations; software engineering applies them practically
- C. Computer science is a branch of software engineering
- D. They have no relationship

Answer: B

Q811. How does the V-model differ from the traditional Waterfall model?

- A. The V-model has no testing phase
- B. The V-model explicitly associates each development phase with a corresponding testing phase
- C. The V-model eliminates design
- D. They are exactly the same

Answer: B

Q812. What is the primary risk addressed by the Spiral model?

- A. Lack of programmers
- B. Project risks identified through repeated risk analysis in each spiral iteration
- C. Poor internet connectivity
- D. Office space constraints

Answer: B

Q813. What is evolutionary prototyping?

- A. A prototype that is immediately discarded
- B. A prototype that is incrementally refined and evolves into the final system
- C. A testing technique
- D. A documentation method

Answer: B

Q814. In RAD (Rapid Application Development), what enables faster development?

- A. Skipping all testing
- B. Using component-based construction and automated tools with intensive user involvement
- C. Writing code without design
- D. Eliminating documentation entirely

Answer: B

Q815. What is the key disadvantage of the Waterfall model?

- A. It is too fast
- B. Difficulty accommodating changes after a phase is complete
- C. It requires no documentation
- D. It has too many iterations

Answer: B

Q816. What is a timeboxed iteration?

- A. An iteration with no deadline
- B. An iteration with a fixed, predetermined duration that cannot be extended
- C. An iteration that lasts forever
- D. A phase that only runs at night

Answer: B

Q817. How does the incremental model handle changing requirements?

- A. It cannot handle changes
- B. Later increments can incorporate new or changed requirements without affecting delivered increments
- C. It restarts from scratch
- D. It ignores changes completely

Answer: B

Q818. What is a process framework in software engineering?

- A. A programming library
- B. A foundation that defines the activities, actions, and tasks for a complete software process
- C. A testing tool
- D. A type of database

Answer: B

Q819. What is the primary advantage of using the Spiral model?

- A. It is the simplest model
- B. It combines iterative development with systematic risk management
- C. It never requires customer involvement
- D. It eliminates all project risks

Answer: B

Q820. What distinguishes a prescriptive process model from an adaptive one?

- A. Prescriptive models have no defined steps
- B. Prescriptive models define a specific set of activities in a defined order; adaptive models allow flexibility
- C. Adaptive models never deliver software
- D. There is no difference

Answer: B

Q821. What is the concept of 'velocity' in Agile?

- A. How fast the team types
- B. The amount of work a team can complete in one iteration, measured in story points or similar units
- C. The speed of the internet connection
- D. The number of meetings per week

Answer: B

Q822. What is pair programming in Extreme Programming (XP)?

- A. Two teams working on separate projects
- B. Two developers working together at one workstation, one writing code and the other reviewing in real time
- C. Two developers in different countries
- D. A code review done after deployment

Answer: B

Q823. What is the purpose of Sprint Planning in Scrum?

- A. To assign blame for past failures
- B. To select items from the Product Backlog and define the work plan for the upcoming Sprint
- C. To write the final project report
- D. To conduct user acceptance testing

Answer: B

Q824. What is a Kanban board used for in Agile?

- A. Drawing architectural diagrams
- B. Visualizing work in progress by showing tasks moving through workflow stages
- C. Tracking employee vacation days
- D. Managing email

Answer: B

Q825. What is continuous refactoring in Agile?

- A. Rewriting the entire system regularly
- B. Ongoing improvement of code structure and design without changing its external behavior
- C. Continuously adding new features
- D. Removing features from the product

Answer: B

Q826. What is the role of the Scrum Master during a Sprint?

- A. Assigning tasks to developers
- B. Facilitating the Scrum process, removing impediments, and shielding the team from external distractions
- C. Writing code exclusively
- D. Approving the final product

Answer: B

Q827. What is an epic in Agile?

- A. A very short user story
- B. A large body of work that can be broken down into smaller user stories
- C. A type of software bug
- D. A deployment milestone

Answer: B

Q828. What is the purpose of a Sprint Backlog?

- A. To list all company projects
- B. To contain the selected Product Backlog items and the plan for delivering them during the Sprint
- C. To track customer complaints
- D. To record team vacation schedules

Answer: B

Q829. What is the MoSCoW prioritization technique?

- A. A Russian programming method
- B. A method categorizing requirements as Must have, Should have, Could have, and Won't have
- C. A testing framework
- D. A deployment strategy

Answer: B

Q830. What is a use case scenario?

- A. A test case for unit testing
- B. A specific sequence of interactions between an actor and the system to achieve a particular goal
- C. A deployment plan
- D. A database schema

Answer: B

Q831. What is the purpose of a requirements traceability matrix?

- A. To trace network packets
- B. To map each requirement to its source, design elements, and test cases
- C. To track employee attendance
- D. To monitor server uptime

Answer: B

Q832. What is requirements ambiguity?

- A. Requirements that are too detailed
- B. Requirements that can be interpreted in more than one way
- C. Requirements that are too short
- D. Requirements written in code

Answer: B

Q833. What is the difference between data flow diagrams (DFDs) and use case diagrams?

- A. They are identical
- B. DFDs show how data moves through the system; use case diagrams show interactions between users and system functions
- C. DFDs are only for databases
- D. Use case diagrams show data flows

Answer: B

Q834. What is a JAD (Joint Application Development) session?

- A. A solo programming activity
- B. A facilitated workshop bringing stakeholders and developers together to define requirements collaboratively
- C. A hardware testing procedure
- D. A deployment meeting

Answer: B

Q835. What is the role of a domain model in requirements engineering?

- A. To model database tables
- B. To represent key concepts, relationships, and rules of the problem domain to facilitate shared understanding
- C. To model network topology
- D. To design user interfaces

Answer: B

Q836. What is the difference between structured and unstructured requirements interviews?

- A. There is no difference
- B. Structured interviews follow a predefined set of questions; unstructured interviews allow open-ended exploration
- C. Structured interviews have no questions prepared
- D. Unstructured interviews use questionnaires

Answer: B

Q837. What is the difference between top-down and bottom-up estimation?

- A. They produce identical results
- B. Top-down estimates the whole project then breaks it down; bottom-up estimates individual tasks then aggregates them
- C. Top-down is always more accurate
- D. Bottom-up ignores individual tasks

Answer: B

Q838. What is Brooks' Law?

- A. Adding more programmers always helps
- B. Adding manpower to a late software project makes it later due to communication overhead and ramp-up time
- C. Projects always finish on time
- D. Larger teams produce fewer bugs

Answer: B

Q839. What is the purpose of a Work Breakdown Structure (WBS)?

- A. To break down the team hierarchy
- B. To decompose the project into smaller, manageable work packages that can be estimated and tracked
- C. To break software into functions
- D. To break hardware into components

Answer: B

Q840. What is analogous estimation?

- A. Estimating by writing code first
- B. Using data from similar past projects to estimate the current project's effort and cost
- C. Random guessing
- D. Estimating based on lines of code only

Answer: B

Q841. What is project tracking and control?

- A. Tracking employee locations
- B. Monitoring project progress against the plan and taking corrective action when deviations occur
- C. Controlling access to the server room
- D. Tracking software downloads

Answer: B

Q842. What is the difference between a Gantt chart and a PERT chart?

- A. They are identical
- B. A Gantt chart shows tasks as horizontal bars on a timeline; a PERT chart shows task dependencies as a network diagram
- C. Gantt charts show only costs
- D. PERT charts cannot show time

Answer: B

Q843. What is the concept of 'slack time' in project scheduling?

- A. Time spent not working
- B. The amount of time an activity can be delayed without delaying the overall project completion date
- C. Time reserved for meetings
- D. Break time for developers

Answer: B

Q844. What is the difference between composition and inheritance in OO design?

- A. They are the same concept
- B. Inheritance creates an 'is-a' relationship; composition creates a 'has-a' relationship by containing objects
- C. Composition is always worse
- D. Inheritance is never used

Answer: B

Q845. What is the Template Method design pattern?

- A. A template for writing documentation
- B. A pattern that defines the skeleton of an algorithm in a superclass, deferring specific steps to subclasses
- C. A method for creating templates
- D. A UI design template

Answer: B

Q846. What is the difference between functional and object-oriented decomposition?

- A. They are identical approaches
- B. Functional decomposition breaks the system into functions; OO decomposition breaks it into objects that encapsulate data and behavior
- C. Functional decomposition is always better
- D. OO decomposition ignores data

Answer: B

Q847. What is the State design pattern?

- A. A pattern for managing US states
- B. A pattern that allows an object to alter its behavior when its internal state changes, appearing to change its class
- C. A pattern for database state management
- D. A pattern for version control

Answer: B

Q848. What is the difference between data-centered and data-flow architecture in design?

- A. They are the same
- B. Data-centered architecture has a central repository accessed by components; data-flow architecture processes data through a sequence of transformations
- C. Data-centered ignores data
- D. Data-flow never processes data

Answer: B

Q849. What is the Proxy design pattern?

- A. A network proxy server
- B. A pattern that provides a surrogate or placeholder for another object to control access to it
- C. A debugging tool
- D. A testing framework

Answer: B

Q850. What is the difference between horizontal and vertical scaling?

- A. They are the same
- B. Horizontal scaling adds more machines; vertical scaling adds more power (CPU, RAM) to existing machines
- C. Horizontal scaling reduces capacity
- D. Vertical scaling always requires new hardware

Answer: B

Q851. What is a message queue in distributed architecture?

- A. A queue for email messages
- B. An asynchronous communication mechanism that stores messages until the receiving component is ready to process them
- C. A task assignment queue
- D. A debugging tool

Answer: B

Q852. What is the Model-View-Controller (MVC) pattern?

- A. A version control model
- B. An architectural pattern that separates an application into Model (data), View (UI), and Controller (input handling)
- C. A testing methodology
- D. A project management framework

Answer: B

Q853. What is an API gateway in microservices architecture?

- A. A physical network gateway
- B. A single entry point that routes client requests to appropriate microservices and handles cross-cutting concerns
- C. A database connector
- D. A code compiler

Answer: B

Q854. What is the difference between SOA and microservices architecture?

- A. They are identical
- B. SOA uses an enterprise service bus and shares data models; microservices use lightweight protocols and each service owns its data
- C. SOA is always better
- D. Microservices never use APIs

Answer: B

Q855. What is a load balancer in software architecture?

- A. A device that weighs servers
- B. A component that distributes incoming network traffic across multiple servers to ensure no single server is overwhelmed
- C. A code optimization tool
- D. A testing utility

Answer: B

Q856. What is a service mesh in microservices?

- A. A physical network mesh
- B. A dedicated infrastructure layer that handles service-to-service communication, including load balancing, encryption, and observability
- C. A type of database
- D. A code generation tool

Answer: B

Q857. What is a containerized architecture?

- A. An architecture stored in shipping containers
- B. An approach where applications are packaged with their dependencies in lightweight containers for consistent deployment across environments
- C. An architecture using physical boxes
- D. A hardware architecture

Answer: B

Q858. What is the difference between a persona and a user profile?

- A. They are the same thing
- B. A persona is a fictional representative character based on research; a user profile is actual data about a specific user
- C. Personas are always inaccurate
- D. User profiles are fictional

Answer: B

Q859. What is the principle of consistency in UI design?

- A. Changing the UI frequently
- B. Using the same design patterns, terminology, and interactions throughout the interface so users can predict behavior
- C. Making each page look different
- D. Randomizing button positions

Answer: B

Q860. What is an A/B test in UI design?

- A. A test with only two answers
- B. A method of comparing two versions of a UI element to determine which performs better based on user metrics
- C. A binary code test
- D. An accessibility test

Answer: B

Q861. What is a card sorting exercise used for?

- A. Sorting playing cards
- B. A user research method where participants organize content into categories to inform information architecture
- C. Sorting database records
- D. Organizing code files

Answer: B

Q862. What is the difference between accessibility and usability?

- A. They are identical
- B. Accessibility ensures the interface is usable by people with disabilities; usability measures how easy and efficient the interface is for all users
- C. Accessibility is optional
- D. Usability is only about speed

Answer: B

Q863. What is the purpose of breadcrumb navigation?

- A. Tracking user cookies
- B. Showing users their current location within the site hierarchy and allowing easy navigation to parent pages
- C. Displaying error logs
- D. Managing user sessions

Answer: B

Q864. What is progressive disclosure in UI design?

- A. Gradually revealing source code
- B. Showing only essential information initially and revealing additional details as the user needs them
- C. Disclosing security vulnerabilities
- D. A type of animation effect

Answer: B

Q865. What are Nielsen's error prevention heuristics about?

- A. Preventing hardware errors
- B. Designing interfaces that prevent errors from occurring in the first place, such as confirmation dialogs and input validation
- C. Preventing network outages
- D. Preventing code compilation errors

Answer: B

Q866. What is code coverage and why is it important?

- A. The number of files in a project
- B. A metric measuring the percentage of code executed during testing, indicating how thoroughly the code is tested
- C. How many developers work on the code
- D. The number of code comments

Answer: B

Q867. What is the difference between static and dynamic code analysis?

- A. They are identical
- B. Static analysis examines code without executing it; dynamic analysis examines code behavior during execution
- C. Static analysis is always slower
- D. Dynamic analysis does not require running the code

Answer: B

Q868. What is dependency injection?

- A. Injecting malicious code
- B. A technique where an object receives its dependencies from external sources rather than creating them internally
- C. Adding new library imports
- D. A database insertion technique

Answer: B

Q869. What is a code smell?

- A. Physical odor from a computer
- B. A surface indication in code that suggests a deeper design problem requiring refactoring
- C. A compilation error
- D. A runtime crash

Answer: B

Q870. What is a linter in software development?

- A. A piece of clothing fabric
- B. A tool that analyzes source code to flag programming errors, bugs, stylistic issues, and suspicious constructs
- C. A compiler optimization
- D. A deployment tool

Answer: B

Q871. What is the YAGNI principle?

- A. You Are Going to Need It
- B. You Aren't Gonna Need It: do not implement functionality until it is actually needed
- C. Yet Another Great New Innovation
- D. Your Application Gets New Issues

Answer: B

Q872. What is technical debt?

- A. Money owed for software licenses
- B. The implied cost of future rework caused by choosing a quick solution now instead of a better approach that would take longer
- C. The cost of hardware
- D. Unpaid developer salaries

Answer: B

Q873. What is a code review and why is it important?

- A. Reviewing the project budget
- B. A systematic examination of source code by peers to identify defects, improve quality, and share knowledge
- C. A customer feedback session
- D. A marketing review

Answer: B

Q874. What is the difference between top-down and bottom-up integration testing?

- A. They are identical approaches
- B. Top-down tests from the main module downward using stubs; bottom-up tests from the lowest modules upward using drivers
- C. Top-down is always better
- D. Bottom-up never requires drivers

Answer: B

Q875. What is statement coverage in testing?

- A. Coverage of requirement statements
- B. A metric measuring the percentage of executable code statements that have been exercised by test cases
- C. The number of comment lines tested
- D. Coverage of variable declarations

Answer: B

Q876. What is branch coverage and how does it differ from statement coverage?

- A. They are the same metric
- B. Branch coverage ensures every branch of every decision point is executed, providing stronger coverage than statement coverage alone
- C. Branch coverage is weaker than statement coverage
- D. Branch coverage only applies to switch statements

Answer: B

Q877. What is the purpose of a test harness?

- A. A horse harness for testing
- B. A framework that provides the infrastructure for automating test execution, including drivers, stubs, and result reporting
- C. A physical testing device
- D. A documentation tool

Answer: B

Q878. What is the difference between smoke testing and sanity testing?

- A. They are the same thing
- B. Smoke testing verifies that critical functions work after a build; sanity testing verifies specific functionality after a minor change
- C. Smoke testing uses fire
- D. Sanity testing tests mental health

Answer: B

Q879. What is the purpose of a test stub?

- A. A ticket stub for testing events
- B. A simplified implementation of a component that a module under test depends on, providing predefined responses
- C. A type of test case
- D. A testing report

Answer: B

Q880. What is the concept of 'test automation'?

- A. Robots performing manual testing
- B. Using software tools and scripts to execute tests automatically, compare results, and report outcomes
- C. Automating the hiring of testers
- D. Automatically generating code

Answer: B

Q881. What is the difference between load testing and stress testing?

- A. They are identical
- B. Load testing measures performance under expected conditions; stress testing pushes the system beyond normal capacity to find its breaking point
- C. Load testing is easier
- D. Stress testing is always done first

Answer: B

Q882. What is impact analysis in software maintenance?

- A. Analyzing the impact of a meteor
- B. Evaluating the consequences of a proposed change on the rest of the system to identify affected components and potential risks
- C. Analyzing market impact
- D. Analyzing team morale

Answer: B

Q883. What is program comprehension in maintenance?

- A. Understanding a TV program
- B. The process by which maintainers understand the existing code's structure, behavior, and intent before making modifications
- C. A reading comprehension test
- D. Understanding user requirements

Answer: B

Q884. What is software migration?

- A. Software moving to another country
- B. The process of transferring a software system from one operating environment to another while preserving functionality
- C. Developers changing jobs
- D. Moving files between folders

Answer: B

Q885. What is the concept of 'maintainability index'?

- A. A measure of building maintenance costs
- B. A composite metric that combines code complexity, volume, and other factors to quantify how easy the software is to maintain
- C. An employee performance metric
- D. A hardware reliability metric

Answer: B

Q886. What is the role of regression testing in maintenance?

- A. It has no role
- B. Ensuring that maintenance changes have not introduced new defects into previously working functionality
- C. Testing only new features
- D. Testing only the user interface

Answer: B

Q887. What is code restructuring?

- A. Moving code files to different folders
- B. Modifying the internal structure of existing code to improve readability and maintainability without changing its external behavior
- C. Rewriting code in a different language
- D. Deleting unused code

Answer: B

Q888. What is the challenge of maintaining undocumented code?

- A. There is no challenge
- B. Maintainers must spend significant time understanding the code's purpose and logic through code reading alone, increasing the risk of introducing errors
- C. Undocumented code runs faster
- D. Undocumented code is easier to maintain

Answer: B

Q889. What is the difference between emergency and scheduled maintenance?

- A. They are the same
- B. Emergency maintenance addresses critical failures requiring immediate fixes; scheduled maintenance is planned changes during designated windows
- C. Emergency maintenance is optional
- D. Scheduled maintenance is never planned

Answer: B

Q890. What is a Service Level Agreement (SLA) in the context of maintenance?

- A. A software license agreement
- B. A contract defining expected maintenance response times, resolution times, and availability targets
- C. A developer employment contract
- D. A hardware warranty

Answer: B

Q891. What is a formal technical review (FTR)?

- A. A review of technical documents by lawyers
- B. A structured review meeting where a team of technical staff evaluates a work product to identify defects and verify quality
- C. A hardware inspection
- D. A user acceptance review

Answer: B

Q892. What is defect density as a quality metric?

- A. The physical density of a computer
- B. The number of confirmed defects found per unit of software size, such as defects per thousand lines of code
- C. The number of defects per developer
- D. The weight of defect reports

Answer: B

Q893. What is the Six Sigma approach in software quality?

- A. A Greek letter used in equations
- B. A data-driven methodology aiming to reduce defects to fewer than 3.4 per million opportunities using DMAIC
- C. A programming framework
- D. A type of testing

Answer: B

Q894. What is root cause analysis in QA?

- A. Analyzing tree roots
- B. A systematic process of identifying the fundamental reason why a defect occurred to prevent similar defects in the future
- C. Analyzing code roots
- D. A type of testing

Answer: B

Q895. What is the purpose of quality gates in a development process?

- A. Physical gates at the office entrance
- B. Checkpoints where work products must meet predefined quality criteria before proceeding to the next phase
- C. Gates in logic circuits
- D. Entrance exams for QA engineers

Answer: B

Q896. What is the Pareto analysis used for in QA?

- A. Analyzing Italian software
- B. Identifying the vital few causes that account for the majority of defects, focusing improvement efforts where they have the most impact
- C. Analyzing parrot behavior
- D. A type of performance testing

Answer: B

Q897. What is the ISO 9001 standard in relation to software quality?

- A. A programming language standard
- B. An international standard for quality management systems that defines requirements for consistent product quality and customer satisfaction
- C. A database standard
- D. A network protocol standard

Answer: B

Q898. What is the role of metrics in quality assurance?

- A. Metrics have no role in QA
- B. Metrics provide objective, quantitative data to measure process effectiveness, product quality, and guide improvement decisions
- C. Metrics only measure project cost
- D. Metrics replace testing

Answer: B

Q899. How is cyclomatic complexity calculated?

- A. By counting lines of code
- B. By counting the number of linearly independent paths through the code, calculated as $V(G) = E - N + 2P$ for the control flow graph
- C. By counting the number of classes
- D. By counting the number of variables

Answer: B

Q900. What does a high cyclomatic complexity value indicate?

- A. Simple, easy-to-test code
- B. Code with many decision paths that is more complex, harder to test, and more error-prone
- C. Fast-running code
- D. Well-documented code

Answer: B

Q901. What is function point analysis used for?

- A. Analyzing mathematical functions
- B. Measuring software size based on the functionality delivered to the user, independent of the programming language
- C. Analyzing function calls in code
- D. Measuring hardware performance

Answer: B

Q902. What are the five components measured in function point analysis?

- A. Classes, methods, variables, loops, conditions
- B. External Inputs, External Outputs, External Inquiries, Internal Logical Files, and External Interface Files
- C. Lines, functions, modules, packages, projects
- D. CPU, RAM, disk, network, display

Answer: B

Q903. What is the Halstead effort metric?

- A. The physical effort of typing code
- B. A measure derived from the number of operators and operands in code that estimates the mental effort required to develop or understand the program
- C. The effort to install software
- D. The effort to run tests

Answer: B

Q904. What is the purpose of a GQM (Goal-Question-Metric) approach in metrics?

- A. A gaming qualification metric
- B. To ensure metrics are aligned with specific goals by defining goals first, then questions to assess goal achievement, then metrics to answer questions
- C. A graphics quality metric
- D. A random metrics selection approach

Answer: B

Q905. What is the difference between subjective and objective metrics?

- A. They are the same
- B. Objective metrics are based on quantifiable measurements; subjective metrics involve human judgment and perception
- C. Subjective metrics are always better
- D. Objective metrics are always inaccurate

Answer: B

Q906. What is a Change Control Board (CCB)?

- A. A circuit board for hardware changes
- B. A group of stakeholders responsible for evaluating, approving, or rejecting proposed changes to configuration items
- C. A board game about changes
- D. A software testing board

Answer: B

Q907. What is the purpose of a build pipeline in configuration management?

- A. Building a physical pipeline
- B. An automated workflow that compiles, tests, and packages software whenever changes are committed to the repository
- C. A water pipeline
- D. An email pipeline

Answer: B

Q908. What is the difference between Git merge and Git rebase?

- A. They are identical operations
- B. Merge creates a merge commit preserving branch history; rebase moves commits to the tip of the target branch creating a linear history
- C. Merge always causes conflicts
- D. Rebase never modifies commit history

Answer: B

Q909. What is semantic versioning (SemVer)?

- A. A versioning system based on how versions sound
- B. A versioning scheme using MAJOR.MINOR.PATCH numbers where each indicates the type and scope of changes made
- C. A random version numbering scheme
- D. Versioning based on release dates

Answer: B

Q910. What is the purpose of a branching strategy?

- A. A gardening technique
- B. A set of rules defining how branches are created, named, and merged to organize parallel development and releases
- C. A marketing strategy
- D. A hiring strategy

Answer: B

Q911. What is GitFlow as a branching strategy?

- A. A water flow monitoring tool
- B. A branching model using main, develop, feature, release, and hotfix branches with defined rules for merging between them
- C. A Git tutorial
- D. A deployment tool

Answer: B

Q912. What is the purpose of configuration status accounting?

- A. Tracking the company's finances
- B. Recording and reporting the status of configuration items and change requests throughout the project lifecycle
- C. Counting lines of code
- D. Accounting for team hours

Answer: B

Q913. What is an artifact repository in CI/CD?

- A. A museum for code artifacts
- B. A centralized storage system for build outputs (compiled code, packages, containers) that can be versioned and deployed
- C. A database for test results
- D. A code hosting platform

Answer: B

Q914. What is a risk probability-impact matrix?

- A. A mathematical equation
- B. A grid that plots risks by their probability and impact to prioritize them and determine appropriate response strategies
- C. A spreadsheet template
- D. A type of decision tree

Answer: B

Q915. What are the common categories of software project risks?

- A. Only technical risks exist
- B. Technology, people, organizational, tools, requirements, and estimation risks among others
- C. Only budget risks
- D. Only schedule risks

Answer: B

Q916. What is risk exposure calculated as?

- A. Risk probability minus risk impact
- B. Risk probability multiplied by risk impact, representing the expected loss from the risk
- C. Risk probability divided by risk impact
- D. Risk probability plus risk impact

Answer: B

Q917. What is a risk trigger (or warning sign)?

- A. A gun trigger
- B. An event or condition that indicates a risk is about to occur or has occurred, prompting activation of the risk response plan
- C. A software trigger
- D. A database trigger

Answer: B

Q918. What is the concept of 'secondary risk'?

- A. A less important risk
- B. A new risk that arises as a direct result of implementing a risk response for a primary risk
- C. The second risk identified
- D. A backup risk plan

Answer: B

Q919. What is the purpose of a risk audit?

- A. Auditing the company's finances
- B. Examining and documenting the effectiveness of risk responses and the accuracy of the risk management process
- C. Auditing employee performance
- D. Auditing code quality

Answer: B

Q920. What is SQL injection and how is it prevented?

- A. Injecting SQL databases with vitamins
- B. An attack where malicious SQL code is inserted into application queries; prevented by using parameterized queries and input validation
- C. A database optimization technique
- D. A method for speeding up queries

Answer: B

Q921. What is the principle of defense in depth?

- A. Building deep bunkers
- B. Implementing multiple layers of security controls so that if one layer fails, others still provide protection
- C. Using only one very strong password
- D. Having one firewall

Answer: B

Q922. What is a session hijacking attack?

- A. Stealing a classroom session plan
- B. An attack where an attacker takes over a legitimate user's active session by stealing or predicting the session identifier
- C. Hijacking a training session
- D. Canceling a meeting

Answer: B

Q923. What is role-based access control (RBAC)?

- A. Controlling roles in a theater play
- B. A security model that restricts system access based on the roles assigned to users rather than individual user identities
- C. A role-playing game
- D. A hiring management system

Answer: B

Q924. What is the purpose of penetration testing?

- A. Testing physical building penetration
- B. Simulating real-world attacks against a system to identify vulnerabilities before malicious attackers can exploit them
- C. Testing code penetration rates
- D. Testing market penetration

Answer: B

Q925. What is a security token?

- A. A subway token
- B. A device or software that generates one-time passwords or cryptographic keys used for authenticating user identity
- C. A casino chip
- D. A gift card

Answer: B

Q926. What is the difference between a product backlog and a sprint backlog in Scrum?

- A. They are the same thing used at different stages of the project
- B. The product backlog contains all desired features for the product while the sprint backlog contains only items selected for the current sprint
- C. The sprint backlog is maintained by management while the product backlog is maintained by developers
- D. The product backlog is created once and never changes while the sprint backlog is updated daily

Answer: B

Q927. What is the Scrum of Scrums technique used for?

- A. Replacing the daily standup meeting in small teams
- B. Coordinating work across multiple Scrum teams working on the same product
- C. Training new Scrum Masters on agile practices
- D. Evaluating individual developer performance within a sprint

Answer: B

Q928. What is the Volere requirements template and what is its primary benefit?

- A. A coding template that generates requirements automatically from source code
- B. A structured template that provides a comprehensive framework for capturing all types of requirements systematically
- C. A UML diagram type used only for non-functional requirements
- D. A testing framework that validates requirements against code

Answer: B

Q929. How do quality function deployment (QFD) techniques apply to requirements engineering?

- A. QFD is only used in manufacturing and has no application in software
- B. QFD translates customer needs into technical requirements by mapping relationships between what customers want and how the system will deliver it
- C. QFD replaces all forms of requirements elicitation with automated tools
- D. QFD is a testing technique that validates requirements after implementation

Answer: B

Q930. What is the purpose of a RACI matrix in project management?

- A. To calculate the project budget using mathematical formulas
- B. To define roles and responsibilities by identifying who is Responsible, Accountable, Consulted, and Informed for each task
- C. To schedule tasks on a calendar view similar to a Gantt chart
- D. To track defects found during software testing

Answer: B

Q931. What is parametric estimation in software project management?

- A. Estimation based entirely on the project manager's intuition and experience
- B. Using statistical relationships between historical data and project parameters to calculate estimates
- C. Asking each developer to guess how long their tasks will take
- D. Copying the schedule from a previous project without any adjustments

Answer: B

Q932. What is the difference between a project baseline and a project plan?

- A. They are the same document used at different project phases
- B. A baseline is the approved version of the plan used to measure project performance against, while the plan is the working document that may be updated
- C. A baseline is created after project completion while a plan is created before
- D. A plan is used only for small projects while a baseline is used for large projects

Answer: B

Q933. What is the Builder design pattern used for?

- A. Ensuring a class has only one instance throughout the application
- B. Separating the construction of a complex object from its representation so the same construction process can create different representations
- C. Defining a one-to-many dependency between objects so that changes in one notify all dependents
- D. Providing a simplified interface to a complex subsystem

Answer: B

Q934. What is the Facade design pattern?

- A. A pattern that adds new responsibilities to an object dynamically
- B. A pattern that provides a simplified, unified interface to a set of interfaces in a subsystem
- C. A pattern that allows incompatible interfaces to work together
- D. A pattern that separates abstraction from implementation

Answer: B

Q935. What is the Interface Segregation Principle (ISP)?

- A. All classes must implement every method defined in any interface they use
- B. Clients should not be forced to depend on interfaces they do not use
- C. Interfaces should contain as many methods as possible for reusability
- D. Only one interface should exist per application module

Answer: B

Q936. What is the Iterator design pattern?

- A. A pattern that converts the interface of a class into another interface clients expect
- B. A pattern that provides a way to access elements of a collection sequentially without exposing its underlying representation
- C. A pattern that defines a family of algorithms and makes them interchangeable
- D. A pattern that composes objects into tree structures to represent part-whole hierarchies

Answer: B

Q937. What is the Broker architectural pattern?

- A. A pattern where all requests go through a single centralized database
- B. A pattern that decouples components by using a broker to coordinate communication between distributed services
- C. A pattern that requires all components to be written in the same programming language
- D. A pattern used exclusively for real-time embedded systems

Answer: B

Q938. What is the blackboard architectural pattern?

- A. A pattern where developers write requirements on a physical whiteboard
- B. A pattern where multiple specialized knowledge sources collaborate by reading from and writing to a shared data structure to solve complex problems
- C. A pattern identical to the client-server model but with a different name
- D. A pattern that restricts access to a single database to one user at a time

Answer: B

Q939. What is the F-pattern layout in UI design?

- A. A layout that arranges all elements in a fixed grid pattern
- B. A layout based on eye-tracking research showing users scan content in an F-shaped pattern, reading horizontally across the top then scanning vertically down the left side
- C. A layout that uses only the letter F as a design element
- D. A layout specifically designed for mobile-first applications

Answer: B

Q940. What is the purpose of user journey mapping?

- A. To map the physical location of users using GPS technology
- B. To visualize the complete experience a user has with a product from initial contact through engagement and into long-term use
- C. To create a route map showing the shortest navigation path in the application
- D. To document the travel expenses of the design team

Answer: B

Q941. What is the difference between a framework and a library in software implementation?

- A. They are the same thing; the terms are used interchangeably
- B. A framework calls your code (inversion of control) while a library is called by your code
- C. A library is always larger than a framework in terms of file size
- D. Frameworks are only used for web development while libraries are used for desktop applications

Answer: B

Q942. What is decision table testing?

- A. A management technique for deciding which tests to skip
- B. A black-box testing technique that uses a table to represent combinations of inputs and their corresponding expected outputs or actions
- C. A method for testing database query performance
- D. A way to schedule testing activities in a project timeline

Answer: B

Q943. What is the difference between a test driver and a test stub?

- A. They are the same thing used at different project phases
- B. A test driver calls the module under test (simulating the caller), while a test stub replaces a module called by the module under test (simulating the dependency)
- C. A test driver is used for black-box testing while a test stub is used for white-box testing
- D. A test driver runs tests automatically while a test stub requires manual execution

Answer: B

Q944. What is Lehman's Law of Continuing Change and its implications for software maintenance?

- A. Software that is never used does not need maintenance
- B. A system used in a real-world environment must be continually adapted or it becomes progressively less satisfactory to users
- C. All software eventually becomes unmaintainable regardless of effort
- D. Software maintenance costs decrease over time as the system stabilizes

Answer: B

Q945. What is the difference between prevention costs and appraisal costs in the cost of quality model?

- A. Prevention costs and appraisal costs are the same category of quality costs
- B. Prevention costs are spent on activities to prevent defects from occurring, while appraisal costs are spent on activities to detect defects that have already been introduced
- C. Appraisal costs prevent defects while prevention costs detect them
- D. Prevention costs only apply to hardware while appraisal costs only apply to software

Answer: B

Q946. How does the ISO/IEC 25010 quality model categorize software quality characteristics?

- A. It defines only two quality characteristics: functionality and performance
- B. It defines eight quality characteristics including functional suitability, performance efficiency, compatibility, usability, reliability, security, maintainability, and portability
- C. It defines quality solely based on the number of defects per thousand lines of code
- D. It categorizes quality only by user satisfaction survey scores

Answer: B

Q947. What is the depth of inheritance tree (DIT) metric and what does a high value indicate?

- A. It measures the number of files in a project; a high value indicates a large project
- B. It measures the maximum length from a class to the root of its inheritance hierarchy; a high value indicates greater complexity and potential for inherited behavior issues
- C. It measures the depth of nested loops in code; a high value indicates faster execution
- D. It measures the number of test cases per class; a high value indicates thorough testing

Answer: B

Q948. What is the fan-in and fan-out metric in software design?

- A. Fan-in counts the number of developers who modified a module while fan-out counts the number of testers who tested it
- B. Fan-in measures the number of modules that call a given module while fan-out measures the number of modules called by a given module
- C. Fan-in measures input file size while fan-out measures output file size
- D. Fan-in and fan-out are terms used only in electrical engineering, not software

Answer: B

Q949. What is the purpose of a deployment pipeline in configuration management?

- A. To physically transport software on USB drives to client locations
- B. To automate the process of building, testing, and deploying software through a series of stages from development to production
- C. To create documentation for each software release
- D. To assign developers to specific configuration management tasks

Answer: B

Q950. What is Infrastructure as Code (IaC) in the context of configuration management?

- A. Writing application code that runs on infrastructure servers
- B. Managing and provisioning computing infrastructure through machine-readable configuration files rather than manual processes
- C. Using hardware components as programming languages
- D. Coding directly on production servers without any version control

Answer: B

Q951. What is the RMMM (Risk Mitigation, Monitoring, and Management) plan?

- A. A financial auditing plan for tracking project expenses
- B. A comprehensive plan that documents risk mitigation strategies, monitoring activities, and management approaches for all identified project risks
- C. A plan for managing team member vacation schedules
- D. A security plan for preventing unauthorized access to the codebase

Answer: B

Q952. What is the difference between a risk and an issue in project management?

- A. There is no difference; the terms are interchangeable
- B. A risk is a potential future event that may or may not occur, while an issue is a current problem that has already materialized and needs immediate action
- C. An issue is less important than a risk and can be ignored
- D. Risks are always technical while issues are always organizational

Answer: B

Q953. What is the purpose of a risk-driven spiral model in software development?

- A. To eliminate all risks before any coding begins
- B. To use risk analysis as the primary driver for iteration planning, addressing the highest risks earliest through prototyping and analysis
- C. To create spiral-shaped project documentation
- D. To assign each developer a specific risk to manage independently

Answer: B

Q954. What are technology risks in software projects?

- A. Risks that only affect non-technical project activities
- B. Risks related to the use of unproven technologies, technical complexity, platform dependencies, or performance uncertainties
- C. Risks that are always impossible to mitigate
- D. Risks that only occur during the maintenance phase of a project

Answer: B

Q955. What is the purpose of threat modeling in software security engineering?

- A. To create 3D models of physical security threats to the building
- B. To systematically identify, enumerate, and prioritize potential security threats to a system so appropriate countermeasures can be designed
- C. To model the performance characteristics of the software under load
- D. To create user interface mockups that look secure to end users

Answer: B

Q956. What is cross-site request forgery (CSRF) and how is it prevented?

- A. An attack where multiple websites share the same database and it cannot be prevented
- B. An attack that tricks a user's browser into making unwanted requests to a trusted site, prevented using anti-CSRF tokens and same-site cookie attributes
- C. An attack that only targets mobile applications and is prevented by app store reviews
- D. An encryption algorithm used to protect data in transit

Answer: B

Q957. What is the CIA triad in information security?

- A. A government intelligence agency classification system for software projects
- B. Three core principles of information security: Confidentiality (data accessible only to authorized parties), Integrity (data is accurate and unaltered), and Availability (data is accessible when needed)
- C. A project management framework with three phases: Create, Implement, Archive
- D. A testing methodology using three levels of test cases

Answer: B

Q958. What is the difference between authentication and authorization in security?

- A. They are the same process and can be used interchangeably
- B. Authentication verifies the identity of a user (who they are), while authorization determines what actions or resources the authenticated user is permitted to access (what they can do)
- C. Authorization happens before authentication in all security systems
- D. Authentication is used only for web applications while authorization is used only for mobile applications

Answer: B

Q959. What is the purpose of an assertion in software implementation?

- A. To display output to the user interface
- B. To verify that a condition the programmer believes to be true actually holds at a specific point in the code, catching logical errors early
- C. To encrypt sensitive data before storing it in the database
- D. To create a backup of the source file before compilation

Answer: B

Q960. What is the Weighted Methods per Class (WMC) metric in object-oriented design?

- A. A metric that counts the total weight of class files in kilobytes
- B. The sum of complexities of all methods defined in a class, indicating the class's overall complexity and effort required to develop and maintain it
- C. The average number of parameters across all methods in a class
- D. A metric that measures how many classes inherit from a given class

Answer: B

Software Engineering — Hard

480 questions

Q961. Which of the following best describes 'socio-technical systems'?

- A. Systems that include both technical and organizational components
- B. Systems that contain only technical components
- C. Web-based social media communication platforms
- D. Internal technical documentation management systems

Answer: A

Q962. What is Lehman's first law of software evolution?

- A. Software systems must continually grow in size or become obsolete in the system context
- B. The inherent complexity of software always decreases over its lifetime during implementation
- C. Well-designed software systems never require updates or modifications at all for quality purposes
- D. A program used in a real-world environment must continually change or become progressively less useful

Answer: D

Q963. In the context of software engineering ethics, what does the ACM/IEEE Code of Ethics emphasize?

- A. Public interest and professional responsibility
- B. Rapid speed of software delivery timelines
- C. Using the minimum documentation possible
- D. Maximizing project profits above all else

Answer: A

Q964. What is the key challenge addressed by Brooks' 'No Silver Bullet' paper?

- A. Network bandwidth and throughput constraints by the organization
- B. A critical lack of modern programming languages as a standard approach
- C. Fundamental hardware performance limitations in practice typically
- D. Essential complexity of software cannot be eliminated by any single technique

Answer: D

Q965. Which concept distinguishes software engineering from ad hoc programming?

- A. Systematic, disciplined, and quantifiable approach
- B. Simply writing more lines of source code
- C. Using a standard keyboard for input tasks
- D. Utilizing more recent hardware component sets

Answer: A

Q966. What is the 'wicked problem' concept in software engineering?

- A. A straightforward and simple coding bug fix across all phases
- B. A problem that is difficult to define and has no definitive solution
- C. A software problem caused by a virus attack at every stage
- D. A recurring hardware malfunction issue type in a systematic way

Answer: B

Q967. What is the primary difference between software engineering and systems engineering?

- A. Software engineering also includes hardware design within the system boundary
- B. Systems engineering exclusively covers testing by the development process
- C. The two disciplines are entirely identical for the project goals according to best practices
- D. Systems engineering covers the entire system including hardware; SE focuses on software

Answer: D

Q968. What does 'accidental complexity' mean in software engineering?

- A. Unexpected failures in hardware components as defined by standards
- B. Complexity introduced by the tools and methods used, not inherent to the problem
- C. Random and unpredictable bugs in source code over the entire lifecycle
- D. Inherent complexity from the problem domain and its related activities

Answer: B

Q969. Which of the following represents a key principle of software engineering according to Sommerville?

- A. Always maximize the total length of source code
- B. Systems should be developed using a managed and understood process
- C. Avoid writing any form of technical documentation
- D. Use only one single programming language always

Answer: B

Q970. What role does abstraction play in software engineering?

- A. It increases the total length of the code
- B. It eliminates the need for software testing
- C. It helps manage complexity by hiding unnecessary details
- D. It completely removes all project documentation

Answer: C

Q971. What is the Rational Unified Process (RUP)?

- A. An iterative framework with four phases: Inception, Elaboration, Construction, Transition
- B. A basic and simple coding technique only as defined by standards for all project stakeholders
- C. A relational database design model type in all development efforts by the project team members
- D. A specialized software testing tool suite within the given constraints

Answer: A

Q972. What distinguishes the concurrent development model from other models?

- A. It only supports a single developer at a time to achieve project objectives
- B. All activities occur simultaneously with states rather than sequential phases
- C. It completely eliminates all development phases and related components
- D. It has absolutely no testing activities at all in the engineering discipline

Answer: B

Q973. In the Spiral model, what happens if a risk cannot be mitigated?

- A. Skip directly to the deployment phase
- B. Continue development regardless of the risk
- C. Ignore the risk and move forward
- D. The project may be terminated

Answer: D

Q974. What is the Personal Software Process (PSP)?

- A. A hardware development process model type as part of the methodology
- B. A standard network communication protocol in the development process
- C. A team management and organizational method for effective project outcomes
- D. A structured framework for individual developers to measure and improve their work

Answer: D

Q975. What is the Team Software Process (TSP)?

- A. A method designed for a single developer only during the software lifecycle
- B. A specific type of software testing process within the project scope
- C. A hardware design and manufacturing method for the development team
- D. A framework that builds on PSP to guide teams in developing software-intensive products

Answer: D

Q976. What does CMMI stand for and what is its purpose?

- A. Capability Maturity Model Integration — a process improvement framework
- B. Central Management Module Index for projects in the system context
- C. Computer Made Machine Interface for hardware throughout the project
- D. Code Maintenance Manual Integration standard during implementation

Answer: A

Q977. What are the five maturity levels of CMMI?

- A. Initial, Managed, Defined, Quantitatively Managed, Optimizing
- B. Code, Test, Deploy, Monitor, Maintain in practice typically
- C. Beginner, Intermediate, Advanced, Expert, Master
- D. Plan, Do, Check, Act, Review for quality purposes

Answer: A

Q978. What is the key difference between prescriptive and adaptive process models?

- A. Prescriptive models do not involve any planning at every stage across all phases
- B. Adaptive models typically require more documentation by the organization
- C. There is absolutely no real difference between them as a standard approach
- D. Prescriptive models define a fixed set of activities; adaptive models adjust based on feedback

Answer: D

Q979. What is the Cleanroom software engineering approach?

- A. A formal methods approach emphasizing defect prevention over defect removal
- B. A hardware testing and validation environment according to best practices
- C. A physical room dedicated to coding activities for the project goals
- D. A model focused on clean code formatting style in a systematic way

Answer: A

Q980. In the context of process models, what is a 'process pattern'?

- A. A software design pattern for code modules within the system boundary
- B. An individual test case for code validation and its related activities
- C. A specific coding style guideline for teams by the development process
- D. A template describing a proven process solution to a recurring problem

Answer: D

Q981. What is the Scaled Agile Framework (SAFe)?

- A. A framework for scaling Agile practices to enterprise level
- B. A framework designed for individual small teams
- C. A comprehensive testing automation framework tool
- D. An established coding standard for all projects

Answer: A

Q982. What is the difference between Scrum and Kanban?

- A. Scrum uses fixed-length sprints; Kanban uses continuous flow with WIP limits
- B. Kanban uses fixed-length sprints; Scrum does not across all phases
- C. The two methodologies are completely identical at every stage
- D. Scrum has no defined roles; Kanban has defined roles in a systematic way

Answer: A

Q983. What is 'technical debt' in Agile context?

- A. Physical hardware equipment purchase expenses according to best practices
- B. The cost of rework caused by choosing quick solutions over better approaches
- C. Annual recurring software licensing fee costs within the system boundary
- D. Financial money owed to the development team for the project goals

Answer: B

Q984. What is the Crystal methodology family?

- A. A family of Agile methodologies sized by team size and criticality
- B. A specialized software testing automation tool by the development process
- C. A general-purpose programming language format and its related activities
- D. A relational database management system design over the entire lifecycle

Answer: A

Q985. What is Feature-Driven Development (FDD)?

- A. A random unstructured development approach only as defined by standards
- B. A comprehensive documentation standard and format within the given constraints
- C. An Agile method that organizes development around building features in two-week iterations
- D. A formal software testing methodology and approach for all project stakeholders

Answer: C

Q986. In Agile estimation, what is 'Planning Poker'?

- A. A recreational team card game activity only in all development efforts
- B. A project planning and scheduling tool suite by the project team members
- C. A formal risk analysis assessment method type and related components
- D. A consensus-based estimation technique using cards with numbers

Answer: D

Q987. What is the purpose of a 'Burndown Chart'?

- A. To track overall project financial expenses
- B. To show remaining work versus time in a sprint or project
- C. To measure and assess overall code quality
- D. To track daily employee attendance records

Answer: B

Q988. What is Dynamic Systems Development Method (DSDM)?

- A. An Agile framework that fixes time and cost, varying features to meet business needs
- B. A standard network communication protocol type for effective project outcomes
- C. A hardware design and manufacturing method only to achieve project objectives
- D. A relational database design methodology approach in the engineering discipline

Answer: A

Q989. What does 'MoSCoW' prioritization stand for?

- A. Monitor, Observe, Scan, Check, Optimize, Watch
- B. Must have, Should have, Could have, Won't have
- C. A Russian city
- D. Manage, Organize, Schedule, Control, Outline, Work

Answer: B

Q990. What is the concept of 'Emergent Design' in Agile?

- A. Designing the entire system completely upfront first
- B. Allowing the design to evolve incrementally as understanding grows
- C. Having absolutely no design process at all ever
- D. Copying designs from existing similar systems only

Answer: B

Q991. What is the Volere requirements template?

- A. A software testing automation template setup within the project scope
- B. A comprehensive template for organizing and documenting requirements with quality gateways
- C. A deployment configuration setup template for the development team throughout the project
- D. A source code implementation template file during the software lifecycle

Answer: B

Q992. What is goal-oriented requirements engineering?

- A. An approach where requirements are derived from stakeholder and system goals
- B. Setting performance and speed benchmarks for quality purposes
- C. Setting rigid project deadline target dates in the system context
- D. Writing only goals instead of actual code during implementation

Answer: A

Q993. What is the KAOS methodology?

- A. A deployment and release methodology guide by the organization
- B. A goal-oriented requirements engineering methodology using formal methods
- C. A specific software testing methodology only as a standard approach
- D. A random unstructured development approach in practice typically

Answer: B

Q994. What is requirements volatility?

- A. Requirements that remain perfectly stable always at every stage
- B. A type of volatile computer memory hardware across all phases
- C. A recurring issue with hardware component sets in a systematic way
- D. The tendency of requirements to change during the development process

Answer: D

Q995. What is viewpoint-oriented requirements engineering?

- A. Collecting requirements from multiple stakeholder viewpoints to ensure completeness
- B. Only considering developer technical viewpoints within the system boundary
- C. Deliberately ignoring all stakeholder views given according to best practices
- D. Using only a single stakeholder perspective view for the project goals

Answer: A

Q996. What is the difference between 'shall' and 'should' in requirements documents?

- A. 'Shall' indicates mandatory requirements; 'should' indicates desirable but optional
- B. The two terms mean exactly the same thing always by the development process
- C. Neither term is used in requirements documents ever over the entire lifecycle
- D. 'Should' is mandatory and 'shall' is always optional and its related activities

Answer: A

Q997. What is formal specification in requirements engineering?

- A. Drawing informal pictures of system requirements
- B. Using mathematical notation to precisely define system requirements
- C. Using programming code as the requirements source
- D. Writing requirements in an informal manner only

Answer: B

Q998. What are cross-cutting concerns in requirements?

- A. Requirements that are very easy to implement fast as defined by standards
- B. Requirements that are considered not important for all project stakeholders
- C. Requirements that affect multiple modules and cannot be cleanly decomposed
- D. Requirements from only one single stakeholder within the given constraints

Answer: C

Q999. What is the i* (i-star) framework?

- A. A numerical system for rating software products in all development efforts
- B. A goal-oriented and agent-oriented modeling framework for requirements engineering
- C. A general-purpose application programming framework by the project team members
- D. A comprehensive software testing automation framework and related components

Answer: B

Q1000. What is the difference between derived and imposed requirements?

- A. Derived requirements are always considered optional in the development process during the software lifecycle
- B. Derived requirements come from higher-level requirements; imposed requirements come from external constraints
- C. They are fundamentally the same concept entirely to achieve project objectives in the engineering discipline
- D. Imposed requirements are always derived from code for effective project outcomes as part of the methodology

Answer: B

Q1001. What is the COCOMO II model's key improvement over COCOMO I?

- A. It uses significantly simpler mathematical formulas within the system boundary
- B. It provides much faster calculation processing speed for the project goals
- C. It produces much better visual graphics output according to best practices
- D. It accounts for modern development practices like reuse and object-oriented development

Answer: D

Q1002. What is the Delphi estimation technique?

- A. A specific programming implementation technique only
- B. A comprehensive software testing methodology type
- C. A production deployment automation technique type
- D. An expert-based estimation method using anonymous iterative rounds

Answer: D

Q1003. What is the Putnam model (SLIM)?

- A. A personal health and diet improvement plan only by the development process
- B. A comprehensive software testing process model type and its related activities
- C. A software architecture and design model for teams over the entire lifecycle
- D. A software estimation model relating effort to delivery time using the Rayleigh curve

Answer: D

Q1004. What is the difference between bottom-up and top-down estimation?

- A. Top-down estimation is always much more accurate in all development efforts by the project team members
- B. Bottom-up estimation is always significantly faster within the given constraints
- C. Bottom-up estimates individual tasks then aggregates; top-down estimates the whole project then decomposes
- D. Both approaches always give identical final results as defined by standards for all project stakeholders

Answer: C

Q1005. What is the concept of 'Brooks' Law'?

- A. Adding people to a late software project makes it later
- B. A specific programming coding law and its rules
- C. Established laws about code formatting standards
- D. More developers always mean faster project work

Answer: A

Q1006. What is the Schedule Performance Index (SPI) in EVM?

- A. A comprehensive software testing progress index in the engineering discipline
- B. A specific source code quality coding metric only to achieve project objectives
- C. The ratio of earned value to planned value, indicating schedule efficiency
- D. A generic system processing speed measurement and related components

Answer: C

Q1007. What is the Cost Performance Index (CPI) in EVM?

- A. A general consumer price comparison index only for effective project outcomes
- B. A specific source code quality coding metric only as part of the methodology
- C. The ratio of earned value to actual cost, indicating cost efficiency
- D. A hardware system performance rating metric in the development process

Answer: C

Q1008. What is Monte Carlo simulation used for in project management?

- A. A recreational casino gambling game activity only during the software lifecycle
- B. Simulating project outcomes using probability distributions to assess risk
- C. A specific software coding implementation technique within the project scope
- D. A software architecture and design method guide for the development team

Answer: B

Q1009. What is the concept of 'crashing' in project scheduling?

- A. An unplanned system failure or complete breakdown
- B. Adding resources to critical path tasks to reduce project duration
- C. Permanently removing tasks from the project plan
- D. Completely cancelling the entire project scope

Answer: B

Q1010. What is fast tracking in project management?

- A. Reducing the total number of project requirements during implementation
- B. Requiring the team to work overtime every day throughout the project
- C. Skipping all software testing activities entirely in the system context
- D. Performing sequential activities in parallel to compress the schedule

Answer: D

Q1011. What is the SOLID acronym in object-oriented design?

- A. Single Responsibility, Open/Closed, Liskov Substitution, Interface Segregation, Dependency Inversion
- B. A comprehensive software testing framework set for effective project outcomes
- C. A specific general-purpose coding language type in the engineering discipline
- D. A classification of physical building materials and related components to achieve project objectives

Answer: A

Q1012. What is the Dependency Inversion Principle?

- A. All dependencies should always point downward only as part of the methodology
- B. All module dependencies should be completely removed in the development process
- C. High-level modules should not depend on low-level modules; both should depend on abstractions
- D. Dependencies between modules should be circular during the software lifecycle

Answer: C

Q1013. What is the Interface Segregation Principle?

- A. Avoid defining all interfaces in the system entirely
- B. Use one single large monolithic interface always
- C. Clients should not be forced to depend on interfaces they don't use
- D. Combine all interfaces into a single large one

Answer: C

Q1014. What is the Decorator design pattern?

- A. A pattern that attaches additional responsibilities to an object dynamically
- B. A user interface visual styling pattern only within the project scope
- C. A relational database schema design pattern type for the development team
- D. A comprehensive software testing design pattern throughout the project

Answer: A

Q1015. What is the difference between composition and inheritance?

- A. The two concepts are completely identical always in the system context
- B. Inheritance is never used in modern source code for quality purposes
- C. Composition is always the wrong approach to use during implementation
- D. Inheritance creates an 'is-a' relationship; composition creates a 'has-a' relationship

Answer: D

Q1016. What is the Adapter design pattern?

- A. A comprehensive software testing adapter tool
- B. A relational database connection driver and tool
- C. A physical electrical power adapter device unit
- D. A pattern that allows incompatible interfaces to work together

Answer: D

Q1017. What is the Command design pattern?

- A. A pattern that encapsulates a request as an object, allowing parameterization and queuing
- B. A production deployment automation command set by the organization at every stage
- C. A command-line interface terminal command only in practice typically
- D. A comprehensive software testing command suite as a standard approach

Answer: A

Q1018. What is the Template Method pattern?

- A. A source code boilerplate template file format across all phases in a systematic way
- B. A software testing test case plan template set according to best practices
- C. A comprehensive project documentation template for the project goals
- D. A pattern that defines the skeleton of an algorithm, deferring some steps to subclasses

Answer: D

Q1019. What is the concept of 'design debt'?

- A. Financial debt for purchasing design tools only within the system boundary
- B. The accumulated cost of design shortcuts that will require future rework
- C. A specific type of design pattern category name by the development process
- D. A concept related only to software testing work and its related activities

Answer: B

Q1020. What is the difference between architectural design and detailed design?

- A. They are fundamentally identical concepts entirely over the entire lifecycle as defined by standards
- B. Architectural design only involves UI layout work in all development efforts by the project team members
- C. Detailed design always comes before architectural for all project stakeholders within the given constraints
- D. Architectural design defines system structure; detailed design specifies algorithms and data structures within components

Answer: D

Q1021. What is the TOGAF framework?

- A. The Open Group Architecture Framework for enterprise architecture
- B. A production deployment automation framework for effective project outcomes
- C. A general-purpose source code framework tool to achieve project objectives
- D. A comprehensive software testing framework type in the engineering discipline

Answer: A

Q1022. What is the Architecture Tradeoff Analysis Method (ATAM)?

- A. A method for evaluating architectural decisions against quality attribute requirements
- B. A mathematical formula for doing calculations as part of the methodology
- C. A comprehensive software testing methodology during the software lifecycle
- D. A specific source code writing methodology only in the development process

Answer: A

Q1023. What is the concept of 'architectural erosion'?

- A. The gradual degradation of architecture as implementation deviates from the intended design
- B. A specific category of reported software defect for the development team
- C. A concept related to software testing only throughout the project in the system context
- D. Physical decay of building materials over time within the project scope

Answer: A

Q1024. What is a Domain-Specific Software Architecture (DSSA)?

- A. A generic all-purpose architecture template type during implementation
- B. A network infrastructure architecture blueprint in practice typically
- C. An architecture tailored to a specific application domain with reusable components
- D. A comprehensive software testing architecture for quality purposes

Answer: C

Q1025. What is the CQRS pattern?

- A. A specific source code implementation pattern by the organization
- B. A standard database query language command type as a standard approach
- C. A comprehensive software testing pattern type at every stage
- D. Command Query Responsibility Segregation — separating read and write operations

Answer: D

Q1026. What is the Strangler Fig pattern?

- A. A migration pattern that incrementally replaces a legacy system by gradually routing traffic to a new system
- B. An application security protection pattern type within the system boundary by the development process
- C. A comprehensive software testing pattern approach for the project goals according to best practices
- D. A botanical plant species classification category across all phases in a systematic way

Answer: A

Q1027. What is the concept of 'fitness functions' in evolutionary architecture?

- A. Physical exercise routines for the team members and its related activities
- B. User interface design metrics and score values as defined by standards
- C. System performance benchmarks and baselines set over the entire lifecycle
- D. Automated tests that verify architectural characteristics are maintained

Answer: D

Q1028. What is the hexagonal architecture (ports and adapters)?

- A. A building shaped with six equal sides design for all project stakeholders
- B. A relational database schema design method type within the given constraints
- C. An architecture that isolates the core logic from external concerns using ports and adapters
- D. A network topology and routing structure layout in all development efforts

Answer: C

Q1029. What is the difference between architectural patterns and architectural styles?

- A. Styles are more specific than patterns overall in the engineering discipline
- B. Patterns are broader in scope than styles are and related components to achieve project objectives
- C. Styles are broader design philosophies; patterns are specific solutions to recurring problems
- D. The two concepts are completely identical overall by the project team members

Answer: C

Q1030. What is the concept of 'Conway's Law' in software architecture?

- A. A specific legal regulatory requirement and law for effective project outcomes
- B. A standard source code formatting law and rule as part of the methodology
- C. A comprehensive software testing principle type in the development process
- D. Organizations design systems that mirror their own communication structure

Answer: D

Q1031. What is the difference between formative and summative usability evaluation?

- A. Formative evaluation only occurs after launch date within the system boundary
- B. The two types of evaluation are the same thing according to best practices
- C. Summative evaluation occurs during development by the development process
- D. Formative evaluates during development to improve; summative evaluates the finished product

Answer: D

Q1032. What is cognitive load theory in UI design?

- A. A database theory about query optimization work within the given constraints in all development efforts
- B. A programming concept about memory management as defined by standards for all project stakeholders
- C. The theory that users have limited mental processing capacity, and designs should minimize unnecessary cognitive effort
- D. Brain scanning technology for development teams and its related activities over the entire lifecycle

Answer: C

Q1033. What is the Gestalt principle of proximity?

- A. Physical distance between computer workstations by the project team members
- B. A security principle about access proximity level to achieve project objectives
- C. A networking concept about node proximity range and related components
- D. Elements placed close together are perceived as belonging to the same group

Answer: D

Q1034. What is WCAG?

- A. Web Content Accessibility Guidelines — standards for making web content accessible
- B. A creational object-oriented design pattern type as part of the methodology
- C. A comprehensive software testing framework tool for effective project outcomes
- D. A general-purpose programming language type only in the engineering discipline

Answer: A

Q1035. What is the concept of 'dark patterns' in UI design?

- A. Security patterns for data protection only within the project scope
- B. Night mode designs for reducing eye strain during the software lifecycle
- C. Dark color themes for user interface screens in the development process
- D. Deceptive design practices that trick users into unintended actions

Answer: D

Q1036. What is card sorting in UX research?

- A. A method where users organize topics into categories to inform information architecture
- B. Sorting records stored in a database table only throughout the project
- C. Sorting a deck of standard playing cards only for the development team
- D. A specific software testing automation technique in the system context

Answer: A

Q1037. What is the System Usability Scale (SUS)?

- A. A performance scale for rating hardware speed in practice typically
- B. A standardized questionnaire for measuring perceived usability of a system
- C. A musical scale for composing melody pieces during implementation
- D. A weight scale for measuring device components for quality purposes

Answer: B

Q1038. What is the difference between learnability and memorability in usability?

- A. Learnability is considered less important overall for the project goals according to best practices
- B. Memorability always comes first in any evaluation at every stage across all phases in a systematic way
- C. The two concepts are fundamentally the same thing as a standard approach by the organization
- D. Learnability is how easy it is to learn initially; memorability is how easy it is to reuse after a period of not using

Answer: D

Q1039. What is the concept of 'information scent' in UI design?

- A. Visual and textual cues that help users predict what they'll find by following a link or path
- B. Data encryption and protection techniques used by the development process
- C. The physical fragrance of a commercial product within the system boundary
- D. File compression and archiving algorithms used and its related activities

Answer: A

Q1040. What is the difference between user-centered design (UCD) and activity-centered design (ACD)?

- A. User-centered design completely ignores all users for all project stakeholders
- B. UCD focuses on user needs and preferences; ACD focuses on the activities users perform
- C. Activity-centered design is entirely obsolete now as defined by standards
- D. The two approaches are completely identical overall over the entire lifecycle

Answer: B

Q1041. What is the concept of 'clean code'?

- A. Code that has absolutely no comments at all as defined by standards
- B. Source code that has been fully minimized down for all project stakeholders
- C. Source code that has been fully obfuscated over within the given constraints
- D. Code that is easy to read, understand, and maintain, following best practices

Answer: D

Q1042. What is the Boy Scout Rule in coding?

- A. A comprehensive software testing protocol type
- B. Write source code in an outdoor setting always
- C. Always leave the code cleaner than you found it
- D. A recreational camping and outdoors activity rule

Answer: C

Q1043. What is the Law of Demeter?

- A. A relational database normalization rule set type and related components
- B. A fundamental law of classical physics principles by the project team members
- C. An ancient Greek legal and civil law code rule in all development efforts
- D. A design guideline stating that an object should only talk to its immediate friends

Answer: D

Q1044. What is technical debt and how does it relate to implementation?

- A. Annual recurring software licensing fee cost totals for effective project outcomes
- B. The cost of shortcuts in implementation that will need to be fixed later
- C. Monetary financial debt owed by a company entity to achieve project objectives
- D. Physical hardware equipment purchase cost amounts in the engineering discipline

Answer: B

Q1045. What is the difference between defensive programming and offensive programming?

- A. Offensive programming is always incorrect to use during the software lifecycle
- B. The two approaches are completely identical overall in the development process
- C. Defensive handles errors gracefully; offensive fails fast to expose bugs early
- D. Sports programming for athletic events coverage as part of the methodology

Answer: C

Q1046. What is code obfuscation?

- A. A specific debugging and tracing technique tool for the development team
- B. A comprehensive software testing method approach throughout the project
- C. Deliberately making code difficult to understand to prevent reverse engineering
- D. Making source code more readable for developers within the project scope

Answer: C

Q1047. What is the concept of 'separation of concerns' in implementation?

- A. Separating test code from production code files during implementation
- B. Organizing code so each section addresses a distinct concern or functionality
- C. Physical separation of developer workstations area in the system context
- D. Separating frontend code from backend code only for quality purposes

Answer: B

Q1048. What is a design smell versus a code smell?

- A. Design smells are always much worse than code smells as a standard approach
- B. The two concepts are completely identical and same in practice typically
- C. Code smells are always much worse than design smells by the organization
- D. Code smells are implementation-level issues; design smells are higher-level structural problems

Answer: D

Q1049. What is the 'broken window theory' in software development?

- A. A user interface visual design theory and concept in a systematic way
- B. A theory about physical broken windows in homes at every stage across all phases
- C. Neglecting small code quality issues leads to progressively worse code quality over time
- D. A comprehensive software testing theory approach for the project goals

Answer: C

Q1050. What is the principle of least astonishment in programming?

- A. Deliberately surprising users with behavior changes according to best practices
- B. Minimizing features for simplicity above all else within the system boundary
- C. A component should behave in a way that most users expect it to behave
- D. Maximizing performance above all other attributes by the development process

Answer: C

Q1051. What is mutation testing?

- A. Performing genetic and biological testing samples
- B. Testing software for computer virus threats only
- C. Introducing small changes to code to check if tests can detect them
- D. Testing for mutations in stored database data only

Answer: C

Q1052. What is the difference between structural and functional testing?

- A. Structural testing completely ignores the code base during implementation for quality purposes
- B. Functional testing examines internal code structure throughout the project in the system context
- C. Structural testing is based on internal code structure; functional testing is based on specifications
- D. The two approaches are fundamentally the same thing for the development team

Answer: C

Q1053. What is path testing?

- A. Testing all possible execution paths through the code
- B. Testing GPS navigation system paths and routes
- C. Testing operating system file paths and path names
- D. Testing network routing paths and protocol flows

Answer: A

Q1054. What is cyclomatic complexity and how does it relate to testing?

- A. A metric related to bicycle complexity and speed in practice typically as a standard approach
- B. A relational database complexity metric and score by the organization at every stage
- C. A code complexity metric that indicates the minimum number of test cases needed for full branch coverage
- D. A network topology complexity metric and rating across all phases in a systematic way

Answer: C

Q1055. What is fuzz testing?

- A. Testing using fuzzy logic-based algorithm systems for the project goals
- B. Providing random, invalid, or unexpected data as input to find vulnerabilities
- C. Testing audio and sound reproduction quality level within the system boundary
- D. Testing with fuzzy and blurred test image inputs according to best practices

Answer: B

Q1056. What is the concept of 'test coverage criteria'?

- A. The total file size of all test script source files
- B. Rules that define what aspects of the software must be tested
- C. The total number of testers on the project team
- D. Simply how many test scripts currently exist total

Answer: B

Q1057. What is property-based testing?

- A. Testing real estate management software systems by the development process
- B. Testing that specified properties or invariants hold for all valid inputs
- C. Testing CSS stylesheet property values and styles and its related activities
- D. Testing object property and getter method values over the entire lifecycle

Answer: B

Q1058. What is the pesticide paradox in software testing?

- A. Testing agricultural management software systems within the given constraints
- B. Running the same tests repeatedly will eventually stop finding new defects
- C. Using pesticides on actual computer bugs and pests as defined by standards
- D. A chemical testing methodology approach and method for all project stakeholders

Answer: B

Q1059. What is the difference between fault, error, and failure?

- A. The three concepts are fundamentally the same thing in all development efforts
- B. An error is always immediately visible to all users to achieve project objectives
- C. A fault is a defect in code; an error is an incorrect state; a failure is observable incorrect behavior
- D. A failure is what always causes a fault to happen by the project team members and related components

Answer: C

Q1060. What is model-based testing?

- A. Testing fashion model runway performance quality in the engineering discipline
- B. Generating test cases automatically from a model of the system's behavior
- C. Testing relational database data model structures for effective project outcomes
- D. Testing three-dimensional rendering model objects as part of the methodology

Answer: B

Q1061. What is Lehman's Law of Self-Regulation?

- A. Software evolution processes are self-regulating with statistically determinable trends
- B. Software automatically regulates itself without help in practice typically
- C. Programs never need any updates or patches applied as a standard approach
- D. Developers regulate themselves without any guidance by the organization

Answer: A

Q1062. What is the concept of 'software archaeology'?

- A. Digging up old physical computer equipment devices at every stage across all phases
- B. A museum exhibit about computing history artifacts in a systematic way for the project goals
- C. Physical hardware recovery and data retrieval work according to best practices
- D. The study and recovery of knowledge about legacy systems through code analysis and documentation

Answer: D

Q1063. What is the difference between wrapping and re-engineering legacy systems?

- A. Re-engineering is always significantly faster overall and its related activities
- B. Wrapping is always significantly more expensive work by the development process
- C. The two approaches are completely identical overall within the system boundary
- D. Wrapping adds a new interface around the legacy system; re-engineering restructures its internals

Answer: D

Q1064. What is the concept of 'feature interaction' in maintenance?

- A. Comprehensive feature testing and validation plans for all project stakeholders
- B. Features physically meeting one another in person over the entire lifecycle
- C. Formal feature documentation and specification docs as defined by standards
- D. Unexpected behavior when independently developed features conflict or interfere

Answer: D

Q1065. What is the ripple effect in software maintenance?

- A. Physical water ripples in a pond surface area within the given constraints
- B. A user interface animation and transition effect by the project team members
- C. Audio sound effects in the user interface design in all development efforts
- D. The propagation of changes through the system, causing unintended side effects

Answer: D

Q1066. What is the Lehman-Belady model of software evolution?

- A. A creational object-oriented design pattern type to achieve project objectives
- B. A comprehensive software testing model and plan in the engineering discipline
- C. A purely mathematical model for calculation work and related components
- D. A set of laws describing the behavior and evolution of large software systems

Answer: D

Q1067. What is technical debt in the context of maintenance?

- A. Accumulated shortcuts and suboptimal decisions that increase future maintenance cost
- B. Financial debt for purchasing hardware equipment for effective project outcomes
- C. Monthly developer salary payments and wage costs as part of the methodology
- D. Annual recurring software licensing fee cost totals in the development process

Answer: A

Q1068. What is the concept of 'design recovery' in maintenance?

- A. Using reverse engineering to reconstruct higher-level design abstractions from code
- B. Recovering from catastrophic design errors quickly within the project scope
- C. Recovering accidentally deleted design documents during the software lifecycle
- D. Creating design backup copies and archive storage for the development team

Answer: A

Q1069. What is software rejuvenation?

- A. Periodically restarting or reinitializing software to prevent degradation from accumulated errors
- B. Making software visually look brand new and fresh throughout the project in the system context
- C. Completely rewriting the entire software system over during implementation
- D. Updating only the user interface visual design look for quality purposes in practice typically

Answer: A

Q1070. What is the relationship between maintainability and other quality attributes?

- A. Maintainability is completely independent of everything at every stage
- B. Maintainability often trades off with performance but supports reliability and portability
- C. There is absolutely no relationship at all between them as a standard approach
- D. Maintainability conflicts with all other quality attributes by the organization

Answer: B

Q1071. What is the GQM (Goal-Question-Metric) approach?

- A. A recreational video gaming strategy method approach in all development efforts by the project team members
- B. A source code implementation coding approach method and related components to achieve project objectives
- C. A systematic approach to defining goals, deriving questions, and identifying metrics for quality measurement
- D. A comprehensive software testing approach method only in the engineering discipline

Answer: C

Q1072. What is Statistical Process Control (SPC) in software QA?

- A. A source code implementation coding technique method as part of the methodology
- B. A software architecture and design method guide book in the development process
- C. Using statistical methods to monitor and control software processes to ensure they operate predictably
- D. A statistical report generated for project managers for effective project outcomes

Answer: C

Q1073. What is the concept of 'process maturity' in QA?

- A. The overall size of the development process and scope for the development team throughout the project
- B. Processes that have become outdated and old over time during the software lifecycle
- C. The execution speed of all development processes running within the project scope
- D. The extent to which an organization's processes are defined, managed, measured, and continuously improved

Answer: D

Q1074. What is Root Cause Analysis (RCA)?

- A. A systematic process for identifying the underlying causes of defects or problems
- B. Finding the roots of a physical plant in soil in the system context
- C. A production deployment automation process method for quality purposes
- D. A source code implementation coding technique only during implementation

Answer: A

Q1075. What is the difference between product quality and process quality?

- A. Product quality is entirely not measurable or assessable by the organization
- B. The two concepts are completely identical to each other in practice typically
- C. Product quality measures the software itself; process quality measures the development process
- D. Process quality is completely irrelevant to all outcomes as a standard approach

Answer: C

Q1076. What is a causal analysis meeting?

- A. An informal casual meeting with team colleagues at every stage
- B. A meeting to analyze the root causes of defects and identify preventive actions
- C. A recreational team building social activity event in a systematic way
- D. A project planning and scheduling meeting session across all phases

Answer: B

Q1077. What is the concept of 'quality gates' in software development?

- A. A type of network firewall configuration and setup according to best practices
- B. Physical security gates at building entrance areas for the project goals
- C. Checkpoints in the development process where quality criteria must be met before proceeding
- D. Security gates for data access control setup only within the system boundary

Answer: C

Q1078. What is the Clean Room approach to quality?

- A. A physically tidy and organized workspace and area by the development process
- B. A hardware manufacturing clean room facility space and its related activities
- C. A software development approach using formal verification to prevent defects rather than test for them
- D. A specific testing environment setup and space area over the entire lifecycle

Answer: C

Q1079. What is the relationship between CMMI and software quality?

- A. CMMI only measures physical hardware quality metrics for all project stakeholders
- B. There is absolutely no relationship between the two as defined by standards
- C. CMMI provides a framework for improving processes, which directly impacts software quality
- D. CMMI completely replaces all software testing types within the given constraints

Answer: C

Q1080. What is the concept of 'zero defect' philosophy?

- A. Deliberately ignoring all reported defect reports and related components
- B. Having absolutely no tests in the entire project by the project team members
- C. A quality approach emphasizing prevention and doing things right the first time
- D. A goal that is impossible to ever achieve at all in all development efforts

Answer: C

Q1081. What is the Chidamber and Kemerer (CK) metrics suite?

- A. A collection of design patterns for modules and code
- B. A comprehensive software testing tool suite and set
- C. An established coding standard and guideline set only
- D. A set of six object-oriented metrics: WMC, DIT, NOC, CBO, RFC, LCOM

Answer: D

Q1082. What is the Response For a Class (RFC) metric?

- A. A network communication protocol metric and measure during the software lifecycle
- B. A relational database performance metric and measure within the project scope
- C. The count of all methods that can be invoked in response to a message to the class
- D. A comprehensive software testing metric and measure for the development team

Answer: C

Q1083. What is the Lack of Cohesion of Methods (LCOM) metric?

- A. A production deployment readiness metric and measure during implementation
- B. A metric measuring how dissimilar methods are in a class based on the instance variables they use
- C. A comprehensive software testing metric and measure in the system context
- D. A measure of total source code line length and size throughout the project

Answer: B

Q1084. What is the Coupling Between Objects (CBO) metric?

- A. A relational database join performance metric and score as a standard approach
- B. A network bandwidth and latency metric and measurement in practice typically
- C. A physical coupling measurement between parts and things for quality purposes
- D. The count of classes to which a class is coupled through method calls, variables, or inheritance

Answer: D

Q1085. What is the concept of 'measurement dysfunction'?

- A. Broken physical measurement tools and device parts by the organization
- B. Errors in hardware measurement instruments and tools across all phases
- C. When metrics are gamed or misused, leading to counterproductive behavior
- D. Dysfunctional software that crashes frequently daily at every stage

Answer: C

Q1086. What is GQM+ Strategies?

- A. A comprehensive software testing strategy and plan according to best practices
- B. An extension of GQM that aligns measurement goals with organizational business goals
- C. A recreational video gaming strategy guide and plan in a systematic way
- D. A source code implementation coding strategy and plan for the project goals

Answer: B

Q1087. What is the concept of 'metric validation' in software engineering?

- A. Validating application source code for errors and bugs by the development process
- B. Testing software requirements for completeness and scope and its related activities
- C. Ensuring that a metric actually measures what it claims to measure and is useful
- D. Testing physical measurement tools for their accuracy within the system boundary

Answer: C

Q1088. What is the Number of Children (NOC) metric?

- A. The total number of comments in the source code base
- B. The total number of source code files in the project
- C. The number of child processes currently running now
- D. The count of immediate subclasses of a class in the inheritance tree

Answer: D

Q1089. What is the difference between size-oriented and function-oriented metrics?

- A. Size-oriented uses LOC as the basis; function-oriented uses function points
- B. Function-oriented metrics use LOC as the main basis as defined by standards
- C. Size-oriented metrics use function points as basis for all project stakeholders
- D. The two metric types are fundamentally identical overall over the entire lifecycle

Answer: A

Q1090. What is the concept of a 'balanced scorecard' in software metrics?

- A. A scorecard for athletic sports competition events within the given constraints
- B. A formal source code peer review evaluation and form by the project team members
- C. A comprehensive software testing checklist form sheet in all development efforts
- D. A strategic management framework measuring performance across multiple perspectives

Answer: D

Q1091. What is the concept of 'configuration identification'?

- A. Identifying individual automated test case scenarios in a systematic way
- B. Selecting and labeling configuration items and documenting their characteristics
- C. Identifying specific software bugs and defect reports across all phases
- D. Identifying individual team member profile records at every stage

Answer: B

Q1092. What is the difference between configuration management and change management?

- A. CM only involves application source code files and code by the development process
- B. Change management is always much broader in its scope within the system boundary
- C. CM manages all configuration items; change management specifically handles the change request process
- D. The two concepts are completely identical in all ways for the project goals according to best practices

Answer: C

Q1093. What is a Software Bill of Materials (SBOM)?

- A. A project budget and financial document and report over the entire lifecycle
- B. A grocery shopping list for the team members only and its related activities
- C. A team member roster and directory listing sheet as defined by standards
- D. A formal record of the components and dependencies in a piece of software

Answer: D

Q1094. What is Infrastructure as Code (IaC)?

- A. Building physical infrastructure roads and buildings for all project stakeholders
- B. A general-purpose coding language definition and spec within the given constraints
- C. A comprehensive software testing framework and tool set in all development efforts
- D. Managing and provisioning computing infrastructure through machine-readable configuration files

Answer: D

Q1095. What is the concept of 'immutable infrastructure'?

- A. An approach where infrastructure components are replaced rather than modified, ensuring consistency
- B. Permanent physical hardware installations and setups and related components
- C. Infrastructure that physically cannot ever change form by the project team members
- D. A specific type of database storage system and service to achieve project objectives

Answer: A

Q1096. What is semantic versioning?

- A. Alphabetical versioning of all software release names for effective project outcomes
- B. Using random numbering for all software version names in the engineering discipline
- C. A versioning scheme using MAJOR.MINOR.PATCH to convey meaning about underlying changes
- D. Date-based versioning of all software release versions as part of the methodology

Answer: C

Q1097. What is a monorepo strategy?

- A. Multiple separate repos per each single project type
- B. Storing multiple projects in a single version control repository
- C. One separate repository per individual project file
- D. Having absolutely no repository at all for the code

Answer: B

Q1098. What is the concept of 'configuration status accounting'?

- A. Recording and reporting the status of configuration items and change requests throughout the project
- B. Financial accounting and bookkeeping tasks and work in the development process
- C. Tracking developer working hours per task and activity within the project scope
- D. Counting total source code lines in the whole project during the software lifecycle

Answer: A

Q1099. What is a feature flag (feature toggle)?

- A. A physical flag or banner for signaling and marking for the development team
- B. A test indicator for pass or fail status result only in the system context
- C. A technique allowing features to be enabled or disabled without deploying new code
- D. A marker for identifying software bugs and issues only throughout the project

Answer: C

Q1100. What is the concept of 'GitOps'?

- A. A comprehensive software testing approach and method only in practice typically as a standard approach
- B. An operational framework where Git is the single source of truth for infrastructure and application deployment
- C. An established coding standard and guideline set for teams by the organization at every stage
- D. Only performing basic Git operations and nothing else during implementation for quality purposes

Answer: B

Q1101. What is the Expected Monetary Value (EMV) in risk analysis?

- A. An employee salary expectation and forecast amount in practice typically
- B. A project budget and financial metric and measurement as a standard approach
- C. A product pricing strategy and approach and methodology by the organization
- D. The product of probability and monetary impact of a risk, used in decision trees

Answer: D

Q1102. What is the concept of 'risk-driven development'?

- A. Deliberately ignoring risks in the development process in a systematic way for the project goals
- B. A comprehensive software testing approach and method only according to best practices
- C. Developing intentionally risky software applications at every stage across all phases
- D. A development approach that uses risk analysis to prioritize and guide architecture and design decisions

Answer: D

Q1103. What is a decision tree in risk management?

- A. A test hierarchy and organization structure and setup and its related activities
- B. A diagram showing decision choices, chance events, probabilities, and outcomes for risk analysis
- C. A physical tree growing in a garden area outdoors within the system boundary
- D. A source code tree and directory structure and layout by the development process

Answer: B

Q1104. What is sensitivity analysis in risk management?

- A. An emotional sensitivity analysis technique approach over the entire lifecycle
- B. A source code analysis technique and tool and method for all project stakeholders
- C. A security vulnerability analysis technique and method as defined by standards
- D. A technique that determines which risks have the greatest potential impact on project outcomes

Answer: D

Q1105. What is the concept of 'risk appetite' vs 'risk tolerance'?

- A. The two concepts are completely identical in every way within the given constraints in all development efforts
- B. Risk appetite is the level of risk an organization is willing to pursue; risk tolerance is the acceptable variation in outcomes
- C. Risk appetite is always smaller than risk tolerance level by the project team members and related components
- D. Risk tolerance always means absolutely zero risk accepted to achieve project objectives in the engineering discipline

Answer: B

Q1106. What is Monte Carlo simulation in risk management?

- A. A technique using random sampling to simulate project outcomes and assess risk probability
- B. A comprehensive software testing simulation and approach in the development process
- C. A source code implementation coding simulation and test as part of the methodology
- D. A recreational casino gambling game activity and event for effective project outcomes

Answer: A

Q1107. What is residual risk?

- A. The risk that remains after risk responses have been implemented
- B. A production deployment issue and concern and problem
- C. No risk at all remaining after response is applied
- D. A specific type of source code error found in the code

Answer: A

Q1108. What is secondary risk?

- A. A risk that is considered less important than others during the software lifecycle
- B. A specific type of reported software bug and defect for the development team
- C. A new risk that arises as a direct result of implementing a risk response
- D. A backup risk for contingency planning purposes only within the project scope

Answer: C

Q1109. What is the concept of 'risk-based testing'?

- A. A testing approach that uses risk analysis to prioritize test cases and allocate testing effort
- B. Testing software without a formal plan or schedule during implementation
- C. Random and unstructured software testing approaches in the system context
- D. Testing software with risky input data and scenarios throughout the project

Answer: A

Q1110. What is the Failure Mode and Effects Analysis (FMEA) in software?

- A. A failure analysis report document summary and overview for quality purposes in practice typically
- B. A systematic technique for identifying potential failure modes, their causes, and effects on the system
- C. A production deployment analysis and report and summary by the organization at every stage
- D. A source code implementation coding technique and method as a standard approach

Answer: B

Q1111. What is threat modeling?

- A. Creating detailed threat assessment reports and summaries for the development team
- B. A systematic approach to identifying, quantifying, and addressing security threats to a system
- C. Modeling threats in three-dimensional space and models within the project scope
- D. A specific type of software testing approach and method throughout the project

Answer: B

Q1112. What is the STRIDE threat model?

- A. A source code implementation coding framework and library for quality purposes in practice typically
- B. A comprehensive software testing model and approach only as a standard approach by the organization
- C. A threat classification model: Spoofing, Tampering, Repudiation, Information Disclosure, DoS, Elevation of Privilege
- D. A physical walking gait analysis technique and method in the system context during implementation

Answer: C

Q1113. What is the concept of 'security by design'?

- A. Integrating security considerations into every phase of the software development lifecycle
- B. Designing security logos and brand imagery and assets at every stage
- C. Adding security features only after deployment to production across all phases
- D. Writing security documentation exclusively and nothing else in a systematic way

Answer: A

Q1114. What is a zero-day vulnerability?

- A. A vulnerability that is only one day old at most today within the system boundary
- B. A previously unknown vulnerability that is exploited before a patch is available
- C. A vulnerability with absolutely zero impact on the system according to best practices
- D. A vulnerability that was fixed on day zero of discovery for the project goals

Answer: B

Q1115. What is the DREAD risk assessment model?

- A. A comprehensive software testing model and approach only and its related activities
- B. An established coding standard and guideline and rule set over the entire lifecycle
- C. An assessment of fear and dread feelings and emotions by the development process
- D. A model rating threats by Damage, Reproducibility, Exploitability, Affected users, Discoverability

Answer: D

Q1116. What is secure software development lifecycle (SSDLC)?

- A. A comprehensive software testing lifecycle and approach for all project stakeholders
- B. A production deployment lifecycle plan and schedule doc within the given constraints
- C. An approach that integrates security practices into each phase of the traditional SDLC
- D. A brand new programming language created from scratch as defined by standards

Answer: C

Q1117. What is the difference between symmetric and asymmetric encryption?

- A. Asymmetric encryption uses only one single key for all by the project team members
- B. The two encryption types are fundamentally the same thing in all development efforts
- C. Symmetric encryption uses two different keys always now and related components
- D. Symmetric uses the same key for encryption and decryption; asymmetric uses a key pair

Answer: D

Q1118. What is a Common Vulnerabilities and Exposures (CVE)?

- A. A standardized system for identifying and naming publicly known cybersecurity vulnerabilities
- B. A comprehensive software testing framework and tool suite for effective project outcomes
- C. An established coding standard and guideline and rule set in the engineering discipline
- D. A specific type of computer virus or malware and threat to achieve project objectives

Answer: A

Q1119. What is the concept of 'attack surface'?

- A. The physical surface of a computer case and housing as part of the methodology
- B. A network configuration and topology setup and structure during the software lifecycle
- C. A specific type of user interface design and layout only in the development process
- D. The total number of points where an unauthorized user can try to enter or extract data

Answer: D

Q1120. What is the principle of 'fail secure' vs 'fail open'?

- A. The two concepts are fundamentally identical in all ways within the project scope
- B. Fail secure allows all access on any failure or error event throughout the project
- C. Fail open is always the more secure approach and method for the development team
- D. Fail secure denies access by default on failure; fail open allows access on failure

Answer: D

Q1121. A company is experiencing frequent project delays and cost overruns. Which fundamental software engineering principle would most directly address this?

- A. Hiring additional developers without changing methodology
- B. Reducing documentation to accelerate coding activities
- C. Adopting structured process models with clear milestones
- D. Purchasing more powerful servers for the development team

Answer: C

Q1122. In Lehman's laws of software evolution, what does the law of continuing change state?

- A. Software quality improves automatically without any effort
- B. Software maintenance effort reduces with each new release
- C. Software complexity decreases naturally over time periods
- D. Software must continually adapt or become less satisfactory

Answer: D

Q1123. A team finds that fixing one bug consistently introduces two new defects. Which concept best explains this phenomenon?

- A. Software entropy and increasing system complexity over time
- B. Effective regression testing preventing all defect propagation
- C. Proper documentation reducing overall maintenance difficulty
- D. Successful modular design isolating component dependencies

Answer: A

Q1124. When applying the 'separation of concerns' principle, what is the primary architectural benefit achieved?

- A. Documentation requirements are completely eliminated from work
- B. Each module addresses a distinct feature independently of others
- C. Testing becomes unnecessary due to simplified code structure
- D. All system functionality is consolidated into a single module

Answer: B

Q1125. A startup must deliver a product quickly with uncertain requirements. Which combination of software engineering principles is most appropriate?

- A. Comprehensive upfront documentation with rigid phase milestones
- B. Iterative development with frequent stakeholder feedback cycles
- C. Big bang integration testing after all modules are completed
- D. Waterfall model with complete requirements before any coding

Answer: B

Q1126. What is Brooks' law and how does it impact late software projects?

- A. Late projects benefit most from adding experienced engineers
- B. Doubling the team size will exactly halve the delivery time
- C. Adding people to a late project makes it even later overall
- D. Project delays are always caused by insufficient team members

Answer: C

Q1127. In the context of sociotechnical systems, why must software engineers consider organizational factors?

- A. Organizational factors have no measurable impact on software quality
- B. Technical solutions must align with organizational processes and culture
- C. Software operates independently of any human or organizational context
- D. Engineering decisions should be based solely on technical merit alone

Answer: B

Q1128. A legacy system handles critical operations but uses obsolete technology. What is the recommended engineering approach?

- A. Ignoring the system until it fails completely in production
- B. Complete system rewrite without referencing existing functionality
- C. Immediate shutdown followed by building from scratch entirely
- D. Incremental modernization while maintaining system availability

Answer: D

Q1129. What distinguishes emergent system properties from designed properties in software engineering?

- A. Emergent properties arise from component interactions not individual parts
- B. Emergent properties can be fully predicted from individual components
- C. Designed properties only appear after system deployment to production
- D. Emergent properties are explicitly defined in requirements documentation

Answer: A

Q1130. How does Conway's law relate software architecture to organizational structure?

- A. Technical architecture solely determines how teams are organized
- B. Software architecture is always independent of team organization
- C. Organizational hierarchy has no influence on system module design
- D. System designs mirror the communication structure of the organization

Answer: D

Q1131. A government agency requires extensive documentation and formal sign-offs at each development stage. Which process model best fits this environment?

- A. Waterfall model with formal reviews and documentation at each phase
- B. Extreme programming with minimal documentation and rapid iterations
- C. Scrum framework with two-week sprints and daily standup meetings
- D. Kanban method with continuous flow and no predefined phase gates

Answer: A

Q1132. A team is building a safety-critical system for an aircraft. Why would the V-model be preferred over agile approaches?

- A. V-model ensures rigorous verification and validation at each level
- B. V-model eliminates the need for any form of requirements analysis
- C. Agile approaches always produce higher quality safety-critical code
- D. Agile methods provide more thorough documentation for regulators

Answer: A

Q1133. In the spiral model, how does the team determine the scope of each iteration?

- A. By implementing all remaining features in every single spiral loop
- B. By evaluating identified risks and prioritizing their mitigation efforts
- C. By following a fixed predetermined schedule regardless of any risks
- D. By randomly selecting features from the product backlog for building

Answer: B

Q1134. A project started with the waterfall model but requirements keep changing. What is the most pragmatic transition strategy?

- A. Shift to an incremental model preserving completed phase artifacts
- B. Continue with waterfall and reject all incoming change requests
- C. Abandon process models entirely and use ad hoc development now
- D. Restart the entire project from scratch using a different language

Answer: A

Q1135. What is the primary challenge of using evolutionary prototyping for large-scale enterprise systems?

- A. Enterprise systems cannot use any iterative development technique
- B. Large systems never benefit from any prototyping approach at all
- C. Prototypes may evolve into poorly architected production systems
- D. Evolutionary prototyping always produces perfectly structured code

Answer: C

Q1136. How does the Rational Unified Process (RUP) handle changing requirements compared to traditional plan-driven approaches?

- A. RUP accommodates changes through iterative development within structured phases
- B. RUP completely prevents any requirement changes after the inception phase ends
- C. RUP requires a complete project restart whenever any requirement changes occur
- D. RUP has no mechanism for handling requirements changes during development work

Answer: A

Q1137. A medical device company must comply with FDA validation requirements. How does this constraint affect process model selection?

- A. It favors models with traceable verification and validation documentation
- B. It mandates that no testing be performed until final system delivery
- C. It eliminates all possible process models from consideration entirely
- D. It requires using only agile methods without any formal documentation

Answer: A

Q1138. What is the fundamental trade-off between plan-driven and agile process models in terms of project governance?

- A. Plan-driven offers more control and predictability; agile offers more flexibility and speed
- B. Plan-driven and agile models provide identical levels of governance and predictability
- C. Plan-driven models are always inferior to agile in every project context and scenario
- D. Agile models provide more formal governance and control than plan-driven approaches do

Answer: A

Q1139. When applying the DevOps model, how does continuous integration change the traditional process model structure?

- A. It separates development and operations into completely independent workflows
- B. It requires all code to be written before any integration testing takes place
- C. It merges development and operations phases into a continuous automated pipeline
- D. It completely eliminates the need for any testing or quality assurance activities

Answer: C

Q1140. A startup pivots its product direction every three months based on market feedback. Which process model characteristic is most essential?

- A. Fixed scope defined entirely at project start without any later changes
- B. Long sequential phases with comprehensive documentation at every stage
- C. Short iteration cycles with frequent delivery and customer feedback loops
- D. Complete architecture designed before any implementation work begins

Answer: C

Q1141. A team's velocity has been declining over the past five sprints despite no personnel changes. What is the most likely root cause to investigate?

- A. The team has become too skilled and experienced for the project work
- B. Increasing technical debt making new development progressively harder
- C. The product backlog has become too small with too few remaining items
- D. Sprint duration has been too short for any meaningful work to occur

Answer: B

Q1142. How should a Scrum team handle a situation where the Product Owner frequently changes priorities mid-sprint?

- A. Remove the Product Owner from the team to prevent further disruption to the process flow
- B. Educate the Product Owner on sprint commitment and defer changes to next sprint planning
- C. Immediately abandon current sprint work and start on the new priorities without question
- D. Ignore all priority changes and continue working on the original plan without communication

Answer: B

Q1143. In SAFe (Scaled Agile Framework), what is the purpose of a Program Increment planning event?

- A. To eliminate the need for individual team sprint planning sessions throughout the increment
- B. To replace all Scrum ceremonies with a single comprehensive monthly planning meeting event
- C. To align multiple agile teams on shared objectives and dependencies for a planning period
- D. To assign specific tasks to individual developers across all teams in the entire organization

Answer: C

Q1144. A distributed agile team across three time zones is struggling with communication. Which practice would most effectively address this challenge?

- A. Requiring all team members to work the same hours regardless of their local time zone
- B. Establishing overlapping core hours and using asynchronous communication tools effectively
- C. Eliminating all meetings and relying exclusively on email for all team communications
- D. Splitting the team into independent groups that never need to communicate with each other

Answer: B

Q1145. What is the relationship between the Agile Manifesto values and the twelve agile principles?

- A. The values replace the principles and make them unnecessary for any agile team practice
- B. The values and principles are completely unrelated documents with different authors entirely
- C. The principles contradict the values and should be followed instead of the manifesto values
- D. The principles provide specific guidance for applying the four abstract manifesto values

Answer: D

Q1146. When transitioning from waterfall to agile, what is the most significant organizational challenge?

- A. Shifting from command-and-control management to servant leadership and team autonomy
- B. Purchasing new software development tools and IDE licenses for all team members now
- C. Reducing the number of meetings from weekly status updates to no meetings at all
- D. Renaming existing project roles to match agile terminology without changing practices

Answer: A

Q1147. How does technical debt impact agile teams differently than plan-driven teams?

- A. Agile teams are immune to technical debt because they refactor code in every single sprint
- B. Technical debt affects plan-driven teams more because they have longer development cycles
- C. Agile teams feel the impact sooner due to shorter iteration cycles and continuous delivery
- D. Plan-driven teams never accumulate technical debt due to their thorough upfront planning

Answer: C

Q1148. A team is practicing Scrum but the sprint goal is never met. The team completes individual stories but lacks cohesion. What is the underlying issue?

- A. The sprint goal is not well-defined or the team is not collaborating toward a shared objective
- B. The Product Owner is providing too few user stories for the team to work on each sprint
- C. The sprint duration is too long and the team runs out of meaningful work before it ends
- D. Individual team members are too skilled and complete work too quickly without coordination

Answer: A

Q1149. In Lean software development, what does 'eliminating waste' specifically refer to?

- A. Removing activities that do not add value to the customer from the development process
- B. Reducing the physical waste produced by the office including paper and plastic materials
- C. Firing underperforming team members to improve the overall productivity of the team now
- D. Deleting all source code comments and documentation to reduce the size of the codebase

Answer: A

Q1150. What is the difference between adaptive and predictive project management approaches in the context of agile?

- A. Both approaches are identical in how they handle project scope and requirement changes fully
- B. Adaptive approaches fix all requirements at start; predictive approaches change constantly now
- C. Adaptive approaches embrace change and evolve plans; predictive approaches fix scope upfront
- D. Predictive approaches are more responsive to change than adaptive approaches in all cases

Answer: C

Q1151. A healthcare system must comply with HIPAA regulations. How should these regulatory constraints be captured in requirements?

- A. As project management constraints affecting only the delivery timeline and budget
- B. As functional requirements describing the user interface layout and color scheme
- C. As non-functional requirements specifying security, privacy, and audit trail standards
- D. As optional features that can be implemented if time and budget allow during work

Answer: C

Q1152. During elicitation, stakeholders express contradictory requirements. What is the most effective resolution approach?

- A. Implement all contradictory requirements simultaneously without resolving the conflicts
- B. Facilitate negotiation sessions to find compromises and document agreed-upon priorities
- C. Ignore the less senior stakeholder's requirements and implement only the manager's view
- D. Postpone the entire project until all stakeholders naturally reach unanimous agreement

Answer: B

Q1153. What is the relationship between requirements volatility and project risk in software engineering?

- A. Requirements volatility has no measurable impact on project risk or delivery outcomes
- B. Requirements stability increases project risk because it prevents creative innovation work
- C. Volatile requirements always reduce project risk by keeping the team adaptable and ready
- D. Higher requirements volatility increases project risk through scope creep and rework costs

Answer: D

Q1154. How should safety-critical requirements be specified differently from standard functional requirements?

- A. Using informal natural language descriptions without any additional precision or detail
- B. With formal notation, rigorous verification criteria, and detailed failure mode analysis
- C. Exactly the same way as any other functional requirement without special treatment given
- D. By omitting them from documentation and relying on developer intuition during coding

Answer: B

Q1155. A team discovers during implementation that a key requirement is technically infeasible. What is the appropriate requirements engineering response?

- A. Conduct impact analysis, negotiate alternatives with stakeholders, and update the SRS document
- B. Cancel the entire project because a single requirement cannot be implemented as specified
- C. Continue attempting to implement the infeasible requirement regardless of technical limitations
- D. Silently remove the requirement from the SRS without informing any stakeholders about it

Answer: A

Q1156. What is the 'requirements gap' and how does it affect software projects?

- A. It is the physical distance between the development team and the client office locations
- B. It is the difference between documented requirements and actual stakeholder needs and expectations
- C. It is the time delay between writing requirements and implementing them in source code only
- D. It is the salary difference between requirements analysts and software development engineers

Answer: B

Q1157. How do goal-oriented requirements engineering approaches differ from traditional approaches?

- A. They focus exclusively on technical implementation details rather than business-level objectives
- B. They eliminate the need for any stakeholder communication during the requirements phase entirely
- C. They focus on stakeholder goals and derive requirements from high-level objectives systematically
- D. They produce identical results to traditional approaches but use different terminology and tools

Answer: C

Q1158. In a large enterprise system with hundreds of requirements, how should traceability be managed effectively?

- A. By assigning one developer to memorize all requirement-to-code mappings without any tools
- B. By avoiding traceability entirely since it adds overhead without providing any project value
- C. By manually maintaining paper documents that cross-reference every requirement to code lines
- D. Using automated traceability tools with bidirectional links between requirements and artifacts

Answer: D

Q1159. What is the impact of cultural differences on requirements elicitation in global software development projects?

- A. Cultural differences affect communication styles, expectations, and interpretation of requirements
- B. Cultural differences have no impact on requirements because software is a universal language
- C. All cultures have identical approaches to expressing needs and defining system requirements
- D. Cultural factors only affect project management and have no relevance to requirements work

Answer: A

Q1160. How does viewpoint-oriented requirements engineering handle the challenge of multiple perspectives?

- A. It combines all viewpoints into one without distinguishing their sources or resolving conflicts
- B. It selects a single viewpoint and ignores all other perspectives throughout the entire project
- C. It systematically captures and reconciles requirements from different stakeholder viewpoints
- D. It eliminates the need for stakeholder involvement by relying solely on developer assumptions

Answer: C

Q1161. A project is showing a Cost Performance Index (CPI) of 0.8 and Schedule Performance Index (SPI) of 0.9. What does this indicate?

- A. The project is under budget by 20% and ahead of schedule by 10% of planned work
- B. The project is over budget by 20% and behind schedule by 10% of planned progress
- C. The project is exactly on budget and on schedule according to the baseline plan
- D. The project has been completed successfully with all deliverables meeting standards

Answer: B

Q1162. How should a project manager handle a situation where the estimated effort exceeds the available budget?

- A. Cancel the project immediately without exploring any alternative approaches or solutions
- B. Proceed with the full scope and hope the team can work faster than the estimates suggest
- C. Add more developers to reduce the schedule without adjusting the budget or scope at all
- D. Negotiate scope reduction with stakeholders while maintaining core functionality intact

Answer: D

Q1163. In Brooks' Mythical Man-Month, why does adding developers to a late project make it later?

- A. New developers instantly become productive and reduce the workload on existing team members
- B. Projects always finish on time regardless of team size changes during the development period
- C. Adding developers has no measurable effect on project schedule in any known situation today
- D. Communication overhead grows quadratically and new members need ramp-up training time

Answer: D

Q1164. A team uses PERT estimation with optimistic=2 weeks, most likely=4 weeks, and pessimistic=12 weeks. What is the expected duration?

- A. Five weeks based on the weighted average formula used in PERT estimation technique
- B. Two weeks because PERT always selects the most optimistic estimate for planning work
- C. Twelve weeks because PERT always uses the pessimistic estimate for safety margin only
- D. Four weeks because the most likely estimate is always the final expected duration value

Answer: A

Q1165. How does the theory of constraints apply to software project management?

- A. It requires constraining all team members to work at exactly the same speed and pace
- B. It identifies and focuses improvement efforts on the bottleneck limiting overall throughput
- C. It suggests adding constraints to every activity to improve project discipline overall
- D. It eliminates all project constraints by removing deadlines and budget limitations now

Answer: B

Q1166. A project manager discovers that two critical team members plan to leave mid-project. What proactive measures should be taken?

- A. Ignore the risk and assume the remaining team can absorb all additional workload easily
- B. Immediately cancel the project since replacing team members is impossible to accomplish
- C. Reduce product quality standards so the remaining team can deliver with less total effort
- D. Document knowledge, cross-train team members, and plan for replacement hiring immediately

Answer: D

Q1167. What is the relationship between project size and defect density in software engineering?

- A. Project size has absolutely no correlation with defect density in any published research data
- B. Defect density remains perfectly constant regardless of project size or complexity levels
- C. Smaller projects always have more defects than larger projects due to less thorough testing
- D. Larger projects tend to have higher defect density due to increased complexity and coupling

Answer: D

Q1168. How should a project manager balance the competing demands of multiple stakeholders with different priorities?

- A. Randomly select which stakeholder requirements to implement without any systematic approach
- B. Ignore all stakeholder input and make decisions based solely on the development team opinion
- C. Use stakeholder analysis to understand influence and interest, then negotiate compromises
- D. Always prioritize the requirements of the most senior stakeholder without any negotiation

Answer: C

Q1169. What is Monte Carlo simulation used for in project management and how does it improve estimation accuracy?

- A. It eliminates all uncertainty from project estimates by using a fixed mathematical formula only
- B. It provides a single deterministic estimate that is always accurate for every project situation
- C. It replaces the need for expert judgment by using random numbers without any project data
- D. It runs thousands of simulations with varying inputs to produce probability-based schedule forecasts

Answer: D

Q1170. In agile project management, how does the concept of 'servant leadership' change the traditional project manager role?

- A. The manager delegates all responsibility and has no involvement in the project activities at all
- B. The manager increases direct control over all individual tasks and decisions made by the team
- C. The manager role is completely eliminated and no leadership exists within agile project teams
- D. The manager focuses on removing impediments and empowering the team rather than directing work

Answer: D

Q1171. A system needs to support multiple database backends (MySQL, PostgreSQL, MongoDB) without changing business logic. Which design approach is most appropriate?

- A. Hardcoding database queries directly into business logic classes for maximum performance
- B. Abstract Factory pattern with database-specific factory implementations for each backend
- C. Using global variables to store database connection strings throughout the application
- D. Writing separate copies of the business logic for each database backend independently

Answer: B

Q1172. How does the SOLID principles violation manifest in a God class anti-pattern?

- A. The class violates SRP by handling too many responsibilities and having many reasons to change
- B. The class follows SRP perfectly by focusing on a single well-defined responsibility only
- C. The class demonstrates ideal OCP by being completely closed to any future extensions work
- D. The class exemplifies LSP by being perfectly substitutable for all of its subtype classes

Answer: A

Q1173. A real-time trading system must process millions of events per second. How should the design balance performance with maintainability?

- A. Sacrifice all code readability and maintainability to achieve maximum raw processing speed
- B. Implement all processing in a single monolithic function for maximum execution throughput
- C. Use the simplest possible design without considering performance requirements at any point
- D. Use event-driven architecture with optimized hot paths while keeping non-critical paths clean

Answer: D

Q1174. What is the relationship between design patterns and anti-patterns in software engineering?

- A. Anti-patterns are always better than design patterns for solving complex engineering problems
- B. Design patterns cause problems while anti-patterns provide effective solutions for all cases
- C. Design patterns solve recurring problems; anti-patterns are commonly used but ineffective solutions
- D. Design patterns and anti-patterns are identical concepts with interchangeable names and usage

Answer: C

Q1175. How does the Hexagonal Architecture (Ports and Adapters) improve testability of business logic?

- A. It isolates business logic from external dependencies through ports allowing easy mock injection
- B. It embeds external dependencies directly into business logic making testing straightforward now
- C. It makes testing harder by creating multiple layers of indirection between all components now
- D. It eliminates the need for any testing by ensuring the business logic is always correct already

Answer: A

Q1176. A development team is debating between using inheritance hierarchies and composition for code reuse. What factors should guide this decision?

- A. Always use composition because inheritance should never be used in any software system today
- B. Use inheritance for genuine is-a relationships; prefer composition for flexible behavior reuse
- C. Always use inheritance because it is fundamentally superior to composition in every scenario
- D. The choice is purely aesthetic and has no impact on system flexibility or maintainability

Answer: B

Q1177. How does the Command pattern enable undo functionality in software applications?

- A. The Command pattern only supports redo operations and cannot implement undo functionality now
- B. Undo is implemented by restarting the entire application from its initial state each time needed
- C. The Command pattern prevents any actions from being undone once they have been executed fully
- D. Each command object encapsulates an action and its inverse allowing execution reversal on demand

Answer: D

Q1178. What is the impact of circular dependencies on software design quality and maintainability?

- A. They simplify the build process by reducing the number of modules that need compilation work
- B. They have no measurable impact on software design quality or long-term maintainability at all
- C. They increase coupling, complicate builds, and make modules impossible to test in isolation easily
- D. They improve design quality by ensuring all modules are tightly integrated with each other now

Answer: C

Q1179. In domain-driven design, how does the Bounded Context pattern help manage complexity in large systems?

- A. It requires a single unified domain model used identically across all parts of the entire system
- B. It eliminates the need for any domain modeling by using generic data structures everywhere now
- C. It defines clear boundaries where specific domain models apply preventing model conflicts globally
- D. It increases complexity by requiring every team to understand all domain models in the system

Answer: C

Q1180. How should a designer handle the tension between the DRY principle and the need for module independence?

- A. Accept controlled duplication when shared abstractions would create inappropriate coupling between modules
- B. Always eliminate every instance of duplication regardless of the coupling it introduces between parts
- C. Force all modules to share a single implementation even when their requirements diverge significantly
- D. Ignore the DRY principle entirely because code duplication has no negative consequences on quality

Answer: A

Q1181. A company needs to migrate a monolithic e-commerce platform to microservices. What is the recommended architectural migration strategy?

- A. Identify bounded contexts, extract services incrementally, and use an API gateway for routing
- B. Rewrite the entire system from scratch in a new language without referencing existing code
- C. Keep the monolith exactly as is and simply rename it as a microservices architecture now
- D. Split the monolith randomly into services without analyzing domain boundaries or dependencies

Answer: A

Q1182. How does the CAP theorem constrain the design of distributed database architectures?

- A. The CAP theorem only applies to relational databases and not to any NoSQL database systems
- B. A distributed system can guarantee at most two of consistency, availability, and partition tolerance
- C. Partition tolerance is optional in distributed systems and can be safely ignored in all designs
- D. A distributed system can easily achieve all three properties simultaneously without any trade-offs

Answer: B

Q1183. A streaming analytics platform must process ten million events per second with sub-second latency. Which architectural pattern is most suitable?

- A. Single-threaded monolithic application processing events sequentially one at a time in order
- B. Batch processing architecture that collects events for hourly processing in large bulk operations
- C. Traditional request-response architecture with synchronous database queries for each single event
- D. Event sourcing with partitioned stream processing and in-memory computation for hot data paths

Answer: D

Q1184. What are the architectural implications of choosing eventual consistency over strong consistency?

- A. Eventual consistency has no implications and behaves identically to strong consistency in all cases
- B. Strong consistency always provides better performance than eventual consistency in every scenario
- C. Eventual consistency eliminates all data inconsistency problems permanently without any trade-offs
- D. The system gains availability and performance but must handle temporary data inconsistencies gracefully

Answer: D

Q1185. How should an architect design a system to be resilient against cascading failures in a microservices environment?

- A. Remove all error handling to reduce code complexity and improve overall system performance metrics
- B. Implement circuit breakers, bulkheads, and graceful degradation patterns between all service boundaries
- C. Use a single shared database for all services to eliminate any communication failures between them
- D. Make every service synchronously dependent on every other service to ensure data consistency always

Answer: B

Q1186. What is the architectural significance of Conway's law when designing large-scale distributed systems?

- A. Conway's law has been disproven and has no relevance to modern software architecture decisions
- B. System architecture should be designed without any consideration of team organization structure
- C. Team structure should align with desired system architecture to avoid communication mismatches
- D. Conway's law only applies to monolithic systems and is irrelevant for distributed architectures

Answer: C

Q1187. A healthcare system requires 99.999% availability. How does this requirement influence the architectural decisions?

- A. It requires a single centralized server to minimize complexity and potential points of failure in system
- B. It mandates redundancy, automatic failover, geographic distribution, and elimination of single points of failure
- C. It means the system should avoid any redundancy to reduce the number of components that can fail
- D. It has no impact on architecture because availability is purely an operational and infrastructure concern

Answer: B

Q1188. How does the saga pattern handle distributed transactions across microservices?

- A. It uses traditional two-phase commit protocol locking all services until the transaction completes fully
- B. It coordinates a sequence of local transactions with compensating actions for failure rollback scenarios
- C. It requires all microservices to share a single database to ensure transactional consistency always
- D. It avoids any form of transaction management and accepts data corruption as an acceptable outcome

Answer: B

Q1189. What architectural considerations are necessary when designing a multi-tenant SaaS application?

- A. Multi-tenant applications should use a separate deployed instance for every individual tenant always
- B. All tenants must share identical configurations with no customization options available for any tenant
- C. Data isolation strategies, per-tenant customization, scalability, and fair resource allocation mechanisms
- D. Multi-tenancy requires no special architectural considerations beyond a standard single-user application

Answer: C

Q1190. How does the sidecar pattern enhance service mesh architectures in cloud-native applications?

- A. It deploys supporting functionality alongside each service without modifying the service code itself
- B. It embeds all infrastructure concerns directly into business logic code of every service instance
- C. It requires every service to implement its own service discovery and load balancing logic fully
- D. It eliminates the need for any networking between services by running everything in one process

Answer: A

Q1191. A healthcare application is used by elderly patients with limited technical experience. Which UI design principles should be prioritized?

- A. Large touch targets, high contrast, simple navigation, clear typography, and minimal cognitive load
- B. Advanced keyboard shortcuts, minimal visual feedback, and hidden navigation behind icon-only menus
- C. Small fonts, complex multi-level menus, gesture-only navigation, and low contrast color schemes
- D. Dense information displays, small buttons, animation-heavy interfaces, and technical terminology use

Answer: A

Q1192. How does cognitive load theory influence the number of options presented in a UI at one time?

- A. The number of options has no measurable impact on user performance or satisfaction in any context
- B. Cognitive load theory suggests options should never be limited because users can process infinite items
- C. Too many options increase cognitive load and decision time, following Hick's law principles closely
- D. Users always prefer having maximum options displayed simultaneously regardless of task complexity

Answer: C

Q1193. A team is designing a dashboard for data analysts who monitor real-time metrics. What design pattern best supports their workflow?

- A. A slideshow that automatically rotates through metrics without any user control or interaction
- B. Full-screen modal dialogs for each individual metric requiring sequential manual navigation
- C. A single fixed text-based report with no interactive elements or visualization capabilities
- D. Overview plus detail pattern with drill-down capabilities and customizable widget arrangements

Answer: D

Q1194. How should error messages be designed to minimize user frustration and support error recovery?

- A. Cryptic error codes without any explanation or guidance for users on how to resolve the problem
- B. Blaming the user for causing the error and demanding they contact technical support immediately
- C. Clear language explaining what went wrong, why it happened, and specific steps to fix the issue
- D. Hiding all errors silently without any notification so the user does not experience any frustration

Answer: C

Q1195. What is the impact of animation and micro-interactions on perceived application performance?

- A. All animations slow down perceived performance and should be completely eliminated from every app
- B. Animations should only be used for decorative purposes and never for functional feedback to users
- C. Animations have no measurable impact on how users perceive the speed of the application interface
- D. Well-designed animations mask loading times and provide continuity making the app feel faster overall

Answer: D

Q1196. A multi-language application serves users across twenty countries. What UI design challenges must be addressed?

- A. Text expansion and contraction, RTL layout support, cultural color meanings, and locale-specific formats
- B. Internationalization only requires translating text strings and has no impact on layout or design work
- C. All languages and cultures use identical UI conventions so no design adaptations are ever necessary
- D. Multi-language support only affects backend code and has no visible impact on the user interface

Answer: A

Q1197. How does the concept of 'dark patterns' in UI design undermine user trust?

- A. Dark patterns refer exclusively to using dark color themes and have no ethical implications at all
- B. Dark patterns manipulate users into unintended actions through deceptive design techniques intentionally
- C. Dark patterns are always beneficial because they guide users toward the best possible decisions
- D. Dark patterns have no impact on user trust because most users cannot detect manipulative design

Answer: B

Q1198. A financial trading application must display rapidly changing data without overwhelming the user. How should the UI handle this?

- A. Use identical visual treatment for all data regardless of importance or rate of change observed
- B. Use visual prioritization, controlled update frequencies, and highlight only significant data changes
- C. Remove all real-time data and only show daily summary reports to reduce information overload
- D. Display every data change immediately with full-screen alerts for each individual price movement

Answer: B

Q1199. What is the relationship between information architecture and UI design in creating effective digital products?

- A. Information architecture defines content structure and organization that UI design then visualizes
- B. UI design determines information architecture since visual layout always precedes content structure
- C. Information architecture only applies to physical spaces and has no relevance to digital products
- D. Information architecture and UI design are completely unrelated disciplines with no overlap at all

Answer: A

Q1200. How should a design team approach creating an inclusive design that works for users with motor, visual, cognitive, and auditory disabilities?

- A. Design only for the average user and provide a separate simplified version for disabled users later
- B. Assistive technologies handle all accessibility needs so designers do not need to consider them now
- C. Follow WCAG guidelines, test with assistive technologies, and involve users with disabilities in testing
- D. Accessibility is a legal requirement only and does not require any changes to the actual design work

Answer: C

Q1201. A team is experiencing frequent merge conflicts in a large codebase. What implementation practices would most effectively reduce this problem?

- A. Larger infrequent commits combining many changes to reduce the total number of merge operations
- B. Smaller frequent commits, feature branching strategies, and modular code with clear ownership areas
- C. Having all developers work on the same branch and file simultaneously without any coordination
- D. Avoiding version control entirely and using shared network drives for all source code storage

Answer: B

Q1202. How does test-driven development (TDD) change the implementation workflow compared to traditional development?

- A. Tests are completely eliminated from the development process to accelerate feature delivery speed
- B. TDD and traditional development follow identical workflows with no differences in approach at all
- C. Implementation is completed fully before any tests are considered or written for the codebase
- D. Tests are written before implementation code, driving design decisions and ensuring test coverage

Answer: D

Q1203. A legacy codebase has no automated tests. What is the safest approach to introduce testing?

- A. Delete all existing code and replace it with test stubs that will be implemented later on
- B. Skip testing entirely since the codebase has worked without tests for its entire lifetime so far
- C. Write characterization tests capturing current behavior before making any code changes to system
- D. Rewrite the entire codebase from scratch with tests rather than adding tests to existing code

Answer: C

Q1204. How does feature flag implementation enable continuous deployment while managing risk?

- A. Feature flags completely prevent any bugs from reaching the production environment automatically
- B. Feature flags are only useful for UI changes and cannot be applied to backend logic or services
- C. Flags allow deploying code to production with features disabled, enabling gradual controlled rollout
- D. Flags eliminate the need for any testing since features can be toggled off if problems occur now

Answer: C

Q1205. What is the impact of choosing a monorepo versus polyrepo strategy on implementation workflow?

- A. Monorepos are only suitable for small projects and cannot scale to large enterprise codebases
- B. Monorepos and polyrepos have identical workflow implications with no meaningful differences at all
- C. Monorepos simplify cross-project changes but require sophisticated build and CI tooling at scale
- D. Polyrepos are always superior because they eliminate all cross-project dependency management needs

Answer: C

Q1206. A performance-critical system component is a bottleneck. How should the team approach optimization?

- A. Apply every known optimization technique to all code regardless of whether it is a bottleneck
- B. Rewrite the entire codebase in assembly language without profiling or measuring performance first
- C. Profile to identify actual bottlenecks, optimize measured hot paths, and benchmark improvements
- D. Avoid any optimization since premature optimization is the root of all evil in every situation

Answer: C

Q1207. How does clean code philosophy address the trade-off between code readability and cleverness?

- A. It treats readability and cleverness as identical qualities that never conflict with each other
- B. It suggests that code readability is irrelevant since compilers handle all code interpretation
- C. It encourages maximally clever code because shorter code is always better regardless of clarity
- D. It prioritizes readability over cleverness because code is read far more often than it is written

Answer: D

Q1208. What are the implementation challenges of handling distributed system failures gracefully?

- A. All distributed failures can be prevented entirely through proper implementation techniques alone
- B. Distributed systems never experience failures because modern networks are perfectly reliable now
- C. Network partitions, partial failures, and timeout handling require careful idempotent retry design
- D. Failure handling is identical in distributed and single-machine systems with no added complexity

Answer: C

Q1209. How does immutable infrastructure change the deployment and implementation practices?

- A. Immutable infrastructure prevents any updates or patches from being applied to the system ever
- B. Immutable infrastructure is identical to traditional mutable server management in all aspects
- C. Servers are continuously patched and modified in place to apply updates and configuration changes
- D. Servers are never modified after creation; changes require building and deploying new instances

Answer: D

Q1210. What is the role of code complexity metrics like cyclomatic complexity in guiding implementation decisions?

- A. Complexity metrics should only be used for documentation purposes and never influence coding work
- B. Higher complexity always means better code because it handles more edge cases and scenarios now
- C. Cyclomatic complexity has no correlation with defect rates or maintenance difficulty in any study
- D. High complexity indicates code that is harder to test, maintain, and more likely to contain defects

Answer: D

Q1211. A team has 90% code coverage but still experiences production bugs. What does this reveal about their testing strategy?

- A. High coverage does not guarantee test quality; tests may lack meaningful assertions and edge cases
- B. Code coverage is meaningless and should be completely ignored as a testing quality metric now
- C. The remaining ten percent of uncovered code contains all bugs and coverage should reach exactly full
- D. Ninety percent coverage guarantees that ninety percent of all possible bugs have been detected

Answer: A

Q1212. How does property-based testing complement example-based testing?

- A. It generates hundreds of random inputs to test properties that should hold for all valid inputs
- B. It is identical to example-based testing and provides no additional value to the test suite now
- C. It only tests a single predefined example and is less thorough than example-based approaches
- D. It replaces example-based testing entirely and provides complete coverage of all possibilities

Answer: A

Q1213. A safety-critical aviation system requires DO-178C compliance. How does this affect the testing approach?

- A. It allows the team to choose any testing approach without following any specific testing standards
- B. It only requires basic unit testing without any structural coverage or formal verification work
- C. It mandates structural coverage analysis, requirements-based testing, and formal verification methods
- D. It eliminates the need for any testing because certification replaces all quality activities now

Answer: C

Q1214. What is the oracle problem in software testing and how does it affect test automation?

- A. The oracle problem has been completely solved and no longer affects any testing activities today
- B. The oracle problem only applies to manual testing and has no impact on automated testing work
- C. Difficulty in determining expected outputs for complex scenarios limits automated test effectiveness
- D. It refers to the challenge of selecting which database system to use for storing test data only

Answer: C

Q1215. How does chaos engineering differ from traditional testing approaches?

- A. It replaces all traditional testing and eliminates the need for unit and integration test suites
- B. It intentionally introduces failures in production to verify system resilience and recovery capability
- C. It follows traditional scripted test cases executed in a controlled staging environment only today
- D. It avoids any deliberate failure injection and relies on monitoring to detect issues passively

Answer: B

Q1216. A microservices architecture has complex service interactions. What testing strategy best addresses integration risks?

- A. Full end-to-end testing of every possible service interaction path for complete coverage goal
- B. Contract testing verifying service interfaces plus consumer-driven contracts between service pairs
- C. Skipping all testing and relying on production monitoring to detect integration failures only
- D. Only unit testing individual services without any integration or contract testing between them

Answer: B

Q1217. What is the relationship between testability and software design quality?

- A. Poorly designed code is always easier to test than well-designed code due to simpler structure
- B. Highly testable code typically exhibits good design qualities like loose coupling and high cohesion
- C. Testability has no relationship with design quality and they are completely independent concerns
- D. Testability should never influence design decisions because it compromises performance always

Answer: B

Q1218. How should a team prioritize test automation efforts when resources are limited?

- A. Automate high-value repetitive tests first: smoke tests, regression suites, and critical path scenarios
- B. Automate only the tests that are easiest to write regardless of their value to the project now
- C. Avoid all automation and rely entirely on manual testing regardless of the available resources
- D. Automate random tests without considering their value or frequency of execution by the team

Answer: A

Q1219. What is the impact of flaky tests on development team productivity and confidence?

- A. Flaky tests are beneficial because they simulate real-world conditions and improve test coverage
- B. Flaky tests improve team productivity by adding unpredictability that keeps developers alert now
- C. Flaky tests erode confidence in the test suite and cause teams to ignore legitimate test failures
- D. Flaky tests have no measurable impact on team productivity or confidence in the testing suite

Answer: C

Q1220. How does shift-left testing change the economics of defect detection in software projects?

- A. Testing earlier increases costs because more comprehensive testing infrastructure is needed upfront
- B. Finding defects earlier in development is significantly cheaper than finding them in production later
- C. Shift-left testing has no impact on the cost of defect detection at any stage of development
- D. Defects found in production are always cheaper to fix than those found during early development

Answer: B

Q1221. A legacy system with no documentation or original developers must be maintained. What is the recommended approach?

- A. Make changes randomly and deploy them immediately without understanding the system structure
- B. Use reverse engineering and program comprehension techniques to recover design knowledge first
- C. Refuse to maintain the system and recommend shutting it down immediately without replacement
- D. Rewrite the entire system from scratch without studying the existing system behavior at all

Answer: B

Q1222. How does Lehman's law of conservation of organizational stability affect maintenance planning?

- A. Adding more developers always proportionally increases the rate of maintenance activity and output
- B. The rate of development activity tends to remain constant over a system's lifetime regardless of resources
- C. Organizational stability has no measurable effect on maintenance planning or resource allocation
- D. Maintenance effort always decreases over time as the system becomes more stable and mature now

Answer: B

Q1223. A system experiences performance degradation every few months requiring manual restarts. What maintenance strategy addresses this systematically?

- A. Add more hardware resources each time degradation occurs without investigating the root cause
- B. Ignore the degradation until the system completely fails and then perform emergency maintenance
- C. Manually restart the system only when users complain about noticeable performance problems now
- D. Implement proactive software rejuvenation with automated scheduled restarts and resource cleanup

Answer: D

Q1224. What is the relationship between software architecture erosion and increasing maintenance costs?

- A. Architecture erosion increases coupling and decreases cohesion making changes progressively harder
- B. Architecture erosion has no impact on maintenance costs because architecture is only relevant initially
- C. Erosion decreases maintenance costs by simplifying the system structure through organic evolution now
- D. Architecture quality always improves over time through maintenance activities without deliberate effort

Answer: A

Q1225. How should a maintenance team prioritize between corrective, adaptive, perfective, and preventive maintenance requests?

- A. Preventive maintenance should always have lowest priority since it addresses only future problems
- B. Prioritize critical corrections first, then assess adaptive, perfective, and preventive by business value
- C. All maintenance types should be treated with equal priority without any systematic prioritization
- D. Always complete all corrective maintenance before considering any other type of maintenance request

Answer: B

Q1226. What is the economic justification for investing in preventive maintenance of software systems?

- A. Preventing defects is significantly cheaper than fixing them after they manifest in production use
- B. There is no economic basis for preventive maintenance since all software eventually needs replacing
- C. Preventive maintenance never provides any return on investment and should always be avoided now
- D. The cost of preventive maintenance always exceeds the cost of fixing defects after they occur

Answer: A

Q1227. How does microservices architecture change the maintenance approach compared to monolithic systems?

- A. Individual services can be maintained, deployed, and scaled independently reducing maintenance scope
- B. Microservices and monolithic systems require identical maintenance approaches without differences
- C. Microservices eliminate all maintenance needs because services are small enough to rewrite quickly
- D. Monolithic systems are always easier to maintain than microservices in every possible scenario now

Answer: A

Q1228. A company wants to decide whether to maintain or replace a legacy system. What factors should drive this decision?

- A. Business value, maintenance cost trajectory, technical risk, and availability of replacement options
- B. Legacy systems should always be maintained indefinitely regardless of costs or technical risks
- C. Only the age of the system should be considered since older systems should always be replaced now
- D. The decision should be based solely on the programming language used regardless of other factors

Answer: A

Q1229. What is the difference between software maintenance and software evolution from a theoretical perspective?

- A. Maintenance implies keeping software functional; evolution implies continuous adaptation and growth
- B. Maintenance and evolution are identical concepts with no theoretical distinction between them now
- C. Maintenance is a proactive process while evolution is purely reactive to environmental changes
- D. Evolution only occurs during initial development while maintenance happens exclusively after release

Answer: A

Q1230. How do maintenance metrics like Mean Time to Repair (MTTR) guide maintenance process improvement?

- A. MTTR measures repair speed, identifying bottlenecks in diagnosis, implementation, and deployment phases
- B. MTTR has no practical value for improving maintenance processes and should not be tracked at all
- C. Lower MTTR values always indicate worse maintenance performance and should be increased over time
- D. MTTR only measures the time to find bugs and does not include the time to fix and deploy them now

Answer: A

Q1231. How does the Deming cycle (PDCA) apply to continuous quality improvement in software?

- A. PDCA requires completing all four steps simultaneously without any sequential ordering of them
- B. The cycle should only be executed once and never repeated for continuous improvement purposes
- C. The Deming cycle only applies to manufacturing and has no relevance to software quality at all
- D. Plan improvements, Do implement them, Check results, Act to standardize or adjust the approach

Answer: D

Q1232. A company at CMMI Level 2 wants to reach Level 3. What fundamental change is required?

- A. Reducing the number of quality activities to simplify the overall process framework at the org
- B. Transitioning from project-level managed processes to organization-wide defined standard processes
- C. Maintaining the same project-level processes without any standardization across the organization
- D. Only purchasing more expensive quality assurance tools without changing any existing processes now

Answer: B

Q1233. What is the relationship between technical debt and software quality degradation over time?

- A. Technical debt accumulates quality compromises that progressively degrade system quality over time
- B. Accumulating technical debt always improves software quality by simplifying the system structure
- C. Technical debt only affects development speed and never impacts the quality of the software now
- D. Technical debt has no relationship with software quality and can be safely ignored by all teams

Answer: A

Q1234. How should an organization measure the return on investment of quality assurance activities?

- A. Calculate ROI based solely on the number of QA engineers hired regardless of their actual output
- B. Only measure the number of defects found without considering the cost of finding or preventing them
- C. ROI of QA cannot be measured because quality improvements are intangible and subjective in all cases
- D. Compare prevention and appraisal costs against reduction in failure costs and customer satisfaction gains

Answer: D

Q1235. A software product consistently fails in production despite passing all quality gates. What systemic issue might explain this?

- A. Quality gates always guarantee production quality and any failures must be caused by hardware now
- B. The quality gates are too strict and should be relaxed to allow more defects through to production
- C. Quality gates may not adequately represent real-world conditions and usage patterns of the system
- D. The development team is too skilled and the quality gates cannot keep up with their code changes

Answer: C

Q1236. How do organizational culture and management commitment affect quality assurance effectiveness?

- A. Management commitment has no impact on QA because quality is solely the responsibility of testers
- B. Strong management commitment and quality culture are essential for effective QA adoption and practice
- C. Excessive management involvement always reduces QA effectiveness by creating bureaucratic overhead
- D. Organizational culture is irrelevant to quality outcomes since only tools and processes matter now

Answer: B

Q1237. What is the Six Sigma approach and how does it apply to software quality management?

- A. It aims to reduce process variation to achieve fewer than 3.4 defects per million opportunities made
- B. Six Sigma only works for hardware manufacturing and has no proven application in software at all
- C. It requires eliminating all defects completely with zero tolerance for any imperfection in product
- D. Six Sigma increases process variation intentionally to improve creativity in software development

Answer: A

Q1238. How does model-based testing support quality assurance for complex software systems?

- A. It eliminates the need for any testing by mathematically proving software correctness from models
- B. It replaces quality assurance entirely by generating defect-free code directly from the models
- C. It only works for simple systems and cannot handle the complexity of real-world applications now
- D. It generates test cases from system models ensuring systematic coverage of modeled behaviors fully

Answer: D

Q1239. What is the impact of shifting quality responsibility to the entire development team rather than a separate QA group?

- A. It doubles the quality effort because both developers and QA testers do identical redundant work
- B. It has no impact on quality outcomes because only dedicated QA teams can effectively ensure quality
- C. It integrates quality thinking into all activities leading to earlier defect detection and prevention
- D. It eliminates quality entirely because developers do not have the skills to perform QA activities

Answer: C

Q1240. How should quality requirements be quantified and measured for a performance-critical real-time system?

- A. Quality requirements for real-time systems cannot be quantified and must remain purely qualitative
- B. Only measure response time and ignore all other quality attributes for real-time system evaluation
- C. Define measurable SLAs for latency, throughput, and availability with continuous monitoring infrastructure
- D. Use the same quality metrics as batch processing systems without any adaptation for real-time needs

Answer: C

Q1241. A team is using LOC as the primary productivity metric. What are the fundamental problems with this approach?

- A. LOC is technology-independent and measures delivered value perfectly for all project contexts
- B. LOC is the most accurate and comprehensive metric for measuring developer productivity always
- C. LOC penalizes efficient code, varies by language, and does not measure delivered business value
- D. There are no problems with using LOC as a productivity metric in any development environment

Answer: C

Q1242. How should metrics be used to avoid Goodhart's law in software development teams?

- A. Set aggressive metric targets and penalize teams that fail to meet them to ensure compliance
- B. Goodhart's law does not apply to software metrics and can be safely ignored by all managers
- C. Use metrics for insight rather than targets because optimizing for metrics distorts the behavior
- D. Use a single metric for all performance evaluations to ensure simplicity and consistency now

Answer: C

Q1243. A project shows decreasing defect density but increasing customer complaints. What does this discrepancy reveal?

- A. Testing may miss requirement-level defects that affect user experience but not code-level quality
- B. Customer complaints are irrelevant to software quality and should be ignored by the team always
- C. The metrics system is working perfectly and no investigation is needed for this discrepancy at all
- D. Decreasing defect density always means improved customer satisfaction without any exceptions now

Answer: A

Q1244. How does the Balanced Scorecard approach apply to software engineering measurement?

- A. It measures performance across financial, customer, process, and learning dimensions holistically
- B. It only measures financial metrics and ignores all other aspects of software project performance
- C. It measures a single dimension of performance and provides a narrow view of project health only
- D. It is exclusively a hardware manufacturing tool with no application to software engineering now

Answer: A

Q1245. What is the relationship between code churn metrics and defect prediction in software projects?

- A. Code churn has no correlation with defect rates and provides no predictive value for quality now
- B. Lower code churn always indicates more defects because unchanged code accumulates hidden bugs
- C. Code churn only measures developer activity and has no relationship to code quality whatsoever
- D. High code churn in specific files correlates with higher defect probability in those areas later

Answer: D

Q1246. How should an organization establish a metrics program without creating a dysfunctional measurement culture?

- A. Publicly rank individual developers by all metrics and reward the highest performers exclusively
- B. Use metrics exclusively for performance reviews and salary decisions without any other purpose
- C. Focus on team-level process improvement metrics rather than individual performance punishment data
- D. Collect as many metrics as possible regardless of their relevance to organizational goals now

Answer: C

Q1247. What is the role of benchmarking in software metrics and what are its limitations?

- A. Benchmarking provides perfectly comparable results across all organizations without any limitations
- B. It is impossible to benchmark software organizations because no two projects are ever similar now
- C. Benchmarking compares performance against peers but context differences limit direct comparisons
- D. Benchmarking only works for hardware measurements and has no application in software contexts

Answer: C

Q1248. How do object-oriented metrics like CK metrics suite differ from traditional complexity metrics?

- A. CK metrics and traditional metrics measure identical properties with identical calculation methods
- B. CK metrics only apply to procedural programming and cannot be used for object-oriented systems
- C. Traditional metrics are always superior to CK metrics for assessing object-oriented code quality
- D. CK metrics measure class-level properties like coupling, cohesion, inheritance, and method complexity

Answer: D

Q1249. A manager wants to compare productivity across teams using different technologies. What measurement approach is most fair?

- A. Compare teams solely on lines of code produced regardless of the programming languages they use
- B. Use technology-independent metrics like function points delivered per unit of effort for comparison
- C. Use identical metrics thresholds for all teams regardless of their technology stack differences
- D. Avoid any cross-team comparison because different technologies make all comparisons meaningless

Answer: B

Q1250. What is the statistical validity concern when using small sample sizes for software metrics?

- A. Software metrics do not follow statistical distributions so sample size is completely irrelevant
- B. Small sample sizes always produce more accurate statistics than large sample sizes in all cases
- C. Sample size has no impact on the statistical validity of software metrics or their conclusions
- D. Small samples produce unreliable statistics with wide confidence intervals limiting conclusion validity

Answer: D

Q1251. A large distributed team experiences frequent integration failures. How should their CM practices be improved?

- A. Eliminate automated testing to speed up the integration process and reduce pipeline execution time
- B. Have each developer maintain a separate repository that is never integrated with other team work
- C. Reduce commit frequency to monthly batches so integration issues can be resolved all at once now
- D. Implement CI with automated builds and tests, smaller commits, and trunk-based development workflow

Answer: D

Q1252. How does infrastructure as code (IaC) extend traditional configuration management principles?

- A. It replaces all software configuration management with manual server setup and configuration work
- B. It applies version control, reproducibility, and automation principles to infrastructure provisioning
- C. It eliminates the need for version control by automating all infrastructure changes automatically
- D. IaC has no relationship with configuration management principles or practices in any context now

Answer: B

Q1253. What are the trade-offs between feature branching and trunk-based development strategies?

- A. Trunk-based development cannot work with more than three developers on any given project team
- B. Feature branching is always superior to trunk-based development in every project context scenario
- C. Feature branches isolate work but delay integration; trunk-based enables faster integration but needs discipline
- D. Feature branching and trunk-based development produce identical outcomes with no trade-offs at all

Answer: C

Q1254. A production incident requires an emergency fix while a major release is in progress. How should configuration management handle this?

- A. Create a hotfix branch from the production release tag, apply the fix, merge to both main and release branches
- B. Apply the fix directly to the production server without any version control tracking or review
- C. Cancel the major release entirely and restart the development from scratch to include the fix
- D. Delay the emergency fix until the major release is complete to avoid any branching complexity

Answer: A

Q1255. How does GitOps extend configuration management principles to infrastructure and deployment operations?

- A. GitOps eliminates the need for any version control in infrastructure management and operations
- B. GitOps requires manual deployment processes without any automation or version control involvement
- C. GitOps only applies to application code and has no relevance to infrastructure or operations work
- D. GitOps uses Git as the single source of truth for infrastructure and deployment configurations now

Answer: D

Q1256. What is the impact of poor configuration management on software project outcomes?

- A. It leads to integration failures, lost changes, deployment errors, and inability to reproduce builds
- B. Poor CM has no measurable impact on project outcomes because CM is purely administrative overhead
- C. Only large projects are affected by poor CM while small projects are completely immune to problems
- D. Poor CM only affects the testing phase and has no impact on development or deployment activities

Answer: A

Q1257. How should database schema changes be managed within a configuration management framework?

- A. Use versioned migration scripts tracked in version control with automated application during deployment
- B. Store schema documentation in email threads rather than in the version control repository system
- C. Database schemas should never change after initial deployment regardless of evolving requirements
- D. Apply schema changes directly to production databases without any tracking or version control work

Answer: A

Q1258. What is the relationship between configuration management and reproducible builds?

- A. CM has no impact on build reproducibility since builds are determined solely by compiler versions
- B. Reproducible builds are impossible regardless of how thorough the configuration management practices
- C. CM ensures all build inputs are versioned and tracked enabling identical build reproduction every time
- D. Reproducible builds only matter for open-source projects and are irrelevant for commercial software

Answer: C

Q1259. How does artifact management complement source code version control in a complete CM strategy?

- A. Artifact management replaces source code version control entirely making repositories unnecessary
- B. Only source code needs version control; built artifacts never need tracking or management at all
- C. Artifact management and source code control serve identical purposes and using both is redundant
- D. It tracks versioned build outputs, dependencies, and deployable packages alongside source code control

Answer: D

Q1260. A team needs to maintain multiple product versions simultaneously for different customers. What CM strategy supports this?

- A. Long-lived release branches with selective backporting of changes between supported version branches
- B. Maintain separate codebases for each customer version without any shared source code between them
- C. Use a single branch and deploy different versions by modifying configuration at deployment time only
- D. Force all customers to use the same version and reject any requests for version-specific support

Answer: A

Q1261. A project manager discovers that the team's key architect is considering leaving. How should this personnel risk be managed?

- A. Ignore the risk because the architect has not officially resigned from the project team yet now
- B. Immediately terminate the architect to control the timing of their departure from the team work
- C. Prevent the architect from communicating with other team members to reduce knowledge dependency
- D. Document critical knowledge, cross-train team members, and develop a succession plan immediately

Answer: D

Q1262. How does Monte Carlo simulation enhance risk assessment compared to deterministic estimation methods?

- A. It models uncertainty by running thousands of simulations to produce probability distributions of outcomes
- B. Monte Carlo simulation only works for financial risk assessment and not for software project risks
- C. It provides a single exact answer that is always more accurate than any probabilistic approach used
- D. It eliminates all uncertainty from risk assessment by computing the exact outcome every single time

Answer: A

Q1263. A project depends on a third-party API that may be deprecated. How should this dependency risk be assessed and managed?

- A. Build the entire application tightly coupled to the API without any abstraction or fallback plans
- B. Ignore the deprecation risk because third-party providers never actually remove their APIs from use
- C. Wait until the API is actually deprecated before taking any action to address the dependency risk
- D. Create abstraction layers, monitor API provider communications, and plan migration alternatives early

Answer: D

Q1264. What is the relationship between risk appetite and risk tolerance in organizational risk management?

- A. Risk appetite and risk tolerance are identical concepts with no difference in their definitions at all
- B. Risk appetite refers to specific measurable thresholds while tolerance is a general attitude toward risk
- C. Risk appetite is the overall willingness to take risk; tolerance is the acceptable variation around objectives
- D. Organizations should have zero risk appetite and zero risk tolerance to ensure project success always

Answer: C

Q1265. How should risk management be adapted for agile projects compared to traditional plan-driven projects?

- A. Agile projects have no risks because the iterative nature eliminates all uncertainty automatically now
- B. Apply identical heavyweight risk management processes regardless of whether the project is agile
- C. Integrate risk identification into sprint planning and retrospectives with continuous lightweight assessment
- D. Perform risk management only once at project start and never revisit risks throughout agile sprints

Answer: C

Q1266. A software product must launch on a fixed date for a regulatory deadline. How does this constraint affect risk management?

- A. It elevates schedule risks to critical priority and may require scope negotiation as the main lever
- B. Schedule constraints only affect the budget and have no impact on scope or quality risk factors
- C. The regulatory deadline should be ignored if the team determines it is not technically achievable
- D. Fixed deadlines eliminate all risks because the team will naturally work harder to meet the date

Answer: A

Q1267. What is the difference between known risks, unknown risks, and unknowable risks in project management?

- A. Known risks are identified; unknown risks are unidentified but identifiable; unknowable risks cannot be anticipated
- B. Unknowable risks are the only category worth managing since known and unknown risks resolve themselves
- C. Unknown risks do not exist because thorough analysis will always identify every possible risk event
- D. All three categories are identical because every risk can be identified with sufficient analysis effort

Answer: A

Q1268. How does the concept of 'risk budgeting' help project managers allocate resources effectively?

- A. Risk budgeting wastes project resources on events that may never occur and should be avoided always
- B. All project budget should be allocated to development activities with nothing reserved for risk events
- C. It reserves specific time and budget for managing identified risks proportional to their exposure values
- D. Risk budgeting only applies to financial institutions and has no relevance to software project work

Answer: C

Q1269. What is the impact of cognitive biases on risk assessment accuracy in software projects?

- A. Cognitive biases always improve risk assessment by providing diverse perspectives on potential risks
- B. Cognitive biases only affect individual contributors and never influence management risk decisions
- C. Project managers are immune to cognitive biases due to their training and professional experience
- D. Biases like optimism, anchoring, and groupthink lead to systematic underestimation of project risks

Answer: D

Q1270. A distributed development team faces geopolitical risks across multiple countries. How should these be incorporated into the risk management framework?

- A. Only consider geopolitical risks if a team member specifically raises a concern during a meeting
- B. Assume all regions have identical risk profiles and apply a single uniform risk assessment globally
- C. Assess regulatory, political, and infrastructure risks per region and create contingency plans for each
- D. Geopolitical risks are outside the scope of software project risk management and should be ignored

Answer: C

Q1271. A web application stores user passwords in plaintext in the database. What is the proper remediation approach?

- A. Encrypt passwords with a symmetric key stored alongside the passwords in the same database table
- B. Store passwords in a separate plaintext file on the server instead of in the database tables
- C. Continue storing passwords in plaintext because hashing adds unnecessary system complexity now
- D. Hash passwords using bcrypt or Argon2 with per-user salts and migrate all existing user records

Answer: D

Q1272. How does the STRIDE threat model categorize security threats?

- A. Scalability, Traceability, Reusability, Interoperability, Durability, and Extensibility of design
- B. Spoofing, Tampering, Repudiation, Information disclosure, Denial of service, Elevation of privilege
- C. Speed, Testing, Reliability, Integration, Deployment, and Efficiency of the system components
- D. Storage, Transmission, Retrieval, Indexing, Deletion, and Encryption of application data now

Answer: B

Q1273. A company discovers that a critical zero-day vulnerability affects their production system. What is the recommended incident response?

- A. Shut down all company systems permanently until the vulnerability is fixed by the software vendor
- B. Post details of the vulnerability publicly to warn other companies before applying any patches
- C. Isolate affected systems, assess impact, apply temporary mitigations, develop and deploy a permanent fix
- D. Ignore the vulnerability because zero-day means no exploit exists yet in the wild for it now

Answer: C

Q1274. How does the principle of 'security by design' differ from 'security by obscurity'?

- A. Security by design has been proven ineffective and all modern systems rely on obscurity instead
- B. Security by design builds protection into architecture; obscurity relies on hiding implementation details
- C. Security by obscurity is more effective because attackers cannot exploit what they cannot find now
- D. Security by design and security by obscurity are identical approaches producing equivalent protection

Answer: B

Q1275. What is the impact of the DevSecOps approach on traditional security engineering practices?

- A. It delays all security testing until after production deployment to avoid slowing down development
- B. It eliminates all security activities because DevOps automation makes applications inherently secure
- C. DevSecOps has no meaningful impact on security practices compared to traditional approaches now
- D. It integrates security into every stage of the CI/CD pipeline rather than treating it as a gate only

Answer: D

Q1276. How should an organization implement a zero-trust security architecture?

- A. Verify every request regardless of source, enforce least privilege, and assume breach at all times
- B. Implement a single perimeter firewall and trust everything inside the corporate network boundary
- C. Trust all internal network traffic and only verify requests coming from external network sources
- D. Remove all authentication and authorization controls to achieve maximum system simplicity now

Answer: A

Q1277. What is the relationship between security requirements and functional requirements in software engineering?

- A. Functional requirements always take priority and security requirements can be ignored if needed
- B. Security requirements are completely independent of functional requirements with no interaction now
- C. Security requirements constrain how functional requirements are implemented to protect system assets
- D. Security requirements only apply to the network layer and never affect application functionality

Answer: C

Q1278. How does supply chain security affect modern software development practices?

- A. Supply chain security is irrelevant because all third-party libraries are thoroughly vetted already
- B. Only first-party code can contain vulnerabilities since third-party libraries are always secure now
- C. Third-party dependencies can introduce vulnerabilities requiring continuous monitoring and auditing
- D. Supply chain attacks only target hardware and have no relevance to software development practices

Answer: C

Q1279. What is the role of security testing throughout the software development lifecycle?

- A. Early security testing is wasteful because most vulnerabilities are introduced during deployment
- B. Security testing should only occur during the final phase before deployment to production servers
- C. Security testing at every phase catches vulnerabilities earlier when they are cheaper to fix overall
- D. Security testing is unnecessary if the development team follows secure coding guidelines properly

Answer: C

Q1280. How should cryptographic key management be implemented in enterprise software systems?

- A. Store encryption keys in plaintext files alongside the encrypted data for easy access and retrieval
- B. Use hardware security modules, automated key rotation, and strict access controls for all key materials
- C. Use a single encryption key for all data across the entire organization without any rotation ever
- D. Hardcode encryption keys directly into the application source code for convenience and reliability

Answer: B

Q1281. What does Lehman's Law of Continuing Growth state?

- A. Software always decreases in size
- B. Software must continuously grow in functionality to remain satisfactory to users
- C. Software never changes after release
- D. Growth always improves quality

Answer: B

Q1282. In the context of socio-technical systems, what is an organizational system?

- A. A system made only of hardware
- B. A broader system that includes social and organizational policies and rules governing how people interact with the technical system
- C. A database management system
- D. An operating system

Answer: B

Q1283. What is the 'no silver bullet' argument by Fred Brooks?

- A. That all software should use silver encryption
- B. That there is no single technique that will dramatically improve software productivity by an order of magnitude
- C. That only one programming language should be used
- D. That testing alone can guarantee quality

Answer: B

Q1284. What distinguishes 'essential complexity' from 'accidental complexity' in software?

- A. Essential complexity can be eliminated; accidental cannot
- B. Essential complexity is inherent to the problem; accidental complexity arises from the chosen solution approach
- C. They are the same concept
- D. Essential complexity relates only to hardware

Answer: B

Q1285. What is the Standish Group's CHAOS report primarily known for?

- A. Evaluating programming languages
- B. Analyzing software project success and failure rates
- C. Rating software companies
- D. Benchmarking hardware performance

Answer: B

Q1286. Which of the following best describes the concept of 'wicked problems' in software engineering?

- A. Problems that involve malicious code
- B. Problems that are inherently ill-defined and have no definitive solution
- C. Simple problems with obvious solutions
- D. Problems caused by poor coding practices

Answer: B

Q1287. What is the principle of 'separation of concerns' in software engineering?

- A. Keeping developers separate from testers
- B. Dividing a program into distinct sections that each address a separate concern or functionality
- C. Separating the budget from the schedule
- D. Using different computers for different tasks

Answer: B

Q1288. What is Lehman's Law of Declining Quality?

- A. Software quality always improves
- B. Software will be perceived as declining in quality unless rigorously maintained and adapted to its operational environment
- C. Quality is unrelated to maintenance
- D. Only hardware quality declines

Answer: B

Q1289. What role does abstraction play in managing software complexity?

- A. It makes software run faster
- B. It hides unnecessary details so developers can focus on relevant aspects at each level
- C. It eliminates all bugs
- D. It replaces the need for testing

Answer: B

Q1290. What is the 'Pareto Principle' as applied to software defects?

- A. All modules have equal defect rates
- B. Roughly 80% of defects are found in approximately 20% of the modules
- C. Defects can never be found
- D. Testing eliminates 100% of defects

Answer: B

Q1291. What are the four phases of the Rational Unified Process (RUP)?

- A. Plan, Code, Test, Deploy
- B. Inception, Elaboration, Construction, Transition
- C. Analysis, Design, Code, Maintain
- D. Gather, Build, Verify, Release

Answer: B

Q1292. What is the cleanroom software engineering process?

- A. Development done in a physically clean room
- B. A process emphasizing formal specification, incremental development, and statistical testing to produce near-zero-defect software
- C. A process that only uses one programmer
- D. A process that avoids all documentation

Answer: B

Q1293. What is the concept of 'process tailoring'?

- A. Designing new programming languages
- B. Adapting a standard process model to fit specific project needs, constraints, and organizational culture
- C. Ignoring all processes
- D. Writing code without any process

Answer: B

Q1294. In the Spiral model, what happens during the 'planning' quadrant?

- A. Code is written
- B. Project scope, schedule, and cost estimates are reviewed and the next iteration is planned based on risk analysis results
- C. The software is deployed
- D. Users are trained

Answer: B

Q1295. What is the difference between verification and validation in process models?

- A. They are the same activity
- B. Verification checks if the product is built correctly per specification; validation checks if the right product is built for the user
- C. Verification is done after deployment; validation is done before coding
- D. Neither involves testing

Answer: B

Q1296. What is a stage-gate process in software development?

- A. A process with no checkpoints
- B. A process where each phase ends with a formal review (gate) that must be passed before proceeding to the next stage
- C. A process that only uses one language
- D. An automated deployment pipeline

Answer: B

Q1297. What is the key criticism of prescriptive process models in modern software development?

- A. They are too flexible
- B. They can be too rigid and document-heavy, failing to adapt to rapidly changing requirements and technologies
- C. They produce software too quickly
- D. They never include testing

Answer: B

Q1298. What is the Capability Maturity Model Integration (CMMI) and how does it relate to process models?

- A. A programming language
- B. A framework for assessing and improving organizational process maturity across five levels
- C. A type of database
- D. A testing tool

Answer: B

Q1299. What is the Cynefin framework and how does it relate to Agile?

- A. A programming framework
- B. A sense-making model that categorizes problems into domains (simple, complicated, complex, chaotic) to guide the appropriate management approach
- C. A database management system
- D. A version control system

Answer: B

Q1300. What is the difference between Scrum and Scrumban?

- A. There is no difference
- B. Scrumban combines Scrum's structure with Kanban's flow-based approach, using WIP limits instead of strict Sprint commitments
- C. Scrumban eliminates all Scrum ceremonies
- D. Scrumban is only for hardware development

Answer: B

Q1301. What is the concept of 'swarming' in Agile?

- A. The team working on many tasks simultaneously
- B. Multiple team members collaborating on a single high-priority item to complete it quickly
- C. A technique for writing documentation
- D. A method for deploying software

Answer: B

Q1302. What is Large-Scale Scrum (LeSS)?

- A. A version of Scrum for small teams
- B. A framework for applying Scrum principles to multiple teams working on a single product with minimal additional roles or artifacts
- C. A testing methodology
- D. A documentation standard

Answer: B

Q1303. What is the 'iron triangle' trade-off in Agile versus traditional project management?

- A. Agile fixes scope and varies time/cost
- B. In traditional management scope is fixed while in Agile time and cost are fixed and scope varies
- C. They handle trade-offs identically
- D. Agile ignores all constraints

Answer: B

Q1304. What is Behavior-Driven Development (BDD) and how does it extend TDD?

- A. BDD replaces all testing
- B. BDD extends TDD by using natural language specifications that describe system behavior from the user's perspective
- C. BDD is unrelated to testing
- D. BDD only tests database operations

Answer: B

Q1305. What is the concept of 'Minimum Viable Product' (MVP) in Agile?

- A. The cheapest product possible
- B. The smallest version of a product that delivers enough value to early adopters and provides validated learning
- C. A product with no features
- D. The final release version

Answer: B

Q1306. What is the Spotify model in Agile?

- A. A music streaming app
- B. An organizational model using Squads, Tribes, Chapters, and Guilds to balance team autonomy with cross-functional alignment
- C. A testing framework
- D. A deployment strategy

Answer: B

Q1307. What is the Agile concept of 'sustainable pace'?

- A. Working overtime every day
- B. Maintaining a work rhythm that the team can sustain indefinitely without burnout
- C. Releasing software once a year
- D. Working only 2 hours per day

Answer: B

Q1308. What are Agile release trains (ARTs) in the Scaled Agile Framework (SAFe)?

- A. Physical trains that deliver software
- B. Long-lived teams of Agile teams that plan, commit, and execute together in Program Increments
- C. A type of version control
- D. A deployment tool

Answer: B

Q1309. What is the KAOS methodology in requirements engineering?

- A. A methodology for causing chaos
- B. A goal-oriented requirements engineering approach that derives requirements from high-level organizational goals
- C. A testing technique
- D. A coding standard

Answer: B

Q1310. What is the difference between safety requirements and security requirements?

- A. They are the same
- B. Safety requirements prevent harm to people or environment; security requirements protect against malicious threats
- C. Safety is about passwords
- D. Security is about physical safety

Answer: B

Q1311. What is the problem of 'requirements creep'?

- A. Requirements that are too small
- B. The uncontrolled expansion of project scope as new requirements are continuously added without proper change management
- C. Requirements that are well-managed
- D. A beneficial increase in features

Answer: B

Q1312. What is the i* (iStar) framework in requirements engineering?

- A. A programming language
- B. An agent-oriented modeling framework that captures stakeholder goals, dependencies, and relationships
- C. A testing framework
- D. A deployment tool

Answer: B

Q1313. How does ethnography contribute to requirements elicitation?

- A. It replaces all other elicitation techniques
- B. It reveals implicit requirements by observing users in their actual work environment to understand real practices
- C. It only studies ancient cultures
- D. It is a testing technique

Answer: B

Q1314. What is the Kano model in requirements analysis?

- A. A database model
- B. A model that classifies features into must-be, one-dimensional, attractive, indifferent, and reverse categories based on customer satisfaction
- C. A process model
- D. A testing model

Answer: B

Q1315. What is the concept of 'requirements engineering for self-adaptive systems'?

- A. Systems that never change
- B. Specifying requirements that allow a system to monitor itself and dynamically adjust its behavior to meet goals
- C. Requirements written by AI
- D. Systems without any requirements

Answer: B

Q1316. What is the role of natural language processing (NLP) in modern requirements engineering?

- A. It has no role
- B. NLP tools help detect ambiguities, inconsistencies, and missing information in requirements documents
- C. NLP replaces all requirements activities
- D. NLP is only used for coding

Answer: B

Q1317. What is the Putnam model (SLIM) used for in software estimation?

- A. Estimating physical weight of hardware
- B. Using a Rayleigh curve to relate project effort, development time, and software size for effort estimation
- C. Estimating employee satisfaction
- D. Predicting market trends

Answer: B

Q1318. What is the difference between COCOMO I and COCOMO II?

- A. There is no difference
- B. COCOMO II accounts for modern development practices including reuse, object-oriented development, and iterative processes
- C. COCOMO I is newer
- D. COCOMO II only estimates hardware costs

Answer: B

Q1319. What is the Wideband Delphi estimation technique?

- A. A single expert making all estimates
- B. A consensus-based estimation method where experts independently estimate, discuss differences, and re-estimate iteratively
- C. A database query optimization method
- D. An automated code generation technique

Answer: B

Q1320. What is the concept of 'planning poker' in Agile estimation?

- A. A card game for entertainment
- B. A consensus-based estimation technique where team members simultaneously reveal their estimates using cards to avoid anchoring bias
- C. A method for assigning tasks randomly
- D. A technique for prioritizing bug fixes

Answer: B

Q1321. What is the concept of 'cone of uncertainty' in estimation?

- A. A physical cone used in meetings
- B. The principle that estimation accuracy improves as the project progresses, with early estimates having 4x variability
- C. Uncertainty never decreases
- D. Estimates are always accurate from the start

Answer: B

Q1322. What is the Theory of Constraints (TOC) as applied to software project management?

- A. A theory about hardware limitations
- B. A management philosophy that identifies the most critical limiting factor (bottleneck) and systematically improves it
- C. A programming paradigm
- D. A database optimization theory

Answer: B

Q1323. What is the concept of 'project velocity' in iterative estimation?

- A. The speed of code compilation
- B. Using historical throughput data from past iterations to forecast how much work can be completed in future iterations
- C. The network speed of the team
- D. How fast code is deployed

Answer: B

Q1324. What is Goldratt's Critical Chain Project Management (CCPM)?

- A. Managing chains of emails
- B. A method that accounts for resource constraints and uses project buffers instead of task-level buffers to protect the schedule
- C. A version control strategy
- D. A code review technique

Answer: B

Q1325. What is the purpose of post-mortem analysis in project management?

- A. Analyzing deceased team members
- B. Reviewing a completed project to identify lessons learned, successes, and areas for improvement
- C. Analyzing legacy code
- D. Testing production systems

Answer: B

Q1326. What is the Visitor design pattern?

- A. A pattern for tracking website visitors
- B. A pattern that separates an algorithm from the object structure it operates on, allowing new operations without modifying the objects
- C. A pattern for user authentication
- D. A pattern for logging

Answer: B

Q1327. What is the Abstract Factory design pattern?

- A. A factory building blueprint
- B. A pattern that provides an interface for creating families of related objects without specifying their concrete classes
- C. An abstract art design tool
- D. A database abstraction layer

Answer: B

Q1328. What is the Law of Demeter and why is it important?

- A. A mathematical law
- B. A design principle stating that a method should only call methods on its own object, parameters, created objects, or direct components, reducing coupling
- C. A law about documentation
- D. A law about data encryption

Answer: B

Q1329. What is the difference between the Bridge and Adapter patterns?

- A. They are identical
- B. The Bridge pattern separates abstraction from implementation by design; the Adapter makes existing incompatible interfaces work together after the fact
- C. Bridge is for databases; Adapter is for files
- D. Neither involves interfaces

Answer: B

Q1330. What is the concept of 'design debt' in software engineering?

- A. The financial cost of design tools
- B. The accumulated cost of shortcuts and suboptimal design decisions that will require future refactoring
- C. Debt incurred by buying design software
- D. The cost of hiring designers

Answer: B

Q1331. What is the Mediator design pattern?

- A. A legal mediation tool
- B. A pattern that defines an object that encapsulates how a set of objects interact, promoting loose coupling by preventing direct references
- C. A middleware framework
- D. A database connector

Answer: B

Q1332. What is the concept of 'design for testability'?

- A. Making the UI look like a test
- B. Structuring software so that individual components can be easily isolated and tested through techniques like dependency injection
- C. Writing tests before design
- D. Testing the design document

Answer: B

Q1333. What is the Chain of Responsibility design pattern?

- A. An organizational hierarchy chart
- B. A pattern that passes a request along a chain of handlers, where each handler decides whether to process it or pass it to the next handler
- C. A supply chain management tool
- D. A blockchain pattern

Answer: B

Q1334. What is the Architecture Tradeoff Analysis Method (ATAM) used for?

- A. Analyzing stock trades
- B. Evaluating software architectures against quality attributes by identifying sensitivity points, tradeoff points, and risks
- C. Analyzing code performance
- D. Trading software licenses

Answer: B

Q1335. What is the concept of 'eventual consistency' in distributed systems?

- A. Systems that are always consistent
- B. A model where, given enough time without new updates, all replicas of data will converge to the same value
- C. Systems that are never consistent
- D. A database locking mechanism

Answer: B

Q1336. What does the CAP theorem state?

- A. All distributed systems are perfect
- B. A distributed system can provide at most two of three guarantees: Consistency, Availability, and Partition tolerance simultaneously
- C. Distributed systems have no limitations
- D. CAP stands for Code, Application, Program

Answer: B

Q1337. What is the Strangler Fig pattern for migrating legacy systems?

- A. A pattern for growing plants
- B. A migration strategy that gradually replaces parts of a legacy system with new components until the old system is completely replaced
- C. A security vulnerability pattern
- D. A testing pattern

Answer: B

Q1338. What is the concept of 'architectural fitness functions'?

- A. Exercise routines for architects
- B. Automated tests that verify whether an architecture meets its intended quality attributes and constraints over time
- C. A measure of code readability
- D. A design pattern

Answer: B

Q1339. What is the Saga pattern in microservices?

- A. A storytelling framework
- B. A pattern for managing distributed transactions across multiple services using a sequence of local transactions with compensating actions
- C. A logging framework
- D. A deployment pattern

Answer: B

Q1340. What is the concept of 'domain-driven design' (DDD) in architecture?

- A. Designing domain names for websites
- B. An approach that models software structure and behavior around the business domain using bounded contexts, aggregates, and ubiquitous language
- C. A database design approach
- D. A UI design methodology

Answer: B

Q1341. What is the Bulkhead pattern in resilient architecture?

- A. A ship construction pattern
- B. A pattern that isolates system components into pools so that failure in one does not cascade and bring down the entire system
- C. A database partitioning strategy
- D. A UI layout pattern

Answer: B

Q1342. What is a cell-based architecture?

- A. An architecture inspired by biological cells
- B. An architecture that groups related microservices into independently deployable and scalable cells, each containing a complete stack
- C. A spreadsheet-based architecture
- D. A hardware design approach

Answer: B

Q1343. What is the difference between summative and formative usability evaluation?

- A. They are the same
- B. Formative evaluation identifies problems during design to improve the product; summative evaluation measures overall usability of the finished product
- C. Summative is done first
- D. Formative evaluation cannot find problems

Answer: B

Q1344. What is the GOMS model in HCI?

- A. A graphics rendering model
- B. A cognitive modeling method that predicts user performance by analyzing Goals, Operators, Methods, and Selection rules
- C. A database model
- D. A project management model

Answer: B

Q1345. What is the concept of 'cognitive walkthrough' in usability evaluation?

- A. Walking while thinking
- B. A usability inspection method where evaluators step through task scenarios to identify learnability problems from a new user's perspective
- C. A physical exercise program
- D. A code review technique

Answer: B

Q1346. What is Hick's Law and its implication for UI design?

- A. A law about screen resolution
- B. The time to make a decision increases logarithmically with the number of choices, implying designers should minimize options
- C. A law about animation speed
- D. A law about data transfer

Answer: B

Q1347. What is the ISO 9241-210 standard about?

- A. Database design
- B. An international standard providing requirements and recommendations for human-centered design of interactive systems
- C. Network protocols
- D. Software licensing

Answer: B

Q1348. What is the concept of 'design thinking' in UI/UX?

- A. Thinking about design while sleeping
- B. A human-centered innovation approach using empathy, ideation, prototyping, and testing to solve complex design problems iteratively
- C. A method for writing code
- D. A database design technique

Answer: B

Q1349. What is the Weber-Fechner Law in the context of UI design?

- A. A law about network latency
- B. A psychophysical law stating that the perceived difference between stimuli is proportional to their ratio, affecting how users notice UI changes
- C. A database normalization rule
- D. A coding standard

Answer: B

Q1350. What is the concept of 'emotional design' as described by Don Norman?

- A. Designing software that makes users cry
- B. A design philosophy addressing three levels of user experience: visceral (appearance), behavioral (usability), and reflective (self-image)
- C. Designing emoticons
- D. A psychological test

Answer: B

Q1351. What is the concept of 'defensive programming'?

- A. Programming while on defense in sports
- B. Writing code that anticipates and handles unexpected inputs, states, and errors gracefully to improve robustness
- C. Programming in a secure facility
- D. Programming with antivirus software

Answer: B

Q1352. What is the difference between optimistic and pessimistic concurrency control?

- A. They produce identical results
- B. Optimistic assumes conflicts are rare and checks at commit time; pessimistic locks resources preemptively to prevent conflicts
- C. Optimistic always uses locks
- D. Pessimistic never uses locks

Answer: B

Q1353. What is the concept of 'immutability' and why is it valuable in implementation?

- A. Code that cannot be read
- B. Objects whose state cannot be modified after creation, reducing bugs from unintended state changes and simplifying concurrent programming
- C. Variables that can change freely
- D. A database locking mechanism

Answer: B

Q1354. What is the 'fail-fast' principle in implementation?

- A. Making software crash as often as possible
- B. Detecting and reporting errors immediately at the point of failure rather than allowing them to propagate and cause obscure failures later
- C. Deploying quickly without testing
- D. Skipping error handling

Answer: B

Q1355. What is aspect-oriented programming (AOP)?

- A. Programming from different angles
- B. A programming paradigm that separates cross-cutting concerns like logging, security, and transactions from business logic
- C. Programming for aspect ratios
- D. A type of object-oriented programming

Answer: B

Q1356. What is the concept of 'continuous delivery' pipeline?

- A. A pipe for delivering water
- B. An automated software release pipeline that ensures code changes are always in a deployable state through build, test, and staging automation
- C. A continuous email notification system
- D. A hardware delivery system

Answer: B

Q1357. What is the concept of 'feature flags' in implementation?

- A. Physical flags on servers
- B. Conditional switches in code that enable or disable features at runtime without deploying new code
- C. Flags in a debugging tool
- D. Country flags in the UI

Answer: B

Q1358. What is the concept of 'contract-based programming'?

- A. Legal contracts for developers
- B. A programming approach where functions define preconditions, postconditions, and invariants as formal contracts
- C. Freelance programming
- D. A licensing scheme

Answer: B

Q1359. What is the strangler pattern applied to codebase modernization?

- A. A gardening technique
- B. Gradually replacing legacy code by wrapping it with new implementations and redirecting calls until the old code can be safely removed
- C. Strangling competition
- D. A compression algorithm

Answer: B

Q1360. What is the concept of 'trunk-based development'?

- A. Developing software on a tree trunk
- B. A branching strategy where all developers commit to a single main branch frequently, using short-lived feature branches
- C. Using elephant trunks as input devices
- D. A storage management technique

Answer: B

Q1361. What is mutation testing and how does it assess test quality?

- A. Testing with mutated viruses
- B. A technique that introduces small changes (mutants) to the source code and checks whether existing tests detect these changes
- C. Genetic testing of software
- D. Testing software mutations over time

Answer: B

Q1362. What is the concept of 'combinatorial testing' (pairwise testing)?

- A. Testing combinations of locks
- B. A technique that tests interactions between parameters by covering all pairwise combinations, significantly reducing the number of test cases needed
- C. Combining two test suites
- D. Testing with two testers

Answer: B

Q1363. What is symbolic execution in testing?

- A. Testing with symbols instead of code
- B. A technique that analyzes programs by using symbolic values instead of concrete inputs to explore multiple execution paths simultaneously
- C. A type of code signing
- D. Writing tests using special symbols

Answer: B

Q1364. What is the concept of 'chaos engineering' in testing?

- A. Creating chaos in the testing team
- B. Deliberately injecting failures into a production system to test its resilience and identify weaknesses before real failures occur
- C. Disorganized testing practices
- D. Testing without any plan

Answer: B

Q1365. What is the oracle problem in software testing?

- A. A problem with Oracle databases
- B. The difficulty of determining the correct expected output for a given test input, especially for complex or novel systems
- C. A licensing issue
- D. A performance problem

Answer: B

Q1366. What is metamorphic testing?

- A. Testing software transformation
- B. A technique that addresses the oracle problem by defining relationships between inputs and outputs across multiple test executions
- C. Testing software evolution
- D. Testing user interface transitions

Answer: B

Q1367. What is the concept of 'shift-left testing'?

- A. Moving test results to the left column
- B. Integrating testing activities earlier in the software development lifecycle to find defects sooner when they are cheaper to fix
- C. Testing only left-aligned text
- D. A specific testing tool

Answer: B

Q1368. What is fuzzing (fuzz testing)?

- A. Making software fuzzy
- B. An automated testing technique that provides invalid, unexpected, or random data as inputs to discover vulnerabilities and crashes
- C. Testing with blurred images
- D. A type of performance test

Answer: B

Q1369. What is Lehman's Law of Increasing Complexity and how does it affect maintenance?

- A. Complexity always decreases
- B. As software evolves, its complexity increases unless deliberate effort is made to reduce it, making maintenance progressively harder
- C. Complexity has no impact on maintenance
- D. Only new software is complex

Answer: B

Q1370. What is the concept of 'software rejuvenation'?

- A. Making old software look new visually
- B. Proactively restarting or refreshing software systems to prevent failures caused by accumulated degradation such as memory leaks
- C. Celebrating a software's anniversary
- D. Installing software updates

Answer: B

Q1371. What is the 'ripple effect' in software maintenance?

- A. An animation effect in the UI
- B. The phenomenon where a change to one part of the system causes unintended changes or failures in other parts
- C. A water-themed design pattern
- D. A network propagation technique

Answer: B

Q1372. What is the concept of 'technical debt management' in maintenance?

- A. Managing financial debt
- B. Strategically identifying, measuring, and prioritizing the repayment of accumulated design and code shortcuts to maintain system health
- C. Managing team debt
- D. Managing hardware debt

Answer: B

Q1373. What is the difference between big-bang and incremental migration for legacy systems?

- A. They are the same approach
- B. Big-bang replaces the entire legacy system at once; incremental migration gradually replaces parts while both systems run in parallel
- C. Big-bang is always safer
- D. Incremental migration is never used

Answer: B

Q1374. What is the concept of 'software archeology' in maintenance?

- A. Digging up old hardware
- B. The practice of investigating and understanding legacy code through code reading, tooling, and historical analysis when documentation is missing
- C. Collecting vintage software
- D. Building a software museum

Answer: B

Q1375. What is the concept of 'planned obsolescence' versus 'unplanned obsolescence' in software?

- A. They are identical
- B. Planned obsolescence is an intentional end-of-life strategy; unplanned obsolescence occurs when external changes make the software unusable unexpectedly
- C. Neither affects maintenance
- D. Obsolescence never happens in software

Answer: B

Q1376. What is the concept of 'evolution-aware software design'?

- A. Designing software inspired by biological evolution
- B. Designing software from the start with the expectation of future changes, making modification points explicit and changes easier
- C. Using evolutionary algorithms
- D. Designing without any plan

Answer: B

Q1377. What is the challenge of maintaining systems with undocumented dependencies?

- A. There is no challenge
- B. Hidden dependencies can cause unexpected cascading failures when changes are made, as maintainers cannot anticipate which components will be affected
- C. Undocumented dependencies improve performance
- D. Dependencies do not affect maintenance

Answer: B

Q1378. What is the concept of 'continuous maintenance' in DevOps?

- A. Maintenance done once a year
- B. An approach where maintenance activities are integrated into daily development workflows through automation, monitoring, and continuous improvement
- C. Maintaining only during business hours
- D. Outsourcing all maintenance

Answer: B

Q1379. What is the GQM (Goal-Question-Metric) paradigm and how is it applied?

- A. A gaming metric
- B. A structured approach where quality goals are defined, questions identify what needs measurement, and metrics provide the answers
- C. A graphics quality metric
- D. A database query method

Answer: B

Q1380. What is the concept of 'quality culture' in an organization?

- A. A culture of testing everything
- B. An organizational environment where quality is valued by everyone, built into processes, and continuously improved as a shared responsibility
- C. A culture of writing documentation
- D. A culture of using expensive tools

Answer: B

Q1381. What is the Rayleigh model used for in defect prediction?

- A. Predicting weather patterns
- B. A statistical model that describes the expected distribution of defect discovery rates over time during development and testing
- C. Predicting stock prices
- D. Predicting network traffic

Answer: B

Q1382. What is the concept of 'defect amplification' in the software process?

- A. Making defects louder
- B. The phenomenon where undetected defects in one phase generate additional defects in subsequent phases, multiplying the total defect count
- C. Amplifying error messages
- D. Increasing test coverage

Answer: B

Q1383. What is Cleanroom software engineering in the context of QA?

- A. Developing in a sterile room
- B. A quality-focused approach using formal specification, structured programming, and statistical usage testing to achieve near-zero-defect software
- C. A method of cleaning code
- D. Developing with air purifiers

Answer: B

Q1384. What is the concept of 'quality attribute scenarios' in quality assessment?

- A. Movie scripts about quality
- B. Structured descriptions that specify quality attribute requirements through source, stimulus, environment, artifact, response, and response measure
- C. Testing scenarios only
- D. User story templates

Answer: B

Q1385. What is the relationship between process maturity and defect rates?

- A. There is no relationship
- B. Higher process maturity levels are generally associated with lower defect rates and more predictable quality outcomes
- C. Higher maturity means more defects
- D. Process maturity is irrelevant to defects

Answer: B

Q1386. What is the concept of 'total quality management' (TQM) applied to software?

- A. A testing management tool
- B. A management philosophy emphasizing continuous improvement, customer focus, and involvement of all employees in quality across the entire organization
- C. A quality certification
- D. A project management methodology

Answer: B

Q1387. What is the Coupling Between Objects (CBO) metric in the CK suite?

- A. The number of physical connections between servers
- B. The count of other classes to which a class is coupled through method calls, variable access, or inheritance, indicating its dependency level
- C. The number of database connections
- D. The number of API calls

Answer: B

Q1388. What is the concept of 'measurement dysfunction' in software metrics?

- A. Broken measurement tools
- B. The phenomenon where measuring a specific attribute causes people to optimize for that metric at the expense of actual goals
- C. Metrics that produce wrong numbers
- D. Dysfunction in the measurement team

Answer: B

Q1389. What is Halstead's concept of 'program vocabulary' and 'program length'?

- A. The number of English words in comments
- B. Vocabulary is the number of distinct operators and operands; length is the total number of operators and operands used in the program
- C. The number of variable names
- D. The number of code files

Answer: B

Q1390. What is the concept of 'technical debt ratio' as a metric?

- A. The company's financial debt ratio
- B. The ratio of the cost to fix all code issues to the total cost of developing the code, indicating the relative amount of accumulated technical shortcuts
- C. The ratio of developers to managers
- D. The ratio of tests to code

Answer: B

Q1391. What is the concept of 'object-oriented design metrics' and how do they differ from traditional metrics?

- A. They are the same as traditional metrics
- B. OO metrics measure class-specific attributes like inheritance depth, coupling, and cohesion that traditional metrics like LOC and cyclomatic complexity do not capture
- C. OO metrics only measure code size
- D. Traditional metrics are always superior

Answer: B

Q1392. What is the concept of 'code churn' as a metric?

- A. The sound of a hard drive
- B. The frequency and volume of code changes over time, where high churn in a module often correlates with higher defect rates
- C. The number of developers who leave
- D. The compilation time

Answer: B

Q1393. What is the purpose of 'software reliability growth models'?

- A. Growing software like plants
- B. Mathematical models that use defect discovery data to predict remaining defects and estimate when software will reach a target reliability level
- C. Models for increasing server capacity
- D. Models for team growth

Answer: B

Q1394. What is the concept of 'metric validation' in software measurement?

- A. Checking if numbers add up correctly
- B. The process of ensuring that a metric actually measures what it claims to measure and is useful for its intended purpose
- C. Validating test results
- D. Checking code syntax

Answer: B

Q1395. What is the concept of 'configuration identification' in SCM?

- A. Identifying configuration files
- B. The process of selecting configuration items, defining their relationships, and establishing a naming and labeling scheme for unique identification
- C. Identifying team members
- D. Identifying hardware components

Answer: B

Q1396. What is the concept of 'infrastructure as code' (IaC) in configuration management?

- A. Building physical infrastructure
- B. Managing and provisioning computing infrastructure through machine-readable definition files rather than manual configuration
- C. Writing code on paper
- D. A type of programming language

Answer: B

Q1397. What is the concept of 'immutable infrastructure' in modern configuration management?

- A. Infrastructure that is impossible to change
- B. A practice where servers are never modified after deployment; any change requires building and deploying new instances from scratch
- C. Infrastructure made of unbreakable materials
- D. A type of hardware

Answer: B

Q1398. What is a Software Bill of Materials (SBOM) and why is it important?

- A. A bill for software purchases
- B. A formal, machine-readable inventory of all components, libraries, and dependencies in a software product for security and compliance tracking
- C. A financial document
- D. An employee roster

Answer: B

Q1399. What is the concept of 'GitOps' as a configuration management approach?

- A. Using Git for social media
- B. An operational model where Git repositories serve as the single source of truth for declarative infrastructure and application configuration
- C. A Git tutorial series
- D. A Git plugin

Answer: B

Q1400. What is the difference between 'feature toggles' and 'feature branches' as configuration strategies?

- A. They are identical
- B. Feature branches isolate work in separate code branches; feature toggles allow incomplete features to exist in the main branch but be disabled at runtime
- C. Feature toggles never work
- D. Feature branches are always better

Answer: B

Q1401. What is the concept of 'configuration drift' and how is it managed?

- A. Physical movement of servers
- B. The gradual divergence of actual system configurations from their intended state, managed through automated monitoring and remediation
- C. Network signal drift
- D. Time zone changes

Answer: B

Q1402. What is the concept of 'dependency management' in configuration management?

- A. Managing employee dependencies
- B. Systematically tracking, resolving, and updating third-party libraries and components that a software project depends on
- C. Managing database dependencies
- D. Managing hardware dependencies

Answer: B

Q1403. What is the concept of 'blue-green deployment' in release management?

- A. Using blue and green color themes
- B. A deployment strategy using two identical production environments where traffic is switched from the current (blue) to the new (green) version
- C. A testing strategy using colors
- D. A UI design technique

Answer: B

Q1404. What is canary deployment and how does it reduce release risk?

- A. Deploying software to canary birds
- B. Gradually rolling out changes to a small subset of users before deploying to the entire infrastructure, allowing early detection of issues
- C. A type of load testing
- D. A security scanning technique

Answer: B

Q1405. What is the Expected Monetary Value (EMV) technique in risk analysis?

- A. The expected salary of the project manager
- B. A quantitative technique that calculates the average outcome by multiplying each possible outcome's value by its probability and summing all results
- C. The expected project budget
- D. The expected revenue from the project

Answer: B

Q1406. What is a decision tree in the context of risk management?

- A. A tree used for making physical decisions
- B. A graphical tool that models decision options, chance events, probabilities, and outcomes to evaluate the expected value of different risk response strategies
- C. A type of data structure
- D. A tree in a garden used for thinking

Answer: B

Q1407. What is Boehm's Top 10 Risk Items approach?

- A. A list of the top 10 most expensive software products
- B. A risk management method that continuously tracks the top 10 project risks, ensuring management attention is focused on the most critical risks
- C. A list of the top 10 programming languages
- D. A ranking of the top 10 developers

Answer: B

Q1408. What is the concept of 'risk appetite' versus 'risk tolerance'?

- A. They are identical concepts
- B. Risk appetite is the broad level of risk an organization is willing to accept to achieve goals; risk tolerance is the specific acceptable variation from a particular objective
- C. Risk appetite is about food preferences
- D. Risk tolerance is about physical pain

Answer: B

Q1409. What is the concept of 'anti-fragile' systems in risk management?

- A. Systems that break easily
- B. Systems that actually improve and become stronger when exposed to stressors, volatility, and disorder
- C. Fragile systems that are not working
- D. Systems with no risk

Answer: B

Q1410. What is the concept of 'risk burndown' in Agile projects?

- A. Burning risk documents
- B. A chart tracking the total risk exposure over time, showing whether the project's overall risk is decreasing as planned
- C. A risk that causes burnout
- D. Deleting risk entries from the backlog

Answer: B

Q1411. What is the Monte Carlo simulation used for in risk analysis?

- A. Simulating casino gambling
- B. Running thousands of simulations with randomly varied risk parameters to produce a probability distribution of possible project outcomes
- C. Simulating weather patterns
- D. Simulating network traffic

Answer: B

Q1412. What is the concept of 'residual risk' in risk management?

- A. Risk that has been completely eliminated
- B. The remaining level of risk after risk response strategies have been applied, which must be documented and monitored
- C. Risk from restaurant software
- D. The first risk identified

Answer: B

Q1413. What is the STRIDE threat model and how is each element addressed?

- A. A walking exercise program
- B. A framework categorizing threats as Spoofing, Tampering, Repudiation, Information Disclosure, Denial of Service, and Elevation of Privilege, each with specific countermeasures
- C. A software development methodology
- D. A testing framework

Answer: B

Q1414. What is the concept of 'zero trust architecture'?

- A. Trusting no one in the office
- B. A security model that never implicitly trusts any entity inside or outside the network and requires continuous verification of every access request
- C. An architecture with zero features
- D. An architecture that trusts everything

Answer: B

Q1415. What is a supply chain attack in software security?

- A. Attacking a shipping company
- B. An attack that compromises a trusted third-party component, library, or tool in the software supply chain to inject malicious code into dependent systems
- C. Attacking a grocery store chain
- D. A logistics attack

Answer: B

Q1416. What is the concept of 'security by design' (Secure SDLC)?

- A. Designing secure buildings
- B. Integrating security practices and considerations into every phase of the software development lifecycle rather than adding them as an afterthought
- C. Designing security guard uniforms
- D. A type of physical security

Answer: B

Q1417. What is the concept of 'least privilege escalation paths' in security analysis?

- A. The shortest path to a promotion
- B. Analyzing how an attacker with minimal initial access could chain together vulnerabilities or misconfigurations to escalate their privileges to a higher level
- C. Walking the shortest path in an office
- D. A career development path

Answer: B

Q1418. What is the concept of 'attack surface' in security engineering?

- A. The physical surface of a computer
- B. The total set of points in a system where an attacker could potentially enter, extract data, or cause damage
- C. A type of computer monitor
- D. A mouse pad surface

Answer: B

Q1419. What is homomorphic encryption?

- A. Encryption that makes all data look the same
- B. A form of encryption that allows computations to be performed on encrypted data without decrypting it first, preserving privacy during processing
- C. Simple symmetric encryption
- D. A type of hash function

Answer: B

Q1420. What is the concept of 'secure coding guidelines' like CERT and CWE?

- A. Guidelines for physical building security
- B. Standardized catalogs of coding vulnerabilities (CWE) and best practices (CERT) that help developers write code resistant to common security flaws
- C. Guidelines for writing fast code
- D. Guidelines for team management

Answer: B

Q1421. What is the concept of 'DevSecOps' and how does it differ from traditional security approaches?

- A. A type of firewall
- B. An approach that integrates security practices directly into the DevOps pipeline, making security automated, continuous, and a shared responsibility rather than a late-stage gate
- C. A developer social network
- D. A security certification

Answer: B

Q1422. What is the Unified Software Development Process and how does it differ from the Rational Unified Process (RUP)?

- A. The Unified Process is a generic framework while RUP is a specific commercial implementation with detailed guidance and tools
- B. RUP is a generic framework while the Unified Process is a commercial product
- C. They are identical processes with different names used by different companies
- D. The Unified Process supports only agile methods while RUP supports only plan-driven methods

Answer: A

Q1423. How does the concept of 'technical debt' influence the selection and adaptation of software process models over time?

- A. Technical debt has no effect on process model selection since it is a coding-level concern
- B. Process models should never be changed once selected regardless of technical debt
- C. Accumulated technical debt may force teams to adopt more structured processes that include explicit refactoring and quality improvement phases
- D. Technical debt only affects waterfall projects and has no impact on iterative processes

Answer: C

Q1424. How does the problem frames approach by Michael Jackson differ from traditional requirements analysis methods?

- A. It focuses on identifying recurring problem structures in the real world rather than modeling solutions directly
- B. It eliminates the need for stakeholder involvement in requirements
- C. It is identical to use case modeling but with different terminology
- D. It only works for web-based applications and cannot be applied to embedded systems

Answer: A

Q1425. What role does requirements prioritization using the Wiegers' method play in managing conflicting stakeholder demands?

- A. It eliminates all conflicts by automatically removing low-priority requirements
- B. It uses relative weighting of benefit, penalty, cost, and risk to create an objective priority score that facilitates negotiation
- C. It assigns priority based solely on the seniority of the stakeholder who proposed the requirement
- D. It defers all prioritization decisions until after implementation is complete

Answer: B

Q1426. How does Little's Law apply to software project management for predicting delivery timelines?

- A. It states that adding more people always reduces delivery time proportionally
- B. It establishes that average cycle time equals work in progress divided by throughput, helping teams predict delivery based on current workflow
- C. It proves that all software projects will exceed their initial estimates by exactly 50%
- D. It only applies to manufacturing and has no relevance to software projects

Answer: B

Q1427. What is the Flyweight design pattern and when is it most beneficial?

- A. A pattern that ensures a class has only one instance with global access
- B. A pattern that uses sharing to support large numbers of fine-grained objects efficiently by separating intrinsic and extrinsic state
- C. A pattern that defines a skeleton algorithm in a base class with steps deferred to subclasses
- D. A pattern for creating families of related objects without specifying concrete classes

Answer: B

Q1428. How does the concept of 'connascence' extend the traditional understanding of coupling in software design?

- A. Connascence is simply another name for coupling with no additional insight
- B. Connascence categorizes coupling into static and dynamic forms with varying strengths, providing a more nuanced framework for evaluating design quality
- C. Connascence eliminates the need to consider cohesion in design evaluation
- D. Connascence only applies to functional programming and not object-oriented design

Answer: B

Q1429. What is the concept of 'architectural quanta' in evolutionary architecture and how does it affect deployment?

- A. It refers to the smallest unit of source code that can be compiled independently
- B. It is an independently deployable artifact with high functional cohesion and synchronous connascence, defining the boundary for architectural characteristics
- C. It is a measurement unit for counting the number of microservices in an architecture
- D. It is a quantum computing concept that has no relevance to software architecture

Answer: B

Q1430. How does the KLM-GOMS (Keystroke-Level Model) predict user task completion time?

- A. It measures the actual time users take and averages the results across a sample group
- B. It sums estimated times for basic operations including keystrokes, pointing, homing, drawing, mental preparation, and system response
- C. It uses machine learning algorithms trained on historical user data
- D. It relies solely on subjective user satisfaction surveys

Answer: B

Q1431. How does the concept of 'satisficing' in UI design differ from 'optimizing' and what are its implications?

- A. Satisficing means users always find the best option while optimizing means they settle for less
- B. Users often choose the first option that meets their minimum criteria rather than evaluating all options for the best one, which means interfaces should make good-enough options easy to find
- C. Satisficing and optimizing are identical concepts in UI design
- D. Satisficing only applies to expert users while novice users always optimize their choices

Answer: B

Q1432. What is concolic testing and how does it combine concrete and symbolic execution?

- A. It is a form of manual testing that combines interviews and observations
- B. It simultaneously performs concrete execution with actual inputs and symbolic execution with symbolic values to systematically explore execution paths
- C. It tests only the user interface using both automated and manual approaches
- D. It is an obsolete testing technique replaced entirely by fuzz testing

Answer: B

Q1433. What is the 'all-pairs' or pairwise testing technique and why is it effective for combinatorial testing?

- A. It requires testing every possible combination of all input parameters exhaustively
- B. It ensures that every pair of input parameter values is tested together at least once, significantly reducing test cases while catching most interaction faults
- C. It tests only two input parameters and ignores all others
- D. It is a technique where two testers review the same test case for accuracy

Answer: B

Q1434. What is the Orthogonal Defect Classification (ODC) system and how does it improve quality analysis?

- A. It is a tool for automatically fixing defects in source code
- B. It classifies defects along independent dimensions (type, trigger, impact) to provide actionable feedback about the development process
- C. It is a testing technique that replaces all other forms of quality assurance
- D. It measures only the number of defects without any classification

Answer: B

Q1435. How does the concept of 'quality attributes trade-off' affect architectural decisions in quality assurance?

- A. Quality attributes never conflict with each other, so trade-offs are unnecessary
- B. Improving one quality attribute (e.g., security) can negatively impact another (e.g., performance), requiring deliberate architectural trade-off analysis
- C. Trade-offs only apply during testing and have no effect on architecture
- D. All quality attributes can be maximized simultaneously with sufficient budget

Answer: B

Q1436. What is the Response for a Class (RFC) metric in the CK metrics suite and what are its implications?

- A. It counts the number of HTTP responses a web class can generate
- B. It counts the set of methods that can potentially be executed in response to a message received by an object of the class, indicating testing complexity
- C. It measures the average response time of a class in milliseconds
- D. It counts the number of comment lines in a class file

Answer: B

Q1437. How does the concept of 'metric entropy' relate to software measurement programs?

- A. Metric entropy means all metrics become more accurate over time as data accumulates
- B. As measurement programs mature, collected metrics can become stale, gamed, or lose their diagnostic power, requiring periodic reassessment of what is measured and why
- C. Metric entropy is a thermodynamic law that has no application to software engineering
- D. It ensures that all metrics remain stable and never need to be updated or replaced

Answer: B

Q1438. How does the concept of 'black swan events' challenge traditional risk management approaches in software projects?

- A. Black swan events are easily predicted by standard risk identification techniques
- B. These are highly improbable, high-impact events that lie outside normal expectations and cannot be predicted by historical data, requiring resilience-based rather than prediction-based risk strategies
- C. Black swan events only affect financial industries and have no relevance to software
- D. Traditional risk management already fully accounts for black swan events through contingency planning

Answer: B

Q1439. What is the difference between inherent risk and residual risk, and why is this distinction critical for risk management?

- A. Inherent and residual risk are synonymous terms used interchangeably in all contexts
- B. Inherent risk is the risk level before any controls or mitigations are applied, while residual risk is the remaining risk after mitigation measures are implemented
- C. Residual risk is always zero after proper mitigation strategies are applied
- D. Inherent risk only applies to new projects while residual risk applies only to maintenance projects

Answer: B

Q1440. What is the DREAD risk assessment model and how does it quantify security threats?

- A. DREAD is a psychological assessment of how much developers fear certain vulnerabilities
- B. DREAD rates threats on five dimensions: Damage potential, Reproducibility, Exploitability, Affected users, and Discoverability, each scored to calculate an overall threat rating
- C. DREAD is a network protocol for secure data transmission
- D. DREAD is a testing framework that automatically detects all security vulnerabilities

Answer: B

SUBJECT

Web Development

1,530 Multiple Choice Questions

Web Development — Easy

510 questions

Q1. What does HTML stand for?

- A. Hyper Transfer Markup Language
- B. High Tech Modern Language
- C. Home Tool Markup Language
- D. Hyper Text Markup Language

Answer: D

Q2. Which language is primarily used for styling web pages?

- A. HyperText Markup
- B. JavaScript Logic
- C. Python Scripting
- D. Cascading Styles

Answer: D

Q3. What is a web browser?

- A. Software that displays web pages
- B. A system for managing databases
- C. A tool for compiling source code
- D. A backend application framework

Answer: A

Q4. Which of the following is a web browser?

- A. Google Chrome
- B. Apache HTTP
- C. MySQL Shell
- D. Node.js CLI

Answer: A

Q5. What does URL stand for?

- A. Uniform Reference Locator
- B. Uniform Resource Locator
- C. Universal Resource Locator
- D. Universal Reference Link

Answer: B

Q6. Which language adds interactivity to web pages?

- A. HTML markup
- B. JavaScript
- C. SQL queries
- D. CSS styling

Answer: B

Q7. What is a website?

- A. A single standalone web page
- B. A program that runs web servers
- C. A plugin installed in browsers
- D. A collection of related web pages

Answer: D

Q8. What is the role of a web server?

- A. To browse and view the internet
- B. To write and compile source code
- C. To design visual page layouts
- D. To store and serve web pages

Answer: D

Q9. Which file extension is used for HTML files?

- A. .py
- B. .html
- C. .js
- D. .css

Answer: B

Q10. What is front-end development?

- A. Configuring the network settings
- B. Managing and querying databases
- C. Building the user-visible interface
- D. Programming logic on the server

Answer: C

Q11. What does HTTP stand for?

- A. Hyper Text Transport Protocol
- B. High Tech Transfer Protocol
- C. Hyper Text Transfer Protocol
- D. Home Transfer Text Protocol

Answer: C

Q12. What is a client in web architecture?

- A. The web server that hosts content
- B. The database that stores records
- C. The network cable carrying signals
- D. The device or browser making requests

Answer: D

Q13. What is the default port for HTTP?

- A. Port 443
- B. Port 21
- C. Port 8080
- D. Port 80

Answer: D

Q14. What does DNS stand for?

- A. Dynamic Network Server
- B. Data Network Service
- C. Digital Naming Standard
- D. Domain Name System

Answer: D

Q15. Which HTTP method is used to retrieve data from a server?

- A. POST
- B. DELETE
- C. PUT
- D. GET

Answer: D

Q16. What is an IP address?

- A. A human-readable domain website name
- B. A server-side scripting language for web
- C. A software application like a web browser
- D. A numerical label assigned to each device

Answer: D

Q17. What does the status code 200 mean in HTTP?

- A. OK (Success)
- B. Not Found Error
- C. Redirect Response
- D. Internal Server Error

Answer: A

Q18. What is the role of a web server in the client-server model?

- A. To store only the user password data
- B. To forward requests to other clients
- C. To process requests and send responses
- D. To display web pages to the end user

Answer: C

Q19. Which protocol is used for sending emails?

- A. SMTP
- B. DNS
- C. HTTP
- D. FTP

Answer: A

Q20. What is a web API?

- A. An interface for applications to communicate
- B. A CSS framework for responsive design
- C. A relational database management system
- D. An HTML element used for page layout

Answer: A

Q21. Which HTML tag is used to create a paragraph?

- A. <pg>
- B. <para>
- C. <p>
- D. <text>

Answer: C

Q22. What is the correct HTML tag for the largest heading?

- A. <head> tag
- B. <heading>
- C. <h6> tag
- D. <h1> tag

Answer: D

Q23. Which HTML tag is used to create a hyperlink?

- A. <a>
- B. <url>
- C. <link>
- D. <href>

Answer: A

Q24. Which attribute specifies the URL in an anchor tag?

- A. href
- B. url
- C. link
- D. src

Answer: A

Q25. What is the correct HTML tag for inserting an image?

- A. <photo>
- B. <image>
- C. <pic>
- D.

Answer: D

Q26. Which HTML tag creates an unordered list?

- A. <list>
- B.
- C.
- D.

Answer: B

**Q27. What does the
 tag do?**

- A. Creates a border
- B. Makes text bigger
- C. Creates bold text
- D. Inserts a line break

Answer: D

Q28. Which HTML tag makes text bold?

- A. Only the tag
- B. Only the tag
- C. The <bold> tag only
- D. Both and

Answer: D

Q29. What is the root element of an HTML document?

- A. <html>
- B. <head>
- C. <doctype>
- D. <body>

Answer: A

Q30. Which tag defines a table row in HTML?

- A. <td> tag
- B. <tr> tag
- C. <th> tag
- D. <table> tag

Answer: B

Q31. What does CSS stand for?

- A. Colorful Style Sheets
- B. Computer Style Sheets
- C. Creative Style System
- D. Cascading Style Sheets

Answer: D

Q32. Which CSS property changes the text color?

- A. text-color
- B. color
- C. font-color
- D. fg-color

Answer: B

Q33. How do you select an element with the ID 'header' in CSS?

- A. #header
- B. header
- C. .header
- D. *header

Answer: A

Q34. Which CSS property sets the background color?

- A. color-property
- B. back-color-css
- C. background-color
- D. bgcolor-attr

Answer: C

Q35. How do you select elements with the class 'active' in CSS?

- A. active
- B. *active
- C. #active
- D. .active

Answer: D

Q36. Which property changes the font size?

- A. text-font
- B. text-size
- C. font-size
- D. text-style

Answer: C

Q37. What is the correct CSS syntax to change font to Arial?

- A. text-font: Arial;
- B. font-family: Arial;
- C. family-font: Arial;
- D. font: Arial;

Answer: B

Q38. Which property adds space inside an element's border?

- A. padding
- B. spacing
- C. margin
- D. gap

Answer: A

Q39. Which property adds space outside an element's border?

- A. spacing
- B. margin
- C. padding
- D. outer-space

Answer: B

Q40. How do you add a comment in CSS?

- A. <!-- comment -->
- B. /* comment */
- C. // comment
- D. # comment

Answer: B

Q41. What is CSS Flexbox?

- A. A JavaScript library for DOM element manipulation
- B. A CSS keyframe animation type for effects
- C. A third-party CSS framework for styling pages
- D. A layout model for one-dimensional rows or columns

Answer: D

Q42. Which CSS property creates a flex container?

- A. display: flex
- B. flex: true
- C. layout: flex
- D. container: flex

Answer: A

Q43. What is CSS Grid?

- A. A JavaScript grid plugin for data tables
- B. A third-party CSS component library
- C. A replacement exclusively for HTML tables
- D. A two-dimensional layout for rows and columns

Answer: D

Q44. Which property defines grid columns?

- A. column-template
- B. grid-template-columns
- C. display-columns
- D. grid-columns

Answer: B

Q45. What is a media query in CSS?

- A. A query for fetching database records
- B. A technique applying styles based on device traits
- C. An HTML element for embedding external content
- D. A JavaScript API for browser feature detection

Answer: B

Q46. Which property controls the direction of flex items?

- A. flex-align
- B. flex-direction
- C. flex-orient
- D. flex-flow

Answer: B

Q47. What does justify-content do in Flexbox?

- A. Distributes space between items along main axis
- B. Sets the font family for text inside items
- C. Aligns items along the cross axis vertically
- D. Changes the text alignment within flex items

Answer: A

Q48. What is a CSS transition?

- A. A JavaScript event triggered by user action
- B. A smooth change between CSS property values
- C. An HTML animation tag for visual effects
- D. A page redirect to a new URL location

Answer: B

Q49. Which property creates a CSS animation?

- A. transition
- B. animation
- C. motion
- D. animate

Answer: B

Q50. What does align-items do in Flexbox?

- A. Sets item width
- B. Aligns items along the main axis
- C. Aligns items along the cross axis
- D. Changes item order

Answer: C

Q51. Which keyword declares a variable in JavaScript?

- A. variable
- B. var
- C. declare
- D. let

Answer: B

Q52. What is the correct way to write a single-line comment in JavaScript?

- A. <!-- comment -->
- B. /* comment */
- C. // comment
- D. # comment

Answer: C

Q53. Which operator is used for strict equality in JavaScript?

- A. ===
- B. ==
- C. =
- D. !=

Answer: A

Q54. What does console.log() do?

- A. Creates a modal dialog box on screen
- B. Outputs a message to the browser console
- C. Logs the user into a website account
- D. Creates a new log file on disk

Answer: B

Q55. Which data type represents true or false?

- A. Number
- B. Undefined
- C. String
- D. Boolean

Answer: D

Q56. How do you create a function in JavaScript?

- A. create myFunc() {}
- B. def myFunc() {}
- C. func myFunc() {}
- D. function myFunc() {}

Answer: D

Q57. What is an array in JavaScript?

- A. An ordered collection of values
- B. A type of reusable function
- C. A single primitive value in memory
- D. A CSS property declaration

Answer: A

Q58. How do you write an if statement in JavaScript?

- A. if (x > 5) {}
- B. if x > 5 {}
- C. when (x > 5) {}
- D. if x > 5 then

Answer: A

Q59. Which method adds an element to the end of an array?

- A. add()
- B. insert()
- C. append()
- D. push()

Answer: D

Q60. What does the typeof operator return?

- A. The CSS type selector for an HTML element
- B. A string indicating the data type of a value
- C. The file type extension of a source file
- D. The element type in an HTML document

Answer: B

Q61. What is a Promise in JavaScript?

- A. A looping construct for iterating over arrays
- B. A type of variable for storing constant values
- C. A guarantee of optimal performance metrics
- D. An object representing an async operation result

Answer: D

Q62. What keyword makes a function asynchronous?

- A. defer
- B. await
- C. promise
- D. async

Answer: D

Q63. What does the await keyword do?

- A. Stops the program from running permanently
- B. Pauses async execution until a Promise settles
- C. Creates a timed delay between code statements
- D. Waits for CSS stylesheets to fully load

Answer: B

Q64. What is the spread operator in JavaScript?

- A. The ... operator expanding iterables into elements
- B. The => operator for defining arrow function bodies
- C. The ** operator for exponentiation calculations
- D. The -> operator for creating arrow function syntax

Answer: A

Q65. What is an arrow function?

- A. A function that draws arrow graphics on canvas
- B. A built-in math function for calculations
- C. A recursive function that calls itself repeatedly
- D. A shorter syntax for writing functions using =>

Answer: D

Q66. What does fetch() do in JavaScript?

- A. Makes HTTP requests to retrieve network resources
- B. Downloads files directly to the local disk
- C. Imports CSS files into a JavaScript module
- D. Fetches DOM elements by their CSS selectors

Answer: A

Q67. What is a JavaScript module?

- A. A debugging tool for inspecting runtime state
- B. A browser plugin for extending functionality
- C. A testing framework for unit test execution
- D. A file that exports and imports reusable code

Answer: D

Q68. Which keyword is used to export from a module?

- A. output
- B. export
- C. module
- D. send

Answer: B

Q69. What is the rest parameter in JavaScript?

- A. A sleep function that pauses code execution
- B. Syntax (...args) collecting remaining arguments
- C. A default parameter value for function calls
- D. A parameter specifically for REST API calls

Answer: B

Q70. What is try...catch used for?

- A. Testing CSS styles in the browser
- B. Importing modules from external files
- C. Handling errors and exceptions in code
- D. Creating loops that iterate over data

Answer: C

Q71. What is React?

- A. A CSS framework for responsive page design
- B. A relational database management system tool
- C. A server-side backend application framework
- D. A JavaScript library for building user interfaces

Answer: D

Q72. What is a component in React?

- A. A reusable piece of UI with its own logic
- B. A server route for handling requests
- C. A CSS class for styling HTML elements
- D. A database table for storing records

Answer: A

Q73. What is Vue.js?

- A. A video player for streaming content
- B. A testing tool for browser automation
- C. A CSS library for animation effects
- D. A progressive framework for building UIs

Answer: D

Q74. What is Angular?

- A. A distributed version control system tool
- B. A CSS grid system for page layout design
- C. A TypeScript-based web framework by Google
- D. A NoSQL document-oriented database system

Answer: C

Q75. What is JSX?

- A. A new standalone programming language
- B. A CSS preprocessor like Sass or Less
- C. A syntax extension for HTML-like code in JS
- D. A JSON alternative data exchange format

Answer: C

Q76. What are props in React?

- A. Database properties for configuring schemas
- B. Browser settings for rendering preferences
- C. Data passed from parent to child components
- D. CSS properties applied to styled elements

Answer: C

Q77. What is state in React?

- A. A geographical US state name string value
- B. An object holding data that triggers re-renders
- C. A CSS state like hover or focus pseudo-class
- D. A server state for managing session data

Answer: B

Q78. What is npm?

- A. A server-side programming language runtime
- B. Node Package Manager for managing dependencies
- C. A CSS framework for responsive web design
- D. A relational database management system

Answer: B

Q79. What does npm install do?

- A. Installs CSS styles into browser cache
- B. Installs dependencies from package.json
- C. Installs a new browser on the system
- D. Installs a new operating system update

Answer: B

Q80. What is a Single Page Application (SPA)?

- A. A mobile app built with native platform-specific tools
- B. A static HTML page with no interactive features at all
- C. A web app that loads once and dynamically updates content
- D. A website containing only a single paragraph of text

Answer: C

Q81. What is a server?

- A. A computer providing services to clients
- B. A type of relational database system
- C. A CSS framework for styling web layouts
- D. A web browser for viewing internet pages

Answer: A

Q82. What is Node.js?

- A. A JavaScript runtime on V8 for server-side code
- B. A CSS preprocessor for advanced styling
- C. A NoSQL document-oriented database system
- D. A frontend JavaScript UI framework library

Answer: A

Q83. What is an API endpoint?

- A. A CSS selector for targeting page elements
- B. A database connection string configuration
- C. The final termination point of an API
- D. A specific URL where an API receives requests

Answer: D

Q84. What is Express.js?

- A. A frontend library for building user interfaces
- B. A CSS framework for responsive page layouts
- C. A minimal web application framework for Node.js
- D. A NoSQL database for document-based storage

Answer: C

Q85. What is a database?

- A. An organized collection of structured data
- B. A CSS file for styling web page elements
- C. A web server for hosting application files
- D. A programming language for writing scripts

Answer: A

Q86. What does CRUD stand for?

- A. Create, Read, Update, Delete
- B. Compile, Run, Upload, Debug
- C. Connect, Retrieve, Upload, Download
- D. Cache, Route, Update, Deploy

Answer: A

Q87. What is middleware in backend development?

- A. Software between request and response for processing
- B. A CSS framework for building page component layouts
- C. A frontend component for rendering user interfaces
- D. A type of relational database engine system

Answer: A

Q88. What is JSON?

- A. A relational database management system tool
- B. JavaScript Object Notation for data interchange
- C. A CSS extension for advanced style features
- D. A server-side programming language runtime

Answer: B

Q89. What is a route in backend development?

- A. A URL pattern mapped to a handler function
- B. A database query for fetching records
- C. A CSS path for selecting nested elements
- D. A physical network cable routing path

Answer: A

Q90. What is the purpose of a package.json file?

- A. To store database records in JSON format
- B. To configure CSS preprocessor build steps
- C. To style web pages with CSS declarations
- D. To define project metadata and dependencies

Answer: D

Q91. What is server-side rendering?

- A. A database operation for aggregating records
- B. A CSS technique for optimizing visual layout
- C. Rendering CSS on the server before delivery
- D. Generating HTML on the server before sending it

Answer: D

Q92. Which language is commonly used for server-side programming?

- A. HTML
- B. Python
- C. CSS
- D. SVG

Answer: B

Q93. What is PHP?

- A. A relational database management system
- B. A CSS framework for responsive web layouts
- C. A server-side scripting language for web dev
- D. A JavaScript library for building interfaces

Answer: C

Q94. What is Django?

- A. A CSS framework for responsive component design
- B. A NoSQL database for storing document collections
- C. A high-level Python web framework for rapid dev
- D. A JavaScript library for building front-end UIs

Answer: C

Q95. What is a template engine?

- A. A database tool for managing schema migrations
- B. A JavaScript framework for reactive UI development
- C. Software combining templates with data for HTML
- D. A CSS generator for building style sheets

Answer: C

Q96. What is the purpose of a controller in server-side MVC?

- A. To control the browser rendering engine
- B. To control CSS styles of page elements
- C. To handle requests, process data, and respond
- D. To control the database schema directly

Answer: C

Q97. What is a server-side session?

- A. Data stored on server to maintain user state
- B. A database backup for disaster recovery
- C. A CSS animation effect with keyframes
- D. A JavaScript variable in the global scope

Answer: A

Q98. What is Flask?

- A. A JavaScript runtime built on V8 engine
- B. A CSS container element for layouts
- C. A NoSQL database for document storage
- D. A lightweight Python web framework for apps

Answer: D

Q99. What is a request object in server-side programming?

- A. A database query for retrieving record sets
- B. A CSS class for styling specific elements
- C. An object containing incoming HTTP request data
- D. A JavaScript event fired by user interactions

Answer: C

Q100. What does npm start typically do?

- A. Starts CSS compilation for the build process
- B. Opens a browser window to the application URL
- C. Creates a new project with boilerplate code
- D. Runs the start script in package.json for server

Answer: D

Q101. What is MySQL?

- A. A CSS framework for responsive web page layouts
- B. A web server for hosting application endpoints
- C. An open-source relational database management system
- D. A JavaScript library for building web interfaces

Answer: C

Q102. What is MongoDB?

- A. A NoSQL document database storing JSON-like data
- B. A CSS framework for building styled components
- C. A web server for serving static content files
- D. A SQL relational database using table schemas

Answer: A

Q103. What does SQL stand for?

- A. Simple Query Language
- B. Structured Query Language
- C. Server Query Logic
- D. Standard Question Language

Answer: B

Q104. Which SQL command retrieves data?

- A. RETRIEVE
- B. FETCH
- C. GET
- D. SELECT

Answer: D

Q105. What is a primary key?

- A. The first column created in a database table
- B. A password used for database authentication
- C. A unique identifier for each record in a table
- D. A foreign key linking to another related table

Answer: C

Q106. What SQL command inserts new data?

- A. CREATE INTO
- B. ADD INTO
- C. INSERT INTO
- D. PUT INTO

Answer: C

Q107. What is a table in a relational database?

- A. A JavaScript object for storing key-value pairs
- B. A CSS grid for laying out page components
- C. An HTML table element for displaying data
- D. A structured collection of rows and columns

Answer: D

Q108. What is PostgreSQL?

- A. A CSS preprocessor for advanced style features
- B. An advanced open-source relational database system
- C. A JavaScript framework for building web apps
- D. A web browser for viewing internet content

Answer: B

Q109. What SQL command deletes data?

- A. DESTROY
- B. DELETE
- C. DROP DATA
- D. REMOVE

Answer: B

Q110. What is a database schema?

- A. The structure defining data organization and relations
- B. An HTML template for rendering dynamic page content
- C. A JavaScript object for storing config data
- D. A CSS layout blueprint for page structure

Answer: A

Q111. What is HTTPS?

- A. HTTP with encryption using TLS/SSL for security
- B. A database protocol for encrypted connections
- C. A faster version of HTTP for static content
- D. A programming language for server applications

Answer: A

Q112. What is a password hash?

- A. A one-way transformation into a fixed-length string
- B. A password hint for helping users remember
- C. A password reset link sent via email notification
- D. An encrypted password stored with a cipher key

Answer: A

Q113. What is XSS (Cross-Site Scripting)?

- A. A JavaScript framework for building secure apps
- B. A CSS extension for adding custom properties
- C. A vulnerability where malicious scripts are injected
- D. A database attack targeting stored procedures

Answer: C

Q114. What is a firewall?

- A. A NoSQL database for storing security records
- B. A JavaScript library for client-side validation
- C. A CSS border property for styling elements
- D. A system monitoring and controlling network traffic

Answer: D

Q115. What is authentication?

- A. A JavaScript function for data validation
- B. The process of verifying a user's identity
- C. A CSS technique for styling form elements
- D. Authorizing access to protected resources

Answer: B

Q116. What is a cookie in web security?

- A. A CSS property for styling decorated elements
- B. A small data piece stored for session management
- C. A browser game for entertainment purposes
- D. A JavaScript variable stored in local memory

Answer: B

Q117. What does SSL stand for?

- A. Secure Socket Layer
- B. Simple Security Layer
- C. Secure Server Link
- D. Server Side Logic

Answer: A

Q118. What is a CAPTCHA?

- A. A test distinguishing humans from automated bots
- B. A type of database for storing user accounts
- C. A JavaScript function for validating form data
- D. A CSS framework for building secure page layouts

Answer: A

Q119. What is two-factor authentication (2FA)?

- A. Using two passwords for login authentication
- B. Using two browsers for testing compatibility
- C. Having two accounts for backup and recovery
- D. A method requiring two forms of identification

Answer: D

Q120. Why should user input be validated?

- A. To improve SEO ranking in search engine results
- B. To improve CSS styling across different browsers
- C. To speed up the database query execution time
- D. To prevent malicious data and protect against attacks

Answer: D

Q121. What is web performance optimization?

- A. A CSS framework for building optimized page layouts
- B. A JavaScript library for animation performance
- C. Making websites look visually better with design
- D. Techniques to make websites load and respond faster

Answer: D

Q122. What is caching?

- A. Storing data copies for faster subsequent access
- B. A type of database for key-value storage
- C. A payment method for online transactions
- D. A CSS property for visual element styling

Answer: A

Q123. What is image compression?

- A. Permanently deleting images from the server
- B. Reducing image file size while keeping quality
- C. A JavaScript function for image manipulation
- D. A CSS filter for applying visual image effects

Answer: B

Q124. What is minification?

- A. A CSS technique for reducing selector specificity
- B. Making fonts smaller on the rendered page
- C. Removing unnecessary characters without changing function
- D. A database optimization for reducing table sizes

Answer: C

Q125. What is lazy loading?

- A. A database query that runs in the background
- B. A CSS animation with a slow timing function
- C. Deferring non-critical resource loading until needed
- D. Loading pages intentionally at a slower speed

Answer: C

Q126. What is a CDN used for in web performance?

- A. A database network for distributed data storage
- B. Creating dynamically generated HTML pages
- C. Distributing content across global servers for speed
- D. A CSS framework for building responsive layouts

Answer: C

Q127. Why are large images bad for web performance?

- A. They look bad on high-resolution displays
- B. They break CSS layouts on mobile devices
- C. They increase page load time and bandwidth use
- D. They cause JavaScript runtime errors in browsers

Answer: C

Q128. What is page load time?

- A. A server uptime metric for availability
- B. The time to write the application code
- C. The total time for a web page to fully load
- D. A CSS duration for transition animations

Answer: C

Q129. What does the async attribute do on a script tag?

- A. Makes CSS load asynchronously in parallel
- B. Downloads script without blocking HTML parsing
- C. Makes the script execute synchronously in order
- D. Delays the script loading permanently from page

Answer: B

Q130. What is browser caching?

- A. Storing resources locally for faster repeat access
- B. A JavaScript method for storing session data
- C. A CSS property for caching computed style values
- D. A browser plugin for extending capabilities

Answer: A

Q131. What is unit testing?

- A. Testing individual functions of code in isolation
- B. Testing the entire application end to end
- C. Testing only the database layer and queries
- D. Testing CSS styles across different browsers

Answer: A

Q132. What is a bug in software?

- A. A planned feature in the software design
- B. A CSS property that is incorrectly applied
- C. An error in code causing unexpected behavior
- D. A computer virus that damages system files

Answer: C

Q133. What is the browser developer console used for?

- A. Writing CSS styles for page elements
- B. Managing database records and schemas
- C. Debugging JavaScript and inspecting elements
- D. Designing web page visual layouts

Answer: C

Q134. What is console.log() used for in debugging?

- A. Creating a user activity log in the database
- B. Outputting values to the console for inspection
- C. Sending logs to the remote server for storage
- D. Creating log files on the server disk

Answer: B

Q135. What is a test case?

- A. Conditions to verify a feature works correctly
- B. A computer case for housing hardware
- C. A CSS container for grouping styled elements
- D. A JavaScript class for creating object types

Answer: A

Q136. What is manual testing?

- A. Testing with reference to user documentation
- B. Testing only the documentation for accuracy
- C. Testing performed using automated tools
- D. Testing software by hand without automation

Answer: D

Q137. What does a breakpoint do in debugging?

- A. Pauses code execution at a specific line for inspection
- B. Breaks the code permanently so it stops running
- C. Breaks the CSS layout causing elements to misalign
- D. Creates a responsive breakpoint for media queries

Answer: A

Q138. What is the Network tab in browser DevTools for?

- A. Managing network configuration settings in browser
- B. Creating network topology diagrams for documentation
- C. Connecting to WiFi network for internet access
- D. Monitoring HTTP requests, responses, and their timing

Answer: D

Q139. What is regression testing?

- A. Re-testing features after changes to check nothing broke
- B. Testing CSS regression across different browsers
- C. Testing database regression after schema updates
- D. Testing a feature for the very first time

Answer: A

Q140. What is a linter?

- A. A debugger for stepping through code execution
- B. A CSS framework for responsive page layouts
- C. A testing framework for automated unit testing
- D. A tool analyzing code for errors and style violations

Answer: D

Q141. What is web hosting?

- A. A service making a website accessible on the internet
- B. A database service for storing application records
- C. A CSS framework for building web page layouts
- D. A JavaScript library for reactive user interfaces

Answer: A

Q142. What is a domain name?

- A. A server name assigned during initial configuration
- B. A database name for identifying data collections
- C. A human-readable address used to access a website
- D. A numerical IP address for server identification

Answer: C

Q143. What is FTP?

- A. A CSS framework for styled web page components
- B. File Transfer Protocol for transferring files over networks
- C. A JavaScript function for processing user form data
- D. A database query for retrieving and filtering records

Answer: B

Q144. What is GitHub Pages?

- A. A database service for cloud-based data storage
- B. A social media platform for sharing content
- C. A free static site hosting service from GitHub
- D. A JavaScript framework for building web apps

Answer: C

Q145. What is a production environment?

- A. A coding environment for writing application code
- B. A testing environment for verifying bug fixes
- C. The live environment where users access the application
- D. A staging environment for pre-release validation

Answer: C

Q146. What does deploy mean in web development?

- A. To write application source code in an IDE
- B. To debug code and find programming errors
- C. To make an application accessible on a server
- D. To test code for correctness and reliability

Answer: C

Q147. What is Netlify?

- A. A JavaScript framework for building web applications
- B. A cloud platform for deploying sites and serverless
- C. A CSS framework for building responsive web layouts
- D. A NoSQL database for storing document collections

Answer: B

Q148. What is a staging environment?

- A. The production server running the live application
- B. A development computer for writing application code
- C. A CSS staging area for previewing style changes
- D. A pre-production environment mirroring production

Answer: D

Q149. What is version control?

- A. A database version for tracking schema migrations
- B. A CSS versioning system for managing style updates
- C. Controlling which browser version users run
- D. A system tracking and managing code changes over time

Answer: D

Q150. What is a build process?

- A. Transforming source code into optimized deploy files
- B. Building HTML pages manually without any tools
- C. A CSS compilation step for preprocessing styles
- D. Building a physical server room infrastructure

Answer: A

Q151. What is a Web API?

- A. A JavaScript variable for storing temporary data
- B. An interface allowing software to communicate via HTTP
- C. A CSS framework for building styled web layouts
- D. A NoSQL database for storing document records

Answer: B

Q152. What is JSON used for in Web APIs?

- A. Storing CSS variables for theming and customization
- B. Styling web pages with CSS declarations and rules
- C. A lightweight data format for client-server exchange
- D. Creating HTML elements for dynamic page rendering

Answer: C

Q153. What is a GET request?

- A. Getting database access for read operations
- B. Getting CSS styles for an HTML element
- C. An HTTP request that retrieves data without modifying
- D. Getting HTML templates from the server cache

Answer: C

Q154. What is a POST request?

- A. An HTTP request sending data to create a resource
- B. Posting database queries for batch execution
- C. Posting CSS updates to a remote style server
- D. Posting content on a social media platform

Answer: A

Q155. What is an API key?

- A. A physical key for server room access control
- B. A CSS key for targeting specific styled elements
- C. A unique identifier used to authenticate API requests
- D. A keyboard shortcut for developer tool actions

Answer: C

Q156. What does REST stand for?

- A. Remote Execution of Server Tasks
- B. Real-time Event Streaming Technology
- C. Responsive State Transition
- D. Representational State Transfer

Answer: D

Q157. What is an API response?

- A. Data sent back by server after processing a request
- B. A CSS file for styling page components
- C. An HTML page always returned from the server
- D. A JavaScript function for handling async events

Answer: A

Q158. What is Postman used for?

- A. Managing database records through a visual tool
- B. Sending physical mail through the postal service
- C. Testing and debugging APIs by sending requests
- D. Writing CSS stylesheets for web page elements

Answer: C

Q159. What is an HTTP status code?

- A. A JavaScript error code thrown during execution
- B. A database error code from failed query operations
- C. A numerical code indicating the result of a request
- D. A CSS class name applied to styled elements

Answer: C

Q160. What is a query parameter in a URL?

- A. A CSS query for selecting styled page elements
- B. A key-value pair after ? to send data with requests
- C. A database query for filtering stored records
- D. A JavaScript query for searching DOM tree nodes

Answer: B

Q161. What is TypeScript?

- A. A CSS preprocessor for advanced style syntax
- B. A typed superset of JavaScript compiling to JS
- C. A new unrelated programming language for web
- D. A database language for structured query access

Answer: B

Q162. What is a package manager?

- A. A database manager for schema and migration control
- B. A CSS framework for responsive component-based design
- C. A tool automating installing and managing dependencies
- D. A delivery service for shipping physical packages

Answer: C

Q163. What is Sass/SCSS?

- A. A database language for writing structured queries
- B. A web server for hosting static and dynamic content
- C. A JavaScript library for building reactive interfaces
- D. A CSS preprocessor adding variables, nesting, mixins

Answer: D

Q164. What is a linter?

- A. A CSS animation for transition visual effects
- B. A database tool for query optimization and profiling
- C. A testing framework for automated unit test execution
- D. A tool analyzing code to flag errors and enforce standards

Answer: D

Q165. What is Git?

- A. A NoSQL database for storing document collections
- B. A CSS framework for building responsive layouts
- C. A web browser for viewing internet content
- D. A distributed version control system for code

Answer: D

Q166. What is a code formatter?

- A. A CSS beautifier for styling elements only
- B. A database formatter for prettifying query output
- C. A JavaScript compiler for transpiling modern syntax
- D. A tool automatically formatting code to consistent style

Answer: D

Q167. What is Tailwind CSS?

- A. A utility-first CSS framework with low-level classes
- B. A database tool for managing schema migrations
- C. A testing framework for end-to-end browser tests
- D. A JavaScript framework for building user interfaces

Answer: A

Q168. What is a component library?

- A. A CSS file for global application style definitions
- B. A book library for reading documentation offline
- C. A collection of pre-built reusable UI components
- D. A JavaScript runtime for executing server code

Answer: C

Q169. What is Markdown?

- A. A JavaScript library for DOM manipulation events
- B. A lightweight markup language for formatted text
- C. A database query language for structured access
- D. A CSS framework for responsive component styling

Answer: B

Q170. What is an environment variable?

- A. A variable set outside the app configuring behavior
- B. A CSS variable defined with custom property syntax
- C. A database field for storing record attribute values
- D. A JavaScript global variable accessible from any scope

Answer: A

Q171. What does CSS stand for?

- A. Colorful Style Sheets
- B. Cascading Style Sheets
- C. Computer Style Sheets
- D. Creative Style Systems

Answer: B

Q172. What is the purpose of a text editor in web development?

- A. To write and edit source code
- B. To compile server-side programs
- C. To manage database table records
- D. To browse and navigate websites

Answer: A

Q173. Which protocol is used to load web pages?

- A. HTTP protocol
- B. FTP protocol
- C. SSH protocol
- D. SMTP protocol

Answer: A

Q174. What is a web page?

- A. A database query statement
- B. A server configuration file
- C. A single document on the web
- D. A network routing protocol

Answer: C

Q175. What does a back-end developer do?

- A. Designs the visual user interface
- B. Creates graphic assets for the site
- C. Writes server-side application logic
- D. Tests the browser compatibility only

Answer: C

Q176. Which tool is used to inspect web page elements?

- A. Terminal emulator
- B. Browser DevTools
- C. Package installer
- D. Database manager

Answer: B

Q177. What is a hyperlink?

- A. A styling rule for text decoration
- B. A script that validates user input
- C. A method for storing session data
- D. A clickable reference to another page

Answer: D

Q178. What is the main function of DNS?

- A. To validate HTML page structure
- B. To encrypt data during transfer
- C. To compile source code on servers
- D. To translate domain names to IPs

Answer: D

Q179. What does IDE stand for?

- A. Internal Data Encryption
- B. Integrated Development Environment
- C. Interactive Design Editor
- D. Internet Development Engine

Answer: B

Q180. Which language is used for web page structure?

- A. JavaScript
- B. HTML markup
- C. Python lang
- D. Cascading SS

Answer: B

Q181. What does HTTP stand for?

- A. Hyper Terminal Text Program
- B. HyperText Transfer Protocol
- C. High Transfer Text Protocol
- D. Home Text Transfer Protocol

Answer: B

Q182. What is an IP address?

- A. A programming language for building web applications
- B. A unique numerical label for devices on a network
- C. A type of database used to store website content
- D. A visual design tool for creating page mockup files

Answer: B

Q183. What is the default port number for HTTP?

- A. Port 443
- B. Port 80
- C. Port 25
- D. Port 21

Answer: B

Q184. What does HTTPS add to HTTP?

- A. Encryption via SSL/TLS
- B. Faster data transfer
- C. Database connectivity
- D. Image optimization

Answer: A

Q185. What is a client in client-server architecture?

- A. The device that stores all website data files
- B. The protocol that encodes data for safe transfer
- C. The cable that connects computers on a network
- D. The software that requests resources from servers

Answer: D

Q186. Which HTTP method is used to retrieve data?

- A. DELETE method
- B. POST method
- C. PUT method
- D. GET method

Answer: D

Q187. What is a web host?

- A. A browser extension for ad blocking
- B. A language for scripting web pages
- C. A service that stores website files
- D. A tool for debugging CSS layouts

Answer: C

Q188. What does a 404 status code mean?

- A. Resource not found
- B. Server internal error
- C. Redirect to new URL
- D. Successful request

Answer: A

Q189. What is bandwidth in web context?

- A. The number of users currently accessing a specific web server
- B. The maximum data transfer rate of a network connection link
- C. The visual width of a web page displayed in the browser window
- D. The total storage space available on a web hosting server disk

Answer: B

Q190. What is latency in networking?

- A. The delay before data transfer begins between two endpoints
- B. The maximum number of connections a server can handle at once
- C. The total storage capacity available on a remote file server
- D. The encryption strength used to protect data during transfers

Answer: A

Q191. Which HTML tag creates a paragraph?

- A.
 tag
- B. tag
- C. <div> tag
- D. <p> tag

Answer: D

Q192. What does the <a> tag create?

- A. An image element
- B. A paragraph block
- C. A table row item
- D. A hyperlink element

Answer: D

Q193. Which attribute specifies the URL in an anchor tag?

- A. type attribute
- B. alt attribute
- C. src attribute
- D. href attribute

Answer: D

Q194. What is the purpose of the <head> element?

- A. To display the main visible content of the page
- B. To contain metadata and links to external resources
- C. To create a navigation menu across the top of site
- D. To define a footer section at the bottom of a page

Answer: B

Q195. Which tag is used for the largest heading?

- A. <h6> tag
- B. <h2> tag
- C. <h3> tag
- D. <h1> tag

Answer: D

Q196. What does the tag do?

- A. Generates a table layout
- B. Inserts an image element
- C. Creates a hyperlink to page
- D. Defines a text paragraph

Answer: B

Q197. Which HTML tag creates a line break?

- A. <lb> tag
- B.
 tag
- C. <hr> tag
- D. <nl> tag

Answer: B

Q198. What is the purpose of the <title> tag?

- A. To add a tooltip on hover state
- B. To set the browser tab title text
- C. To display a heading on the page
- D. To create bold text in the body

Answer: B

Q199. Which tag creates an unordered list?

- A. tag
- B. tag
- C. tag
- D. <dl> tag

Answer: B

Q200. What does the <body> tag contain?

- A. Visible content of the page
- B. Page metadata and style links
- C. Server configuration settings
- D. External script file references

Answer: A

Q201. Which CSS property changes text color?

- A. background-color
- B. border
- C. color
- D. font-size

Answer: C

Q202. How do you select an element by its ID in CSS?

- A. Using a hash prefix
- B. Using a dot prefix
- C. Using an at symbol
- D. Using a tilde prefix

Answer: A

Q203. Which property sets the background color?

- A. font-color
- B. text-background
- C. color property
- D. background-color

Answer: D

Q204. What does the font-size property control?

- A. The size of the text characters shown
- B. The weight or boldness of the text
- C. The spacing between lines of text
- D. The font family used for the element

Answer: A

Q205. How do you add a CSS comment?

- A. /* comment text */
- B. # comment text here
- C. <!-- comment here -->
- D. // comment text here

Answer: A

Q206. Which property controls element width?

- A. width
- B. size
- C. length
- D. height

Answer: A

Q207. What value makes text bold?

- A. text-weight: bold
- B. text-style: bold
- C. font-weight: bold
- D. font-style: bold

Answer: C

Q208. Which property adds space inside an element's border?

- A. border property
- B. padding property
- C. spacing property
- D. margin property

Answer: B

Q209. How do you select all paragraph elements?

- A. #p selector
- B. p selector
- C. *p selector
- D. .p selector

Answer: B

Q210. Which property controls text alignment?

- A. text-position
- B. align-text
- C. font-align
- D. text-align

Answer: D

Q211. What does a media query do in CSS?

- A. Embeds media files like videos
- B. Creates multimedia animations
- C. Applies styles based on conditions
- D. Queries a media database server

Answer: C

Q212. What CSS property creates a flexible container?

- A. display: flex
- B. display: table
- C. display: inline
- D. display: block

Answer: A

Q213. What unit is relative to the viewport width?

- A. cm unit
- B. em unit
- C. vw unit
- D. px unit

Answer: C

Q214. What does the flex-direction property control?

- A. The direction of flex items
- B. The color of flex items
- C. The opacity of container
- D. The border of flex items

Answer: A

Q215. What is CSS Grid used for?

- A. Adding animations to pages
- B. Creating two-dimensional layouts
- C. Validating form input fields
- D. Styling text fonts and colors

Answer: B

Q216. What does the gap property do in Grid?

- A. Sets the grid background color
- B. Sets the font size of grid text
- C. Sets spacing between grid items
- D. Sets the border of grid items

Answer: C

Q217. What is a breakpoint in responsive design?

- A. A screen width where the layout changes its presentation style
- B. A point where the code stops executing during debugging mode
- C. A network point where data packets are inspected for threats
- D. A JavaScript error that breaks the page rendering completely

Answer: A

Q218. What does justify-content do in flexbox?

- A. Removes all padding and margin from the flex container element
- B. Aligns items vertically along the cross axis of the container
- C. Changes the font justification to monospaced character spacing
- D. Aligns items horizontally along the main axis of the container

Answer: D

Q219. What is the rem unit relative to?

- A. Root element font size
- B. Device pixel density
- C. Parent element font size
- D. Viewport total height

Answer: A

Q220. What does align-items do in flexbox?

- A. It sets the text alignment for content inside each flex item child
- B. It aligns flex items along the cross axis of the flex container
- C. It defines the number of items that can fit on a single flex row
- D. It adds animated alignment transitions between flex item positions

Answer: B

Q221. What keyword declares a variable that can be reassigned?

- A. const keyword
- B. final keyword
- C. var keyword
- D. let keyword

Answer: D

Q222. What is the result of typeof null in JavaScript?

- A. 'boolean' val
- B. 'object' type
- C. 'undefined'
- D. 'null' string

Answer: B

Q223. Which operator checks strict equality?

- A. === operator
- B. != operator
- C. = operator
- D. == operator

Answer: A

Q224. What does console.log() do?

- A. Outputs to the console
- B. Sends data to server
- C. Writes to a log file
- D. Displays an alert box

Answer: A

Q225. How do you write a single-line comment?

- A. <!-- comment -->
- B. # comment here
- C. // comment here
- D. /* comment */

Answer: C

Q226. What is an array in JavaScript?

- A. A CSS property for arranging page elements
- B. A function that processes text strings only
- C. A single numeric value stored in memory
- D. An ordered collection of values in a list

Answer: D

Q227. What does the push() method do to an array?

- A. Adds an element to end
- B. Removes the last element
- C. Sorts elements in order
- D. Reverses array elements

Answer: A

Q228. What is a function in JavaScript?

- A. A reusable block of code that performs a task
- B. A network protocol for exchanging data between nodes
- C. A database query that retrieves stored information
- D. A CSS selector that targets specific HTML elements

Answer: A

Q229. What value does a variable have before assignment?

- A. zero value
- B. undefined
- C. null value
- D. empty string

Answer: B

Q230. Which method converts a string to a number?

- A. parseInt() method
- B. toString() method
- C. concat() method
- D. stringify() method

Answer: A

Q231. What is a Promise in JavaScript?

- A. A guaranteed synchronous operation result value
- B. A CSS animation timing function for elements
- C. A method for declaring global variables only
- D. An object representing a future async outcome

Answer: D

Q232. What keyword makes a function asynchronous?

- A. defer keyword
- B. await keyword
- C. yield keyword
- D. async keyword

Answer: D

Q233. What does the await keyword do?

- A. Pauses async function execution until a Promise settles
- B. Creates a new thread for parallel processing of data
- C. Stops all JavaScript execution on the current web page
- D. Delays code execution by a specific number of milliseconds

Answer: A

Q234. What is the DOM?

- A. A database object model for SQL queries
- B. A design optimization module for CSS
- C. A disk operating mechanism for files
- D. A document object model for web pages

Answer: D

Q235. What does JSON stand for?

- A. JavaScript Object Notation
- B. Joint System Object Network
- C. JavaScript Oriented Namespace
- D. Java Standard Object Naming

Answer: A

Q236. What is an arrow function?

- A. A shorter syntax for writing function expressions
- B. A function that draws arrows on a canvas element
- C. A special function that handles keyboard arrow keys
- D. A function that only works with array operations

Answer: A

Q237. What is the purpose of try-catch?

- A. To test network speed before loading page resources
- B. To test CSS styles before applying them to elements
- C. To handle errors gracefully without crashing the code
- D. To catch mouse click events on specific HTML elements

Answer: C

Q238. What does the fetch API do?

- A. Makes network requests to get or send data
- B. Retrieves CSS styles from a stylesheet file
- C. Downloads and installs browser plugin extensions
- D. Fetches DOM elements by their class name value

Answer: A

Q239. What is localStorage?

- A. A CSS property for element positioning
- B. Client-side persistent key-value storage
- C. A method for caching DNS query results
- D. Server-side database for user data

Answer: B

Q240. What is an event listener?

- A. A function that waits for and responds to a specific event trigger
- B. A CSS property that detects hover and focus states on an element
- C. A server that listens for incoming network connections from clients
- D. A debugging tool that records all console output to a log file

Answer: A

Q241. What is React primarily used for?

- A. Server-side database management
- B. Operating system development work
- C. Network protocol implementation
- D. Building user interface components

Answer: D

Q242. What is a component in frontend frameworks?

- A. A database table configuration file
- B. A reusable piece of user interface
- C. A CSS file that styles a single page
- D. A server route handler for requests

Answer: B

Q243. What is Vue.js?

- A. A mobile operating system app
- B. A database query language spec
- C. A backend server framework
- D. A progressive frontend framework

Answer: D

Q244. What does Angular use for templates?

- A. Python template syntax
- B. Ruby ERB template files
- C. Enhanced HTML with directives
- D. Plain JavaScript strings

Answer: C

Q245. What is JSX in React?

- A. A JSON schema extension for validation
- B. A CSS preprocessor for React styles
- C. A JavaScript testing framework tool
- D. A syntax extension mixing HTML with JS

Answer: D

Q246. What is state in a frontend framework?

- A. Data that determines how a component renders and behaves
- B. The geographic location of the server hosting the application
- C. The CSS styling applied to an element at any given moment
- D. The HTTP status code returned by the server for each request

Answer: A

Q247. What is npm used for?

- A. Managing JavaScript packages
- B. Running database migrations
- C. Compiling TypeScript to JS
- D. Styling web page elements

Answer: A

Q248. What is a Single Page Application?

- A. An app that works only on single monitor setup
- B. A website with only one paragraph
- C. A web page that disables all user scrolling
- D. An app that updates without full page reloads

Answer: D

Q249. What is the virtual DOM?

- A. A database for storing HTML template files
- B. A lightweight copy of the real DOM tree
- C. A physical server that stores DOM data
- D. A CSS framework for virtual reality pages

Answer: B

Q250. What is Svelte?

- A. A compiler-based frontend framework
- B. A backend REST API framework tool
- C. A CSS animation library for web pages
- D. A database migration utility script

Answer: A

Q251. What is a server in web development?

- A. A computer that serves resources
- B. A browser extension for debugging
- C. A CSS framework for styling
- D. A front-end layout component

Answer: A

Q252. What is an API?

- A. A database management system tool
- B. An interface for software communication
- C. A visual design tool for web pages
- D. A type of web browser extension

Answer: B

Q253. What does REST stand for?

- A. Remote Execution Server Tool
- B. Real-time Event Stream Technology
- C. Representational State Transfer
- D. Responsive Element Style Template

Answer: C

Q254. What is a database used for?

- A. Creating user interface layouts
- B. Compiling source code files
- C. Storing and retrieving data
- D. Styling web page elements

Answer: C

Q255. What is Node.js?

- A. A CSS preprocessor tool
- B. A browser testing framework
- C. A JavaScript server runtime
- D. A database query language

Answer: C

Q256. What is a route in backend development?

- A. A URL pattern mapped to a handler
- B. A database table relationship link
- C. A network cable connecting servers
- D. A CSS path for animation movement

Answer: A

Q257. What does CRUD stand for?

- A. Compile, Run, Upload, Deploy
- B. Cache, Route, Update, Debug
- C. Create, Read, Update, Delete
- D. Connect, Render, Update, Drop

Answer: C

Q258. What is middleware in backend development?

- A. The physical hardware between client and server machines in the network
- B. Software that processes requests between receiving and sending response
- C. The CSS styles applied between two animation keyframes in a transition
- D. A database layer that sits between the raw data files and the queries

Answer: B

Q259. What is Express.js?

- A. A database query builder
- B. A CSS animation library
- C. A frontend testing suite
- D. A Node.js web framework

Answer: D

Q260. What format is commonly used for API responses?

- A. HTML format
- B. XML format
- C. CSV format
- D. JSON format

Answer: D

Q261. What is server-side rendering?

- A. Compiling JavaScript into machine code on the server for faster speed
- B. Rendering CSS animations on the server before sending them to client
- C. Processing database queries on the client side instead of the server
- D. Generating HTML on the server before sending it to the browser client

Answer: D

Q262. What is PHP primarily used for?

- A. Mobile app development
- B. Desktop GUI applications
- C. Server-side web scripting
- D. Database engine creation

Answer: C

Q263. What is a session in web development?

- A. A JavaScript testing framework function
- B. A CSS animation sequence that plays once
- C. A way to store user data across requests
- D. A type of database join operation query

Answer: C

Q264. What is a cookie?

- A. A small data piece stored in the browser
- B. A CSS property for styling form inputs
- C. A server configuration file for caching
- D. A database table for user preferences

Answer: A

Q265. What does MVC stand for?

- A. Managed-Virtual-Cache
- B. Main-View-Component
- C. Model-View-Controller
- D. Module-Variable-Class

Answer: C

Q266. What is a template engine?

- A. A tool that generates CSS from design files automatically
- B. A tool that combines templates with data to produce HTML output
- C. A tool that creates database schemas from visual diagram models
- D. A tool that compiles TypeScript into standard JavaScript code

Answer: B

Q267. What is Django?

- A. A database management system
- B. A CSS preprocessor language
- C. A JavaScript runtime environment
- D. A Python-based web framework

Answer: D

Q268. What is Ruby on Rails?

- A. A database visualization platform
- B. A Ruby-based web framework tool
- C. A CSS grid framework for layouts
- D. A JavaScript package manager app

Answer: B

Q269. What is a request-response cycle?

- A. The process of compiling source code and deploying it to a production box
- B. The cycle of backing up and restoring database records on a fixed schedule
- C. The time it takes CSS to render on screen after being parsed by browser
- D. A client sends a request and the server returns a corresponding response

Answer: D

Q270. What does a web framework provide?

- A. Only visual design tools for creating page mockups and wireframes easily
- B. Hardware drivers for connecting web servers to network switch devices
- C. Pre-built tools and structure for building web applications efficiently
- D. Browser extensions that enhance the user experience of any website

Answer: C

Q271. What is SQL?

- A. A server configuration format
- B. A JavaScript testing framework
- C. A styling language for web pages
- D. A language for managing databases

Answer: D

Q272. What is a primary key in a database?

- A. The main index for text search
- B. The first column in every table
- C. The password for database access
- D. A unique identifier for each row

Answer: D

Q273. What does SELECT do in SQL?

- A. Creates new database tables
- B. Deletes specified records
- C. Updates existing row values
- D. Retrieves data from tables

Answer: D

Q274. What is MongoDB?

- A. A CSS styling framework
- B. A JavaScript build tool
- C. A relational SQL database
- D. A NoSQL document database

Answer: D

Q275. What is a table in a relational database?

- A. A JavaScript data structure for key-value pairs used in code
- B. A CSS property that creates grid-like layouts on web page elements
- C. A collection of related data organized in rows and columns format
- D. A server configuration file that maps URLs to handler functions

Answer: C

Q276. What does INSERT do in SQL?

- A. Removes rows from a table
- B. Adds new rows to a table
- C. Searches for text patterns
- D. Modifies existing row data

Answer: B

Q277. What is a foreign key?

- A. A temporary key generated during database backup process
- B. A field that references a primary key in another table
- C. A key used to encrypt database files for security
- D. A key that allows foreign users to access the database

Answer: B

Q278. What is MySQL?

- A. A relational SQL database
- B. A CSS preprocessor tool
- C. A JavaScript framework
- D. A NoSQL document store

Answer: A

Q279. What does UPDATE do in SQL?

- A. Modifies existing data in the database table rows
- B. Deletes the entire database from the server storage
- C. Creates a new table in the database schema file
- D. Retrieves all records matching a specified condition

Answer: A

Q280. What is PostgreSQL?

- A. A command-line tool for managing server deployment pipelines
- B. An advanced open-source relational database management system
- C. A frontend CSS framework for web layout design styling
- D. A JavaScript library for building user interface components

Answer: B

Q281. What is HTTPS?

- A. A CSS styling standard
- B. HTTP with TLS encryption
- C. A faster version of HTML
- D. A database access protocol

Answer: B

Q282. What is a password hash?

- A. A one-way encrypted form of password
- B. A method for generating passwords
- C. A visible password shown on screen
- D. A type of password recovery email

Answer: A

Q283. What is XSS?

- A. A database query optimizer
- B. A server-side language tool
- C. Cross-site scripting attack
- D. A CSS framework for styling

Answer: C

Q284. What is a firewall?

- A. A JavaScript error handling method
- B. A CSS property for styling borders
- C. A database backup recovery process
- D. A system that filters network traffic

Answer: D

Q285. What does SSL stand for?

- A. System Security Level
- B. Server Side Language
- C. Secure Sockets Layer
- D. Standard Style Library

Answer: C

Q286. What is two-factor authentication?

- A. Logging in from two different devices at the same time for redundancy
- B. Using two different browsers to log in to the same account simultaneously
- C. Creating two separate passwords for a single user account on the site
- D. Requiring two forms of identity verification for accessing an account

Answer: D

Q287. What is a CAPTCHA used for?

- A. To style form input fields
- B. To verify a user is human
- C. To encrypt form submissions
- D. To speed up page loading

Answer: B

Q288. What is malware?

- A. A CSS animation framework
- B. A type of web server software
- C. A database storage format
- D. Software designed to cause harm

Answer: D

Q289. What is phishing?

- A. A process for optimizing database queries for faster execution time
- B. A performance testing method for measuring page load speed metrics
- C. A technique for compressing images to reduce file size on servers
- D. A fraudulent attempt to obtain sensitive information by deception

Answer: D

Q290. What is encryption?

- A. Creating backup copies of important files on the server
- B. Deleting sensitive data from the database permanently
- C. Compressing files to reduce their storage size on disk
- D. Converting data into an unreadable format for protection

Answer: D

Q291. What is page load time?

- A. Time to write the page code
- B. Time to deploy the page code
- C. Time for a page to fully display
- D. Time for DNS to resolve domain

Answer: C

Q292. What does minification do?

- A. Encrypts the source code files
- B. Adds comments to the code
- C. Compiles code to machine code
- D. Removes unnecessary characters

Answer: D

Q293. What is caching?

- A. Encrypting data in the database
- B. Storing data for faster future access
- C. Deleting old files from server
- D. Compiling code before deployment

Answer: B

Q294. What does image compression do?

- A. Converts images to HTML format
- B. Adds metadata to image files
- C. Increases image resolution quality
- D. Reduces image file size for web

Answer: D

Q295. What is lazy loading?

- A. Slowing down the server to reduce power usage
- B. Loading resources only on the server side cache
- C. Deferring resource loading until they are needed
- D. Loading all resources immediately on page start

Answer: C

Q296. What does a CDN help with?

- A. Writing server code
- B. Database management
- C. Code compilation speed
- D. Faster content delivery

Answer: D

Q297. What is browser caching?

- A. Encrypting browser data for privacy and security use
- B. Compiling web pages into native apps for offline use
- C. Deleting browser history to free up disk space now
- D. Storing web resources locally in the browser for reuse

Answer: D

Q298. What does the defer attribute do on a script tag?

- A. It loads the script after HTML parsing is complete first
- B. It delays the script execution by a fixed time in seconds
- C. It defers CSS loading until after JavaScript execution
- D. It prevents the script from running in production builds

Answer: A

Q299. What is bandwidth?

- A. The maximum file size a server can store on disk
- B. The number of pages a website has in its sitemap
- C. The visual width of a web page layout in pixels
- D. The data transfer capacity of a network connection

Answer: D

Q300. What is a performance bottleneck?

- A. A CSS property that controls animation speed on a page
- B. A feature that limits the overall performance of a system
- C. A tool used to measure the speed of a web application
- D. A server configuration for limiting maximum user connections

Answer: B

Q301. What is unit testing?

- A. Testing network connectivity and speed
- B. Testing individual components in isolation
- C. Testing the user interface visual layout
- D. Testing the entire application end to end

Answer: B

Q302. What is a bug in software?

- A. A planned software enhancement
- B. A CSS styling improvement idea
- C. An error or defect in the code
- D. A feature request from the user

Answer: C

Q303. What is the purpose of the browser console?

- A. To design database table structures
- B. To view logs, errors, and run JavaScript
- C. To create new web pages from templates
- D. To manage web server configurations

Answer: B

Q304. What does a linter do?

- A. It removes lint from clothing images
- B. It generates test cases from source
- C. It compresses files for deployment
- D. It checks code for style and errors

Answer: D

Q305. What is debugging?

- A. Deploying code to production
- B. Finding and fixing code errors
- C. Writing documentation for code
- D. Adding new features to software

Answer: B

Q306. What is a test case?

- A. A special case statement in JavaScript switch expressions
- B. A container for storing test data in the database table
- C. A set of conditions to verify expected software behavior
- D. A type of briefcase used by quality assurance engineers

Answer: C

Q307. What tool is commonly used for JavaScript testing?

- A. Nginx server
- B. MySQL server
- C. Photoshop
- D. Jest framework

Answer: D

Q308. What is a breakpoint in debugging?

- A. A network point where data packets are split for transfer
- B. A point in code where execution pauses for inspection work
- C. A CSS property that breaks text into multiple lines on page
- D. A screen width where the responsive layout changes its form

Answer: B

Q309. What is integration testing?

- A. Testing individual functions completely isolated from others
- B. Testing the application under extreme load conditions only
- C. Testing how different modules or components work together
- D. Testing CSS integration with HTML elements on the web page

Answer: C

Q310. What is a test suite?

- A. A single test case for one function
- B. A collection of related test cases
- C. A CSS framework for test pages
- D. A hotel suite for QA team members

Answer: B

Q311. What is deployment in web development?

- A. Writing the initial source code for a new feature in the application
- B. Designing the user interface mockups in a visual design tool app
- C. Making a website available on a server for users to access online
- D. Testing the application for bugs before releasing it to the public

Answer: C

Q312. What is a domain name?

- A. A server-side scripting language
- B. A database management system tool
- C. A human-readable web address name
- D. A CSS layout framework for grids

Answer: C

Q313. What does FTP stand for?

- A. File Transfer Protocol
- B. Front Template Page
- C. Full Type Platform
- D. Fast Text Processing

Answer: A

Q314. What is shared hosting?

- A. Hosting where you own the entire physical server exclusively for your site
- B. Hosting where multiple websites share a single server and its resources
- C. Hosting where files are shared across multiple CDN edge server locations
- D. Hosting where developers share code through version control repositories

Answer: B

Q315. What is a SSL certificate?

- A. A certificate showing a website passed all accessibility compliance tests
- B. A certification for web developers from an accredited training institution
- C. A license required to operate a web hosting business in certain countries
- D. A digital certificate that enables HTTPS encryption for a website domain

Answer: D

Q316. What is GitHub Pages?

- A. A backend API framework tool
- B. A static site hosting service
- C. A CSS component library app
- D. A database hosting service

Answer: B

Q317. What is a VPS?

- A. Virtual Private Server
- B. Variable Processing System
- C. Visual Page Styling
- D. Virtual Protocol Service

Answer: A

Q318. What does DNS do?

- A. Compiles JavaScript for production use
- B. Translates domain names to IP addresses
- C. Styles web page elements with colors
- D. Creates database tables automatically

Answer: B

Q319. What is cloud hosting?

- A. Hosting in a weather station
- B. Hosting through email services
- C. Hosting files on a local computer
- D. Hosting on virtual cloud servers

Answer: D

Q320. What is Netlify used for?

- A. Hosting and deploying web applications
- B. Writing server-side Python programs
- C. Managing SQL databases for web apps
- D. Creating mobile applications for iOS

Answer: A

Q321. What is a REST API?

- A. A database query language spec
- B. A CSS framework for styling
- C. A JavaScript testing framework
- D. An API following REST principles

Answer: D

Q322. What HTTP method retrieves data?

- A. DELETE method
- B. PUT method
- C. POST method
- D. GET method

Answer: D

Q323. What is JSON used for in APIs?

- A. Exchanging data between client and server systems
- B. Styling web page elements with colors and fonts
- C. Compiling JavaScript code for production deployment
- D. Creating database tables and defining their schemas

Answer: A

Q324. What does an API endpoint represent?

- A. A JavaScript variable type name
- B. A CSS selector targeting elements
- C. A database table column value
- D. A specific URL for a resource

Answer: D

Q325. What HTTP method creates a new resource?

- A. GET method
- B. DELETE method
- C. PUT method
- D. POST method

Answer: D

Q326. What is an API key?

- A. A database encryption algorithm name
- B. A CSS property for keyboard events
- C. A token for authenticating API requests
- D. A physical key for server room access

Answer: C

Q327. What does HTTP status 200 mean?

- A. OK / Success
- B. Unauthorized
- C. Not found
- D. Server error

Answer: A

Q328. What is a request header?

- A. A database table column header
- B. The first line of HTML on a page
- C. A CSS header styling property
- D. Metadata sent with an HTTP request

Answer: D

Q329. What does CORS stand for?

- A. Cache Optimized Resource Store
- B. Cross-Origin Resource Sharing
- C. Central Origin Response System
- D. Client Object Request Service

Answer: B

Q330. What is a webhook?

- A. A CSS animation hook for elements
- B. An automated HTTP callback notification
- C. A JavaScript debugging hook method
- D. A hook for hanging server cables

Answer: B

Q331. What is a Progressive Web App?

- A. A web app that requires progressive download of all assets before loading
- B. A web app that only works on desktop computers and not on mobile devices
- C. A web app with native-like features such as offline support and installability
- D. A web app that progressively removes features as more users access the site

Answer: C

Q332. What is TypeScript?

- A. A CSS preprocessor for styling
- B. A database query language spec
- C. A server hosting platform tool
- D. A typed superset of JavaScript

Answer: D

Q333. What is a static site generator?

- A. A tool that generates dynamic server-rendered pages on every request received
- B. A tool that creates static CSS styles without any responsive design support
- C. A tool that builds HTML pages at build time from templates and data sources
- D. A tool that generates static IP addresses for web server hosting purposes

Answer: C

Q334. What is a headless CMS?

- A. A CMS with no user interface for content editors or administrators to use
- B. A CMS that only supports headless browser testing and not content editing
- C. A CMS that runs without a server and stores content in browser storage
- D. A CMS that provides content via APIs without a fixed presentation layer

Answer: D

Q335. What is serverless computing?

- A. Computing without any servers involved in the entire application stack
- B. Computing that only works when the server is completely turned off today
- C. Running code in cloud functions without managing server infrastructure
- D. A deployment model where code runs directly in the user browser client

Answer: C

Q336. What is WebAssembly?

- A. A JavaScript framework for building web applications with assembly patterns
- B. A binary format for running compiled code in the browser at near-native speed
- C. A new CSS layout system for assembling web page components together
- D. A server-side assembly language for writing high-performance API endpoints

Answer: B

Q337. What is a monorepo?

- A. A server configuration where all services run on one physical machine only
- B. A single repository containing multiple related projects or packages together
- C. A database that stores all data in a single table without any relationships
- D. A single repository containing only one small microservice application code

Answer: B

Q338. What is Jamstack?

- A. Java, Angular, and MongoDB for building enterprise web applications stack
- B. JavaScript, APIs, and Markup for building modern web applications stack
- C. jQuery, Apache, and MariaDB for building traditional web application stack
- D. JSON, AJAX, and MySQL for building data-driven web applications stack

Answer: B

Q339. What is a design system?

- A. A system for automatically generating web page designs from text descriptions
- B. A tool for designing database schemas using visual drag-and-drop interfaces
- C. A CSS framework that generates design tokens from Figma files automatically
- D. A collection of reusable components and guidelines for consistent user interfaces

Answer: D

Q340. What are web components?

- A. Server-side components that generate HTML on the server before sending
- B. Components built exclusively using React or Vue frontend framework libraries
- C. CSS animation components that add visual effects to web page elements
- D. Custom reusable HTML elements using native browser APIs and standards

Answer: D

Q341. What is a domain name in web development?

- A. A human-readable address for a website
- B. A programming language for servers
- C. A type of database management system
- D. A tool for editing HTML files

Answer: A

Q342. Which of the following is an example of a markup language?

- A. Python
- B. Java
- C. HTML
- D. C++

Answer: C

Q343. What is the purpose of JavaScript in web development?

- A. To structure web page content
- B. To define database schemas
- C. To add interactivity and dynamic behavior
- D. To manage server hardware

Answer: C

Q344. What is the purpose of an SSL certificate on a website?

- A. To speed up JavaScript execution
- B. To encrypt data transmitted between the browser and the server
- C. To store user passwords
- D. To compress CSS files

Answer: B

Q345. What is a search engine?

- A. A tool that finds and displays web pages based on user queries
- B. A server that stores website files
- C. A programming language for building websites
- D. A device that connects computers to the internet

Answer: A

Q346. Which file extension is commonly used for CSS files?

- A. .html
- B. .js
- C. .css
- D. .php

Answer: C

Q347. What is the World Wide Web?

- A. A system of interlinked hypertext documents accessed via the internet
- B. A type of programming language
- C. A hardware component inside computers
- D. A local network connecting two computers

Answer: A

Q348. What does a text editor do in web development?

- A. Compiles code into machine language
- B. Allows developers to write and edit code
- C. Hosts websites on the internet
- D. Manages database records

Answer: B

Q349. What is an internet browser's primary function?

- A. To compile programming languages
- B. To render and display web pages
- C. To store website databases
- D. To write server-side code

Answer: B

Q350. What is the difference between the internet and the World Wide Web?

- A. They are the same thing
- B. The internet is the network infrastructure while the web is a service that runs on it
- C. The web is a type of hardware device
- D. The internet only supports email

Answer: B

Q351. What is a server in web architecture?

- A. A computer that provides resources or services to other computers
- B. A type of web browser
- C. A CSS framework
- D. A programming language

Answer: A

Q352. What is a URL used for?

- A. To style web pages
- B. To locate and access a specific resource on the web
- C. To compile JavaScript
- D. To manage database tables

Answer: B

Q353. What does the status code 500 indicate in HTTP?

- A. Successful request
- B. Page not found
- C. Internal server error
- D. Redirect to another page

Answer: C

Q354. What is a request in the client-server model?

- A. Data sent from the client to the server asking for a resource or action
- B. A CSS rule applied to a web page
- C. A JavaScript function call
- D. A database table definition

Answer: A

Q355. What is a response in the client-server model?

- A. A CSS animation effect
- B. Data sent from the server back to the client after processing a request
- C. A type of JavaScript variable
- D. A method of creating HTML elements

Answer: B

Q356. What protocol does HTTPS use for encryption?

- A. FTP
- B. TLS/SSL
- C. SMTP
- D. DNS

Answer: B

Q357. What is the default port for HTTPS?

- A. 80
- B. 21
- C. 443
- D. 8080

Answer: C

Q358. What does the HTTP DELETE method do?

- A. Creates a new resource
- B. Retrieves data from the server
- C. Removes a specified resource
- D. Updates part of a resource

Answer: C

Q359. What is a domain name system (DNS) record?

- A. An entry in a DNS database that maps a domain to information like an IP address
- B. A JavaScript variable declaration
- C. A CSS property for fonts
- D. An HTML form element

Answer: A

Q360. What is the purpose of an HTTP header?

- A. To display visible content on the page
- B. To provide metadata about the request or response
- C. To create database connections
- D. To compile source code

Answer: B

Q361. What does the `<!DOCTYPE html>` declaration do?

- A. It creates a new paragraph
- B. It declares the document type as HTML5
- C. It adds a comment to the page
- D. It links a CSS stylesheet

Answer: B

Q362. Which HTML tag is used for creating a numbered list?

- A. `<ul>`
- B. `<ol>`
- C. `<li>`
- D. `<dl>`

Answer: B

Q363. What is the purpose of the `<hr>` tag in HTML?

- A. To create a heading
- B. To insert a horizontal rule (line)
- C. To make text italic
- D. To add a hyperlink

Answer: B

Q364. Which HTML tag is used to display an image?

- A. `<picture>`
- B. `<photo>`
- C. `<img>`
- D. `<image>`

Answer: C

Q365. What does the `<em>` tag do in HTML?

- A. Creates a line break
- B. Emphasizes text, typically displayed in italics
- C. Makes text bold
- D. Adds a hyperlink

Answer: B

Q366. Which attribute specifies the source of an image in the `<img>` tag?

- A. href
- B. link
- C. src
- D. url

Answer: C

Q367. What is the purpose of the `<footer>` tag in HTML5?

- A. To create a navigation menu
- B. To define the footer section of a document or section
- C. To insert a form
- D. To add a heading

Answer: B

Q368. Which HTML element is used to define a table header cell?

- A. <td>
- B. <th>
- C. <tr>
- D. <thead>

Answer: B

Q369. What does the <header> tag represent in HTML5?

- A. The main content area
- B. Introductory content or a group of navigational aids
- C. A form input field
- D. A table column

Answer: B

Q370. Which tag is used to create a text input field in a form?

- A. <textarea>
- B. <input>
- C. <select>
- D. <button>

Answer: B

Q371. Which CSS property controls the size of text?

- A. text-size
- B. font-weight
- C. font-size
- D. text-height

Answer: C

Q372. What does the display: block property do?

- A. Hides the element
- B. Makes the element take up the full width available
- C. Makes the element inline
- D. Removes the element from the page

Answer: B

Q373. How do you apply CSS styles directly to an HTML element?

- A. Using an external stylesheet
- B. Using the style attribute on the element
- C. Using a JavaScript file
- D. Using a server-side script

Answer: B

Q374. Which property changes the background image of an element?

- A. bg-image
- B. background-image
- C. image-bg
- D. back-img

Answer: B

Q375. What does the border property do in CSS?

- A. Adds space inside the element
- B. Creates a line around the element
- C. Changes the text color
- D. Sets the element's width

Answer: B

Q376. Which CSS property makes text italic?

- A. text-decoration
- B. font-weight
- C. font-style
- D. text-transform

Answer: C

Q377. What does the color property in CSS affect?

- A. The background color of an element
- B. The text color of an element
- C. The border color
- D. The shadow color

Answer: B

Q378. How do you link an external CSS file to an HTML document?

- A. Using the <style> tag
- B. Using the <link> tag in the <head> section
- C. Using the <script> tag
- D. Using the <css> tag

Answer: B

Q379. Which CSS property sets the height of an element?

- A. size
- B. height
- C. tall
- D. length

Answer: B

Q380. What is a CSS class selector?

- A. A selector that targets elements by their tag name
- B. A selector that targets elements by a class name using a dot prefix
- C. A selector that targets elements by their ID
- D. A selector that targets all elements

Answer: B

Q381. What is a CSS animation?

- A. A way to gradually change CSS property values over time
- B. A method for compressing images
- C. A server-side rendering technique
- D. A type of HTML element

Answer: A

Q382. What does the flex-grow property do?

- A. Shrinks a flex item
- B. Specifies how much a flex item should grow relative to other items
- C. Sets the item's fixed width
- D. Changes the flex direction

Answer: B

Q383. What does the grid-gap property do in CSS Grid?

- A. Adds padding inside grid items
- B. Sets the space between grid rows and columns
- C. Changes the grid background color
- D. Defines the number of grid columns

Answer: B

Q384. What is a CSS transform?

- A. A way to change an element's shape, size, or position visually
- B. A method for loading CSS files
- C. A type of CSS variable
- D. A database operation

Answer: A

Q385. What does the overflow: hidden property do?

- A. Shows all content regardless of size
- B. Clips content that exceeds the element's boundaries
- C. Adds a scrollbar to the element
- D. Makes the element transparent

Answer: B

Q386. What does the vh unit represent in CSS?

- A. Virtual height of the parent element
- B. One percent of the viewport height
- C. The height of the HTML element
- D. A fixed pixel value

Answer: B

Q387. What is the purpose of the z-index property in CSS?

- A. To set the font size
- B. To control the stacking order of overlapping elements
- C. To change the text color
- D. To set the element width

Answer: B

Q388. What does the flex-shrink property control?

- A. How much a flex item grows
- B. How much a flex item shrinks relative to other items when space is limited
- C. The direction of flex items
- D. The gap between flex items

Answer: B

Q389. What does the grid-template-rows property define?

- A. The number of columns in a grid
- B. The sizes and structure of the grid rows
- C. The gap between grid items
- D. The alignment of grid items

Answer: B

Q390. What is a CSS gradient?

- A. A type of HTML element
- B. A smooth transition between two or more colors used as a background
- C. A JavaScript animation technique
- D. A method for compressing CSS files

Answer: B

Q391. What is a string in JavaScript?

- A. A number value
- B. A sequence of characters enclosed in quotes
- C. A boolean value
- D. An HTML element

Answer: B

Q392. What does the + operator do with two numbers in JavaScript?

- A. Subtracts them
- B. Concatenates them as strings
- C. Adds them together
- D. Multiplies them

Answer: C

Q393. What is a for loop used for in JavaScript?

- A. Declaring variables
- B. Repeating a block of code a specified number of times
- C. Defining functions
- D. Creating objects

Answer: B

Q394. What does the length property return on an array?

- A. The first element
- B. The last element
- C. The number of elements in the array
- D. The data type of the array

Answer: C

Q395. What is an object in JavaScript?

- A. A type of loop
- B. A collection of key-value pairs
- C. A CSS selector
- D. A type of comment

Answer: B

Q396. What does the return statement do in a function?

- A. Logs a message to the console
- B. Ends the function execution and returns a value to the caller
- C. Creates a new variable
- D. Starts a loop

Answer: B

Q397. What is the purpose of the else keyword in JavaScript?

- A. To declare a variable
- B. To provide an alternative block of code when the if condition is false
- C. To create a loop
- D. To define a function

Answer: B

Q398. How do you declare a constant in JavaScript?

- A. Using the var keyword
- B. Using the let keyword
- C. Using the const keyword
- D. Using the constant keyword

Answer: C

Q399. What does the pop() method do to an array?

- A. Adds an element to the beginning
- B. Removes and returns the last element
- C. Sorts the array
- D. Reverses the array

Answer: B

Q400. What is a boolean value in JavaScript?

- A. A text string
- B. A decimal number
- C. A value that is either true or false
- D. An array of numbers

Answer: C

Q401. What is the purpose of the setTimeout() function?

- A. To repeat code indefinitely
- B. To execute a function after a specified delay in milliseconds
- C. To stop all JavaScript execution
- D. To create a new variable

Answer: B

Q402. What is the setInterval() function used for?

- A. To run a function only once
- B. To repeatedly execute a function at specified time intervals
- C. To clear all timers
- D. To pause JavaScript execution

Answer: B

Q403. What does the map object store in JavaScript?

- A. Only string values
- B. Key-value pairs where keys can be any type
- C. Only number values
- D. Only boolean values

Answer: B

Q404. What is the purpose of the catch block in error handling?

- A. To throw new errors
- B. To handle and respond to errors that occur in the try block
- C. To prevent all errors
- D. To log messages to the console

Answer: B

Q405. What does the finally block do in error handling?

- A. It replaces the catch block
- B. It executes code after try and catch regardless of whether an error occurred
- C. It prevents errors from happening
- D. It only runs when there is no error

Answer: B

Q406. What is a callback function?

- A. A function that is never called
- B. A function passed as an argument to another function to be executed later
- C. A function that returns HTML
- D. A function that only works with arrays

Answer: B

Q407. What is the difference between let and const?

- A. They are identical keywords
- B. let allows reassignment while const does not allow reassignment after initialization
- C. const is faster than let
- D. let cannot be used in loops

Answer: B

Q408. What does the typeof operator return for an array?

- A. 'array'
- B. 'object'
- C. 'list'
- D. 'collection'

Answer: B

Q409. What is template literal syntax used for?

- A. Creating HTML templates
- B. Embedding expressions inside strings using backticks and `${}` syntax
- C. Defining CSS templates
- D. Creating database templates

Answer: B

Q410. What does the spread operator (...) do when used with an array?

- A. Deletes the array
- B. Expands the array elements into individual values
- C. Sorts the array
- D. Reverses the array

Answer: B

Q411. What is the purpose of a frontend framework?

- A. To manage database connections
- B. To provide structure and tools for building user interfaces efficiently
- C. To handle server-side routing
- D. To compress images

Answer: B

Q412. What is a prop in frontend frameworks?

- A. A CSS property
- B. Data passed from a parent component to a child component
- C. A database query
- D. A server configuration

Answer: B

Q413. What is component-based architecture?

- A. A server deployment strategy
- B. Building UIs from small, reusable, self-contained pieces called components
- C. A database schema design
- D. A CSS layout method

Answer: B

Q414. What is Next.js used for?

- A. Managing databases
- B. Building React applications with server-side rendering and static generation
- C. Styling web pages
- D. Testing JavaScript code

Answer: B

Q415. What is the purpose of routing in a frontend framework?

- A. To style elements
- B. To navigate between different pages or views within the application
- C. To manage database connections
- D. To compile JavaScript

Answer: B

Q416. What is Nuxt.js?

- A. A CSS framework
- B. A framework built on top of Vue.js for server-side rendering and static generation
- C. A database management tool
- D. A testing library

Answer: B

Q417. What is a template in frontend frameworks?

- A. A database schema
- B. The markup that defines the visual structure of a component
- C. A server configuration file
- D. A build tool

Answer: B

Q418. What is event handling in frontend frameworks?

- A. Sending emails
- B. Responding to user interactions like clicks, input changes, or form submissions
- C. Managing server events
- D. Processing database queries

Answer: B

Q419. What is conditional rendering in frontend frameworks?

- A. Rendering CSS conditions
- B. Showing or hiding UI elements based on conditions in the application state
- C. Conditionally loading databases
- D. Applying conditional server rules

Answer: B

Q420. What is a lifecycle method in a component?

- A. A CSS animation phase
- B. A function that runs at specific stages of a component's existence such as mounting, updating, or unmounting
- C. A database migration step
- D. A server restart procedure

Answer: B

Q421. What is an API route in a backend application?

- A. A CSS selector path
- B. A URL pattern that maps to a specific handler function on the server
- C. A JavaScript import path
- D. A database table name

Answer: B

Q422. What is the purpose of environment variables in backend development?

- A. To style web pages
- B. To store configuration values like API keys and database URLs outside the codebase
- C. To create HTML elements
- D. To define CSS variables

Answer: B

Q423. What is a database connection string?

- A. A CSS selector
- B. A string containing the information needed to connect to a database, including host, port, and credentials
- C. An HTML attribute
- D. A JavaScript string method

Answer: B

Q424. What is the purpose of a controller in the MVC pattern?

- A. To render CSS styles
- B. To handle incoming requests, process them, and return appropriate responses
- C. To store data in the database
- D. To create HTML templates

Answer: B

Q425. What is a model in the MVC pattern?

- A. A CSS layout model
- B. The component responsible for data management, business logic, and database interaction
- C. An HTML template
- D. A routing mechanism

Answer: B

Q426. What is the purpose of a view in the MVC pattern?

- A. To manage the database
- B. To present data to the user in a formatted output
- C. To handle routing
- D. To process API requests

Answer: B

Q427. What is a backend framework?

- A. A CSS grid system
- B. A software tool that provides structure and utilities for building server-side applications
- C. A browser extension
- D. A frontend testing library

Answer: B

Q428. What does the HTTP POST method do?

- A. Retrieves a resource
- B. Submits data to the server to create a new resource
- C. Deletes a resource
- D. Updates part of a resource

Answer: B

Q429. What is a status code in the 4xx range?

- A. A successful response
- B. A server error
- C. A client error indicating the request was invalid or unauthorized
- D. A redirect response

Answer: C

Q430. What is npm used for in backend development?

- A. To style web pages
- B. To manage packages and dependencies for Node.js projects
- C. To create databases
- D. To design user interfaces

Answer: B

Q431. What is a request handler in server-side programming?

- A. A CSS class
- B. A function that processes incoming HTTP requests and sends back responses
- C. An HTML element
- D. A database table

Answer: B

Q432. What is middleware in Express.js?

- A. A CSS preprocessor
- B. Functions that execute during the request-response cycle with access to request and response objects
- C. A database middleware layer
- D. A frontend component

Answer: B

Q433. What is the purpose of a response status code?

- A. To style the response
- B. To indicate the result of the server's attempt to process the request
- C. To create database entries
- D. To compile JavaScript

Answer: B

Q434. What is a REST endpoint?

- A. A CSS property
- B. A specific URL path on a server that accepts HTTP requests for a particular resource
- C. A JavaScript function
- D. A database connection point

Answer: B

Q435. What is the purpose of the response body in HTTP?

- A. To set HTTP headers
- B. To contain the actual data being sent back to the client, such as HTML, JSON, or files
- C. To define the URL
- D. To establish database connections

Answer: B

Q436. What is Node.js primarily used for?

- A. Styling web pages
- B. Running JavaScript on the server side to build scalable network applications
- C. Creating CSS animations
- D. Managing databases visually

Answer: B

Q437. What is the purpose of routing in server-side applications?

- A. To style navigation menus
- B. To map incoming request URLs and methods to specific handler functions
- C. To route database queries
- D. To manage CSS imports

Answer: B

Q438. What is a JSON response?

- A. A CSS format
- B. Data sent from the server formatted as a JavaScript Object Notation string
- C. An HTML template
- D. A database query result only

Answer: B

Q439. What is the purpose of a static file server?

- A. To generate dynamic content
- B. To serve unchanging files like images, CSS, and JavaScript directly to clients
- C. To manage database connections
- D. To process form submissions

Answer: B

Q440. What is an error handler in server-side applications?

- A. A CSS error page
- B. A function that catches and processes errors, returning appropriate error responses to clients
- C. A JavaScript lint rule
- D. A database recovery tool

Answer: B

Q441. What is a database record?

- A. A CSS rule
- B. A single entry or row in a database table containing related data
- C. An HTML element
- D. A JavaScript variable

Answer: B

Q442. What is the purpose of the WHERE clause in SQL?

- A. To create a new table
- B. To filter results based on specified conditions
- C. To sort query results
- D. To join two tables

Answer: B

Q443. What is a collection in MongoDB?

- A. A CSS class collection
- B. A group of related documents, similar to a table in relational databases
- C. An array in JavaScript
- D. A set of HTML elements

Answer: B

Q444. What does the ORDER BY clause do in SQL?

- A. Creates a new order
- B. Sorts the query results based on one or more columns
- C. Orders database tables
- D. Reorders table columns

Answer: B

Q445. What is the purpose of the LIMIT clause in SQL?

- A. To limit database size
- B. To restrict the number of rows returned by a query
- C. To limit the number of tables
- D. To limit column types

Answer: B

Q446. What is a document in MongoDB?

- A. A text file
- B. A JSON-like record that stores data as field-value pairs
- C. A CSS stylesheet
- D. An HTML page

Answer: B

Q447. What does the DELETE statement do in SQL?

- A. Removes the entire database
- B. Removes one or more rows from a table based on a condition
- C. Deletes a column from a table
- D. Removes an index

Answer: B

Q448. What is the purpose of the GROUP BY clause in SQL?

- A. To group CSS styles
- B. To group rows with the same values in specified columns, often used with aggregate functions
- C. To group database tables
- D. To create column groups

Answer: B

Q449. What is SQLite?

- A. A CSS framework
- B. A lightweight, file-based relational database that requires no separate server process
- C. A JavaScript library
- D. An HTML template engine

Answer: B

Q450. What does the UPDATE statement do in SQL?

- A. Creates a new table
- B. Modifies existing data in one or more rows of a table
- C. Updates the database software
- D. Changes the table name

Answer: B

Q451. What is the purpose of a strong password?

- A. To make logging in faster
- B. To make unauthorized access difficult by using a complex combination of characters
- C. To improve website performance
- D. To style user profiles

Answer: B

Q452. What is a security vulnerability?

- A. A CSS layout issue
- B. A weakness in a system that can be exploited by attackers to gain unauthorized access or cause harm
- C. A JavaScript error
- D. A slow database query

Answer: B

Q453. What does the HTTPS padlock icon in a browser indicate?

- A. The website loads faster
- B. The connection between the browser and server is encrypted
- C. The website has no bugs
- D. The website uses CSS

Answer: B

Q454. What is a security patch?

- A. A CSS design pattern
- B. A software update that fixes a known security vulnerability
- C. A JavaScript library
- D. A database backup

Answer: B

Q455. What is social engineering in the context of security?

- A. A web development technique
- B. Manipulating people into divulging confidential information or performing actions that compromise security
- C. A CSS framework
- D. A database design method

Answer: B

Q456. What is the purpose of logging out of a web application?

- A. To clear the CSS cache
- B. To end the user session and prevent unauthorized access to the account
- C. To delete the user account
- D. To clear the browser history

Answer: B

Q457. What is a denial-of-service (DoS) attack?

- A. A CSS rendering technique
- B. An attack that overwhelms a server with excessive traffic to make it unavailable to legitimate users
- C. A JavaScript debugging method
- D. A database indexing technique

Answer: B

Q458. Why should software and libraries be kept up to date?

- A. To change the website design
- B. To patch known security vulnerabilities and protect against recently discovered threats
- C. To improve CSS performance
- D. To add more JavaScript features

Answer: B

Q459. What is the purpose of an access control list (ACL)?

- A. To control CSS access
- B. To define which users or systems are allowed or denied access to specific resources
- C. To manage database access speed
- D. To control HTML rendering

Answer: B

Q460. What is data backup important for in web security?

- A. To style web pages
- B. To ensure data can be recovered in case of attacks, hardware failures, or accidental deletion
- C. To speed up the website
- D. To create new databases

Answer: B

Q461. What is website performance?

- A. The visual design of a website
- B. How quickly a website loads and responds to user interactions
- C. The number of pages on a website
- D. The programming language used

Answer: B

Q462. Why do large files slow down a website?

- A. Large files improve performance
- B. Large files take longer to download over the network, increasing page load time
- C. Large files are more secure
- D. Large files use less bandwidth

Answer: B

Q463. What is the purpose of compressing images for the web?

- A. To increase image quality
- B. To reduce file size so images load faster without significantly impacting visual quality
- C. To change the image format
- D. To add filters to images

Answer: B

Q464. What is a content delivery network (CDN) used for?

- A. Creating website content
- B. Delivering website content from servers geographically closer to users for faster loading
- C. Managing databases
- D. Writing CSS styles

Answer: B

Q465. What does code minification do?

- A. Makes code easier to read
- B. Removes whitespace, comments, and shortens variable names to reduce file size
- C. Adds comments to code
- D. Encrypts the code

Answer: B

Q466. What is render time in web performance?

- A. The time to write code
- B. The time the browser takes to display content on the screen after receiving the data
- C. The time to deploy a website
- D. The time to create a database

Answer: B

Q467. Why is reducing HTTP requests important for performance?

- A. More requests make a site faster
- B. Each request adds network overhead, so fewer requests mean faster page loads
- C. Requests improve security
- D. Requests have no impact on speed

Answer: B

Q468. What is the purpose of using responsive images?

- A. To make images move on the page
- B. To serve appropriately sized images based on the device screen size, avoiding unnecessarily large downloads
- C. To add image filters
- D. To create image slideshows

Answer: B

Q469. What does the term 'above the fold' mean in web performance?

- A. A CSS folding technique
- B. The content visible on screen without scrolling, which should load as quickly as possible
- C. A JavaScript function
- D. A database technique

Answer: B

Q470. What is the impact of web fonts on page performance?

- A. Web fonts have no impact
- B. Custom web fonts add extra file downloads that can delay text rendering
- C. Web fonts make pages load faster
- D. Web fonts reduce file size

Answer: B

Q471. What is automated testing?

- A. Testing done by a human tester
- B. Using software to run tests automatically and compare results to expected outcomes
- C. A type of CSS animation
- D. A database backup method

Answer: B

Q472. What is a test assertion?

- A. A CSS declaration
- B. A statement that checks whether a specific condition is true in a test
- C. A JavaScript variable
- D. A database constraint

Answer: B

Q473. What is a test runner?

- A. A CSS preprocessor
- B. A tool that executes test files and reports the results
- C. A database migration tool
- D. An HTML validator

Answer: B

Q474. What is the purpose of the describe block in testing?

- A. To describe CSS properties
- B. To group related test cases together under a descriptive label
- C. To describe database tables
- D. To describe HTML elements

Answer: B

Q475. What is the purpose of the it or test function in testing?

- A. To iterate over arrays
- B. To define an individual test case with a description and test logic
- C. To create CSS iterations
- D. To iterate database records

Answer: B

Q476. What is an expected value in testing?

- A. A CSS expected property
- B. The correct result that a test verifies the actual output against
- C. A database expected row
- D. An HTML expected element

Answer: B

Q477. What is a test failure?

- A. A CSS rendering error
- B. When the actual output of code does not match the expected result in a test assertion
- C. A server crash
- D. A database timeout

Answer: B

Q478. What does the term 'green test' mean?

- A. A test with green CSS styling
- B. A test that passes, indicating the code behaves as expected
- C. A test for environmental websites
- D. A test that runs slowly

Answer: B

Q479. What is the Elements panel in browser DevTools used for?

- A. To write JavaScript code
- B. To inspect and modify the HTML structure and CSS styles of a web page in real time
- C. To manage databases
- D. To run automated tests

Answer: B

Q480. What is a smoke test?

- A. A test for browser smoke effects
- B. A basic test that verifies the most critical functionality works, ensuring the application is stable enough for further testing
- C. A performance stress test
- D. A CSS animation test

Answer: B

Q481. What is continuous integration (CI)?

- A. A CSS integration method
- B. The practice of automatically building and testing code every time changes are pushed to the repository
- C. A database integration technique
- D. An HTML template system

Answer: B

Q482. What is continuous deployment (CD)?

- A. Continuously designing CSS
- B. Automatically deploying code to production after it passes all tests in the CI pipeline
- C. Continuously deleting files
- D. Continuously creating databases

Answer: B

Q483. What is a build process in web development?

- A. Building HTML elements manually
- B. The automated process of transforming source code into optimized production-ready files
- C. Building database tables
- D. Building CSS manually

Answer: B

Q484. What is the difference between development, staging, and production environments?

- A. They are all the same
- B. Development is for coding, staging is for pre-release testing, and production serves real users
- C. Only production is important
- D. Staging replaces production

Answer: B

Q485. What is Vercel used for?

- A. Writing CSS
- B. Deploying and hosting web applications, particularly those built with Next.js and other frontend frameworks
- C. Managing databases
- D. Testing JavaScript

Answer: B

Q486. What is the purpose of a .gitignore file?

- A. To ignore CSS files
- B. To specify files and directories that Git should not track, like node_modules and environment files
- C. To ignore HTML files
- D. To ignore all JavaScript files

Answer: B

Q487. What is a deployment pipeline in web development?

- A. A physical pipe for network cables
- B. An automated sequence of steps that builds, tests, and deploys code to production
- C. A tool for managing CSS stylesheets
- D. A method for compressing image files

Answer: B

Q488. What does the npm build command typically do?

- A. Builds a database
- B. Compiles and bundles the application's source code into optimized production-ready files
- C. Builds CSS from scratch
- D. Builds HTML templates

Answer: B

Q489. What happens during a typical deployment workflow?

- A. CSS files are manually uploaded
- B. Code is built, tested, and published to a server so users can access the updated application
- C. Databases are deleted and recreated
- D. HTML files are emailed to users

Answer: B

Q490. What is the purpose of a CDN in deployment?

- A. To create databases
- B. To distribute static assets across global servers for faster delivery to users worldwide
- C. To compile code
- D. To manage version control

Answer: B

Q491. What is an API in simple terms?

- A. A CSS framework
- B. A set of rules that allows different software applications to communicate with each other
- C. A JavaScript library
- D. A database management tool

Answer: B

Q492. What is a request body in an API call?

- A. The visual body of a web page
- B. The data sent along with an API request, typically in JSON format
- C. The HTML body element
- D. A CSS body selector

Answer: B

Q493. What does an API response contain?

- A. Only HTML content
- B. The data returned by the server, typically including a status code, headers, and a response body
- C. Only CSS styles
- D. Only JavaScript code

Answer: B

Q494. What is the purpose of the Authorization header in API requests?

- A. To authorize CSS styles
- B. To send authentication credentials like tokens or API keys to prove the requester's identity
- C. To authorize HTML elements
- D. To authorize database access

Answer: B

Q495. What is a path parameter in a REST API URL?

- A. A CSS path property
- B. A variable embedded in the URL path that identifies a specific resource, like /users/123
- C. A file system path
- D. A database path

Answer: B

Q496. What does the HTTP PUT method do in a REST API?

- A. Retrieves a resource
- B. Replaces an entire existing resource with the provided data
- C. Creates a new database
- D. Deletes a resource

Answer: B

Q497. What is the Content-Type header used for?

- A. To set the page title
- B. To specify the format of the data in the request or response body, such as application/json
- C. To define CSS content
- D. To specify HTML content

Answer: B

Q498. What is the purpose of the HTTP status code 201?

- A. Resource not found
- B. A new resource was successfully created
- C. Server error
- D. Redirect to another page

Answer: B

Q499. What tool is commonly used to test API endpoints?

- A. A CSS editor
- B. Postman or similar API testing tools
- C. An image editor
- D. A database viewer

Answer: B

Q500. What is the difference between a query parameter and a path parameter?

- A. They are identical
- B. Path parameters identify specific resources in the URL path while query parameters filter or modify the response and appear after a question mark
- C. Query parameters are faster
- D. Path parameters are optional

Answer: B

Q501. What is Vite used for in modern web development?

- A. Managing databases
- B. A fast build tool and development server that uses native ES modules for quick startup
- C. Styling web pages
- D. Testing JavaScript code

Answer: B

Q502. What is ESLint used for?

- A. Styling web pages
- B. Analyzing JavaScript code for potential errors and enforcing coding standards
- C. Managing databases
- D. Creating HTML templates

Answer: B

Q503. What is the purpose of Prettier in modern development?

- A. To make images prettier
- B. To automatically format code according to consistent style rules, eliminating formatting debates
- C. To improve database performance
- D. To create prettier HTML

Answer: B

Q504. What is a headless CMS in modern web development?

- A. A CMS without a database backend
- B. A content management system that provides content via APIs without a built-in frontend
- C. A CMS that only supports text content
- D. A CMS that runs without a server

Answer: B

Q505. What is a package.json file used for?

- A. To package CSS files
- B. To define project metadata, dependencies, scripts, and configuration for Node.js projects
- C. To package HTML files
- D. To define database packages

Answer: B

Q506. What is GitHub Actions used for?

- A. Styling GitHub pages
- B. Automating workflows like testing, building, and deploying code directly from a GitHub repository
- C. Managing databases on GitHub
- D. Creating GitHub themes

Answer: B

Q507. What is a .env file used for?

- A. To style environments with CSS
- B. To store environment-specific configuration variables like API keys and database URLs
- C. To define HTML environments
- D. To create JavaScript environments

Answer: B

Q508. What is Storybook used for in modern web development?

- A. Creating digital storybooks
- B. Developing, documenting, and testing UI components in isolation outside the main application
- C. Managing server stories
- D. Creating database stories

Answer: B

Q509. What is the purpose of a code formatter?

- A. To format database records
- B. To automatically apply consistent formatting rules to code, ensuring uniform style across a project
- C. To format CSS colors
- D. To format HTML templates

Answer: B

Q510. What is pnpm?

- A. A CSS preprocessor
- B. A fast, disk-space-efficient package manager that uses a content-addressable store to share packages
- C. A JavaScript framework
- D. A database tool

Answer: B

Web Development — Medium

510 questions

Q511. What is the difference between static and dynamic websites?

- A. Static sites load faster but always have fewer total features
- B. Static sites rely on databases while dynamic ones do not
- C. Dynamic sites are completely unable to use HTML markup
- D. Static sites have fixed content while dynamic sites generate it

Answer: D

Q512. Which technology stack is referred to by the term MERN?

- A. MongoDB, Express, React, Node.js
- B. MySQL, Express, React, Node
- C. MariaDB, Ember, React, Nginx
- D. MongoDB, Electron, Redux, Node.js

Answer: A

Q513. What is the purpose of a domain name?

- A. To compile JavaScript code for production
- B. To store website files on a server
- C. To provide a human-readable web address
- D. To encrypt website data during transfer

Answer: C

Q514. What does IDE stand for in web development?

- A. Internal Data Exchange
- B. Integrated Development Environment
- C. Internet Development Engine
- D. Interactive Design Editor

Answer: B

Q515. Which of the following is a version control system?

- A. Docker
- B. Jenkins
- C. Webpack
- D. Git

Answer: D

Q516. What is responsive web design?

- A. Designing websites for one single device
- B. Designing websites that adapt to screen sizes
- C. Designing websites that include animations
- D. Designing websites that respond very quickly

Answer: B

Q517. What is the role of a back-end developer?

- A. Designing the visual user interface
- B. Creating front-end CSS animations
- C. Testing cross-browser compatibility
- D. Managing server-side logic and data

Answer: D

Q518. What is a CMS?

- A. Central Monitoring Service
- B. Client Module Setup
- C. Code Management System
- D. Content Management System

Answer: D

Q519. Which protocol is used for secure communication on the web?

- A. HTTPS
- B. FTP
- C. HTTP
- D. SMTP

Answer: A

Q520. What is a wireframe in web development?

- A. A type of client-side JavaScript framework
- B. A server-side configuration settings file
- C. A basic visual guide for web page layout
- D. A CSS grid system for responsive design

Answer: C

Q521. What is the difference between HTTP and HTTPS?

- A. HTTPS delivers content more quickly
- B. HTTPS encrypts data using TLS/SSL
- C. HTTP provides stronger security features
- D. There is no practical difference at all

Answer: B

Q522. Which HTTP status code indicates a resource was not found?

- A. 200
- B. 500
- C. 404
- D. 301

Answer: C

Q523. What is a RESTful API?

- A. An API built solely with the React library
- B. An API that always requires authentication
- C. An API that exclusively uses GET requests
- D. An API that follows REST design constraints

Answer: D

Q524. What is the purpose of the HTTP POST method?

- A. To retrieve existing data from a resource
- B. To submit data for processing by the server
- C. To delete a specific resource from storage
- D. To update only the HTTP request headers

Answer: B

Q525. What is a proxy server?

- A. A specialized type of relational database
- B. A redundant backup server for data recovery
- C. An intermediary server between client and target
- D. A domain name resolution DNS server system

Answer: C

Q526. What is latency in web architecture?

- A. The total download size of a given web page
- B. The processing speed of the server CPU
- C. The number of concurrent users currently online
- D. The delay between a request and a response

Answer: D

Q527. Which HTTP header specifies the content type of a response?

- A. Cache-Control
- B. Content-Type
- C. Accept-Type
- D. Authorization

Answer: B

Q528. What does a CDN do?

- A. Manages relational database connections
- B. Compiles source code into machine language
- C. Distributes content across global servers
- D. Designs web page layouts automatically

Answer: C

Q529. What is WebSocket used for?

- A. Managing relational database schemas and queries
- B. Full-duplex real-time client-server communication
- C. Applying visual styles to web page elements
- D. Compressing files to reduce transfer size

Answer: B

Q530. What is the purpose of the HTTP PUT method?

- A. To update or replace a resource at a URL
- B. To delete an existing resource completely
- C. To retrieve existing data from a resource
- D. To establish a new connection to a server

Answer: A

Q531. What is the purpose of the <meta> tag?

- A. To provide metadata about the document to browser
- B. To link external CSS files to the HTML document
- C. To display visible metadata on the page directly
- D. To create a navigation menu within the header area

Answer: A

Q532. Which HTML5 tag is used for navigation links?

- A. <links> element
- B. <nav> element
- C. <navigation> tag
- D. <menu> element

Answer: B

Q533. What is the difference between <div> and ?

- A. They are identical in every behavior
- B. is not able to contain any text
- C. <div> is inline, is block-level
- D. <div> is block-level, is inline

Answer: D

Q534. Which attribute provides alternative text for an image?

- A. label
- B. desc
- C. title
- D. alt

Answer: D

Q535. What is the purpose of the <form> tag?

- A. To create a form for user input
- B. To format and style text content
- C. To create a page footer section
- D. To add mathematical formulas

Answer: A

Q536. Which input type creates a checkbox?

- A. type='checkbox'
- B. type='box'
- C. type='check'
- D. type='tick'

Answer: A

Q537. What is the purpose of the action attribute in a form?

- A. To apply CSS styles to the form
- B. To specify where form data is sent
- C. To validate the form input fields
- D. To add a JavaScript event handler

Answer: B

Q538. Which HTML5 element is used for self-contained content like a blog post?

- A. <article>
- B. <div>
- C. <section>
- D. <aside>

Answer: A

Q539. What does target='_blank' do in an anchor tag?

- A. Closes the current browser tab
- B. Opens the link in a new tab
- C. Downloads the linked file
- D. Opens the link in the same tab

Answer: B

Q540. Which tag embeds a video in HTML5?

- A. <embed>
- B. <movie>
- C. <video>
- D. <media>

Answer: C

Q541. What is the CSS Box Model?

- A. A responsive layout framework for components
- B. A model of content, padding, border, margin
- C. A 3D modeling technique for page effects
- D. A CSS animation model for element transitions

Answer: B

Q542. What is the difference between display: none and visibility: hidden?

- A. display: none only hides text content but keeps the element
- B. visibility: hidden removes elements from the document flow
- C. display: none removes from flow; visibility: hidden keeps space
- D. They behave identically in every rendering situation

Answer: C

Q543. Which CSS property is used to float an element?

- A. position
- B. align
- C. float
- D. display

Answer: C

Q544. What does position: relative do?

- A. Makes element fixed on scroll
- B. Positions element relative to the viewport
- C. Positions element relative to its normal position
- D. Removes element from document flow

Answer: C

Q545. What is specificity in CSS?

- A. The algorithm determining which CSS rule applies when multiple match
- B. The order in which CSS properties are declared
- C. The speed at which the browser renders CSS styles
- D. The total number of CSS files loaded on the page

Answer: A

Q546. Which CSS pseudo-class applies when the user hovers over an element?

- A. :focus
- B. ::hover
- C. :hover
- D. :active

Answer: C

Q547. What is the z-index property used for?

- A. Setting the zoom level of an element
- B. Setting the total width of an element
- C. Creating 3D visual effects on elements
- D. Controlling the stack order of elements

Answer: D

Q548. What does box-sizing: border-box do?

- A. Sets the box model back to the content-box
- B. Adds a shadow effect around the element box
- C. Includes padding and border in total dimensions
- D. Removes the border from around the element

Answer: C

Q549. Which selector targets the first child element?

- A. :first
- B. ::first-child
- C. :first-child
- D. :child(1)

Answer: C

Q550. What is the inherit value in CSS?

- A. A reserved keyword in the JavaScript language
- B. Takes the same computed value as parent element
- C. A CSS variable declared with a custom property
- D. Resets the property to its initial default value

Answer: B

Q551. What is the difference between Flexbox and Grid?

- A. They behave identically in all layout situations
- B. Flexbox was introduced after Grid in the spec
- C. Flexbox is one-dimensional, Grid is two-dimensional
- D. Grid is designed exclusively for table layouts

Answer: C

Q552. What does grid-template-areas do?

- A. Creates content for individual grid items
- B. Defines named grid areas for placing items
- C. Creates decorative borders around grid cells
- D. Sets the color scheme for the grid layout

Answer: B

Q553. What is a CSS pseudo-element?

- A. An HTML5 semantic element like header or footer
- B. A fake HTML element added via JavaScript
- C. A keyword styling a specific part of an element
- D. A JavaScript component for building interfaces

Answer: C

Q554. What is the fr unit in CSS Grid?

- A. A frequency measurement unit for media queries
- B. A font ratio measurement unit for sizing
- C. A fraction of available space in the container
- D. A frame rate measurement unit for animations

Answer: C

Q555. How do you create a responsive breakpoint for tablets?

- A. @responsive (tablet)
- B. @tablet { max-width }
- C. @breakpoint (tablet)
- D. @media (max-width: 768px)

Answer: D

Q556. What does flex-wrap: wrap do?

- A. Creates a decorative border around all items
- B. Prevents items from wrapping to new lines
- C. Wraps the text content within each flex item
- D. Allows flex items to wrap onto multiple lines

Answer: D

Q557. What is the gap property in CSS Grid?

- A. A gap in browser support for older versions
- B. A grid rendering error in certain browsers
- C. Sets the spacing between grid rows and columns
- D. The outer margin space around the entire grid

Answer: C

Q558. What does @keyframes define in CSS?

- A. Key events handled in JavaScript listeners
- B. The stages of a CSS animation sequence
- C. Keyboard shortcut bindings for the browser
- D. Database primary and foreign key definitions

Answer: B

Q559. What is the object-fit CSS property?

- A. A flexbox property for controlling item alignment
- B. Specifies how an image resizes to fit its container
- C. A grid alignment property for placing grid items
- D. Fits text content into its parent containers

Answer: B

Q560. What is a mobile-first design approach?

- A. Using only mobile-specific CSS frameworks
- B. Designing websites for desktop screens first
- C. Designing for small screens first then enhancing
- D. Designing websites for mobile devices only

Answer: C

Q561. What is the difference between let and var?

- A. var has block scope; let has function scope
- B. let has block scope; var has function scope
- C. They are identical in scoping behavior
- D. let is significantly slower in performance

Answer: B

Q562. What is a closure in JavaScript?

- A. A type of loop for iterating over arrays
- B. A built-in method to end a running program
- C. A method for closing the browser window
- D. A function accessing its outer scope variables

Answer: D

Q563. What is the purpose of the 'this' keyword?

- A. Refers to the previously declared variable value
- B. Refers to the current CSS style being applied
- C. Refers to the object executing current function
- D. Creates a new object instance in current scope

Answer: C

Q564. What is hoisting in JavaScript?

- A. Moving declarations to the top of their scope
- B. Uploading files from the client to the server
- C. Raising and throwing runtime error exceptions
- D. Lifting DOM elements higher in the page layout

Answer: A

Q565. What is the difference between null and undefined?

- A. undefined always indicates a runtime error occurred
- B. They are identical in meaning and behavior
- C. null is intentional absence; undefined means not assigned
- D. null is for numbers; undefined is for string values

Answer: C

Q566. What does the map() method do on an array?

- A. Creates a geographic map data visualization
- B. Sorts the array elements in ascending order
- C. Creates a new array by applying a function to each item
- D. Maps array indices to their corresponding values

Answer: C

Q567. What is event bubbling in JavaScript?

- A. A CSS animation effect using keyframe transitions
- B. A memory leak issue caused by circular references
- C. Creating animated bubble visual effects on screen
- D. When an event on a child propagates up through parents

Answer: D

Q568. What is the purpose of JSON.parse()?

- A. Compresses JSON data to reduce its file size
- B. Creates a new JSON file on the server disk
- C. Converts a JSON string into a JavaScript object
- D. Validates JSON syntax for correctness and errors

Answer: C

Q569. What is template literal syntax in JavaScript?

- A. Using single quotes around string literals
- B. Using double quotes around string values
- C. Using the + operator for concatenation
- D. Using backticks with \${} for interpolation

Answer: D

Q570. What is destructuring in JavaScript?

- A. Permanently deleting JavaScript objects from memory
- B. Unpacking values from arrays or objects into variables
- C. Breaking code into smaller modules and files
- D. Removing DOM elements from the document tree

Answer: B

Q571. What is Promise.all() used for?

- A. Running multiple promises and waiting for all
- B. Creating a sequential promise chain
- C. Resolving only a single individual promise
- D. Canceling all pending promises at once

Answer: A

Q572. What is the difference between Promise.all() and Promise.allSettled()?

- A. Promise.allSettled() also fails immediately on rejection
- B. Promise.all() is always slower in execution speed
- C. Promise.all() fails on rejection; allSettled() waits for all
- D. They are identical in their behavior

Answer: C

Q573. What is optional chaining (?.) in JavaScript?

- A. Safely accesses nested properties without errors
- B. A ternary operator for conditional expressions
- C. A loop syntax for iterating over object entries
- D. A regex operator for pattern matching in strings

Answer: A

Q574. What is the nullish coalescing operator (??)?

- A. A ternary for inline conditional expressions
- B. Returns right operand when left is null/undefined
- C. A logical OR for boolean expressions
- D. A strict equality check comparing type and value

Answer: B

Q575. What is a Proxy in JavaScript?

- A. A web proxy server for routing network traffic
- B. An object wrapping another to intercept operations
- C. A network library for making HTTP request calls
- D. A design pattern used only in OOP code contexts

Answer: B

Q576. What is the purpose of Object.entries()?

- A. To merge multiple objects into a combined one
- B. To return an array of [key, value] pairs
- C. To create new entries within an object literal
- D. To check if a specific object reference exists

Answer: B

Q577. What is a Set in JavaScript?

- A. A type of map for key-value pair storage
- B. An array method for modifying element order
- C. A function to set values on object properties
- D. A collection of unique values with no duplicates

Answer: D

Q578. What is dynamic import in JavaScript?

- A. Using import() to load modules on demand at runtime
- B. Importing CSS styles dynamically at runtime
- C. Importing all application modules simultaneously
- D. A feature exclusive to the Webpack bundler tool

Answer: A

Q579. What is the for...of loop used for?

- A. Iterating over the keys of an object
- B. Executing queries against a database
- C. A C-style for loop with index counters
- D. Iterating over values of iterable objects

Answer: D

Q580. What does Object.assign() do?

- A. Compares two objects for deep structural equality
- B. Assigns type annotations to JavaScript objects
- C. Creates a deep copy of all nested objects
- D. Copies enumerable own properties to a target (shallow)

Answer: D

Q581. What is the Virtual DOM in React?

- A. A browser API for accessing device hardware and sensors
- B. A lightweight copy of the real DOM for efficient updates
- C. A virtual reality feature for immersive web experiences
- D. A hidden HTML element that is invisible to the end user

Answer: B

Q582. What is React's useState hook?

- A. A database hook for managing query results
- B. A function adding state to functional components
- C. A CSS hook for applying dynamic styles
- D. A routing hook for handling URL navigation

Answer: B

Q583. What is the useEffect hook used for?

- A. Creating CSS effects through JavaScript manipulation
- B. Performing side effects like data fetching in components
- C. Managing component state with reactive data binding
- D. Adding visual effects like shadows and gradients

Answer: B

Q584. What is Vue's reactivity system based on?

- A. Manual DOM manipulation with vanilla JavaScript
- B. JavaScript Proxies that track and trigger updates
- C. Web Workers running in separate background threads
- D. CSS variables with automatic browser recalculation

Answer: B

Q585. What is Angular's two-way data binding?

- A. CSS binding between stylesheets and HTML elements
- B. A database connection pool for managing active queries
- C. Changes in model automatically update the view and vice versa
- D. Server-client binding for real-time data synchronization

Answer: C

Q586. What is React Router?

- A. A server framework for building backend APIs
- B. A library for handling navigation in React apps
- C. A CSS router for switching between stylesheets
- D. A physical network router for internet traffic

Answer: B

Q587. What is a Higher-Order Component (HOC) in React?

- A. A component rendering at the top of a page
- B. A function that takes and returns an enhanced component
- C. A CSS wrapper element for styling child elements
- D. A parent component wrapping all child components

Answer: B

Q588. What is Vuex/Pinia used for in Vue?

- A. Managing CSS styles across Vue applications
- B. Centralized state management for Vue applications
- C. Managing database connections in Vue projects
- D. Server-side routing configuration for Vue apps

Answer: B

Q589. What is the purpose of Angular's NgModule?

- A. A CSS module for scoped component styling
- B. A testing module for automated browser tests
- C. A class organizing components, directives, and services
- D. A database module for connection management

Answer: C

Q590. What is Next.js?

- A. A testing tool for automated end-to-end browser tests
- B. A NoSQL document-oriented database management tool
- C. A CSS framework for responsive web design layouts
- D. A React framework supporting SSR and static generation

Answer: D

Q591. What is the difference between SQL and NoSQL databases?

- A. SQL is designed exclusively for web apps only
- B. They are identical in data model and usage
- C. NoSQL is always faster in every situation
- D. SQL uses tables with schemas; NoSQL uses flexible models

Answer: D

Q592. What is RESTful routing?

- A. A CSS routing system for navigating style sheets
- B. Organizing endpoints around resources via HTTP methods
- C. A database query method for retrieving records
- D. Random URL patterns without any convention

Answer: B

Q593. What is authentication vs authorization?

- A. Authentication verifies identity; authorization controls access
- B. Authentication controls access; authorization verifies identity
- C. Neither is required for modern web application security
- D. They are the same process with different names

Answer: A

Q594. What is an ORM?

- A. A testing tool for automated browser tests
- B. Object-Relational Mapping for querying via objects
- C. A CSS framework for responsive page layouts
- D. A type of relational database engine

Answer: B

Q595. What is the purpose of environment variables?

- A. To store config values like API keys outside code
- B. To create and render HTML elements dynamically
- C. To style web pages with dynamic CSS rules
- D. To define CSS variables for theming support

Answer: A

Q596. What is a session in web development?

- A. A server-side mechanism storing user data across requests
- B. A CSS class applied to styled HTML elements
- C. A database table for persisting application records
- D. A JavaScript function for handling user interactions

Answer: A

Q597. What is rate limiting?

- A. Limiting the speed of CSS animations
- B. Limiting the page load speed for mobile users
- C. Restricting API requests a client can make per period
- D. Limiting the total size of a database instance

Answer: C

Q598. What is the MVC pattern?

- A. A JavaScript library for DOM manipulation tasks
- B. A type of database for NoSQL document storage
- C. A CSS framework for responsive styling layouts
- D. Model-View-Controller, separating data and logic

Answer: D

Q599. What is a webhook?

- A. A CSS hook for attaching dynamic styles
- B. A database trigger running stored procedures
- C. A JavaScript hook for managing component state
- D. An HTTP callback sending data on events

Answer: D

Q600. What is the purpose of CORS middleware?

- A. To style HTTP responses with CSS formatting
- B. To handle CORS headers allowing cross-domain requests
- C. To log errors for debugging server-side issues
- D. To compress data for faster network transfer

Answer: B

Q601. What is the difference between SSR and CSR?

- A. They are identical in every rendering aspect
- B. SSR generates HTML on server; CSR generates in browser
- C. CSR is always faster in initial page load performance
- D. SSR is unable to use any JavaScript on the page

Answer: B

Q602. What is a REST controller?

- A. A component handling HTTP requests for REST endpoints
- B. A database controller for managing query execution
- C. A sleep controller for pausing execution
- D. A CSS component for building styled page layouts

Answer: A

Q603. What is server-side validation?

- A. Validating HTML syntax for standards compliance
- B. Validating CSS on the server for correctness
- C. Checking browser compatibility for CSS features
- D. Checking data on server to ensure requirements

Answer: D

Q604. What is the purpose of a .env file?

- A. To define HTML entities for special character display
- B. To style web pages with dynamic CSS variables
- C. To store environment-specific config like API keys
- D. To configure CSS preprocessor compilation options

Answer: C

Q605. What is a cron job?

- A. A scheduled task running automatically at intervals
- B. A database backup run at a specific time point
- C. A JavaScript timeout for delaying code execution
- D. A CSS animation timer for scheduling effects

Answer: A

Q606. What is the purpose of logging in server-side applications?

- A. Displaying logs visually in the browser console
- B. Recording events and errors for debugging and monitoring
- C. Creating log files that track CSS style changes
- D. Logging into user accounts through authentication

Answer: B

Q607. What is SSR hydration?

- A. A database operation for replicating table data
- B. A CSS technique for fluid responsive layout design
- C. Attaching JS event handlers to server-rendered HTML
- D. Adding water to servers for cooling systems

Answer: C

Q608. What is a reverse proxy in server-side programming?

- A. A database proxy for managing connection pools
- B. A backwards web server that inverts requests
- C. A CSS preprocessor for compiling style files
- D. A server forwarding requests to backend servers

Answer: D

Q609. What is the purpose of body parsing middleware?

- A. To parse CSS files for style extraction
- B. To parse HTML body elements for rendering content
- C. To parse database queries for execution planning
- D. To extract and parse HTTP request bodies into objects

Answer: D

Q610. What is the difference between stateless and stateful server-side apps?

- A. Stateless is always better for every use case scenario
- B. Stateless apps store no session data; stateful apps maintain it
- C. Stateful apps are completely unable to scale horizontally
- D. They are identical in session management behavior

Answer: B

Q611. What is a foreign key?

- A. A column linking two tables via another primary key
- B. A key imported from another country or locale
- C. An encrypted key used for securing data at rest
- D. A backup key for recovering lost database records

Answer: A

Q612. What is database indexing?

- A. Creating structures that speed up data retrieval
- B. Numbering database rows in sequential order
- C. Organizing files on disk for better storage
- D. A CSS technique for ordering layout elements

Answer: A

Q613. What is a JOIN in SQL?

- A. A JavaScript method for merging arrays together
- B. Combining multiple CSS files into a single bundle
- C. A string concatenation operator for joining text
- D. A clause combining rows from tables based on columns

Answer: D

Q614. What is the difference between INNER JOIN and LEFT JOIN?

- A. LEFT JOIN returns only the matching rows from both sources
- B. They are identical in their result sets
- C. INNER JOIN includes NULL values from non-matching entries
- D. INNER returns matches only; LEFT returns all left rows plus matches

Answer: D

Q615. What is database normalization?

- A. Normalizing CSS values for cross-browser consistency
- B. Formatting JSON data for human-readable pretty output
- C. Making data look normal in the display output
- D. Organizing data to reduce redundancy via normal forms

Answer: D

Q616. What is an ODM?

- A. A JavaScript operator for object comparison
- B. A CSS framework for responsive layout design
- C. A type of relational database engine
- D. Object-Document Mapper for objects and NoSQL docs

Answer: D

Q617. What is a transaction in databases?

- A. A CSS transition for animating property changes
- B. A data transfer between server and client machines
- C. A sequence of operations treated as a single unit
- D. A financial payment between two bank accounts

Answer: C

Q618. What are ACID properties?

- A. A type of document-oriented database engine system for storing data
- B. A security protocol for encrypting network traffic between services
- C. Atomicity, Consistency, Isolation, Durability for reliable transactions
- D. A chemistry concept related to pH levels

Answer: C

Q619. What is Redis commonly used for?

- A. As a primary database for all application data
- B. As an in-memory store for caching and sessions
- C. As a CSS preprocessor for compiling styles
- D. As a web server for hosting static content

Answer: B

Q620. What is SQL injection?

- A. A vulnerability where malicious SQL is inserted into queries
- B. A database optimization technique for faster lookups
- C. A JavaScript technique for dynamically building queries
- D. Injecting CSS styles into SQL query statements

Answer: A

Q621. What is CSRF (Cross-Site Request Forgery)?

- A. An attack tricking users into actions on authenticated sites
- B. A CSS framework for building styled web components
- C. A JavaScript error caused by invalid syntax in code
- D. A database attack targeting table join vulnerabilities

Answer: A

Q622. How is CSRF typically prevented?

- A. Disabling JavaScript execution in the browser
- B. Using anti-CSRF tokens to verify request origin
- C. Using CSS styles to hide form submission buttons
- D. Using longer and more complex user passwords

Answer: B

Q623. What is Content Security Policy (CSP)?

- A. A database policy for query execution permissions
- B. A JavaScript policy for variable access control
- C. A CSS policy for restricting style inheritance
- D. An HTTP header restricting which resources load

Answer: D

Q624. What is JWT (JSON Web Token)?

- A. A database token for authenticating query access
- B. A compact self-contained token for secure data transfer
- C. A CSS token for custom property values
- D. A JavaScript variable for storing session data

Answer: B

Q625. What is the principle of least privilege?

- A. Disabling all non-essential application features
- B. Giving admin access to every application user
- C. Granting users only the minimum needed permissions
- D. Using the simplest possible password for accounts

Answer: C

Q626. What is input sanitization?

- A. Compressing HTML to reduce the total page size
- B. Formatting JSON data for consistent key ordering
- C. Removing or encoding dangerous characters from input
- D. Cleaning CSS code by removing unused selectors

Answer: C

Q627. What is CORS and why is it a security concern?

- A. A JavaScript library for making HTTP request calls
- B. Cross-Origin Resource Sharing; misconfiguration exposes data
- C. A CSS framework for cross-browser styling consistency
- D. A NoSQL database system for document-based storage

Answer: B

Q628. What is a brute force attack?

- A. A CSS technique for overriding existing style rules
- B. Systematically trying all possible password combinations
- C. A physical attack on server hardware infrastructure
- D. A network protocol for secure data transmission

Answer: B

Q629. What is the purpose of the HttpOnly cookie flag?

- A. To encrypt the cookie data using symmetric encryption
- B. To prevent JavaScript from accessing the cookie
- C. To make cookies only work over HTTP connections
- D. To make the cookie persistent across browser sessions

Answer: B

Q630. What is OWASP?

- A. A web framework for building server applications
- B. A CSS standard for defining stylesheet specifications
- C. A relational database system for enterprise apps
- D. The Open Web Application Security Project guidelines

Answer: D

Q631. What is the difference between FCP and LCP?

- A. LCP measures only the first pixel painted on screen
- B. FCP is always slower than LCP in page loading time
- C. FCP is when first content appears; LCP is when largest renders
- D. They are identical metrics measuring the same event

Answer: C

Q632. What is code splitting?

- A. Breaking JavaScript bundles into on-demand chunks
- B. Dividing CSS files into smaller partial imports
- C. Database partitioning for distributing table rows
- D. Splitting HTML files across different server endpoints

Answer: A

Q633. What is the purpose of loading='lazy' on images?

- A. To defer loading off-screen images until scrolled to
- B. To load images using CSS background properties
- C. To compress images for smaller file sizes
- D. To slow down image rendering on the page

Answer: A

Q634. What is GZIP compression?

- A. A JavaScript minifier for reducing bundle sizes
- B. A CSS compression tool for minifying stylesheets
- C. Image compression for reducing file sizes
- D. A server-side method reducing transferred file sizes

Answer: D

Q635. What are Core Web Vitals?

- A. A CSS framework for building optimized page layouts
- B. A testing tool for measuring code quality metrics
- C. Google metrics for UX: LCP, INP, and CLS scores
- D. Core JavaScript functions built into the language

Answer: C

Q636. What is render-blocking CSS?

- A. CSS that blocks images from loading on page
- B. CSS that blocks JavaScript from executing code
- C. CSS preventing rendering until fully downloaded
- D. CSS that blocks individual element visibility

Answer: C

Q637. What is prefetching?

- A. A CSS property for preloading font resources
- B. A JavaScript fetch method for making API requests
- C. Fetching data from databases in advance
- D. Loading resources in advance the user will likely need

Answer: D

Q638. What is Cumulative Layout Shift (CLS)?

- A. A CSS layout property for element positioning
- B. A database migration for moving records between tables
- C. A metric measuring unexpected visual shifts during loading
- D. A JavaScript shift operation for bitwise manipulation

Answer: C

Q639. What is the purpose of a service worker in performance?

- A. To minify JavaScript bundles for production
- B. To manage CSS stylesheets and style rules
- C. To compress images for reduced file sizes
- D. To cache assets and serve them offline for speed

Answer: D

Q640. What is tree shaking?

- A. A CSS technique for optimizing selector performance
- B. Removing tree images from web page content
- C. Eliminating unused code from JavaScript bundles
- D. A database optimization for removing unused indexes

Answer: C

Q641. What is integration testing?

- A. Testing CSS integration across page components
- B. Testing only the user interface visual appearance
- C. Testing individual functions in complete isolation
- D. Testing how multiple modules work together as a group

Answer: D

Q642. What is end-to-end (E2E) testing?

- A. Testing from start to finish of the codebase
- B. Testing exclusively the backend server logic only
- C. Testing the complete application from user perspective
- D. Testing only a single API endpoint in isolation

Answer: C

Q643. What is Jest?

- A. A JavaScript testing framework for unit and integration tests
- B. A web server for hosting test environments and staging
- C. A database tool for testing query performance and accuracy
- D. A CSS framework for responsive web design

Answer: A

Q644. What is Cypress used for?

- A. API testing only without any browser interaction
- B. End-to-end testing of web apps in a real browser
- C. CSS processing and compilation of style files
- D. Database testing for query performance and accuracy

Answer: B

Q645. What is mocking in testing?

- A. Making fun of poorly written source code
- B. Creating fake implementations to isolate tested code
- C. A debugging method tracking variable value changes
- D. A CSS technique simulating dynamic state changes

Answer: B

Q646. What is code coverage?

- A. Covering code with explanatory inline comments
- B. A security feature for protecting sensitive source code files
- C. A metric measuring what percentage of code tests execute
- D. A CSS overlay property for stacking visual layer elements

Answer: C

Q647. What is Test-Driven Development (TDD)?

- A. Writing tests before code, then implementing to pass
- B. Developing tests after all code is fully written
- C. A debugging technique for tracing execution flow
- D. A testing framework for running automated test suites

Answer: A

Q648. What is a test suite?

- A. A hotel suite used for testing software
- B. A debugging tool for inspecting test output data
- C. A collection of related test cases grouped together
- D. A CSS test framework for visual regression testing

Answer: C

Q649. What is snapshot testing?

- A. Comparing current output against a stored reference
- B. A database backup for point-in-time recovery
- C. A version control feature for branching code
- D. Taking screenshots of the application pages

Answer: A

Q650. What is the purpose of assertion libraries?

- A. To assert database values match expected outputs
- B. To create assertions within HTML for validation
- C. To provide methods verifying expected code results
- D. To assert CSS values match the design mockups

Answer: C

Q651. What is CI/CD?

- A. A JavaScript library for building interactive web interfaces
- B. A CSS framework for rapid web development workflow
- C. Continuous Integration and Deployment with automated testing
- D. A relational database system for web application data

Answer: C

Q652. What is Docker?

- A. A platform for running apps in isolated containers
- B. A database tool for managing data migrations
- C. A JavaScript runtime for server-side execution
- D. A CSS framework for building grid-based layouts

Answer: A

Q653. What is a container in deployment?

- A. A CSS container element for grouping content
- B. A JavaScript object for storing config values
- C. A database container for grouping related tables
- D. A standalone package with everything to run an app

Answer: D

Q654. What is the purpose of a Dockerfile?

- A. To document CSS style guidelines for the team
- B. To configure a database connection and schema
- C. To define instructions for building a Docker image
- D. To create documentation for the project codebase

Answer: C

Q655. What is AWS?

- A. Amazon Web Services, a cloud platform for hosting
- B. A CSS framework for responsive web design layouts
- C. A JavaScript library for building user interfaces
- D. A web browser for viewing internet content

Answer: A

Q656. What is the difference between PaaS and IaaS?

- A. PaaS provides a platform; IaaS provides raw infrastructure
- B. PaaS provides infrastructure; IaaS provides platforms
- C. They are identical in their service model
- D. IaaS provides a platform for application deployment

Answer: A

Q657. What is an SSL/TLS certificate?

- A. A JavaScript license for commercial software distribution
- B. A CSS certification for web design professionals
- C. A certificate enabling HTTPS by authenticating the server
- D. A database certificate for verifying data integrity

Answer: C

Q658. What are environment variables used for in deployment?

- A. Defining CSS variables for different color schemes
- B. Styling web pages with environment-specific themes
- C. Creating HTML elements for different layout contexts
- D. Storing config values like API keys per environment

Answer: D

Q659. What is a reverse proxy in deployment?

- A. A testing tool for simulating server-side behavior
- B. A CSS technique for reversing element rendering order
- C. A backwards server that inverts incoming requests
- D. A server forwarding requests with load balancing and SSL

Answer: D

Q660. What is a rollback in deployment?

- A. Reverting to a previous version when deploy fails
- B. Rolling back CSS changes to a previous version
- C. A database operation for undoing transaction changes
- D. A JavaScript method for reverting variable values

Answer: A

Q661. What is the difference between REST and GraphQL?

- A. REST uses fixed endpoints; GraphQL uses one for exact data
- B. They are identical in data fetching approach
- C. GraphQL is always faster than REST in every scenario
- D. REST is a newer technology than GraphQL overall

Answer: A

Q662. What is API rate limiting?

- A. Restricting requests a client can make per time period
- B. Limiting the available API features and endpoints
- C. Limiting the number of available API endpoints
- D. Limiting the maximum API response payload size

Answer: A

Q663. What is API versioning?

- A. A JavaScript version number for runtime environments
- B. Updating CSS versions for browser compatibility
- C. A database version for tracking schema migrations
- D. Managing API versions for backward compatibility

Answer: D

Q664. What is OAuth used for with Web APIs?

- A. Delegated authorization for third-party app access
- B. A CSS framework for responsive web page layouts
- C. A JavaScript library for building authentication forms
- D. A database protocol for encrypted connection management

Answer: A

Q665. What is a webhook?

- A. A CSS hook for attaching dynamic style rules
- B. A JavaScript event listener for DOM interactions
- C. A database trigger firing on record changes
- D. A callback URL receiving HTTP POST on events

Answer: D

Q666. What is pagination in APIs?

- A. A JavaScript page event for tracking navigation
- B. Dividing large result sets into smaller pages
- C. Page layout and formatting for print media
- D. A CSS page break controlling print output layout

Answer: B

Q667. What is HATEOAS in REST?

- A. A database concept for relating tables together
- B. A CSS framework for building grid layouts
- C. Responses include links to related resources
- D. A JavaScript pattern for managing app state

Answer: C

Q668. What is API documentation?

- A. Writing CSS comments for code documentation
- B. Database documentation for schema design decisions
- C. Detailed descriptions of endpoints and parameters
- D. HTML documentation for browser compatibility

Answer: C

Q669. What is the purpose of request headers in API calls?

- A. Database headers for defining column schema types
- B. Providing metadata like auth tokens and content type
- C. Styling the request with CSS formatting rules
- D. Creating HTML headers for page title and navigation

Answer: B

Q670. What is a RESTful resource?

- A. A server resource like CPU or memory allocation
- B. A CSS resource like a stylesheet or font file
- C. A JavaScript resource loaded via script elements
- D. An entity identified by URL manipulated via HTTP

Answer: D

Q671. What is a monorepo?

- A. A type of database for centralized data storage
- B. A CSS methodology for organizing style declarations
- C. Multiple related projects stored in one repository
- D. A single repository for one individual project

Answer: C

Q672. What is Jamstack?

- A. A music stack for audio processing applications
- B. An architecture using JavaScript, APIs, and Markup
- C. A CSS framework for building styled web components
- D. A database stack for managing data storage layers

Answer: B

Q673. What is a headless CMS?

- A. A CMS providing content via API without a frontend
- B. A CSS framework for headless browser rendering
- C. A CMS without advanced content management features
- D. A CMS without a header section in the layout

Answer: A

Q674. What is Web Assembly (WASM)?

- A. A binary format enabling near-native browser performance
- B. A JavaScript engine for parsing and executing scripts
- C. A web framework for server-side rendering pages
- D. A CSS assembler for compiling stylesheet syntax

Answer: A

Q675. What is the purpose of ESLint?

- A. Managing databases through a graphical admin interface
- B. Compiling JavaScript into optimized machine code
- C. Minifying CSS files for reduced production bundle sizes
- D. Analyzing JavaScript to find problems and enforce standards

Answer: D

Q676. What is Prettier?

- A. A testing framework for automated end-to-end testing
- B. A CSS framework for responsive web design layouts
- C. An opinionated code formatter enforcing consistent style
- D. A database tool for formatting query output results

Answer: C

Q677. What is a design system?

- A. A testing framework for running automated test suites
- B. A CSS file for global application styling needs
- C. A database schema for normalized table structure design
- D. A comprehensive set of reusable components and guidelines

Answer: D

Q678. What is Storybook?

- A. A CSS framework for styling web page component layouts
- B. A tool for building and documenting UI components in isolation
- C. A book reading app for digital publications
- D. A database tool for visual query builder access and management

Answer: B

Q679. What is a micro-frontend architecture?

- A. A component library for sharing reusable UI element sets
- B. A small frontend with limited feature scope overall
- C. Independently deployable micro apps composing the frontend
- D. A CSS grid layout for building responsive page designs

Answer: C

Q680. What is Deno?

- A. A modern JS/TS runtime with built-in security and TypeScript
- B. A dinosaur-themed web application framework library
- C. A CSS preprocessor for compiling advanced style syntax
- D. A NoSQL database for document-based data storage systems

Answer: A

Q681. What distinguishes a full-stack developer from a front-end developer?

- A. Full-stack only handles database design and server administration
- B. Front-end developers also write all the server-side application code
- C. Full-stack developers exclusively focus on mobile app development work
- D. Full-stack works on both client and server sides of applications

Answer: D

Q682. What is the purpose of version control in web development?

- A. To track changes and enable collaboration among developers
- B. To validate HTML and CSS against web standard guidelines
- C. To minify and compress files for better site performance
- D. To automatically deploy code to production servers quickly

Answer: A

Q683. Which of the following best describes a Content Management System?

- A. A tool for writing and compiling programming language code
- B. A framework for building real-time communication applications
- C. A platform for creating and managing digital content easily
- D. A service for monitoring web server performance and uptime

Answer: C

Q684. What is the role of a package manager in web development?

- A. To host and serve static web pages on a remote cloud server
- B. To encrypt and secure data transmitted between client nodes
- C. To install, update, and manage project dependency libraries
- D. To design responsive layouts for mobile devices and tablets

Answer: C

Q685. What does the term responsive web design mean?

- A. Websites that only function properly on the latest browser versions
- B. Websites that respond to voice commands from the user quickly
- C. Websites that automatically adapt layout to different screen sizes
- D. Websites that load within two seconds on any connection speed

Answer: C

Q686. What is the purpose of a wireframe in web development?

- A. To outline the basic structure and layout of a page visually
- B. To define the database schema for back-end data persistence
- C. To compress images for faster loading times on web servers
- D. To encrypt user authentication tokens stored in the browser

Answer: A

Q687. Which LAMP stack component serves as the web server?

- A. Linux operating system
- B. PHP scripting lang
- C. MySQL data storage
- D. Apache HTTP server

Answer: D

Q688. What does the term single-page application refer to?

- A. An app that dynamically rewrites the page without full reloads
- B. A website that contains only one static HTML landing page total
- C. A web app that can only be accessed on a single device at once
- D. A site designed exclusively for mobile phone portrait view only

Answer: A

Q689. What is the main benefit of using a CSS preprocessor?

- A. It replaces JavaScript entirely for adding interactivity logic
- B. It automatically generates accessible HTML from plain text input
- C. It eliminates the need for any media queries in responsive work
- D. It adds variables, nesting, and functions to standard CSS code

Answer: D

Q690. What is web accessibility primarily concerned with?

- A. Ensuring websites are usable by people with various disabilities
- B. Optimizing search engine rankings through keyword meta tag usage
- C. Preventing unauthorized access to web server admin panel areas
- D. Making websites load faster on slow network connections today

Answer: A

Q691. What is the difference between HTTP GET and POST methods?

- A. GET encrypts data automatically while POST sends it in plain text
- B. GET retrieves data via URL while POST submits data in the body
- C. GET is used only for images while POST is used only for text data
- D. GET sends data in the body while POST appends it to the URL path

Answer: B

Q692. What is the purpose of an HTTP status code in the 3xx range?

- A. To report that the server encountered an internal processing error
- B. To signal that the requested resource has moved or needs redirect
- C. To confirm that the client request was malformed or unauthorized
- D. To indicate a successful operation completed without any errors

Answer: B

Q693. What role does TCP play in web communication?

- A. It ensures reliable ordered delivery of data between two hosts
- B. It encrypts all data packets before they leave the sender device
- C. It compresses image files to reduce bandwidth usage on transfers
- D. It assigns human-readable names to numerical server IP addresses

Answer: A

Q694. What is the purpose of a CDN in web architecture?

- A. To write server-side code for dynamic page generation logic
- B. To provide database replication between primary and backup nodes
- C. To manage version control repositories for development teams
- D. To distribute content across servers closer to end users globally

Answer: D

Q695. What does a reverse proxy server do?

- A. It sits before origin servers and forwards client requests to them
- B. It directly serves web pages to the user browser without caching
- C. It translates domain names into IP addresses for DNS resolution
- D. It compiles source code into executable binaries on the server

Answer: A

Q696. What is the purpose of the HTTP OPTIONS method?

- A. To describe the communication options available for a resource
- B. To upload a new file or document to the server storage location
- C. To establish a persistent WebSocket connection with the server
- D. To delete a specified resource from the server permanently today

Answer: A

Q697. What does stateless mean in the context of HTTP?

- A. The client must re-authenticate for every single page navigation
- B. Each request is independent and carries no memory of prior ones
- C. The server cannot store any files on its local disk permanently
- D. The connection between client and server is never encrypted once

Answer: B

Q698. What is the role of a load balancer in web architecture?

- A. To encrypt all stored database records using advanced hash algorithms
- B. To minify JavaScript files before serving them to client browsers
- C. To distribute incoming traffic across multiple backend server nodes
- D. To automatically generate SSL certificates for domain name hosting

Answer: C

Q699. What is the difference between HTTP/1.1 and HTTP/2?

- A. HTTP/1.1 supports server push while HTTP/2 completely removed it
- B. HTTP/2 removes support for HTTPS and only allows plain HTTP usage
- C. HTTP/2 requires a different port number than HTTP/1.1 for access
- D. HTTP/2 introduces multiplexing allowing multiple parallel requests

Answer: D

Q700. What is a WebSocket and how does it differ from HTTP?

- A. WebSocket is a one-way protocol that only allows server to client data
- B. WebSocket replaces TCP entirely with a new custom transport protocol
- C. WebSocket enables persistent two-way communication between endpoints
- D. WebSocket only works with static HTML pages and cannot handle JSON

Answer: C

Q701. What is the difference between the <div> and elements?

- A. <div> can only contain text while can contain any HTML element
- B. <div> is deprecated in HTML5 while is the recommended replacement
- C. <div> is block-level and is inline in default display behavior
- D. <div> is inline and is block-level in default display behavior

Answer: C

Q702. What is the purpose of semantic HTML elements like <article> and <section>?

- A. They enable JavaScript functionality without writing any script code
- B. They provide meaning to content structure for accessibility and SEO
- C. They replace all div elements and are required for valid HTML5 pages
- D. They apply automatic CSS styling without needing any stylesheet rules

Answer: B

Q703. What does the required attribute do on a form input element?

- A. It automatically formats the input data into proper currency notation
- B. It sets a default placeholder text value inside the input field area
- C. It prevents the form from submitting until the field has been filled
- D. It hides the input field until the user clicks a reveal toggle button

Answer: C

Q704. How does the <picture> element differ from a standard tag?

- A. It automatically applies image filters like blur and grayscale effects
- B. It creates an image gallery carousel with built-in navigation controls
- C. It provides multiple sources for responsive image selection by browser
- D. It only supports SVG format images and cannot display any raster type

Answer: C

Q705. What is the purpose of the data-* attribute in HTML5?

- A. To automatically bind JavaScript event handlers to specific elements
- B. To define database connection strings directly within HTML elements
- C. To store custom data private to the page or application on elements
- D. To encrypt sensitive user information within the DOM tree structure

Answer: C

Q706. What does the <template> element do in HTML5?

- A. It defines email templates for automated notification sending tasks
- B. It holds client-side content not rendered until activated by scripts
- C. It applies a predefined visual theme to the entire web page document
- D. It creates a reusable server-side component for backend rendering

Answer: B

Q707. What is the role of the <meta charset='UTF-8'> tag?

- A. It sets the background color of the web page to a default value
- B. It links an external font library to the current page for styling
- C. It defines the viewport width for mobile responsive design layouts
- D. It specifies the character encoding used for the HTML document text

Answer: D

Q708. What is the difference between the and tags?

- A. indicates semantic importance while is only visual bold
- B. Both tags are completely identical with no semantic differences at all
- C. is deprecated in HTML5 and is the recommended replacement
- D. indicates importance while is only for visual bold style

Answer: A

Q709. What attribute makes an input field read-only but still submittable?

- A. readonly attribute
- B. disabled attribute
- C. hidden attribute
- D. required attribute

Answer: A

Q710. What is the purpose of the <fieldset> and <legend> elements in forms?

- A. They group related form controls and provide a caption for the group
- B. They add animated transitions between different form sections on page
- C. They create dropdown menus and label each option with descriptive text
- D. They define hidden input fields that store session token values only

Answer: A

Q711. What is the difference between margin and padding?

- A. Padding is inside the border while margin is outside the border area
- B. Padding is only for horizontal spacing while margin is only vertical
- C. Margin only applies to inline elements while padding applies to blocks
- D. Margin is inside the border while padding is outside the border area

Answer: A

Q712. What does the z-index property control in CSS?

- A. The stacking order of positioned elements along the depth axis shown
- B. The rotation angle of an element around its central transformation point
- C. The zoom level of the element when displayed in the browser viewport
- D. The horizontal position of an element relative to its parent container

Answer: A

Q713. How does CSS specificity determine which styles are applied?

- A. Only important declarations are considered and all others are ignored
- B. Styles are applied in the order they appear in the HTML source code
- C. Inline styles, IDs, classes, and elements are weighted in that order
- D. The last stylesheet linked in the head always overrides all previous

Answer: C

Q714. What is the CSS box model composed of?

- A. Font, color, background, and shadow properties of text elements
- B. Header, content, sidebar, and footer sections of a webpage layout
- C. Flexbox, grid, float, and position layout mechanisms for styling
- D. Content, padding, border, and margin layers around every element

Answer: D

Q715. What does the position: absolute property do?

- A. Positions element relative to its nearest positioned ancestor element
- B. Positions element relative to the browser viewport scrolling along
- C. Positions element in normal document flow with offset adjustments
- D. Positions element at a fixed spot that ignores scrolling movements

Answer: A

Q716. What is the difference between display: none and visibility: hidden?

- A. display: none removes from flow while visibility: hidden keeps the space
- B. display: none only works on block elements while visibility works on all
- C. display: none hides but keeps space while visibility: hidden removes it
- D. Both properties work identically and there is no difference between them

Answer: A

Q717. How do CSS transitions work?

- A. They animate property changes smoothly over a specified time duration
- B. They only work with JavaScript and cannot be triggered by CSS alone
- C. They require keyframe definitions for every single property animation
- D. They instantly change property values without any animation effect shown

Answer: A

Q718. What is the purpose of the CSS calc() function?

- A. To compute the specificity score of a given CSS selector expression
- B. To count the number of child elements inside a parent container block
- C. To perform mathematical calculations to determine CSS property values
- D. To calculate page load time and display it inside a status bar area

Answer: C

Q719. What does the overflow property control?

- A. The order in which overlapping positioned elements stack on the page
- B. The speed at which CSS transitions animate between two style states
- C. The maximum number of characters shown in a text input field element
- D. How content that exceeds its container dimensions is handled and shown

Answer: D

Q720. What is a CSS pseudo-class?

- A. A CSS property that generates content before or after element content
- B. A keyword added to selectors specifying a special state of an element
- C. A class that is defined in JavaScript rather than in the CSS stylesheet
- D. A method for importing styles from one stylesheet into another one used

Answer: B

Q721. How does the minmax() function work in CSS Grid?

- A. It sets the minimum and maximum font size for responsive text typography
- B. It calculates the minimum and maximum viewport size for media queries
- C. It computes the smallest and largest margin values for element spacing
- D. It defines a size range with minimum and maximum values for grid tracks

Answer: D

Q722. What is the purpose of the grid-template-areas property?

- A. It automatically generates grid lines based on the content size of items
- B. It defines the total number of grid items that can be placed in a grid
- C. It sets background images for specific areas within the grid container
- D. It creates named grid areas for intuitive layout placement of elements

Answer: D

Q723. How does the flex shorthand property work?

- A. It sets the flexibility of the container rather than individual flex items
- B. It combines flex-wrap, flex-flow, and flex-direction into one property
- C. It combines flex-grow, flex-shrink, and flex-basis into one declaration
- D. It controls the visual order of flex items without changing DOM sequence

Answer: C

Q724. What is the difference between auto-fit and auto-fill in CSS Grid?

- A. auto-fill expands items to fill space while auto-fit keeps empty tracks
- B. auto-fill requires explicit row counts while auto-fit calculates them all
- C. auto-fit only works with fixed pixel values while auto-fill uses percents
- D. auto-fit expands items to fill space while auto-fill keeps empty tracks

Answer: D

Q725. How do CSS transforms affect the document flow?

- A. Transforms cause all sibling elements to reposition around the changed one
- B. Transforms visually change elements without affecting document flow at all
- C. Transforms remove the element entirely from the normal document flow
- D. Transforms only work when the element has position absolute or fixed set

Answer: B

Q726. What is the purpose of the clamp() function in CSS?

- A. It clamps scroll position to prevent users from scrolling past content
- B. It limits the number of CSS animations that can run at the same time
- C. It restricts the number of flex items that appear on a single row line
- D. It sets a value that scales between a minimum and maximum with preferred

Answer: D

Q727. How does the order property work in flexbox?

- A. It defines the rendering priority determining which items paint on top
- B. It sets alphabetical sorting order for text content within flex items
- C. It specifies the order in which CSS transitions are applied to each item
- D. It changes the visual order of flex items without altering the DOM order

Answer: D

Q728. What is the aspect-ratio property used for?

- A. To control the animation speed ratio between different keyframe points
- B. To set the ratio of padding to margin for consistent element spacing
- C. To define the ratio between font size and line height in text blocks
- D. To maintain a width-to-height ratio for an element automatically in CSS

Answer: D

Q729. How does the place-items shorthand work in Grid?

- A. It sets the placement algorithm to either sparse or dense packing mode
- B. It combines align-items and justify-items into a single declaration used
- C. It positions grid items at specific pixel coordinates on the web page
- D. It places items outside the grid container in an overflow hidden region

Answer: B

Q730. What is the difference between grid-column and grid-area?

- A. grid-area is deprecated and grid-column is the recommended replacement
- B. grid-area only works with named areas while grid-column works with lines
- C. grid-column sets column placement while grid-area sets full area placement
- D. grid-column is for flexbox while grid-area is exclusively for CSS Grid

Answer: C

Q731. What is the difference between let and var in JavaScript?

- A. let is used for numbers only while var works with all data type values
- B. let cannot be reassigned while var allows reassignment of the variable
- C. let is function-scoped while var is block-scoped in JavaScript code
- D. let is block-scoped while var is function-scoped in JavaScript code

Answer: D

Q732. What is hoisting in JavaScript?

- A. Moving DOM elements to the top of the page during initial rendering
- B. Moving declarations to the top of their scope during the compile phase
- C. Raising the priority of event listeners above other queued callbacks
- D. Uploading JavaScript files to a remote server for production hosting

Answer: B

Q733. What is the difference between == and === in JavaScript?

- A. == coerces types before comparison while === checks type and value both
- B. == is used for strings while === is used for numbers and boolean values
- C. == compares only types while === compares only values of the operands
- D. == is deprecated in modern JavaScript while === is the only valid option

Answer: A

Q734. What is a callback function in JavaScript?

- A. A function that is called automatically when the page finishes loading
- B. A function passed as argument to another function for later execution
- C. A function that calls itself recursively until a condition is satisfied
- D. A function that returns another function as its computed return value

Answer: B

Q735. How does the ternary operator work?

- A. It creates three variables at once using a single declaration statement
- B. It evaluates a condition and returns one of two values based on result
- C. It performs three separate operations on an array simultaneously today
- D. It loops through an array exactly three times for processing elements

Answer: B

Q736. What is destructuring assignment in JavaScript?

- A. Deleting properties from an object permanently from the memory storage
- B. Removing elements from arrays by overwriting them with undefined values
- C. Extracting values from arrays or objects into distinct variables easily
- D. Converting complex objects into simple primitive values automatically

Answer: C

Q737. What does the spread operator do?

- A. It expands an iterable into individual elements in function calls or arrays
- B. It distributes network requests evenly across multiple server endpoints
- C. It removes duplicate values from an array and returns a new unique set
- D. It creates a deep copy of nested objects including all inner references

Answer: A

Q738. What is the purpose of template literals in JavaScript?

- A. To allow embedded expressions and multiline strings using backtick syntax
- B. To define strict type templates for validating function parameter values
- C. To create reusable HTML templates that are rendered on the server side
- D. To generate PDF templates from HTML content for download by site users

Answer: A

Q739. How does the map() array method work?

- A. It filters array elements and returns only those that match a condition
- B. It modifies the original array in place by replacing each element value
- C. It creates a geographical map visualization from array coordinate data
- D. It creates a new array by applying a function to each element of source

Answer: D

Q740. What is the difference between null and undefined?

- A. null is an intentional empty value while undefined means no value assigned
- B. null is used only for objects while undefined is used only for primitives
- C. null triggers an error when accessed while undefined is silently ignored
- D. null is a declared variable with no value while undefined is not declared

Answer: A

Q741. What is the difference between Promise.all and Promise.allSettled?

- A. Promise.all waits for all but allSettled only waits for the first result
- B. Promise.all runs promises sequentially while allSettled runs them parallel
- C. Promise.all rejects on first failure while allSettled waits for all outcomes
- D. Promise.all only works with two promises while allSettled handles any count

Answer: C

Q742. How does event delegation work in JavaScript?

- A. It assigns separate event handlers to every individual child element node
- B. It delegates event processing to a separate web worker thread for speed
- C. It attaches a single handler to a parent to handle events from children
- D. It automatically prevents default browser behavior for all fired events

Answer: C

Q743. What is the purpose of the Symbol primitive type?

- A. To generate Unicode emoji characters for display in user interface text
- B. To create mathematical symbols for complex equation calculations in code
- C. To create unique identifiers that avoid property name collisions in objects
- D. To define currency symbols for formatting monetary values on web pages

Answer: C

Q744. How does the async iterator protocol work with for-await-of?

- A. It limits the iteration count based on available system memory resources
- B. It runs multiple for loops in parallel across different threads at once
- C. It allows asynchronous iteration over streams of data with await support
- D. It converts synchronous iterators into cached versions for faster access

Answer: C

Q745. What is the Proxy object used for in JavaScript?

- A. To create a network proxy for routing HTTP requests through a server
- B. To create backup copies of objects for undo functionality in the editor
- C. To define custom behavior for fundamental object operations like get/set
- D. To convert JavaScript objects into XML format for SOAP API integration

Answer: C

Q746. What is the difference between Object.freeze and Object.seal?

- A. freeze makes objects immutable forever while seal makes them temporary
- B. freeze is for arrays only while seal is for plain objects exclusively used
- C. freeze prevents all changes while seal allows modifying existing properties
- D. freeze works in strict mode only while seal works in any execution mode

Answer: C

Q747. How does the AbortController work with fetch requests?

- A. It converts fetch responses into readable streams for large file handling
- B. It automatically retries failed fetch requests up to three times in a row
- C. It batches multiple fetch requests into a single network call for speed
- D. It provides a way to cancel in-flight fetch requests using abort signals

Answer: D

Q748. What is the purpose of generators in JavaScript?

- A. They automatically generate documentation from JavaScript source code
- B. They compile JavaScript into machine code for better runtime performance
- C. They generate random numbers for use in cryptographic security functions
- D. They create functions that can pause and resume execution using yield

Answer: D

Q749. How does the MutationObserver API work?

- A. It monitors CSS style mutations and prevents unauthorized modifications
- B. It observes network requests and mutates their headers before sending
- C. It tracks variable mutations and logs them for debugging purposes only
- D. It watches for changes to the DOM tree and triggers callbacks for them

Answer: D

Q750. What is a WeakRef and when would you use it?

- A. A weak reference to an object that does not prevent garbage collection
- B. A reference to a CSS variable that resolves at runtime when accessed
- C. A reference to a DOM element that automatically updates when it moves
- D. A reference counted pointer similar to those in C++ or Rust languages

Answer: A

Q751. How does React's reconciliation algorithm work?

- A. It replaces the complete DOM with a new version on each state update call
- B. It diffs the virtual DOM against the previous version and patches changes
- C. It uses web workers to rebuild the DOM tree in a background thread process
- D. It compares the entire DOM tree node by node on every single render cycle

Answer: B

Q752. What is the purpose of React hooks?

- A. To hook into browser APIs that are not available in standard JavaScript
- B. To add state and lifecycle features to functional components without classes
- C. To create HTTP endpoint hooks for RESTful API communication with servers
- D. To connect React components to external CSS stylesheet files for styling

Answer: B

Q753. How does two-way data binding work in Angular?

- A. Binding only works for string values and not numbers or boolean data types
- B. Data flows only from the template to the component class using event calls
- C. Data is bound at compile time and cannot be updated after initial render
- D. Changes in model automatically reflect in view and vice versa in real time

Answer: D

Q754. What is Vue's reactivity system based on?

- A. Manual DOM manipulation with document query selector calls in components
- B. Server-sent events that push data changes from backend to the frontend
- C. Proxy-based observation that tracks dependencies and triggers re-renders
- D. Polling intervals that check for data changes every fixed time interval

Answer: C

Q755. What is the difference between controlled and uncontrolled components in React?

- A. Controlled components are server-rendered while uncontrolled are client-only
- B. Controlled components are class-based while uncontrolled are function-based
- C. Controlled components manage form data via state while uncontrolled use DOM
- D. Controlled components use Redux while uncontrolled use local component state

Answer: C

Q756. What is the purpose of a state management library like Redux?

- A. To manage the state of network connections and retry failed API call requests
- B. To handle database state synchronization between client and server layers
- C. To manage CSS state transitions and animations across all the components
- D. To manage shared application state in a predictable centralized store object

Answer: D

Q757. How does server-side rendering differ from client-side rendering in frameworks?

- A. SSR requires WebSocket connections while CSR uses standard HTTP request calls
- B. SSR only works with Angular while CSR is exclusive to React and Vue frameworks
- C. SSR generates HTML on the server for faster initial load and better SEO results
- D. SSR generates HTML on the client while CSR generates it on the server side

Answer: C

Q758. What is lazy loading in the context of frontend frameworks?

- A. Deferring component or module loading until they are actually needed on page
- B. Loading CSS styles lazily after all JavaScript has finished executing first
- C. Slowing down animations intentionally to create a relaxed user experience
- D. Delaying all network requests until the user explicitly clicks a load button

Answer: A

Q759. What is the purpose of Angular's dependency injection system?

- A. To automatically inject error handling code into every component method call
- B. To inject HTML content from external files into Angular template directives
- C. To inject CSS styles directly into component templates without stylesheets
- D. To provide services and dependencies to components without manual creation

Answer: D

Q760. What are React Server Components and how do they differ from client components?

- A. Server components only work in development mode and are removed in production
- B. Server components run on client with server-like APIs for database access
- C. Server components are deprecated and replaced by client-only rendering now
- D. Server components render on server and send serialized output to the client

Answer: D

Q761. What is the difference between SQL and NoSQL databases?

- A. NoSQL databases cannot handle relational data while SQL handles everything
- B. SQL is only for small datasets while NoSQL is only for large datasets used
- C. SQL databases are always faster than NoSQL for every possible query type
- D. SQL databases use tables with schemas while NoSQL uses flexible data models

Answer: D

Q762. What is the purpose of an ORM in backend development?

- A. To map database records to objects in code without writing raw SQL queries
- B. To manage real-time messaging between connected WebSocket client users
- C. To optimize CSS rendering in server-side rendered pages for performance
- D. To orchestrate multiple microservices into a single deployment pipeline

Answer: A

Q763. How does JWT authentication work?

- A. The client sends username and password with every single request it makes
- B. The server generates a new unique URL for each authenticated user session
- C. The server issues a signed token containing claims that the client stores
- D. The server stores session data in a database and sends a session cookie

Answer: C

Q764. What is the purpose of environment variables in backend applications?

- A. To define CSS custom properties that change based on the user environment
- B. To determine the operating system type for conditional CSS style loading
- C. To store configuration values that change between deployment environments
- D. To set the browser viewport dimensions for responsive testing of layouts

Answer: C

Q765. What is rate limiting and why is it important?

- A. Controlling the speed at which database queries execute on the server side
- B. Setting a maximum file upload size for requests sent to the server endpoint
- C. Limiting the CSS animation frame rate to improve browser performance speed
- D. Restricting the number of API requests a client can make in a time period

Answer: D

Q766. How does connection pooling improve database performance?

- A. It compresses database query results before sending them to the application
- B. It distributes database queries equally across all available server hardware
- C. It reuses existing database connections instead of creating new ones each time
- D. It pools CSS connections to share stylesheets across multiple page views

Answer: C

Q767. What is the purpose of input validation on the server side?

- A. To verify that client-side JavaScript is properly minified before caching
- B. To ensure incoming data meets expected formats and prevents malicious input
- C. To check that image files are in the correct resolution for display on page
- D. To validate that CSS selectors match existing HTML elements in the document

Answer: B

Q768. What is the difference between authentication and authorization?

- A. Authentication verifies identity while authorization checks access permissions
- B. Authentication checks permissions while authorization verifies identity first
- C. Authentication is for APIs only while authorization is for web pages exclusively
- D. Authentication uses tokens while authorization exclusively uses session cookies

Answer: A

Q769. What is GraphQL and how does it differ from REST?

- A. GraphQL requires WebSocket connections while REST uses standard HTTP methods
- B. GraphQL lets clients request exactly the data they need with a single query
- C. GraphQL only supports read operations while REST supports full CRUD actions
- D. GraphQL is a database language while REST is an API architecture style used

Answer: B

Q770. What is the purpose of logging in backend applications?

- A. To measure the length of CSS animations and report their duration in logs
- B. To create visual log-shaped decorations on the user interface of the website
- C. To record application events and errors for monitoring and debugging needs
- D. To generate automatic documentation from source code comment annotations

Answer: C

Q771. How does server-side caching improve application performance?

- A. It caches user passwords on the server to speed up authentication checks
- B. It stores JavaScript bundle files on CDN edge nodes for faster delivery
- C. It caches CSS files in the browser to avoid downloading them every visit
- D. It stores computed results in memory to avoid repeated expensive operations

Answer: D

Q772. What is the difference between stateful and stateless server architectures?

- A. Stateful servers retain client data between requests while stateless do not
- B. Stateful servers are always faster because they remember all previous data
- C. Stateless servers cannot use databases while stateful servers require them
- D. Stateful is for REST APIs only while stateless is for GraphQL APIs exclusively

Answer: A

Q773. How does server-side form validation complement client-side validation?

- A. Server validation runs first and client validation is applied after server
- B. Server validation replaces client validation entirely making it unnecessary
- C. Server validation ensures security since client validation can be bypassed
- D. Server validation is only needed for file uploads and not for text inputs

Answer: C

Q774. What is the purpose of an application server versus a web server?

- A. Application servers are for mobile apps while web servers are for websites only
- B. Web servers are deprecated and application servers have replaced them entirely
- C. Web servers handle HTTP requests while app servers execute business logic also
- D. Application servers only serve static files while web servers handle all logic

Answer: C

Q775. How do server-side redirects differ from client-side redirects?

- A. Server redirects happen via HTTP status codes while client uses JavaScript
- B. Client redirects are more secure because they happen within the browser
- C. Server redirects are slower because they require a full page reload always
- D. Server redirects only work with POST requests while client works with GET

Answer: A

Q776. What is the purpose of a task queue in server-side applications?

- A. To handle time-consuming tasks asynchronously outside the request cycle
- B. To queue CSS animations for smooth rendering on the server-side components
- C. To organize database migration tasks in the order they should be executed
- D. To schedule automated browser tests to run sequentially on a CI pipeline

Answer: A

Q777. How does server-side pagination work for large datasets?

- A. It loads all records into memory then splits them into pages on the server
- B. It caches the entire dataset in the browser and paginates using JavaScript
- C. It fetches only a subset of records per request using offset and limit values
- D. It creates separate database tables for each page of results automatically

Answer: C

Q778. What is the role of a reverse proxy like Nginx in server deployments?

- A. It provides a graphical user interface for managing database tables easily
- B. It converts server-side code into client-side JavaScript for the browser
- C. It reverses the order of HTTP headers to confuse potential attackers today
- D. It sits in front of servers handling load balancing, SSL, and static files

Answer: D

Q779. What is the purpose of database migrations in server-side development?

- A. To move the database server to a different physical location or data center
- B. To manage incremental changes to database schema in a version-controlled way
- C. To migrate data from SQL databases to NoSQL databases for better performance
- D. To create daily backup copies of the database for disaster recovery purposes

Answer: B

Q780. How do background jobs differ from synchronous request handling?

- A. Background jobs process tasks outside the request cycle without blocking it
- B. Background jobs run on the client side while synchronous handlers run on server
- C. Background jobs run immediately within the request cycle blocking the response
- D. Background jobs only handle database operations and not any other task types

Answer: A

Q781. What is database normalization and why is it important?

- A. It normalizes CSS values to consistent units across all stylesheet files
- B. It organizes data to reduce redundancy and improve data integrity in tables
- C. It converts all database text fields to normalized Unicode encoding formats
- D. It equalizes query execution times so all queries run at the same speed

Answer: B

Q782. How do database indexes improve query performance?

- A. They create visual index pages for the database administrator interface
- B. They automatically optimize SQL queries by rewriting them more efficiently
- C. They create data structures that allow the database to find rows faster
- D. They compress database tables to reduce the overall storage requirements

Answer: C

Q783. What is the difference between INNER JOIN and LEFT JOIN?

- A. INNER JOIN returns matching rows while LEFT JOIN includes unmatched left rows
- B. INNER JOIN is faster but less accurate while LEFT JOIN is slower but precise
- C. INNER JOIN returns all rows while LEFT JOIN returns only matching ones
- D. INNER JOIN is for two tables while LEFT JOIN works with three or more tables

Answer: A

Q784. What is a database transaction and what properties does it have?

- A. A financial transaction record stored in the database for accounting use
- B. A scheduled task that runs database maintenance operations periodically
- C. A log entry recording each time a user accesses the database connection
- D. A unit of work with ACID properties ensuring reliability and consistency

Answer: D

Q785. How does connection pooling benefit database-driven web applications?

- A. It distributes connections equally across all tables in the database schema
- B. It pools multiple databases into a single virtual database for unified access
- C. It creates a pool of pre-compiled SQL queries for faster execution speed
- D. It maintains reusable connections reducing overhead of creating new ones each time

Answer: D

Q786. What is the purpose of database views?

- A. To create virtual tables from queries for simplified data access and security
- B. To display database tables in a graphical user interface for administrators
- C. To store compressed backups of database tables for disaster recovery use
- D. To provide real-time visualization charts of database performance metrics

Answer: A

Q787. What is the difference between SQL and NoSQL databases?

- A. SQL databases are free while NoSQL databases always require paid commercial licenses
- B. SQL uses structured schemas and tables while NoSQL uses flexible data models
- C. SQL databases cannot scale while NoSQL databases scale infinitely with ease
- D. SQL is only for small applications while NoSQL is only for enterprise projects

Answer: B

Q788. What is an ORM and how does it simplify database interaction?

- A. It maps database tables to code objects enabling queries without raw SQL usage
- B. It organizes REST methods into a standardized naming convention for APIs
- C. It manages database backups by scheduling them at regular time intervals daily
- D. It monitors database performance and reports optimization recommendations to devs

Answer: A

Q789. What is the purpose of stored procedures in SQL databases?

- A. To save precompiled SQL code on the server for reuse and performance gains
- B. To store JavaScript procedures that run in the browser before page loading
- C. To store CSS procedures for server-side rendering of styled web page content
- D. To archive old database records in compressed format for long-term storage

Answer: A

Q790. How does Redis differ from traditional relational databases?

- A. Redis only supports read operations while relational databases support writes
- B. Redis is an in-memory key-value store optimized for speed and caching needs
- C. Redis requires strict schemas while relational databases are schema-flexible
- D. Redis stores all data on disk while relational databases use in-memory only

Answer: B

Q791. How does CSRF attack work and how can it be prevented?

- A. CSRF intercepts network traffic between client and server to modify data
- B. CSRF brute-forces user passwords by trying common combinations repeatedly
- C. CSRF exploits stored XSS to steal cookies from authenticated user sessions
- D. CSRF tricks authenticated users into making unintended requests using tokens

Answer: D

Q792. What is SQL injection and how is it prevented?

- A. SQL injection inserts malicious SQL through input fields prevented by parameterization
- B. SQL injection overloads the database server with excessive query requests at once
- C. SQL injection exploits weak passwords to gain direct database admin access rights
- D. SQL injection inserts CSS into database queries to alter page styling output

Answer: A

Q793. What is the purpose of Content Security Policy headers?

- A. To restrict which sources of content the browser is allowed to load safely
- B. To define the content type of the HTTP response body sent to the client
- C. To set the content encoding for compressing response data before transfer
- D. To specify the content language for automatic translation of web page text

Answer: A

Q794. How does the SameSite cookie attribute improve security?

- A. It encrypts cookie values so they cannot be read by client-side JavaScript
- B. It forces cookies to expire after a single use preventing session replay attacks
- C. It restricts cookies to same-site requests reducing CSRF attack vulnerability
- D. It limits cookie storage to a maximum of one kilobyte for security purposes

Answer: C

Q795. What is the principle of least privilege in web security?

- A. Giving administrators the highest level of access to all system resources
- B. Providing equal access levels to all users regardless of their actual role
- C. Granting only the minimum permissions necessary to perform required tasks
- D. Restricting all users from accessing any system resources by default only

Answer: C

Q796. How does CORS work and why is it important?

- A. CORS controls which origins can access resources preventing unauthorized access
- B. CORS caches origin responses to reduce the number of server requests needed
- C. CORS encrypts all cross-origin requests to protect data during transmission
- D. CORS compresses resources to reduce bandwidth and improve page load speed

Answer: A

Q797. What is the purpose of input sanitization?

- A. To clean user input by removing or escaping potentially harmful characters
- B. To format input data for visually appealing display in the user interface
- C. To validate that input matches a specific data type before database storage
- D. To compress input data to reduce the payload size of form submission data

Answer: A

Q798. What is a brute force attack?

- A. An attack that tries all possible combinations to guess passwords or keys
- B. An attack that floods the server with traffic to make it unavailable to users
- C. An attack that intercepts data in transit between the client and the server
- D. An attack that exploits a known vulnerability in the server operating system

Answer: A

Q799. How does rate limiting help prevent security attacks?

- A. It limits the size of HTTP response bodies to prevent data exfiltration attacks
- B. It limits the CSS animation frame rate to prevent browser resource exhaustion
- C. It restricts the number of database columns that can be queried simultaneously
- D. It restricts request frequency from a single source to prevent abuse and attacks

Answer: D

Q800. What is the purpose of security headers like X-Frame-Options?

- A. To specify the maximum number of iframes allowed on a single web page view
- B. To set the visual frame and border styling for the web page layout design
- C. To define the frame rate for CSS animations and JavaScript timer intervals
- D. To prevent clickjacking by controlling whether a page can be framed by others

Answer: D

Q801. What are Core Web Vitals and why are they important?

- A. They are Google metrics measuring loading, interactivity, and visual stability
- B. They are CSS properties that control the visual appearance of web page text
- C. They are server health metrics measuring CPU and memory usage of web servers
- D. They are JavaScript functions that validate form data before submission occurs

Answer: A

Q802. How does code splitting improve web application performance?

- A. It splits the database into multiple tables for faster query execution speed
- B. It divides server code into microservices for independent deployment of each
- C. It splits CSS into multiple files that load sequentially one after another
- D. It breaks code into smaller chunks loaded on demand reducing initial bundle size

Answer: D

Q803. What is the critical rendering path?

- A. The sequence of steps the browser takes to convert HTML, CSS, JS to pixels
- B. The path data takes through the network from server to client browser device
- C. The path a user follows through a website to complete a conversion goal task
- D. The route HTTP requests take through load balancers to reach the origin server

Answer: A

Q804. How does tree shaking reduce bundle size?

- A. It removes duplicate database records to reduce storage and query overhead
- B. It prunes old log entries from the server to free up storage disk space use
- C. It eliminates dead code and unused exports from the final JavaScript bundle
- D. It removes unused CSS classes from stylesheets during the build compilation

Answer: C

Q805. What is the purpose of preloading and prefetching resources?

- A. To load critical resources early or fetch future resources in idle time ahead
- B. To download and install browser plugins required by the web page for function
- C. To pre-validate form data on the server before the user submits the form
- D. To preview web pages before navigating to them in the browser tab interface

Answer: A

Q806. How does HTTP caching work with Cache-Control headers?

- A. Cache-Control headers define which HTTP methods are allowed for each resource
- B. Cache-Control headers control the maximum number of concurrent HTTP connections
- C. Cache-Control headers specify how and how long browsers and CDNs cache resources
- D. Cache-Control headers encrypt cached data to prevent unauthorized access today

Answer: C

Q807. What is render-blocking and how can it be avoided?

- A. Database locks that block writes avoided by using optimistic concurrency control patterns
- B. Elements that block user scrolling avoided by removing fixed position CSS properties
- C. Resources that block initial render avoided by async scripts and critical CSS inlining
- D. Server processes that block response avoided by using asynchronous database queries

Answer: C

Q808. How does image format selection affect web performance?

- A. SVG is the best format for photographs while JPEG is best for vector graphics use
- B. PNG is always the smallest format and should be used for all web image content
- C. Modern formats like WebP and AVIF offer better compression than JPEG and PNG today
- D. All image formats perform identically with no difference in file size or quality

Answer: C

Q809. What is the purpose of a service worker for performance optimization?

- A. To compile JavaScript code into machine code for faster execution on the client
- B. To cache resources and serve them offline improving load times on repeat visits
- C. To run CSS animations in a separate thread for smoother visual performance output
- D. To distribute server workload across multiple threads for parallel request handling

Answer: B

Q810. How does DNS prefetching improve page load performance?

- A. It caches the entire web page at the DNS level for instant loading on next visit
- B. It compresses DNS queries to reduce the bandwidth used during name resolution
- C. It prefetches all images on a page before the user scrolls down to view them
- D. It resolves domain names to IP addresses early reducing DNS lookup latency overall

Answer: D

Q811. What is the difference between unit testing and end-to-end testing?

- A. Unit tests are for backend code only while end-to-end tests are for frontend code only
- B. Unit tests verify individual components while E2E tests verify complete user workflows
- C. Unit tests run in production while end-to-end tests run only in development environments
- D. Unit tests are automated while end-to-end tests can only be performed manually today

Answer: B

Q812. What is test-driven development and how does it work?

- A. Tests are driven by user feedback collected through surveys after each deployment
- B. Tests are written after the feature is complete to verify correctness of the code
- C. Tests are written before the code and the code is written to make them pass first
- D. Tests are automatically generated from documentation without developer involvement

Answer: C

Q813. What is mocking in testing and why is it useful?

- A. Creating fake visual mockups of web page designs for client presentation review
- B. Mocking the server response time to test how the frontend handles slow networks
- C. Creating duplicate databases for testing queries without affecting production data
- D. Replacing real dependencies with simulated objects for isolated unit testing work

Answer: D

Q814. How does code coverage measurement help in testing?

- A. It calculates the physical storage space required to store all test result files
- B. It counts the number of developers who have reviewed the code before deployment
- C. It measures the percentage of code executed during tests identifying untested areas
- D. It measures how much of the codebase is covered by documentation comments today

Answer: C

Q815. What is snapshot testing and when is it useful?

- A. Recording network request snapshots for offline testing of API integrations work
- B. Creating snapshots of database state before and after running migration scripts
- C. Taking screenshots of web pages and comparing them manually for visual changes
- D. Capturing component output and comparing against stored reference for regressions

Answer: D

Q816. What is the purpose of continuous integration in testing?

- A. To continuously monitor production servers and restart them when tests fail today
- B. To continuously integrate new CSS styles into the existing stylesheet files daily
- C. To automatically run tests on every code change ensuring quality before merging
- D. To integrate third-party testing tools into the development environment manually

Answer: C

Q817. How does Cypress differ from Selenium for web testing?

- A. Cypress runs in the browser with native access while Selenium uses remote drivers
- B. Cypress only supports unit testing while Selenium supports all types of testing
- C. Cypress requires Python while Selenium only works with Java programming language
- D. Cypress is a server-side testing tool while Selenium tests only client-side code

Answer: A

Q818. What is regression testing?

- A. Testing new features that are being developed for the first time in the sprint
- B. Reducing the number of test cases by removing redundant or obsolete ones today
- C. Verifying that existing functionality still works after code changes are made
- D. Testing the application by reverting to a previous version of the codebase

Answer: C

Q819. What is the testing pyramid and what does it recommend?

- A. Many UI tests at top, fewer integration in middle, fewest unit tests at bottom
- B. Only end-to-end tests are needed because they cover all lower levels automatically
- C. Equal numbers of unit, integration, and end-to-end tests across all three levels
- D. Many unit tests at the base, fewer integration in the middle, fewest E2E on top

Answer: D

Q820. What is the purpose of accessibility testing?

- A. To verify that the website is usable by people with various disabilities needs
- B. To verify that the website source code is accessible through view-source option
- C. To test whether the website can be accessed from different geographic regions
- D. To test how fast the website loads on different network connection speed levels

Answer: A

Q821. What is the difference between continuous integration and continuous deployment?

- A. CI is for frontend code only while CD handles backend deployment pipeline exclusively
- B. CI deploys code to production while CD runs tests and merges code into the main branch
- C. CI requires manual approval while CD runs automatically without any human intervention
- D. CI merges code and runs tests while CD automatically deploys passing code to production

Answer: D

Q822. What is a Docker container and how does it help deployment?

- A. A CSS container element that wraps other elements for layout and styling purposes
- B. A physical container that stores server hardware in a data center facility rack
- C. A lightweight isolated environment that packages an app with all its dependencies
- D. A database container that stores related tables together for organized data access

Answer: C

Q823. What is the purpose of a reverse proxy in deployment?

- A. To reverse engineer the deployment process for rollback to a previous version
- B. To sit before application servers handling load balancing, SSL, and caching tasks
- C. To create a proxy copy of the production database for development testing use
- D. To reverse the order of CSS rules applied to elements for specificity overrides

Answer: B

Q824. What is a deployment pipeline?

- A. A pipeline that transfers database records from staging to production environments
- B. A physical network cable that connects development machines to production servers
- C. An automated process that builds, tests, and deploys code through defined stages
- D. A logging pipeline that streams application logs to a monitoring dashboard view

Answer: C

Q825. How does blue-green deployment reduce downtime?

- A. It colors the deployment logs blue for success and green for errors encountered
- B. It deploys to both environments simultaneously and picks the faster one to serve
- C. It maintains two identical environments switching traffic to reduce deployment risk
- D. It uses blue and green colored servers to visually distinguish environments easily

Answer: C

Q826. What is infrastructure as code and why is it important?

- A. Creating infrastructure diagrams using code for documentation purposes exclusively
- B. Building physical server infrastructure using 3D printing from code specifications
- C. Writing CSS that defines the visual infrastructure layout of web page elements used
- D. Managing infrastructure through code files enabling version control and repeatability

Answer: D

Q827. What is the purpose of environment variables in deployment?

- A. To define the physical environment like temperature where servers are housed today
- B. To set environment-specific font sizes for development versus production display
- C. To store configuration that varies between environments without code changes needed
- D. To set CSS custom properties that change based on the user browser environment

Answer: C

Q828. What is a container registry?

- A. A registry of physical shipping containers used for server hardware delivery tasks
- B. A repository for storing and distributing container images like Docker Hub service
- C. A system that registers new CSS container queries for responsive component design
- D. A database that registers all running containers on a single host machine system

Answer: B

Q829. How does canary deployment work?

- A. It uses canary-colored warning messages to alert users about upcoming changes ahead
- B. It gradually rolls out changes to a small subset of users before full deployment
- C. It deploys to a canary server that only team members can access for final testing
- D. It deploys the entire application to all users simultaneously with no phased rollout

Answer: B

Q830. What is the purpose of health checks in deployed applications?

- A. To monitor the physical health of server hardware like CPU temperature and disk state
- B. To verify that application components are running correctly and able to serve traffic
- C. To check the health of developers working on the application project regularly
- D. To validate that HTML passes W3C health standards for accessibility compliance test

Answer: B

Q831. What is the difference between PUT and PATCH HTTP methods?

- A. PUT is idempotent while PATCH is not and cannot be safely retried on failure
- B. PUT replaces the entire resource while PATCH partially updates specific fields
- C. PUT creates a new resource while PATCH deletes an existing resource from server
- D. PUT is for files only while PATCH works with JSON data responses exclusively

Answer: B

Q832. How does API versioning help maintain backward compatibility?

- A. It tracks how many times each API endpoint has been called for billing usage
- B. It creates version backups of the API database for disaster recovery purposes
- C. It allows multiple API versions to coexist so old clients continue to function
- D. It versions the CSS stylesheets used in API documentation pages for consistency

Answer: C

Q833. What is the purpose of API rate limiting?

- A. To restrict API responses to a fixed number of records per database query
- B. To control request frequency preventing abuse and ensuring fair resource usage
- C. To limit the number of API endpoints that can be created in a single project
- D. To limit the CSS animation frame rate for API loading spinner components

Answer: B

Q834. How does pagination work in REST APIs?

- A. It splits API request payloads across multiple HTTP connections for parallelism
- B. It creates separate HTML pages for each API response sent to the client browser
- C. It automatically paginates database tables into fixed-size partitions for storage
- D. It returns data in smaller chunks using offset, limit, or cursor-based approaches

Answer: D

Q835. What is the purpose of API documentation tools like Swagger?

- A. To automatically generate CSS styles for API response rendering in browsers
- B. To provide interactive documentation for exploring and testing API endpoints
- C. To generate client-side JavaScript code from server-side API handler functions
- D. To create database documentation from SQL schema definitions automatically

Answer: B

Q836. How do WebSockets differ from REST APIs for real-time communication?

- A. WebSockets require API keys while REST APIs use session cookies for authentication
- B. WebSockets only support text data while REST APIs support all content type formats
- C. WebSockets maintain persistent bidirectional connections unlike REST request-response
- D. WebSockets use HTTP request-response while REST uses persistent connections only

Answer: C

Q837. What is the purpose of HATEOAS in REST API design?

- A. To hash API tokens using advanced encryption standards for security purposes
- B. To include hypermedia links in responses enabling clients to discover actions
- C. To hate on APIs that don't follow REST principles by logging warning messages
- D. To handle asynchronous task execution and observe API server health metrics

Answer: B

Q838. What is an API gateway and what functions does it serve?

- A. A single entry point that handles routing, authentication, and rate limiting tasks
- B. A tool that generates API code from visual gateway diagrams designed by teams
- C. A physical gateway device that connects API servers to the internet backbone
- D. A database gateway that translates API calls into SQL queries automatically used

Answer: A

Q839. How does OAuth 2.0 work for API authorization?

- A. It requires users to share their passwords directly with third-party applications
- B. It delegates authorization using access tokens without sharing user credentials
- C. It only works for social media APIs and cannot be used for other API services
- D. It encrypts all API responses using a shared secret key between both parties

Answer: B

Q840. What is the purpose of idempotency in API design?

- A. To ensure that making the same request multiple times produces the same result
- B. To identify duplicate API endpoints and remove them from the server routing table
- C. To verify that API request parameters match the expected data types precisely
- D. To ensure API responses always contain identical formatting and structure layout

Answer: A

Q841. How does server-side rendering differ from static site generation?

- A. SSR generates static files while SSG renders pages dynamically per request
- B. SSR only works with React while SSG works with all frontend frameworks used
- C. SSR and SSG are identical concepts with no practical differences between them
- D. SSR renders pages on each request while SSG pre-builds pages at build time

Answer: D

Q842. What is the purpose of a build tool like Vite or Webpack?

- A. To bundle, transform, and optimize source code for production deployment use
- B. To build physical server infrastructure from configuration code specifications
- C. To build database schemas from visual entity relationship diagram models only
- D. To construct CSS layouts from wireframe sketches using AI image recognition

Answer: A

Q843. What is the purpose of feature flags in modern web development?

- A. To add visual flag icons to features that are new or recently updated for users
- B. To mark CSS features that are not supported in older browsers for fallback use
- C. To enable or disable features dynamically without deploying new code versions
- D. To flag code that contains bugs for prioritized fixing by the development team

Answer: C

Q844. How does edge computing benefit modern web applications?

- A. It sharpens the visual edges of images displayed on the web page for clarity
- B. It caches database queries at the edge of the database cluster for faster reads
- C. It processes requests at network edge locations closer to users reducing latency
- D. It adds decorative edge borders to web page components for visual enhancement

Answer: C

Q845. What is the purpose of a design token in design systems?

- A. A physical token given to designers to access the design tool application suite
- B. A database token that authorizes access to the design system API documentation
- C. A cryptocurrency token used to purchase premium design assets and templates
- D. A named value storing design decisions like colors and spacing for consistency

Answer: D

Q846. How do monorepos benefit large-scale web development teams?

- A. They automatically generate documentation for all projects in the repository
- B. They force all projects to use the same programming language and framework only
- C. They enable code sharing, atomic changes, and unified tooling across projects
- D. They eliminate the need for any version control branching strategies entirely

Answer: C

Q847. What is the difference between a library and a framework?

- A. Libraries are free and open source while frameworks always require paid licenses
- B. A library is larger than a framework and includes more features out of the box
- C. A library is called by your code while a framework calls your code via inversion
- D. Libraries only work in browsers while frameworks work on both server and client

Answer: C

Q848. What is the purpose of a CI/CD pipeline in modern web development?

- A. To create interactive and creative designs for web page user interface layouts
- B. To continuously inspect and clean database records for data quality maintenance
- C. To continuously integrate CSS changes with HTML changes for consistent styling
- D. To automate building, testing, and deploying code changes reliably and quickly

Answer: D

Q849. How does incremental static regeneration work?

- A. It regenerates static pages in the background after a revalidation period expires
- B. It incrementally adds new CSS styles to static pages without rebuilding them all
- C. It regenerates all static pages from scratch on every single user request received
- D. It generates static pages incrementally one at a time during the build process

Answer: A

Q850. What is the purpose of a component library like Shadcn or Material UI?

- A. To store compiled JavaScript components in a shared library for runtime loading
- B. To create a library catalog of all API endpoints for developer documentation use
- C. To provide pre-built accessible UI components for consistent application styling
- D. To compile custom HTML elements into web components for cross-framework sharing

Answer: C

Q851. What is the role of a web framework in development?

- A. It replaces the need for HTML entirely
- B. It provides pre-built tools and structure to speed up development
- C. It is only used for mobile app development
- D. It manages internet service provider connections

Answer: B

Q852. What is the purpose of developer tools in modern browsers?

- A. To create new websites automatically
- B. To inspect, debug, and profile web pages in real time
- C. To replace the need for a code editor
- D. To manage web hosting accounts

Answer: B

Q853. What is a web application compared to a traditional website?

- A. A web application is always built with Python
- B. A web application provides interactive functionality beyond static content
- C. A web application cannot use HTML
- D. A web application runs without a browser

Answer: B

Q854. What is the MEAN stack?

- A. MongoDB, Express, Angular, Node.js
- B. MySQL, Ember, Apache, Nginx
- C. Markdown, ESLint, AWS, Netlify
- D. MongoDB, Electron, Ansible, NPM

Answer: A

Q855. What is cross-browser compatibility?

- A. A tool for compressing images
- B. Ensuring a website works correctly across different web browsers
- C. A method for encrypting web traffic
- D. A framework for building mobile apps

Answer: B

Q856. What is the purpose of a favicon on a website?

- A. To improve server performance
- B. To display a small icon in the browser tab representing the site
- C. To encrypt user data
- D. To validate form inputs

Answer: B

Q857. What does the term 'deployment' mean in web development?

- A. Writing the initial HTML code
- B. Making a website available on a live server for public access
- C. Installing a text editor
- D. Creating a database backup

Answer: B

Q858. What is a progressive web application?

- A. A desktop-only application
- B. A web app that uses modern APIs to deliver native-like experiences
- C. A website that only works in Chrome
- D. A server-side rendering technique

Answer: B

Q859. What is the purpose of semantic versioning in web projects?

- A. To style web pages
- B. To communicate the nature of changes in software releases using major.minor.patch format
- C. To compress JavaScript files
- D. To manage DNS records

Answer: B

Q860. What is the difference between a library and a framework in web development?

- A. They are identical concepts
- B. A library provides specific functions you call, while a framework controls the overall flow of your application
- C. A framework is always smaller than a library
- D. Libraries only work on the server side

Answer: B

Q861. What is the purpose of a load balancer in web architecture?

- A. To compress images on the server
- B. To distribute incoming network traffic across multiple servers
- C. To write JavaScript code
- D. To manage DNS records

Answer: B

Q862. What is the difference between a forward proxy and a reverse proxy?

- A. They are the same thing
- B. A forward proxy acts on behalf of clients while a reverse proxy acts on behalf of servers
- C. Forward proxies are faster
- D. Reverse proxies are only used for caching

Answer: B

Q863. What is the role of caching in web architecture?

- A. To permanently delete old data
- B. To store frequently accessed data closer to the user to reduce latency and server load
- C. To encrypt data in transit
- D. To create new database tables

Answer: B

Q864. What is the role of DNS in web architecture?

- A. To style web pages
- B. To translate human-readable domain names into IP addresses that computers use to locate servers
- C. To compress HTTP responses
- D. To manage database connections

Answer: B

Q865. What is the purpose of the HTTP PATCH method?

- A. To delete a resource entirely
- B. To retrieve all resources
- C. To apply partial modifications to a resource
- D. To create a new resource from scratch

Answer: C

Q866. What is a content delivery network (CDN)?

- A. A tool for writing HTML code
- B. A geographically distributed network of servers that delivers content from the nearest location to the user
- C. A type of programming language
- D. A database management system

Answer: B

Q867. What is the difference between synchronous and asynchronous communication in web architecture?

- A. They are identical approaches
- B. Synchronous blocks until a response is received while asynchronous allows other operations to continue
- C. Synchronous is always faster
- D. Asynchronous cannot be used with HTTP

Answer: B

Q868. What does a status code in the 2xx range indicate?

- A. Client error
- B. Server error
- C. Successful request
- D. Redirect

Answer: C

Q869. What is the role of an API gateway in a distributed system?

- A. To store static files
- B. To act as a single entry point that routes requests to appropriate backend services
- C. To write CSS styles
- D. To manage version control

Answer: B

Q870. What is the purpose of the HTTP Cache-Control header?

- A. To define the page title
- B. To specify caching directives for both requests and responses
- C. To set the content language
- D. To redirect users to a different page

Answer: B

Q871. What is the difference between the <section> and <article> elements?

- A. They are identical elements
- B. A section groups related content while an article represents self-contained, independently distributable content
- C. Section is for navigation only
- D. Article cannot contain headings

Answer: B

Q872. What is the purpose of the autocomplete attribute on form inputs?

- A. To validate the input value
- B. To allow browsers to suggest previously entered values for the field
- C. To make the input read-only
- D. To add placeholder text

Answer: B

Q873. What does the <datalist> element do in HTML5?

- A. Creates a data table
- B. Provides a predefined list of suggestions for an input element
- C. Defines a database connection
- D. Generates a numbered list

Answer: B

Q874. What is the difference between the <script> tag placed in <head> versus at the end of <body>?

- A. There is no difference
- B. Scripts in <head> may block HTML parsing while scripts before </body> allow the page to render first
- C. Scripts in <head> execute slower
- D. Scripts in <body> cannot access the DOM

Answer: B

Q875. What is the purpose of the <main> element in HTML5?

- A. To create a sidebar
- B. To represent the dominant content of the body, unique to the document
- C. To define a footer
- D. To embed external content

Answer: B

Q876. What does the pattern attribute do on an input element?

- A. It sets the input's visual style
- B. It specifies a regular expression that the input value must match for validation
- C. It creates a placeholder text
- D. It disables the input field

Answer: B

Q877. What is the difference between the <link> and <a> tags?

- A. They are interchangeable
- B. <link> defines relationships between the document and external resources while <a> creates clickable hyperlinks
- C. <link> is used for images
- D. <a> is only for email links

Answer: B

Q878. What is the role of the <noscript> element in HTML?

- A. To remove JavaScript from the page
- B. To provide alternative content for users whose browsers do not support or have disabled JavaScript
- C. To add comments to code
- D. To optimize images

Answer: B

Q879. What does the loading attribute on images and iframes control?

- A. The image resolution
- B. Whether the resource is loaded immediately or deferred until it nears the viewport
- C. The image format
- D. The image border style

Answer: B

Q880. What is the purpose of the <output> element in HTML5 forms?

- A. To create a text input
- B. To display the result of a calculation or user action
- C. To submit the form
- D. To define form labels

Answer: B

Q881. What is the difference between relative and absolute units in CSS?

- A. There is no difference
- B. Relative units like em and % scale based on context while absolute units like px have a fixed size
- C. Absolute units are always larger
- D. Relative units only work in flexbox

Answer: B

Q882. What does the position: fixed property do?

- A. Positions the element relative to its parent
- B. Positions the element relative to the browser viewport, staying in place during scrolling
- C. Removes the element from the document flow
- D. Makes the element invisible

Answer: B

Q883. What is the purpose of the CSS opacity property?

- A. To set the element's border width
- B. To control the transparency level of an element
- C. To change the font weight
- D. To set the z-index

Answer: B

Q884. How does the CSS cascade determine which styles are applied?

- A. Styles are applied randomly
- B. Styles are resolved based on specificity, source order, and importance
- C. Only the last stylesheet wins
- D. Only inline styles are applied

Answer: B

Q885. What does the text-decoration property control?

- A. The font size of text
- B. Decorative lines on text such as underline, overline, or line-through
- C. The text alignment
- D. The letter spacing

Answer: B

Q886. What is the purpose of the CSS transform property?

- A. To change an element's color
- B. To apply 2D or 3D transformations like rotate, scale, skew, or translate to an element
- C. To set the element's position
- D. To add a border

Answer: B

Q887. What is the difference between position: sticky and position: fixed?

- A. They behave identically
- B. Sticky toggles between relative and fixed based on scroll position while fixed always stays relative to the viewport
- C. Sticky only works on images
- D. Fixed does not work in modern browsers

Answer: B

Q888. What does the CSS line-height property control?

- A. The width of a line
- B. The vertical space between lines of text
- C. The number of lines visible
- D. The horizontal spacing of characters

Answer: B

Q889. How do CSS pseudo-elements like ::before and ::after work?

- A. They create new HTML elements in the DOM
- B. They insert generated content before or after an element's actual content using CSS alone
- C. They modify JavaScript events
- D. They add attributes to HTML tags

Answer: B

Q890. What is the purpose of the CSS outline property?

- A. To set the background color
- B. To draw a line outside the element's border without affecting layout
- C. To add inner padding
- D. To change the font style

Answer: B

Q891. What is the difference between align-items and align-content in Flexbox?

- A. They are identical properties
- B. align-items aligns items within each line while align-content aligns the lines themselves within the container
- C. align-items only works horizontally
- D. align-content is deprecated

Answer: B

Q892. How does the CSS grid auto-placement algorithm work?

- A. Items are always placed manually
- B. The browser automatically places grid items in the next available cell following row or column order
- C. Items are placed randomly
- D. Auto-placement only works with flexbox

Answer: B

Q893. What does the CSS will-change property do?

- A. Forces an immediate style change
- B. Hints to the browser about which properties will change, allowing it to optimize rendering in advance
- C. Changes the element's content
- D. Prevents any style changes

Answer: B

Q894. How does the CSS repeat() function work in Grid layouts?

- A. It repeats the entire grid layout
- B. It repeats a track pattern a specified number of times, simplifying grid-template definitions
- C. It creates animation loops
- D. It duplicates HTML elements

Answer: B

Q895. What is the difference between a CSS transition and a CSS animation?

- A. They are the same thing
- B. Transitions require a trigger and animate between two states while animations can run independently with multiple keyframes
- C. Transitions are more complex
- D. Animations cannot loop

Answer: B

Q896. What does the CSS filter property allow?

- A. Filtering HTML elements from the DOM
- B. Applying visual effects like blur, brightness, contrast, and grayscale to elements
- C. Filtering network requests
- D. Sorting CSS selectors

Answer: B

Q897. How does the CSS place-items shorthand work?

- A. It positions items using absolute coordinates
- B. It sets both align-items and justify-items in a single declaration
- C. It creates new grid items
- D. It removes items from the layout

Answer: B

Q898. What is the purpose of the CSS shape-outside property?

- A. To clip an element's visible area
- B. To define a shape around which inline content wraps
- C. To create border radius
- D. To define grid areas

Answer: B

Q899. What does the CSS writing-mode property control?

- A. The font family
- B. Whether text flows horizontally or vertically and the direction of block flow
- C. The text color
- D. The font size

Answer: B

Q900. How does the CSS backdrop-filter property differ from filter?

- A. They are identical
- B. backdrop-filter applies effects to the area behind the element while filter applies effects to the element itself
- C. backdrop-filter only works on images
- D. filter is deprecated

Answer: B

Q901. What is the difference between a function declaration and a function expression?

- A. They are identical
- B. Function declarations are hoisted and can be called before their definition while function expressions are not hoisted
- C. Function expressions are faster
- D. Function declarations cannot accept parameters

Answer: B

Q902. What does the filter() method do on an array?

- A. Sorts the array elements
- B. Creates a new array with elements that pass a test provided by a callback function
- C. Removes all elements
- D. Reverses the array order

Answer: B

Q903. What is scope in JavaScript?

- A. The visual appearance of code
- B. The region of code where a variable is accessible and can be referenced
- C. The speed at which code executes
- D. The number of lines in a file

Answer: B

Q904. What is the purpose of the reduce() method on arrays?

- A. To reduce the array's length by one
- B. To execute a reducer function on each element, accumulating a single result value
- C. To delete duplicate elements
- D. To sort elements in descending order

Answer: B

Q905. What is type coercion in JavaScript?

- A. Manually changing a variable's type
- B. The automatic conversion of values from one type to another during operations
- C. A method for creating new data types
- D. A way to validate form inputs

Answer: B

Q906. What is the difference between the forEach() and map() array methods?

- A. They are identical
- B. forEach() executes a function on each element without returning a new array while map() returns a new array of transformed elements
- C. forEach() is faster
- D. map() modifies the original array

Answer: B

Q907. What is the switch statement used for in JavaScript?

- A. Declaring variables
- B. Evaluating an expression against multiple case values and executing the matching block
- C. Creating loops
- D. Defining functions

Answer: B

Q908. What is the purpose of the Object.keys() method?

- A. To lock an object's properties
- B. To return an array of an object's own enumerable property names
- C. To delete all properties from an object
- D. To merge two objects together

Answer: B

Q909. What does the concat() method do with arrays?

- A. Removes elements from an array
- B. Merges two or more arrays into a new array without modifying the originals
- C. Sorts the array
- D. Finds a specific element

Answer: B

Q910. What is short-circuit evaluation in JavaScript?

- A. A way to speed up code execution
- B. When logical operators stop evaluating as soon as the result is determined
- C. A method for handling errors
- D. A type of loop optimization

Answer: B

Q911. What is the purpose of the async iterator protocol in JavaScript?

- A. To iterate over synchronous arrays faster
- B. To enable asynchronous iteration with for-await-of loops over data that arrives over time
- C. To create CSS animations
- D. To manage database cursors

Answer: B

Q912. What is the difference between a shallow copy and a deep copy in JavaScript?

- A. They are the same operation
- B. A shallow copy copies only top-level properties while a deep copy recursively copies all nested objects
- C. Shallow copies are always faster
- D. Deep copies only work with arrays

Answer: B

Q913. How does the nullish coalescing assignment operator (??=) work?

- A. It always assigns the right-hand value
- B. It assigns the right-hand value only if the left-hand variable is null or undefined
- C. It checks for falsy values
- D. It throws an error if the variable is null

Answer: B

Q914. What is the purpose of the queueMicrotask() function?

- A. To create a new thread
- B. To schedule a callback to run as a microtask, executing before the next macrotask
- C. To queue network requests
- D. To manage CSS transitions

Answer: B

Q915. What is the purpose of the Promise.race() method?

- A. To run promises sequentially
- B. To return a promise that resolves or rejects as soon as the first promise in the iterable settles
- C. To cancel all pending promises
- D. To merge promise results

Answer: B

Q916. How do private class fields declared with # differ from underscore convention?

- A. They are the same approach
- B. Private fields with # are truly inaccessible from outside the class while underscore is just a naming convention
- C. Underscore fields are more private
- D. Hash fields are slower

Answer: B

Q917. What is the purpose of the globalThis object?

- A. To create global CSS styles
- B. To provide a standard way to access the global object across different JavaScript environments
- C. To manage global state in React
- D. To define global HTML templates

Answer: B

Q918. What does the Promise.any() method do?

- A. Waits for all promises to reject
- B. Returns a promise that resolves with the first fulfilled promise, ignoring rejections unless all reject
- C. Cancels all promises
- D. Returns all results in order

Answer: B

Q919. What is the purpose of the structuredClone() function?

- A. To clone HTML elements
- B. To create a deep copy of a value using the structured clone algorithm, supporting complex types like Map, Set, and Date
- C. To clone CSS styles
- D. To duplicate event listeners

Answer: B

Q920. What is optional chaining (?.) and how does it prevent errors?

- A. It creates optional function parameters
- B. It safely accesses nested properties by returning undefined instead of throwing an error if a reference is nullish
- C. It chains CSS selectors
- D. It makes all variables optional

Answer: B

Q921. What is the difference between client-side and server-side state management?

- A. They are identical
- B. Client-side state is managed in the browser while server-side state persists data across sessions and users on the server
- C. Client-side state is always faster
- D. Server-side state cannot use databases

Answer: B

Q922. How does the virtual DOM optimization work in React?

- A. It replaces the real DOM entirely
- B. React compares a virtual representation with the previous version and applies only the minimal necessary changes to the real DOM
- C. It creates a copy of the CSS styles
- D. It preloads all images

Answer: B

Q923. What is the purpose of the useRef hook in React?

- A. To manage CSS references
- B. To hold a mutable reference that persists across renders without causing re-renders when changed
- C. To create database references
- D. To reference external APIs

Answer: B

Q924. What is the Composition API in Vue 3 and how does it improve code organization?

- A. It creates CSS compositions
- B. It allows grouping related logic by concern using composable functions instead of scattering it across options
- C. It composes database queries
- D. It manages file compositions

Answer: B

Q925. What is the purpose of code splitting in frontend frameworks?

- A. To split CSS files
- B. To break application code into smaller bundles loaded on demand, improving initial load performance
- C. To split database tables
- D. To divide server instances

Answer: B

Q926. What are React context and when should it be used?

- A. A CSS context property
- B. A way to pass data through the component tree without prop drilling, best for infrequently changing global data
- C. A database context manager
- D. A server execution context

Answer: B

Q927. How does Angular's template-driven forms approach differ from reactive forms?

- A. They are identical approaches
- B. Template-driven forms use directives in templates for simple forms while reactive forms use programmatic control for complex validation
- C. Template-driven forms are always better
- D. Reactive forms cannot validate input

Answer: B

Q928. What is the purpose of error boundaries in React?

- A. To create CSS borders
- B. To catch JavaScript errors in child components and display fallback UI instead of crashing the whole application
- C. To handle network errors
- D. To manage database errors

Answer: B

Q929. What is hydration in the context of server-side rendered applications?

- A. Adding water effects to the UI
- B. The process of attaching event handlers and state to server-rendered HTML on the client side
- C. Compressing server responses
- D. Loading CSS stylesheets

Answer: B

Q930. What is the difference between static site generation (SSG) and server-side rendering (SSR)?

- A. They are identical
- B. SSG generates HTML at build time while SSR generates HTML on each request at runtime
- C. SSG requires more server resources
- D. SSR cannot use caching

Answer: B

Q931. What is the purpose of input validation on the backend?

- A. To style form elements
- B. To ensure incoming data meets expected formats and constraints, preventing malicious or invalid data from being processed
- C. To create HTML forms
- D. To manage CSS classes

Answer: B

Q932. How do JWT tokens work for stateless authentication?

- A. They store session data on the server
- B. They encode user claims into a signed token that the server validates without storing session state
- C. They replace passwords entirely
- D. They only work with SQL databases

Answer: B

Q933. How does middleware chaining work in Express.js?

- A. Middleware functions run in parallel
- B. Each middleware calls next() to pass control to the next function in the chain, allowing sequential processing of requests
- C. Middleware replaces routing
- D. Middleware only works with databases

Answer: B

Q934. What is rate limiting in backend applications?

- A. Limiting CSS animations
- B. Restricting the number of requests a client can make within a time period to prevent abuse
- C. Limiting database size
- D. Restricting the number of HTML elements

Answer: B

Q935. What is the purpose of pagination in API responses?

- A. To style page numbers
- B. To divide large result sets into smaller pages, reducing response size and improving performance
- C. To create PDF documents
- D. To manage page routing

Answer: B

Q936. What is the difference between a relational and a document database?

- A. They store data identically
- B. Relational databases use structured tables with schemas while document databases store flexible JSON-like documents
- C. Relational databases are always faster
- D. Document databases cannot be queried

Answer: B

Q937. What is the purpose of hashing passwords before storing them?

- A. To make passwords shorter
- B. To convert passwords into irreversible fixed-length strings so the original password cannot be recovered if the database is compromised
- C. To encrypt database connections
- D. To compress storage space

Answer: B

Q938. What is the purpose of CORS middleware in a backend application?

- A. To compress responses
- B. To configure which origins, methods, and headers are allowed for cross-origin requests to the API
- C. To create database connections
- D. To manage server logging

Answer: B

Q939. What is the purpose of database seeding?

- A. To delete all data
- B. To populate a database with initial or test data for development or testing purposes
- C. To encrypt database records
- D. To compress database tables

Answer: B

Q940. What is the difference between structured and unstructured logging in backend systems?

- A. They are the same approach
- B. Structured logging outputs logs as parseable data formats like JSON while unstructured logging outputs freeform text
- C. Structured logging is slower
- D. Unstructured logging is more searchable

Answer: B

Q941. What is the purpose of request body parsing in server-side applications?

- A. To parse CSS files
- B. To extract and convert data from the HTTP request body into a usable format like JSON or form data
- C. To parse HTML templates
- D. To parse database queries

Answer: B

Q942. How do server-side sessions maintain user state across requests?

- A. They store all data in the URL
- B. They store a session ID in a cookie and keep session data on the server, associating requests with stored state
- C. They modify HTML files
- D. They use CSS custom properties

Answer: B

Q943. What is the difference between a web server and an application server?

- A. They are identical
- B. A web server handles static content and HTTP while an application server executes business logic and dynamic content
- C. Web servers are faster
- D. Application servers cannot serve HTML

Answer: B

Q944. What is the purpose of request rate limiting middleware?

- A. To limit CSS rendering speed
- B. To restrict the number of requests per client within a time window to prevent abuse and protect server resources
- C. To limit database query speed
- D. To slow down page rendering

Answer: B

Q945. How does server-side caching reduce database load?

- A. It deletes old database records
- B. It stores frequently accessed data in memory so repeated requests are served from cache instead of querying the database
- C. It compresses database files
- D. It indexes all database columns

Answer: B

Q946. How do background job processors like BullMQ work in Node.js applications?

- A. They replace all server routing
- B. They use Redis-backed queues to offload time-consuming tasks to separate worker processes for asynchronous execution
- C. They only process CSS files
- D. They replace database queries

Answer: B

Q947. What is the difference between cookie-based and token-based authentication?

- A. They are identical
- B. Cookie-based auth stores session IDs in cookies managed by the browser while token-based auth sends tokens in request headers
- C. Cookie-based is always more secure
- D. Token-based cannot expire

Answer: B

Q948. What is the purpose of webhook endpoints in server-side applications?

- A. To create CSS hooks
- B. To receive real-time event notifications from external services via HTTP POST requests
- C. To hook into database events
- D. To manage JavaScript event listeners

Answer: B

Q949. What is the purpose of file upload handling in server-side applications?

- A. To upload CSS files
- B. To process multipart form data containing files, validating and storing them securely
- C. To upload database backups
- D. To transfer HTML templates

Answer: B

Q950. What is the role of an ORM like Prisma or Sequelize in server-side Node.js applications?

- A. To render CSS on the server
- B. To provide a type-safe API for database operations, mapping database tables to JavaScript objects
- C. To manage HTTP routing
- D. To compress server responses

Answer: B

Q951. What is the purpose of database indexing and when should indexes be used?

- A. Indexes slow down all queries
- B. Indexes create data structures that speed up read queries on specific columns but add overhead to write operations
- C. Indexes replace primary keys
- D. Indexes only work with MongoDB

Answer: B

Q952. What is database normalization and what problem does it solve?

- A. It normalizes CSS values
- B. It organizes data to minimize redundancy and dependency by dividing large tables into related smaller tables
- C. It normalizes JSON formatting
- D. It standardizes SQL syntax

Answer: B

Q953. What is the difference between a LEFT JOIN and a RIGHT JOIN?

- A. They return the same results
- B. LEFT JOIN returns all rows from the left table with matching right rows, while RIGHT JOIN returns all rows from the right table
- C. LEFT JOIN is faster
- D. RIGHT JOIN is deprecated

Answer: B

Q954. What are ACID properties in database transactions?

- A. A type of database query
- B. Atomicity, Consistency, Isolation, and Durability - guarantees that ensure reliable transaction processing
- C. A chemical database property
- D. A MongoDB-specific feature

Answer: B

Q955. What is the purpose of a composite index in databases?

- A. To compose CSS styles
- B. To index multiple columns together, optimizing queries that filter or sort by those columns in combination
- C. To compose database schemas
- D. To combine multiple databases

Answer: B

Q956. What is the difference between HAVING and WHERE in SQL?

- A. They are interchangeable
- B. WHERE filters rows before grouping while HAVING filters groups after aggregation with GROUP BY
- C. HAVING is faster
- D. WHERE works with aggregate functions

Answer: B

Q957. What is a database trigger?

- A. A CSS event
- B. A stored procedure that automatically executes when a specified event occurs in the database such as INSERT, UPDATE, or DELETE
- C. A JavaScript event listener
- D. An HTML form trigger

Answer: B

Q958. What is the purpose of database connection string parameters like pool size?

- A. To set CSS pool values
- B. To configure how many simultaneous connections the application maintains to the database
- C. To define HTML element pools
- D. To set JavaScript thread pool size

Answer: B

Q959. What is a subquery in SQL?

- A. A CSS sub-selector
- B. A query nested inside another query that provides results used by the outer query
- C. A sub-database
- D. A JavaScript sub-function

Answer: B

Q960. How do database views simplify complex data access?

- A. They replace all tables
- B. They create reusable virtual tables from stored queries, providing abstraction over complex joins and restricting column visibility
- C. They only work with MongoDB
- D. They delete unused data

Answer: B

Q961. What is the purpose of the Secure flag on cookies?

- A. To secure CSS styles
- B. To ensure the cookie is only sent over encrypted HTTPS connections, preventing interception on unencrypted channels
- C. To make cookies larger
- D. To speed up cookie transmission

Answer: B

Q962. How does parameterized queries prevent SQL injection?

- A. They encrypt SQL statements
- B. They separate SQL logic from user data, ensuring user input is treated as data rather than executable SQL code
- C. They make queries faster
- D. They compress query results

Answer: B

Q963. What is the purpose of the X-Content-Type-Options header?

- A. To set the page content type
- B. To prevent browsers from MIME-type sniffing responses away from the declared content type
- C. To compress content
- D. To encrypt content

Answer: B

Q964. What is a cross-site scripting (XSS) stored attack?

- A. A CSS storage technique
- B. An attack where malicious scripts are permanently stored on the server and served to users who view the affected content
- C. A JavaScript storage method
- D. A database optimization

Answer: B

Q965. What is the purpose of multi-factor authentication (MFA)?

- A. To make login faster
- B. To require multiple forms of verification, making unauthorized access significantly harder even if one factor is compromised
- C. To style login pages
- D. To manage multiple user accounts

Answer: B

Q966. What is the purpose of encoding user output in web applications?

- A. To compress HTML
- B. To convert special characters to their safe HTML entities, preventing injected scripts from executing in the browser
- C. To encode CSS styles
- D. To compress JavaScript

Answer: B

Q967. What is session fixation and how is it prevented?

- A. Fixing CSS session styles
- B. An attack where an attacker sets a user's session ID before login; prevented by regenerating the session ID after authentication
- C. A database session repair
- D. A JavaScript debugging technique

Answer: B

Q968. What is the purpose of the Referrer-Policy header?

- A. To set CSS referrer styles
- B. To control how much referrer information is included when navigating from the site, protecting user privacy
- C. To manage database references
- D. To speed up page loads

Answer: B

Q969. What is clickjacking and how is it prevented?

- A. A CSS clicking animation
- B. An attack that tricks users into clicking hidden elements by overlaying transparent frames; prevented with X-Frame-Options or CSP frame-ancestors
- C. A JavaScript click handler
- D. A database click tracking method

Answer: B

Q970. What is the principle of defense in depth?

- A. A CSS depth property
- B. Implementing multiple layers of security controls so that if one layer fails, others still protect the system
- C. A JavaScript nesting technique
- D. A database depth indexing method

Answer: B

Q971. What is the purpose of the font-display CSS property for performance?

- A. To set the font size
- B. To control how custom fonts are displayed during loading, reducing invisible text periods
- C. To display font metadata
- D. To create font animations

Answer: B

Q972. How does HTTP/2 multiplexing improve web performance?

- A. It makes HTTP slower
- B. It allows multiple requests and responses to be sent simultaneously over a single connection, eliminating head-of-line blocking
- C. It replaces HTTPS
- D. It compresses all images

Answer: B

Q973. What is the purpose of critical CSS?

- A. CSS that is hard to write
- B. Inlining the CSS needed for above-the-fold content to eliminate render-blocking stylesheet downloads
- C. Critical CSS errors
- D. CSS for error pages

Answer: B

Q974. What is the Time to First Byte (TTFB) metric?

- A. Time to write the first line of code
- B. The time between the browser requesting a page and receiving the first byte of the response from the server
- C. The time to display the first pixel
- D. The time to first user click

Answer: B

Q975. How does image lazy loading improve performance?

- A. It makes images load in higher quality
- B. It defers loading of off-screen images until the user scrolls near them, reducing initial page load time
- C. It compresses images automatically
- D. It converts all images to SVG

Answer: B

Q976. What is the purpose of preconnect resource hint?

- A. To preconnect CSS styles
- B. To establish early connections (DNS, TCP, TLS) to important third-party domains before they are needed
- C. To connect database pools
- D. To preconnect JavaScript modules

Answer: B

Q977. What is the Cumulative Layout Shift (CLS) metric and why is it important?

- A. A CSS shift property
- B. A Core Web Vital measuring unexpected layout shifts that disrupt user experience during page load
- C. A JavaScript shift operation
- D. A server shift metric

Answer: B

Q978. How does bundling JavaScript files improve performance?

- A. It makes files larger
- B. It combines multiple files into fewer bundles, reducing the number of HTTP requests needed
- C. It encrypts JavaScript
- D. It compresses the server

Answer: B

Q979. What is the purpose of the Largest Contentful Paint (LCP) metric?

- A. To measure the largest CSS file
- B. To measure the time until the largest visible content element is fully rendered, indicating perceived load speed
- C. To measure the largest JavaScript file
- D. To measure database size

Answer: B

Q980. What is the purpose of using a service worker for caching?

- A. To style web pages
- B. To intercept network requests and serve cached responses, enabling offline access and faster repeat visits
- C. To manage CSS workers
- D. To create database workers

Answer: B

Q981. What is the purpose of a mock in testing?

- A. To mock CSS styles
- B. To simulate the behavior of real objects or functions, isolating the code under test from its dependencies
- C. To mock HTML elements
- D. To mock database tables

Answer: B

Q982. What is the difference between a spy and a mock in testing?

- A. They are identical concepts
- B. A spy observes and records calls to a real function while a mock replaces the function entirely with controlled behavior
- C. Spies are faster
- D. Mocks are deprecated

Answer: B

Q983. What is code coverage and what are its limitations?

- A. It measures CSS coverage
- B. It measures the percentage of code executed by tests, but high coverage does not guarantee all edge cases or logic paths are properly tested
- C. It measures database coverage
- D. It has no limitations

Answer: B

Q984. How does snapshot testing work and when is it most useful?

- A. It takes screenshots of the browser
- B. It serializes output to a file and compares future runs against it, most useful for detecting unintended changes in UI components
- C. It creates database snapshots
- D. It only works with React

Answer: B

Q985. What is the purpose of test fixtures?

- A. To fix broken CSS
- B. To set up known test data and state before tests run, ensuring consistent and repeatable test conditions
- C. To fix HTML errors
- D. To fix database bugs

Answer: B

Q986. What is the difference between Cypress and Playwright for end-to-end testing?

- A. They are identical tools
- B. Cypress runs tests in-browser with automatic waiting while Playwright supports multiple browsers and runs outside the browser with better parallel execution
- C. Cypress is always faster
- D. Playwright only supports Chrome

Answer: B

Q987. What is the purpose of test isolation?

- A. To isolate CSS styles
- B. To ensure each test runs independently without being affected by or affecting other tests
- C. To isolate database tables
- D. To isolate server instances

Answer: B

Q988. What is the role of continuous integration in testing?

- A. To integrate CSS continuously
- B. To automatically run tests on every code change, catching bugs early before they reach production
- C. To integrate databases
- D. To integrate HTML templates

Answer: B

Q989. What is the purpose of test doubles and what types exist?

- A. To double CSS styles
- B. To replace real dependencies in tests; types include stubs, mocks, spies, fakes, and dummies
- C. To double database records
- D. To duplicate HTML elements

Answer: B

Q990. What is the purpose of load testing?

- A. To load CSS files
- B. To simulate high traffic to evaluate how the application performs under expected and peak loads
- C. To load database records
- D. To load HTML templates

Answer: B

Q991. What is the purpose of Docker in modern deployment workflows?

- A. To style web pages
- B. To package applications with all dependencies into portable containers that run consistently across environments
- C. To manage CSS variables
- D. To create HTML templates

Answer: B

Q992. What is the difference between a Dockerfile and a docker-compose.yml file?

- A. They are identical files
- B. A Dockerfile defines how to build a single container image while docker-compose.yml orchestrates multiple containers as a service stack
- C. Dockerfile is for CSS
- D. docker-compose is deprecated

Answer: B

Q993. What is the purpose of environment-specific configuration in deployment?

- A. To change CSS per environment
- B. To adjust settings like database URLs, API keys, and feature flags for different environments without code changes
- C. To create environment variables in CSS
- D. To manage HTML per environment

Answer: B

Q994. What is a rollback strategy in deployment?

- A. Rolling back CSS changes
- B. A plan for reverting to a previous stable version of the application when a deployment introduces issues
- C. Rolling back database schemas
- D. Rolling back HTML templates

Answer: B

Q995. What is the purpose of a container registry?

- A. To register CSS containers
- B. To store and distribute Docker container images, providing version management and access control
- C. To register HTML elements
- D. To register database containers

Answer: B

Q996. How does blue-green deployment minimize downtime?

- A. By using blue and green CSS themes
- B. By maintaining two identical production environments and switching traffic from the current (blue) to the new (green) after verification
- C. By deploying only during blue or green light hours
- D. By coloring server logs

Answer: B

Q997. What is the purpose of deployment previews?

- A. To preview CSS changes
- B. To automatically deploy each pull request to a unique URL for review and testing before merging
- C. To preview database migrations
- D. To preview HTML layouts

Answer: B

Q998. What is the purpose of a load balancer in deployment architecture?

- A. To balance CSS loads
- B. To distribute incoming traffic across multiple server instances for reliability and performance
- C. To balance database loads only
- D. To balance JavaScript execution

Answer: B

Q999. What is the purpose of monitoring and alerting in production deployments?

- A. To monitor CSS file sizes
- B. To track application health, performance, and errors, and notify teams when issues are detected
- C. To monitor HTML rendering
- D. To monitor database schemas

Answer: B

Q1000. What is an artifact in the context of deployment?

- A. An ancient CSS technique
- B. A built and versioned package of the application ready for deployment to an environment
- C. A database artifact
- D. An HTML artifact

Answer: B

Q1001. What is the purpose of API versioning and what strategies exist?

- A. To version CSS files
- B. To maintain backward compatibility when making breaking changes; strategies include URL path, query parameter, and header-based versioning
- C. To version HTML files
- D. To version database schemas

Answer: B

Q1002. What is the difference between REST and GraphQL APIs?

- A. They are identical
- B. REST uses fixed endpoints returning predefined data while GraphQL uses a single endpoint where clients specify exactly what data they need
- C. REST is always faster
- D. GraphQL cannot handle mutations

Answer: B

Q1003. What is the purpose of API documentation tools like Swagger/OpenAPI?

- A. To document CSS styles
- B. To provide interactive documentation that describes endpoints, request/response formats, and allows testing APIs directly
- C. To document HTML structures
- D. To document database schemas only

Answer: B

Q1004. What is the purpose of cursor-based pagination compared to offset-based pagination?

- A. To style cursor in CSS
- B. Cursor-based pagination uses an opaque pointer to the last item, avoiding the performance and consistency issues of offset-based pagination on large datasets
- C. To manage database cursors
- D. To paginate CSS styles

Answer: B

Q1005. What is the purpose of HTTP caching headers in API responses?

- A. To cache CSS styles
- B. To control whether and how long API responses can be cached by browsers and intermediaries, reducing server load and latency
- C. To cache HTML pages only
- D. To cache database queries

Answer: B

Q1006. What is the purpose of request validation in API development?

- A. To validate CSS syntax
- B. To verify that incoming API requests contain valid data types, required fields, and acceptable values before processing
- C. To validate HTML structure
- D. To validate database schemas

Answer: B

Q1007. What is the purpose of webhook retry mechanisms?

- A. To retry CSS loading
- B. To automatically reattempt failed webhook deliveries with exponential backoff, ensuring event delivery despite temporary failures
- C. To retry database connections
- D. To retry HTML rendering

Answer: B

Q1008. What is the purpose of API error response standardization?

- A. To standardize CSS errors
- B. To return consistent error formats with status codes, error codes, and messages so clients can handle errors programmatically
- C. To standardize HTML errors
- D. To standardize database errors

Answer: B

Q1009. What is the purpose of API throttling versus rate limiting?

- A. They are identical concepts
- B. Rate limiting blocks requests exceeding a limit while throttling queues or slows them, allowing eventual processing
- C. Throttling is always faster
- D. Rate limiting queues requests

Answer: B

Q1010. How does hypermedia in REST API responses guide client navigation?

- A. It adds CSS styles to responses
- B. It includes links to related resources and available actions, allowing clients to navigate the API without hardcoded URLs
- C. It embeds HTML pages in JSON
- D. It manages database hyperlinks

Answer: B

Q1011. What is the purpose of design tokens in modern design systems?

- A. Tokens for accessing CSS APIs
- B. Atomic design decisions (colors, spacing, typography) stored as platform-agnostic variables that ensure consistency across products
- C. Database access tokens
- D. JavaScript authentication tokens

Answer: B

Q1012. How does monorepo tooling like Turborepo improve development workflows?

- A. It creates CSS monorepos
- B. It manages multiple packages in one repository with intelligent caching, parallel execution, and dependency-aware task orchestration
- C. It only manages databases
- D. It replaces Git

Answer: B

Q1013. How do web components achieve framework-agnostic reusability?

- A. They only work inside React applications
- B. They use browser-native APIs like Custom Elements, Shadow DOM, and HTML Templates to create encapsulated components usable in any framework
- C. They require a special bundler plugin
- D. They depend on a shared CSS library

Answer: B

Q1014. What is the concept of incremental adoption in modern web frameworks?

- A. Installing the entire framework at once
- B. Gradually integrating a framework into an existing application page by page instead of rewriting everything
- C. Adopting only the CSS parts of a framework
- D. Using multiple frameworks simultaneously

Answer: B

Q1015. What is the purpose of TypeScript in modern web development?

- A. To replace JavaScript entirely
- B. To add static type checking to JavaScript, catching type errors at compile time and improving developer experience
- C. To type CSS styles
- D. To type database queries

Answer: B

Q1016. What is the difference between optimistic and pessimistic UI updates?

- A. Optimistic updates are slower but more accurate
- B. Optimistic updates assume success and update the UI immediately before server confirmation, while pessimistic updates wait for the server response
- C. Pessimistic updates never send data to servers
- D. Optimistic updates only work with REST APIs

Answer: B

Q1017. What is the purpose of Bun as a JavaScript runtime?

- A. To bundle CSS files
- B. To provide a fast all-in-one JavaScript runtime, bundler, test runner, and package manager as an alternative to Node.js
- C. To create database bundles
- D. To manage HTML bundles

Answer: B

Q1018. What is the purpose of Web Components standards?

- A. Standard CSS components
- B. A set of browser APIs (Custom Elements, Shadow DOM, HTML Templates) for creating reusable, encapsulated UI components
- C. Standard database components
- D. Standard server components

Answer: B

Q1019. What is the purpose of a build tool like Webpack in modern development?

- A. To build CSS from scratch
- B. To bundle, transform, and optimize JavaScript modules and assets for production deployment
- C. To build database schemas
- D. To build HTML templates

Answer: B

Q1020. What is the purpose of containerization in modern web development?

- A. To contain CSS styles
- B. To package applications with all dependencies into isolated containers that run consistently across any environment
- C. To contain database records
- D. To contain HTML elements

Answer: B

Web Development — Hard

510 questions

Q1021. Which rendering strategy delivers HTML generated at request time from the server?

- A. Static Site Generation
- B. Server-Side Rendering
- C. Client-Side Rendering
- D. Incremental Static Regeneration

Answer: B

Q1022. What is the main advantage of a Single Page Application (SPA)?

- A. It avoids full page reloads by dynamically updating
- B. It does not need to communicate with any APIs
- C. It renders only a single HTML element on screen
- D. It can function with no JavaScript whatsoever

Answer: A

Q1023. What is the JAMstack architecture?

- A. jQuery, Angular, MongoDB
- B. Java, Apache, MySQL stack
- C. JSON, AJAX, MVC
- D. JavaScript, APIs, and Markup

Answer: D

Q1024. Which tool is used for bundling JavaScript modules?

- A. Git
- B. Postman
- C. Parcel
- D. Webpack

Answer: D

Q1025. What is the purpose of a service worker?

- A. To compile TypeScript into JavaScript
- B. To manage and process database queries
- C. To enable offline functionality and sync
- D. To minify and bundle CSS file assets

Answer: C

Q1026. What is Progressive Enhancement?

- A. Starting with a basic experience and adding enhancements for capable ones
- B. Starting with advanced features and removing them for old browsers
- C. Optimizing progressive image loading formats for slower connections
- D. Restricting support exclusively to modern browsers with new features

Answer: A

Q1027. What distinguishes a Progressive Web App (PWA) from a regular web app?

- A. PWAs do not use standard HTML markup or CSS styling
- B. PWAs require React as a mandatory framework dependency
- C. PWAs work offline, can be installed, and send notifications
- D. PWAs only function on mobile devices and tablets

Answer: C

Q1028. In web development, what is hydration?

- A. Applying water-themed CSS styling to elements
- B. Loading and parsing external CSS style sheets
- C. Attaching JavaScript event handlers to server HTML
- D. Compressing and optimizing image file formats

Answer: C

Q1029. What is tree shaking in modern web development?

- A. Balancing binary search trees using JavaScript
- B. Reorganizing the DOM tree for better performance
- C. Clearing the browser cache for fresh content
- D. Removing unused code from the final build output

Answer: D

Q1030. What is the critical rendering path?

- A. The network path from the server to the browser
- B. The route a URL takes through the DNS resolution
- C. The order in which CSS files are loaded and parsed
- D. The sequence of steps the browser takes to render

Answer: D

Q1031. What is the difference between HTTP/1.1 and HTTP/2?

- A. HTTP/2 does not support any form of HTTPS encryption
- B. HTTP/1.1 is significantly faster in performance than HTTP/2
- C. HTTP/2 supports multiplexing requests over one connection
- D. HTTP/2 is text-based while HTTP/1.1 uses binary framing

Answer: C

Q1032. What is CORS in web architecture?

- A. Cross-Origin Resource Sharing, allowing cross-domain requests
- B. A NoSQL document-oriented database for storing data
- C. A lightweight JavaScript library for DOM manipulation
- D. A popular CSS utility-first framework for styling apps

Answer: A

Q1033. What is the purpose of the OPTIONS HTTP method?

- A. To retrieve existing data from a server resource
- B. To update existing data on the specified server
- C. To permanently delete a resource from the server
- D. To describe the communication options for a resource

Answer: D

Q1034. What does the HTTP status code 304 indicate?

- A. Bad request due to client error
- B. Not modified, use the cached version
- C. A permanent redirect to another URL
- D. Internal server processing error

Answer: B

Q1035. What is HTTP/3 based on?

- A. FTP protocol
- B. TCP protocol
- C. WebSocket API
- D. UDP via QUIC

Answer: D

Q1036. What is the Same-Origin Policy?

- A. A security measure restricting cross-origin script access
- B. A policy requiring all files on a single web server
- C. A CSS naming convention for component-scoped styling
- D. A database access rule governing read-write permissions

Answer: A

Q1037. What is a reverse proxy?

- A. A client-side mechanism for caching web assets
- B. A type of VPN for securing remote connections
- C. A proxy system that blocks all incoming traffic
- D. A server that forwards client requests to backends

Answer: D

Q1038. What is content negotiation in HTTP?

- A. Negotiating the pricing for a website project
- B. The process of agreeing on the response format
- C. A handshake protocol for secure connections
- D. A type of compression for reducing file sizes

Answer: B

Q1039. What is the purpose of an ETag header?

- A. To specify the content encoding type
- B. To set and manage session cookies
- C. To encrypt transmitted data in transit
- D. To provide a version identifier for cache

Answer: D

Q1040. What is the difference between stateful and stateless protocols?

- A. Stateful protocols are always significantly faster
- B. Stateless protocols store all data on the server
- C. Stateless protocols do not retain session information
- D. Stateful protocols are unable to use any cookies

Answer: C

Q1041. What is the purpose of the data-* attribute in HTML5?

- A. To define strict data types for form fields
- B. To store custom data accessible via JavaScript
- C. To store CSS data in external stylesheets
- D. To create connections to external databases

Answer: B

Q1042. What is the Shadow DOM?

- A. A hidden HTML element in the page structure
- B. An encapsulated DOM tree that isolates styles
- C. A dark mode feature provided by CSS natively
- D. A deprecated feature from earlier HTML versions

Answer: B

Q1043. What is the purpose of the <template> tag in HTML5?

- A. To create reusable email template layouts
- B. To define HTML content not rendered until activated
- C. To create reusable CSS template style blocks
- D. To import external HTML files into a document

Answer: B

Q1044. Which attribute makes an HTML element editable?

- A. editable='true'
- B. edit='true'
- C. writable='true'
- D. contenteditable='true'

Answer: D

Q1045. What is the purpose of the srcset attribute in ?

- A. To set the image dimensions and size
- B. To provide multiple sources for responsive images
- C. To add a link to the image source code
- D. To set the primary image source URL

Answer: B

Q1046. What is an HTML Web Component?

- A. A standard built-in HTML tag element
- B. A server-side rendering technique for pages
- C. A type of CSS framework for web styling
- D. A reusable custom element with encapsulated code

Answer: D

Q1047. What does the defer attribute do on a <script> tag?

- A. Prevents the script from ever running on page
- B. Downloads script in parallel and runs after parsing
- C. Makes the script load and execute asynchronously
- D. Defers the loading of CSS stylesheets

Answer: B

Q1048. What is the difference between defer and async in script loading?

- A. async preserves execution order; defer does not order them
- B. defer blocks page rendering; async does not block rendering
- C. defer runs after parsing in order; async runs on download
- D. They behave the same in all situations

Answer: C

Q1049. What is ARIA in HTML?

- A. A type of interactive HTML form input element
- B. A client-side JavaScript rendering framework
- C. A CSS animation and transition effects library
- D. Accessible Rich Internet Applications attributes

Answer: D

Q1050. What is the purpose of the <picture> element in HTML5?

- A. To embed inline SVG graphics within the page
- B. To provide multiple sources for responsive images
- C. To create a drawing canvas for vector graphics
- D. To display a photo gallery with thumbnails

Answer: B

Q1051. What is the difference between em and rem units?

- A. em is an absolute unit measured in exact pixels
- B. em is relative to parent font; rem to the root
- C. rem is relative to parent; em is relative to root
- D. They are identical in how they calculate sizing

Answer: B

Q1052. What are CSS custom properties (variables)?

- A. A type of CSS selector for targeting elements
- B. JavaScript variables referenced from within CSS
- C. User-defined values with -- prefix and var() access
- D. Predefined CSS values built into the browser

Answer: C

Q1053. What is the will-change CSS property?

- A. An animation delay property for scheduling effects
- B. A transition property controlling animation timing
- C. A property used to change element content
- D. A hint about which properties will change for optimization

Answer: D

Q1054. How does the CSS calc() function work?

- A. It calculates the total page loading time
- B. It performs math with mixed units in CSS values
- C. It calculates JavaScript expressions in CSS
- D. It counts the number of DOM elements on page

Answer: B

Q1055. What is the stacking context in CSS?

- A. A method of combining multiple CSS stylesheets
- B. A CSS Grid feature for placing items in cells
- C. A method for stacking div elements vertically
- D. A 3D conceptualization of elements along z-axis

Answer: D

Q1056. What does the @supports rule do in CSS?

- A. Provides support for external web fonts
- B. Checks if a browser supports a CSS feature
- C. Enables support for older browser versions
- D. Adds browser support for legacy features

Answer: B

Q1057. What is the difference between :nth-child() and :nth-of-type()?

- A. :nth-child only counts elements sharing the same type
- B. They are identical in their selection behavior
- C. :nth-of-type counts all siblings regardless of their type
- D. :nth-child counts all siblings; :nth-of-type counts same type

Answer: D

Q1058. What is a CSS combinator?

- A. A color combination and palette tool
- B. A CSS framework for rapid prototyping
- C. A character explaining selector relationships
- D. A CSS minification and optimization tool

Answer: C

Q1059. What is the contain property in CSS?

- A. A flexbox property for aligning child items
- B. A property indicating subtree independence for optimization
- C. A property to contain overflowing text content
- D. A text wrapping property for long strings

Answer: B

Q1060. What is the CSS clamp() function?

- A. A function constraining a value between min and max
- B. A function to lock the scroll position on a page
- C. A function to fix element positions in a container
- D. A function to group elements together visually

Answer: A

Q1061. What is CSS Subgrid?

- A. A third-party CSS grid component library
- B. A sub-layout variation of the Flexbox model
- C. A Grid feature letting nested grids adopt parent tracks
- D. A standalone CSS framework for page layouts

Answer: C

Q1062. What is the CSS aspect-ratio property?

- A. A media query feature for responsive design
- B. The ratio of CSS file sizes in a project
- C. A JavaScript API for measuring dimensions
- D. Sets a preferred aspect ratio for an element

Answer: D

Q1063. What is CSS Container Queries?

- A. Docker container related CSS configuration
- B. Media queries designed for flex containers
- C. Queries about the size of CSS containers
- D. Rules applying styles based on container size

Answer: D

Q1064. What is the CSS isolation property?

- A. Removes an element completely from the DOM tree
- B. Creates a new stacking context for the element
- C. Isolates an element from the internet connection
- D. Blocks all external CSS from affecting the element

Answer: B

Q1065. What is the difference between transform and position for moving elements?

- A. transform skips layout recalculation; position triggers it
- B. transform only works within CSS Grid containers
- C. position is faster for rendering than transform
- D. They produce the same visual result always

Answer: A

Q1066. What are CSS Logical Properties?

- A. Properties with conditional logic based on breakpoints
- B. Properties that adapt to writing direction like inline-start
- C. Boolean CSS properties that toggle features on or off
- D. Properties that use logic gates for conditional values

Answer: B

Q1067. What is the prefers-reduced-motion media query?

- A. A query detecting slow network connection speeds
- B. A query specifically targeting mobile touch devices
- C. A query detecting if user prefers less animation
- D. A query for applying print-specific style rules

Answer: C

Q1068. What is CSS Houdini?

- A. A popular CSS framework for responsive layouts
- B. Low-level APIs giving access to the CSS engine
- C. A browser debugging tool for inspecting CSS
- D. A magic trick animation library for CSS effects

Answer: B

Q1069. What is the content-visibility CSS property?

- A. Enables browser to skip rendering off-screen elements
- B. Controls how text content is displayed on screen
- C. An alternative to z-index for controlling layer order
- D. A toggle switching element visibility on and off

Answer: A

Q1070. How does the CSS cascade layer (@layer) work?

- A. It creates z-index layers for positioned elements
- B. It layers multiple background images on a single element
- C. It creates visible visual layers stacked on the page
- D. It defines explicit layers giving control over precedence

Answer: D

Q1071. What is the event loop in JavaScript?

- A. A mechanism handling async callbacks via call stack
- B. An event listener looping through all DOM elements
- C. A standard for loop iterating over event arrays
- D. A DOM traversal method for walking element tree

Answer: A

Q1072. What is the difference between == and === in terms of type coercion?

- A. Both operators perform type coercion
- B. Neither operator performs type coercion
- C. == performs type coercion; === does not
- D. === performs type coercion; == does not

Answer: C

Q1073. What is the Temporal Dead Zone (TDZ)?

- A. A time-related API for scheduling delayed operations
- B. A concept related to memory management and GC
- C. A garbage collection phase reclaiming unused memory
- D. The period between scope entry and let/const declaration

Answer: D

Q1074. What is the difference between call(), apply(), and bind()?

- A. Only bind invokes the function immediately when called
- B. call and apply invoke immediately; bind returns a new function
- C. They are all identical in functionality
- D. call is for arrays; apply is for objects arguments

Answer: B

Q1075. What is a WeakMap in JavaScript?

- A. A collection with weakly referenced object keys for GC
- B. A smaller and more limited version of Map class
- C. A map collection with weak security encryption
- D. A map storing only string keys and string values

Answer: A

Q1076. What is the prototype chain in JavaScript?

- A. A CSS inheritance model for cascading styles
- B. The mechanism objects use to inherit properties
- C. A linked list data structure for sequential data
- D. A blockchain implementation for distributed apps

Answer: B

Q1077. What is tail call optimization in JavaScript?

- A. A DOM rendering optimization for reducing paint operations
- B. A CSS optimization technique for faster style calculations
- C. Optimizing the last element of an array for access
- D. Reusing the stack frame for a function call in tail position

Answer: D

Q1078. What is the difference between Object.freeze() and Object.seal()?

- A. freeze prevents all changes; seal allows modifying values
- B. freeze only prevents deletion of properties
- C. seal prevents all changes including modifications
- D. They are identical in behavior on objects

Answer: A

Q1079. What is a Symbol in JavaScript?

- A. A unique immutable primitive for property identifiers
- B. A string formatting token for template processing
- C. A CSS icon or symbol for decorative elements
- D. A mathematical operator for advanced calculations

Answer: A

Q1080. What is the purpose of generator functions?

- A. To generate HTML markup from template files
- B. To create iterators that pause and resume via yield
- C. To create class instances from constructor functions
- D. To generate random numbers for cryptographic use

Answer: B

Q1081. What is the difference between microtasks and macrotasks?

- A. Microtasks run on a separate background thread
- B. Macrotasks always have higher execution priority
- C. They are identical in execution priority
- D. Microtasks execute before macrotasks in the loop

Answer: D

Q1082. What is a SharedArrayBuffer?

- A. A binary data buffer shared between workers
- B. A database buffer for caching query results
- C. A shared CSS buffer for component styling
- D. An array with elements shared across files

Answer: A

Q1083. What is the Reflect API in JavaScript?

- A. A browser debugging tool for inspecting elements
- B. A library for creating mirror image visual effects
- C. A CSS reflection API for reflecting element visuals
- D. A built-in object with interceptable operation methods

Answer: D

Q1084. What is structural typing in TypeScript?

- A. TypeScript uses nominal typing with explicit declarations
- B. JavaScript has stronger typing than TypeScript overall
- C. They use the identical type system without differences
- D. TypeScript checks compatibility based on type structure

Answer: D

Q1085. What is the AbortController API?

- A. A process manager for handling background server tasks
- B. An error handling mechanism for catching thrown errors
- C. An API to cancel fetch requests and async operations
- D. A method to stop all JavaScript execution immediately

Answer: C

Q1086. What are Decorators in JavaScript/TypeScript?

- A. Functions that modify classes and members via @ syntax
- B. String formatting tools for template processing
- C. CSS decorations like border styles and shadows
- D. HTML attributes added to elements for metadata

Answer: A

Q1087. What is the structuredClone() function?

- A. A function creating deep copies via structured clone
- B. A CSS cloning function for duplicating style rules
- C. A function to clone the DOM structure of elements
- D. A shallow copy method copying only top-level values

Answer: A

Q1088. What is the difference between for...in and for...of?

- A. for...in only works with arrays not objects
- B. for...of iterates over property keys of objects
- C. They are identical in their iteration behavior
- D. for...in iterates keys; for...of iterates values

Answer: D

Q1089. What is a FinalizationRegistry in JavaScript?

- A. An API to register a callback when an object is GC'd
- B. A module registry for tracking installed packages
- C. A service worker registry for managing active workers
- D. A software design pattern for object creation

Answer: A

Q1090. What is the Intl object in JavaScript?

- A. An interface definition for declaring TypeScript contracts
- B. An interval timer for scheduling repeated function calls
- C. An internal object for private variable storage
- D. A namespace for Internationalization with locale formatting

Answer: D

Q1091. What is React's Concurrent Mode?

- A. Parallel CSS processing across multiple browser threads
- B. Multi-threaded JavaScript execution using Web Workers
- C. Running multiple separate React applications at once
- D. Features helping React stay responsive via background rendering

Answer: D

Q1092. What is the purpose of React.memo()?

- A. Creating in-app memo or notification features
- B. A HOC memoizing a component to prevent re-renders
- C. Saving developer notes within the application code
- D. Memory management and garbage collection control

Answer: B

Q1093. What is Angular's Change Detection strategy OnPush?

- A. A Git branching strategy for version control
- B. Detecting CSS property changes in stylesheets
- C. Push notifications sent to mobile device users
- D. A strategy checking only when input references change

Answer: D

Q1094. What is the Composition API in Vue 3?

- A. A music composition tool for creating audio
- B. A CSS composition method for combining styles
- C. A database query API for composing SQL queries
- D. An API organizing component logic using setup()

Answer: D

Q1095. What is code splitting in frontend frameworks?

- A. Splitting CSS files into smaller partial files
- B. Dividing server code into separate microservices
- C. Splitting HTML files across multiple server routes
- D. Breaking the bundle into smaller on-demand chunks

Answer: D

Q1096. What is Server-Side Rendering (SSR) in React/Next.js?

- A. Rendering CSS on the server before page delivery
- B. Running JavaScript only on server, never on client
- C. Generating HTML on the server for each client request
- D. A database rendering technique for cached query results

Answer: C

Q1097. What is React's useCallback hook?

- A. A callback URL handler for authentication flows
- B. An event listener for handling DOM user events
- C. A hook returning a memoized callback function
- D. A state management hook for complex data flows

Answer: C

Q1098. What is the purpose of useMemo in React?

- A. Memoizing expensive computations until deps change
- B. Creating memoized components that never re-render
- C. Memory allocation and management control utilities
- D. Managing in-app memo and notification features

Answer: A

Q1099. What is Angular's Dependency Injection?

- A. A testing technique for isolating component behavior
- B. Importing npm packages into project configuration
- C. A class receiving dependencies from external sources
- D. Injecting CSS dependencies into component styles

Answer: C

Q1100. What is Incremental Static Regeneration (ISR) in Next.js?

- A. Updating static pages after deploy without full rebuild
- B. Regenerating CSS incrementally during development
- C. Incremental database updates during data migration
- D. A version control strategy for incremental releases

Answer: A

Q1101. What is the difference between horizontal and vertical scaling?

- A. Horizontal scaling is always cheaper in every scenario
- B. Vertical adds power to existing servers; horizontal adds servers
- C. Vertical scaling is only applicable to database servers
- D. They are identical in their scaling approach

Answer: B

Q1102. What is a message queue in backend architecture?

- A. A CSS animation queue for scheduling transitions
- B. A database query queue for sequential execution
- C. An async mechanism storing messages for consumers
- D. A chat feature for sending user messages

Answer: C

Q1103. What is GraphQL and how does it differ from REST?

- A. A JavaScript charting tool for rendering bar graphs
- B. A graph database for storing relationship data
- C. A query language where clients specify exact data needed
- D. A CSS graph library for data visualization charts

Answer: C

Q1104. What is the N+1 query problem?

- A. Fetching a list triggers one query plus N for related data
- B. A database naming convention for table schemas
- C. A CSS specificity issue with deeply nested selectors
- D. A JavaScript error caused by off-by-one indexing

Answer: A

Q1105. What is a microservices architecture?

- A. A small lightweight embedded database system
- B. A small application with limited features overall
- C. Independent services communicating via defined APIs
- D. A CSS micro-framework for minimal page styling

Answer: C

Q1106. What is database connection pooling?

- A. A database backup method for disaster recovery plans
- B. Maintaining reusable database connections to reduce overhead
- C. A swimming pool database for aquatic data
- D. A CSS pooling technique for shared style resources

Answer: B

Q1107. What is idempotency in API design?

- A. A security feature for encrypting API payloads
- B. A mathematical concept from pure number theory
- C. A caching mechanism for storing API responses locally
- D. Making the same request multiple times yields same result

Answer: D

Q1108. What is the CAP theorem?

- A. A capping limit on the number of API requests
- B. A CSS theorem about cascade priority order
- C. A JavaScript design pattern for state management
- D. A distributed system can guarantee only two of C, A, P

Answer: D

Q1109. What is event-driven architecture?

- A. A database event log for auditing record changes
- B. CSS event handling through pseudo-class selectors
- C. A React-specific pattern for component lifecycle
- D. A pattern where flow is determined by produced events

Answer: D

Q1110. What is the purpose of a reverse proxy like Nginx?

- A. To reverse engineer API endpoints for testing
- B. To handle load balancing, SSL, caching, and routing
- C. To reverse the order of database query results
- D. To revert code changes to a previous version

Answer: B

Q1111. What is the difference between process-based and event-driven server architectures?

- A. Process-based architectures cannot handle concurrent requests
- B. They are identical in concurrency approach
- C. Event-driven is always slower in handling concurrent requests
- D. Process-based creates threads per request; event-driven uses async I/O

Answer: D

Q1112. What is the purpose of a load balancer?

- A. To balance database queries across shards
- B. To balance CSS style specificity weights
- C. To load JavaScript files in parallel
- D. To distribute traffic across multiple servers

Answer: D

Q1113. What is serverless computing?

- A. A client-side only architecture without any backend logic
- B. Computing without any servers whatsoever
- C. A static site generator for building HTML pages
- D. A cloud model with provider-managed infrastructure on demand

Answer: D

Q1114. What is the difference between monolithic and microservices architecture?

- A. They are identical in application structure
- B. Monolithic is a single unit; microservices are independent
- C. Microservices are always faster in every scenario
- D. Monolithic applications are unable to scale at all

Answer: B

Q1115. What is gRPC?

- A. A CSS framework for responsive web design
- B. A JavaScript library for building reactive interfaces
- C. A high-performance RPC framework using Protocol Buffers
- D. A relational database management system for apps

Answer: C

Q1116. What is the circuit breaker pattern?

- A. A database locking mechanism for transaction isolation levels
- B. An electrical safety feature for server power
- C. A CSS layout pattern for organizing grid-based page elements
- D. A pattern preventing cascading failures by stopping failing requests

Answer: D

Q1117. What is the difference between synchronous and asynchronous processing?

- A. Synchronous processing cannot handle any I/O operations
- B. Synchronous blocks until complete; asynchronous allows other ops
- C. They are identical in execution flow handling
- D. Asynchronous is always slower than synchronous code processing

Answer: B

Q1118. What is server-sent events (SSE)?

- A. A technology for pushing real-time updates over HTTP
- B. Server-side CSS events triggered by style changes
- C. A database trigger that fires on record modifications
- D. A JavaScript event system for handling user interactions

Answer: A

Q1119. What is the purpose of health check endpoints?

- A. Checking only database health and connection status
- B. Checking user health data in medical applications
- C. Checking CSS health for validation errors
- D. Providing a way for load balancers to verify server health

Answer: D

Q1120. What is the Twelve-Factor App methodology?

- A. A database design pattern for normalized tables
- B. A CSS framework following twelve core principles
- C. A methodology for building scalable SaaS apps
- D. A 12-step coding tutorial for learning web dev

Answer: C

Q1121. What is database sharding?

- A. Encrypting data for security and privacy compliance
- B. Compressing tables to reduce disk storage space
- C. Splitting a database into smaller distributed pieces
- D. Breaking a database beyond repair permanently

Answer: C

Q1122. What is the difference between horizontal and vertical partitioning?

- A. Horizontal splits columns; vertical splits rows across servers
- B. Neither approach has any effect on query performance
- C. Horizontal splits rows across servers; vertical splits columns
- D. They are identical in partitioning strategy

Answer: C

Q1123. What is eventual consistency?

- A. A CSS consistency model for cross-browser rendering
- B. Data is always immediately consistent everywhere
- C. All replicas converge to the same value given enough time
- D. Data is never consistent across any of the replicas

Answer: C

Q1124. What is a database migration?

- A. A data backup process for disaster recovery planning
- B. A CSS migration from one framework to another tool
- C. Version-controlled changes to schema that can be rolled back
- D. Moving a database to a new physical data center

Answer: C

Q1125. What is the N+1 query problem and how is it solved?

- A. A JavaScript error caused by off-by-one array access
- B. A CSS specificity issue with deeply nested selectors
- C. A mathematical formula for calculating query cost
- D. N extra queries for related data; solved by eager loading

Answer: D

Q1126. What is a materialized view?

- A. A database object storing query results physically
- B. A CSS view for presenting styled HTML content
- C. A template engine view for generating dynamic pages
- D. A JavaScript view for rendering reactive UI components

Answer: A

Q1127. What is connection pooling in database management?

- A. A swimming pool database for aquatic records
- B. A CSS pooling technique for shared style rule resources
- C. A JavaScript object pool for reusing DOM element references
- D. Maintaining reusable connections to avoid per-request overhead

Answer: D

Q1128. What is a write-ahead log (WAL)?

- A. A CSS log file for tracking style computation results
- B. A blogging technique for writing web content
- C. Changes are written to a log before applying to the DB
- D. A JavaScript console log for outputting debug messages

Answer: C

Q1129. What is database replication?

- A. A backup tool for creating periodic data archive snapshots
- B. A version control system for tracking file change history
- C. Copying and maintaining data across servers for redundancy
- D. Copying CSS files across multiple project directories

Answer: C

Q1130. What is the difference between optimistic and pessimistic locking?

- A. They only apply to NoSQL databases not relational ones
- B. They are identical in concurrency control approach
- C. Optimistic always locks; pessimistic never locks anything
- D. Optimistic checks at commit; pessimistic locks immediately

Answer: D

Q1131. What is a SQL injection attack and how is it prevented?

- A. A CSS vulnerability from untrusted stylesheet loading
- B. A database feature for optimizing query performance
- C. Inserting malicious SQL; prevented by parameterized queries
- D. A JavaScript error from uncaught promise rejections

Answer: C

Q1132. What is the difference between symmetric and asymmetric encryption?

- A. They are identical in their encryption approach
- B. Asymmetric is always faster than symmetric encryption
- C. Symmetric uses one key; asymmetric uses a key pair
- D. Symmetric encryption requires two separate keys

Answer: C

Q1133. What is a man-in-the-middle (MITM) attack?

- A. A database corruption from concurrent write conflicts
- B. An attack intercepting communication between two parties
- C. A network administrator managing server infrastructure
- D. A CSS hack for overriding conflicting style declarations

Answer: B

Q1134. What is the SameSite cookie attribute?

- A. An attribute controlling cookies in cross-site requests
- B. A site name attribute used in HTML anchor tags
- C. A CSS property for styling site-specific page elements
- D. A database attribute for defining column data constraints

Answer: A

Q1135. What is a zero-day vulnerability?

- A. A CSS bug affecting zero-width element rendering
- B. A vulnerability lasting zero days before being fixed
- C. A database error caused by zero-division in queries
- D. A security flaw exploited before the vendor has a patch

Answer: D

Q1136. What is OAuth 2.0?

- A. An authorization framework for limited third-party access
- B. A JavaScript library for building authentication flows
- C. A CSS standard for defining cross-browser specifications
- D. A database protocol for encrypted data transmission

Answer: A

Q1137. What is Subresource Integrity (SRI)?

- A. A database integrity check for verifying data consistency
- B. A CSS property for defining element integrity
- C. A feature verifying fetched resources are not tampered with
- D. A JavaScript framework for building secure web apps

Answer: C

Q1138. What is the purpose of rate limiting in security?

- A. Limiting the speed of CSS transition animations
- B. Restricting request frequency to prevent brute force
- C. Limiting the number of page loads per session
- D. Limiting the number of database queries per session

Answer: B

Q1139. What is X-Frame-Options used for?

- A. A CSS frame property for creating bordered layouts
- B. Preventing iframe embedding to stop clickjacking
- C. A JavaScript frame option for window management
- D. Framing images within decorative CSS borders

Answer: B

Q1140. What is the difference between encryption at rest and in transit?

- A. At rest encryption is only for database systems
- B. At rest encrypts stored data; in transit encrypts transmitted
- C. They are identical in their security approach
- D. In transit encryption is only for local file systems

Answer: B

Q1141. What is the difference between preload, prefetch, and preconnect?

- A. They are identical in resource loading behavior
- B. Prefetch is exclusively used for CSS stylesheet resources only
- C. Preconnect downloads files from the server in advance before needed
- D. Preload fetches current-page resources; prefetch loads future; preconnect opens connections

Answer: D

Q1142. What is the PRPL pattern?

- A. A CSS color pattern for theming web pages
- B. Push, Render, Pre-cache, Lazy-load for delivery
- C. A database pattern for query optimization strategies
- D. A JavaScript design pattern for state management

Answer: B

Q1143. What is the purpose of requestAnimationFrame?

- A. To request data from REST API server endpoints
- B. To request images from a remote CDN server
- C. To schedule animations synced with the repaint cycle
- D. To create CSS animations using keyframe declarations

Answer: C

Q1144. What is the impact of the critical rendering path on performance?

- A. No impact on web application performance at all
- B. It determines how quickly a page becomes visible
- C. It only affects how CSS styles are calculated
- D. It only affects how images are decoded and shown

Answer: B

Q1145. What is HTTP/2 Server Push?

- A. A WebSocket feature for real-time data streaming
- B. A database push operation for replication updates
- C. Proactively sending resources before they are requested
- D. Pushing CSS stylesheets to the client browser

Answer: C

Q1146. What is the difference between synthetic and real user monitoring?

- A. They are identical monitoring approaches
- B. Synthetic uses simulated tests; RUM collects real user data
- C. Synthetic monitoring collects data from actual users
- D. RUM uses simulated automated tests in lab environments

Answer: B

Q1147. What is layout thrashing?

- A. A JavaScript error from undefined variable references
- B. A CSS bug causing layout calculation errors
- C. Forcing repeated layout recalculations by reading/writing DOM
- D. A database deadlock from concurrent transaction conflicts

Answer: C

Q1148. What is the purpose of resource hints in HTML?

- A. To add HTML comments for developer documentation
- B. To create resource files for internationalization support
- C. To inform the browser about resources to preload or prefetch
- D. To add tooltip text to page elements on hover

Answer: C

Q1149. What is the benefit of WebP or AVIF image formats?

- A. Better CSS support across all web browsers
- B. Better database storage efficiency for images
- C. Better JavaScript integration with image APIs
- D. Significantly smaller file sizes than JPEG/PNG

Answer: D

Q1150. What is Time to Interactive (TTI)?

- A. A CSS animation duration for transition effects
- B. The time until the page is fully interactive
- C. A server response time for processing API requests
- D. The time to write the application source code

Answer: B

Q1151. What is the testing pyramid?

- A. A database structure for hierarchical data storage
- B. A pyramid-shaped test data structure for storage
- C. Many unit tests, fewer integration, fewest E2E tests
- D. A CSS shape rendered using clip-path properties

Answer: C

Q1152. What is mutation testing?

- A. Testing DNA sequences in bioinformatics research
- B. Testing code changes before merging into main branch
- C. Database mutation testing for verifying data integrity
- D. Introducing small code changes to verify tests detect them

Answer: D

Q1153. What is the difference between stubs and mocks?

- A. Mocks are simpler and less capable than stub versions
- B. They are identical in their testing approach
- C. Stubs give predetermined responses; mocks verify interactions
- D. Stubs verify that specific method interactions occurred

Answer: C

Q1154. What is visual regression testing?

- A. Testing CSS variables across different components
- B. Testing image quality and resolution across screen sizes
- C. Testing color contrast for accessibility compliance
- D. Comparing screenshots across versions for visual changes

Answer: D

Q1155. What is contract testing in microservices?

- A. Legal testing of software license agreements
- B. CSS contract testing for cross-browser consistency
- C. Testing employment contracts for HR compliance
- D. Verifying services communicate per agreed interfaces

Answer: D

Q1156. What is fuzz testing?

- A. Testing with blurry images for visual quality checks
- B. Providing random invalid data to find vulnerabilities
- C. Testing with fuzzy logic algorithms for AI systems
- D. CSS filter testing for verifying blur effect output

Answer: B

Q1157. What is the purpose of Selenium?

- A. A JavaScript runtime for executing server-side scripts
- B. A framework for automating browser testing interactions
- C. A CSS preprocessor for compiling advanced style syntax
- D. A chemical element found in nature supplements

Answer: B

Q1158. What is Playwright?

- A. A CSS animation tool for creating transition effects
- B. A theater tool for managing stage performances
- C. A modern testing framework with auto-wait and isolation
- D. A database testing tool for verifying query accuracy

Answer: C

Q1159. What is chaos engineering?

- A. A random testing approach without defined test cases
- B. Random coding without any structure or planning
- C. Deliberately introducing failures to test resilience
- D. A messy code style ignoring best practices rules

Answer: C

Q1160. What is the difference between black-box and white-box testing?

- A. Black-box tests without code knowledge; white-box tests with
- B. They are identical in their testing methodology
- C. Black-box testing requires full access to source code
- D. White-box testing ignores internal implementation details

Answer: A

Q1161. What is blue-green deployment?

- A. A strategy using two environments with instant switching
- B. A color-coded database for organizing data categories
- C. Deploying blue and green CSS color themes
- D. A testing framework for comparing two UI versions

Answer: A

Q1162. What is canary deployment?

- A. A testing technique for incremental feature validation
- B. Deploying software for bird tracking applications
- C. A CSS animation strategy for progressive effects
- D. Gradually rolling out to a small subset of users

Answer: D

Q1163. What is Infrastructure as Code (IaC)?

- A. Managing infrastructure through configuration files
- B. A JavaScript pattern for structuring backend code
- C. Writing HTML for server configuration pages
- D. A CSS framework for building server admin interfaces

Answer: A

Q1164. What is Kubernetes?

- A. A NoSQL database for storing application document data
- B. A JavaScript library for building reactive user interfaces
- C. A CSS framework for responsive web design layouts
- D. A container orchestration platform for deployment and scaling

Answer: D

Q1165. What is the difference between horizontal and vertical scaling in deployment?

- A. Vertical adds more server machines to the cluster
- B. They are identical in scaling approach
- C. Horizontal adds more RAM to the current server
- D. Horizontal adds servers; vertical adds resources to existing

Answer: D

Q1166. What is a CDN edge server?

- A. A CDN server near end users for fast delivery
- B. A development server for local test and debug
- C. A database server for write-heavy workloads
- D. A CSS edge property for element border styling

Answer: A

Q1167. What is serverless deployment?

- A. Deploying functions that run on demand in the cloud
- B. Deploying without any server infrastructure at all
- C. A local deployment for development and testing only
- D. A CSS deployment method for updating style sheets

Answer: A

Q1168. What is the purpose of a health check in deployment?

- A. Automated checks verifying app availability for recovery
- B. A code review process for ensuring quality standards
- C. Checking the physical health of development staff
- D. A CSS check for validating style rule correctness

Answer: A

Q1169. What is GitOps?

- A. A CSS framework for building responsive web layouts
- B. A database tool for managing schema version migrations
- C. A Git client application for version control
- D. A framework where Git repos are truth for infrastructure

Answer: D

Q1170. What is the difference between a monorepo and polyrepo?

- A. They are identical in repository management
- B. Polyrepo is exclusively for backend server application code
- C. Monorepo stores all services in one repo; polyrepo uses separate
- D. Monorepo is exclusively for frontend application code only

Answer: C

Q1171. What is GraphQL's N+1 problem and how is it solved?

- A. A math problem related to computational complexity
- B. A REST limitation with fixed endpoint over-fetching
- C. A CSS specificity issue with deeply nested selectors
- D. Nested field queries solved by DataLoader batching

Answer: D

Q1172. What is an API gateway?

- A. A single entry point routing requests to microservices
- B. A JavaScript gateway for module import resolution
- C. A database gateway for managing connection pooling
- D. A CSS gate for controlling style cascade flow

Answer: A

Q1173. What is the difference between idempotent and non-idempotent methods?

- A. They are identical in request handling behavior
- B. Idempotent methods produce same result on repetition; POST may differ
- C. All HTTP methods are non-idempotent in their behavior always
- D. All HTTP methods are idempotent without any exceptions at all

Answer: B

Q1174. What is content negotiation in APIs?

- A. A CSS negotiation for determining style precedence
- B. A database negotiation for connection pool allocation
- C. Negotiating API pricing for commercial usage plans
- D. Client and server agree on response format via headers

Answer: D

Q1175. What is gRPC and when is it preferred over REST?

- A. A high-performance RPC framework using Protocol Buffers
- B. A JavaScript library for real-time data streaming
- C. A database protocol for distributed query execution
- D. A CSS framework for responsive web design layouts

Answer: A

Q1176. What is the purpose of API throttling?

- A. Controlling request rates to prevent abuse and ensure fairness
- B. A CSS technique for limiting animation frame rates on pages
- C. A database optimization for throttling query execution speed
- D. Speeding up API response times with caching

Answer: A

Q1177. What is the difference between polling and SSE?

- A. Polling is more efficient than SSE for real-time data delivery
- B. SSE requires WebSocket protocol to function properly always
- C. They are identical in data delivery approach
- D. Polling makes repeated requests; SSE maintains persistent connection

Answer: D

Q1178. What are API error handling best practices?

- A. Use appropriate status codes, consistent format, and messages
- B. Return 200 for all responses regardless of outcome
- C. Ignore errors and let the client handle all failures
- D. Use CSS error styles to display formatted error pages

Answer: A

Q1179. What is hypermedia-driven API design?

- A. A CSS animation API for creating interactive effects
- B. A multimedia API for playing audio and video
- C. An API where responses include links guiding transitions
- D. A video streaming API for delivering live content

Answer: C

Q1180. What is the Backend for Frontend (BFF) pattern?

- A. A best friend API for social networking features
- B. A CSS pattern for frontend-specific responsive styling
- C. Each frontend has its own backend aggregating services
- D. A database pattern for frontend-optimized query caching

Answer: C

Q1181. What is Islands Architecture?

- A. A database architecture for isolated schema partitioning
- B. An architecture for building island-themed websites
- C. Interactive components hydrated independently in static HTML
- D. A CSS layout technique for island-shaped page elements

Answer: C

Q1182. What is Edge Computing in web development?

- A. A database optimization for edge-case query handling
- B. Computing at the physical edge of a desk surface
- C. A CSS technique for styling elements at container edges
- D. Running code on servers geographically close to users

Answer: D

Q1183. What is the difference between SSR, SSG, and ISR?

- A. They are identical in rendering strategy and output
- B. ISR renders only once and never updates content again
- C. SSR renders per request; SSG at build; ISR updates after deploy
- D. SSG renders per request at the time of each user visit

Answer: C

Q1184. What is Module Federation in Webpack?

- A. A plugin allowing independent builds to share modules
- B. A country federation of web development standards
- C. A database federation distributing queries across nodes
- D. A CSS module system for scoping styles to components

Answer: A

Q1185. What is Streaming SSR?

- A. Streaming CSS stylesheets for progressive loading
- B. Progressively sending HTML chunks as they are rendered
- C. Streaming video content to the end user browser
- D. Streaming database query results for batch processing

Answer: B

Q1186. What are React Server Components?

- A. Components rendering exclusively on the server side
- B. CSS server-side components for pre-rendering styles
- C. Database components for managing query execution
- D. Components that look like physical server hardware

Answer: A

Q1187. What is the purpose of an Abstract Syntax Tree (AST) in web tooling?

- A. A CSS tree layout for hierarchical navigation menus
- B. A tree representation of code used by compilers and linters
- C. A database index structure for optimizing query lookups
- D. A tree visualization for displaying data hierarchies

Answer: B

Q1188. What is Partial Hydration?

- A. Partial CSS loading for progressive style application
- B. Partial database loading for lazy data fetching
- C. Partially loading water into server cooling systems
- D. Selectively hydrating only interactive components

Answer: D

Q1189. What is the concept of Zero-JS pages?

- A. Pages delivered without any visible content at all
- B. Pages with zero interactive features for static content
- C. Blank pages served as placeholder loading screen views
- D. Pages delivered with no JavaScript, relying on HTML/CSS

Answer: D

Q1190. What is Resumability in frameworks like Qwik?

- A. Continuing execution on client from where server left off
- B. A database feature for resuming interrupted transactions
- C. A CSS property for resuming paused animation effects
- D. Pausing a video player for later playback resumption

Answer: A

Q1191. A developer must choose between an SPA and MPA for a content-heavy news site. What is a key drawback of using an SPA?

- A. SPAs generally have worse initial SEO without server rendering
- B. SPAs cannot serve any static assets like images or stylesheets
- C. SPAs require a completely different programming language overall
- D. SPAs are unable to make any asynchronous API calls to servers

Answer: A

Q1192. Which strategy best addresses the challenge of maintaining large-scale CSS in a team environment?

- A. Adopting a methodology like BEM for consistent naming conventions
- B. Writing all styles inline directly on each individual HTML element
- C. Placing every style rule inside a single monolithic CSS file only
- D. Avoiding CSS entirely and using only JavaScript for all styling

Answer: A

Q1193. What is the primary purpose of a service worker in progressive web apps?

- A. To compile TypeScript source files into standard JavaScript code
- B. To create database connections between client and SQL databases
- C. To enable offline functionality and background sync capabilities
- D. To generate server-side rendered HTML pages from template files

Answer: C

Q1194. In a JAMstack architecture, what does the A stand for and what role does it play?

- A. APIs: handles dynamic functionality through reusable interfaces
- B. Authentication: manages user identity verification for all pages
- C. Applications: provides full server-side rendering capabilities
- D. Automation: runs continuous integration pipelines for deployment

Answer: A

Q1195. A team is debating between a monolithic and microservices architecture. Which concern most favors microservices?

- A. The application only serves static content with no dynamic functionality
- B. The application needs independent scaling of different service components
- C. The project has a single developer with no plans to add more team members
- D. The project has very tight deadlines and needs the simplest possible setup

Answer: B

Q1196. What is the main advantage of using WebAssembly alongside JavaScript?

- A. It runs computationally intensive code at near-native speed levels
- B. It provides built-in database functionality without any server need
- C. It eliminates the need for HTML and CSS in modern web applications
- D. It completely replaces JavaScript for all web development tasks

Answer: A

Q1197. Which factor most impacts the choice between SSR and CSR for a web application?

- A. The programming language used for writing automated test scripts
- B. The color scheme and visual design preferences of the stakeholders
- C. The specific brand of web server hardware used for hosting files
- D. The initial load performance and search engine optimization needs

Answer: D

Q1198. What challenge does micro-frontend architecture primarily address?

- A. Eliminating the need for any shared state management in the app
- B. Reducing the total file size of individual CSS stylesheets loaded
- C. Enabling independent development and deployment of UI components
- D. Removing the requirement for build tools in the development flow

Answer: C

Q1199. In a headless CMS architecture, what is the primary role of the front end?

- A. To perform automated backups of the content management system data
- B. To store and manage all the database records for content authors
- C. To consume content via APIs and render it in any presentation layer
- D. To handle user authentication and authorization for CMS admin users

Answer: C

Q1200. Why might a team choose an isomorphic JavaScript approach for their application?

- A. It eliminates the need for testing on multiple browser environments
- B. It shares code between server and client for consistent rendering
- C. It guarantees zero accessibility issues across all device platforms
- D. It removes all server infrastructure requirements for the project

Answer: B

Q1201. In HTTP/2 server push, what is a potential drawback if not implemented carefully?

- A. It can waste bandwidth by pushing resources the client already cached
- B. It completely disables client-side caching for all website resources
- C. It prevents the use of content delivery networks for static assets
- D. It forces the browser to reload the entire page for every navigation

Answer: A

Q1202. What is the primary benefit of HTTP/3 using QUIC over TCP?

- A. QUIC reduces connection setup latency and handles packet loss better
- B. QUIC allows unlimited concurrent connections without any throttling
- C. QUIC replaces DNS resolution entirely with a new naming system used
- D. QUIC eliminates the need for TLS encryption on all connections

Answer: A

Q1203. A site experiences head-of-line blocking. Which protocol upgrade best addresses this?

- A. Switching from HTTPS back to plain HTTP for simpler data transfer
- B. Upgrading from HTTP/1.1 to HTTP/2 which multiplexes streams within
- C. Replacing the web server with a simple FTP file transfer server app
- D. Downgrading from HTTP/2 to HTTP/1.0 for better compatibility range

Answer: B

Q1204. What is the purpose of the HSTS header in web security architecture?

- A. To set the maximum upload file size allowed by the server endpoint
- B. To define which IP addresses are permitted to access the admin panel
- C. To force browsers to only connect via HTTPS for a specified period
- D. To compress response bodies using gzip or brotli encoding methods

Answer: C

Q1205. In a microservices architecture, what problem does an API gateway solve?

- A. It replaces all microservices with a single monolithic application
- B. It eliminates the need for any inter-service communication protocols
- C. It provides a single entry point handling routing and cross-cutting concerns
- D. It automatically generates the front-end UI from backend data models

Answer: C

Q1206. What is the CAP theorem and which two properties can a distributed system guarantee?

- A. Caching, Authentication, Performance — all three must always be met
- B. Concurrency, Accessibility, Portability — only two can be optimized
- C. Compression, Authorization, Persistence — systems must pick one only
- D. Consistency, Availability, Partition tolerance — pick any two of three

Answer: D

Q1207. What is the role of gRPC compared to REST in service-to-service communication?

- A. gRPC uses JSON over HTTP/1.1 while REST uses binary protocol buffers
- B. gRPC uses Protocol Buffers and HTTP/2 for efficient typed communication
- C. gRPC only supports unidirectional calls while REST supports streaming
- D. gRPC requires no schema definition while REST mandates strict schemas

Answer: B

Q1208. What problem does consistent hashing solve in distributed web systems?

- A. It minimizes key redistribution when nodes are added or removed from clusters
- B. It encrypts all data at rest using a consistent algorithm everywhere
- C. It ensures that all web pages are cached for exactly the same duration
- D. It guarantees that all database queries return results in constant time

Answer: A

Q1209. What is the purpose of circuit breaker pattern in distributed architectures?

- A. To physically disconnect failing hardware servers from the data network
- B. To automatically scale up server resources when traffic increases sharply
- C. To encrypt inter-service communication using rotating TLS certificates
- D. To prevent cascading failures by stopping calls to unresponsive services

Answer: D

Q1210. What is eventual consistency and when is it acceptable in web systems?

- A. Data is never consistent and requires manual synchronization by admins
- B. Data consistency is guaranteed only during server maintenance windows
- C. Data is always immediately consistent across all distributed system nodes
- D. Data will become consistent over time and suits high-availability needs

Answer: D

Q1211. A page has multiple forms. How does the form attribute on an input element help?

- A. It prevents the input from being submitted with any form on the page
- B. It applies CSS styles from a specific form to the input field element
- C. It automatically validates the input using the parent form validation
- D. It associates the input with a form it is not nested inside of by ID

Answer: D

Q1212. What is the Shadow DOM and how does it relate to HTML web components?

- A. A caching technique that stores DOM elements in browser memory shadow
- B. A deprecated API for manipulating DOM elements without using JavaScript
- C. An encapsulation mechanism providing scoped styles and DOM for components
- D. A server-side rendering method that generates HTML in a shadow process

Answer: C

Q1213. What is the purpose of the <slot> element in HTML web components?

- A. It assigns a memory slot for caching frequently accessed DOM elements
- B. It marks a position in the DOM where server-rendered content is placed
- C. It creates a placeholder for user-provided content in shadow DOM trees
- D. It defines a time slot for scheduling lazy-loaded resource downloads

Answer: C

Q1214. How does the Content Security Policy meta tag protect against XSS attacks?

- A. It encrypts all HTML content before sending it from server to browser
- B. It restricts which sources of scripts and resources the browser loads
- C. It blocks all JavaScript execution and relies solely on server actions
- D. It automatically sanitizes all user input before inserting into the DOM

Answer: B

Q1215. What accessibility issue arises from using div and span elements for interactive controls?

- A. They cause slower page rendering compared to using native button elements
- B. They cannot have click event listeners attached through JavaScript code
- C. They lack built-in keyboard support, roles, and accessibility semantics
- D. They are invisible to all users including those not using screen readers

Answer: C

Q1216. What is the purpose of the inert attribute introduced in HTML?

- A. It converts a dynamic element into a static image for print stylesheet
- B. It permanently deletes the element and its children from the DOM tree
- C. It creates an animation that slowly fades the element out of the page
- D. It makes an element and descendants non-interactive and hidden from AT

Answer: D

Q1217. How does the <dialog> element improve over custom modal implementations?

- A. It removes the need for any CSS styling on the modal overlay component
- B. It provides native focus trapping, backdrop, and accessibility features
- C. It enables modals to persist across page navigations without JavaScript
- D. It automatically translates modal content into multiple user languages

Answer: B

Q1218. What is the Intersection Observer API commonly used for in HTML pages?

- A. To detect when elements enter or exit the viewport for lazy loading use
- B. To watch for changes in the browser URL hash fragment for navigation
- C. To observe and report CSS syntax errors found within style sheet rules
- D. To monitor server response times and report performance metrics data

Answer: A

Q1219. What is the purpose of the popover attribute in modern HTML?

- A. It enables server-sent events to push notifications to the browser tab
- B. It creates a toggleable popup without JavaScript for tooltip-like behavior
- C. It converts any element into a fixed-position navigation sidebar panel
- D. It defines print-specific styles that appear only during page printing

Answer: B

Q1220. How does the fetchpriority attribute affect resource loading in HTML?

- A. It controls whether the resource is loaded synchronously or not at all
- B. It hints to the browser the relative priority for fetching a resource
- C. It sets the maximum timeout duration for network requests to complete
- D. It determines the order in which CSS rules are applied to the document

Answer: B

Q1221. How does the CSS contain property improve rendering performance?

- A. It prevents any external stylesheets from overriding the element styles
- B. It compresses CSS files automatically during the browser parsing process
- C. It caches the rendered output of an element permanently in the browser
- D. It limits an element's scope so the browser can optimize layout and paint

Answer: D

Q1222. What is the stacking context and how is it created in CSS?

- A. It is the order CSS rules are processed based on their file positions
- B. It is the order in which HTML files are loaded into the browser engine
- C. It is a method for stacking multiple background images on one element
- D. It is an element layer formed by z-index, opacity, or transform properties

Answer: D

Q1223. What is the difference between min-content and max-content sizing?

- A. max-content uses the largest word width while min-content avoids wrapping
- B. min-content forces all content onto one line without any line breaks used
- C. max-content collapses content to zero width and expands on user hover
- D. min-content uses the largest word width while max-content avoids wrapping

Answer: D

Q1224. How does the CSS subgrid feature extend CSS Grid capabilities?

- A. It creates nested grid systems that are completely independent of parent
- B. It automatically converts flexbox layouts into grid layouts for the page
- C. It enables grid items to span across multiple separate grid containers
- D. It lets child grid items inherit track sizing from their parent grid lines

Answer: D

Q1225. What problem does the CSS cascade layer feature (@layer) solve?

- A. It automatically generates responsive breakpoints based on content size
- B. It creates visual layering effects like parallax scrolling on web pages
- C. It provides explicit control over which origin of styles takes precedence
- D. It enables animations that transition between multiple keyframe states

Answer: C

Q1226. How do CSS custom properties differ from preprocessor variables?

- A. Custom properties require JavaScript to read while preprocessor does not
- B. Custom properties only store color values while preprocessor stores any type
- C. Custom properties are resolved at compile time while preprocessor runs live
- D. Custom properties cascade and update at runtime while preprocessor is static

Answer: D

Q1227. What is the purpose of the CSS will-change property and its potential pitfall?

- A. It changes element properties automatically based on user scroll position
- B. It schedules future DOM changes and blocks rendering until they complete
- C. It hints to the browser about upcoming changes but overuse wastes memory
- D. It defines fallback values for properties not supported in older browsers

Answer: C

Q1228. How does the :has() pseudo-class change CSS selector capabilities?

- A. It verifies if a CSS variable has been defined before using it in a rule
- B. It enables parent selection based on the presence of matching child elements
- C. It determines whether the browser supports a given CSS property or not
- D. It checks whether an element has a specific inline style attribute applied

Answer: B

Q1229. What is the CSSOM and how does it relate to the DOM?

- A. The CSSOM is a testing framework for validating CSS syntax and structure
- B. The CSSOM is the CSS object model representing styles paired with the DOM
- C. The CSSOM is a server-side cache for storing precomputed CSS style values
- D. The CSSOM replaces the DOM entirely for pages that use CSS stylesheets

Answer: B

Q1230. How does container query sizing differ from viewport-based media queries?

- A. Container queries are a JavaScript API and not available in pure CSS code
- B. Container queries use the viewport width while media queries use elements
- C. Container queries only work with grid layouts and not flexbox or floats
- D. Container queries respond to the parent container size not the viewport

Answer: D

Q1231. How do CSS container queries change the approach to component-based responsive design?

- A. They remove the need for any CSS Grid or flexbox layout properties entirely
- B. They let components adapt based on their container size rather than viewport
- C. They require JavaScript to detect container dimensions and apply styles
- D. They only work with elements that have a fixed pixel width explicitly set

Answer: B

Q1232. What is the purpose of the @supports rule and how does it enable progressive enhancement?

- A. It verifies that the server supports a specific HTTP protocol version used
- B. It conditionally applies styles based on browser support for CSS features
- C. It checks whether the user's browser supports JavaScript execution enabled
- D. It tests whether external font files have loaded successfully in the page

Answer: B

Q1233. How does the CSS scroll-snap property improve scroll-based user experiences?

- A. It automatically lazy-loads content as the user scrolls down the web page
- B. It creates defined snap points that the viewport locks to while scrolling
- C. It prevents all scrolling and forces users to use navigation buttons only
- D. It changes the scroll speed based on the amount of content on each page

Answer: B

Q1234. What is the difference between implicit and explicit grid tracks in CSS Grid?

- A. Implicit tracks only appear in development mode and vanish in production
- B. Explicit tracks are defined by template properties while implicit are auto
- C. Implicit tracks are defined by the developer while explicit are automatic
- D. Explicit tracks use percentage widths while implicit tracks use pixel only

Answer: B

Q1235. How does the isolation property interact with stacking contexts and blend modes?

- A. It isolates network requests made by elements inside the isolated container
- B. It prevents the element from inheriting any styles from parent containers
- C. It disables all CSS animations within the element to improve performance
- D. It creates a new stacking context limiting blend mode effects to the group

Answer: D

Q1236. What challenges arise when combining CSS Grid and flexbox in complex layouts?

- A. Using both requires doubling the number of media queries for responsiveness
- B. Nested contexts may cause unexpected sizing due to different algorithms used
- C. Combining them automatically disables all CSS transitions and animations
- D. Grid and flexbox cannot be used in the same document or stylesheet at all

Answer: B

Q1237. How does the CSS :is() pseudo-class improve selector maintainability?

- A. It checks if an element exists in the DOM before applying styles to it
- B. It creates conditional styles similar to if-else statements in JavaScript
- C. It validates that CSS property values are within acceptable valid ranges
- D. It groups selectors reducing repetition and uses highest specificity match

Answer: D

Q1238. What is the logical properties model in CSS and why is it important?

- A. It defines logical conditions that determine when styles are applied now
- B. It creates property dependencies where changing one updates others too
- C. It uses flow-relative properties like inline-start for internationalization
- D. It applies mathematical logic to calculate responsive font sizes for text

Answer: C

Q1239. How does the CSS accent-color property simplify form styling?

- A. It applies a color tint to form validation error messages shown to users
- B. It adds a colored border accent around the entire form container element
- C. It sets the accent color for form controls like checkboxes and radio buttons
- D. It changes the text color of all form labels to match the brand palette

Answer: C

Q1240. What is the purpose of the @property rule for CSS custom properties?

- A. It registers custom properties with type checking and initial value support
- B. It creates new CSS properties that override all existing browser defaults
- C. It generates property documentation automatically for developer tooling
- D. It defines JavaScript properties that can be accessed from CSS stylesheets

Answer: A

Q1241. What is a closure and why is it useful in JavaScript?

- A. A function that retains access to its outer scope variables after return
- B. A design pattern for preventing memory leaks in event listener callbacks
- C. A syntax for immediately terminating asynchronous operations mid-flight
- D. A method for closing browser windows programmatically using JavaScript

Answer: A

Q1242. How does the event loop handle microtasks versus macrotasks?

- A. Microtasks execute after current task before the next macrotask is taken
- B. Microtasks and macrotasks are processed in the exact same queue together
- C. Macrotasks always have higher priority and execute before any microtasks
- D. Microtasks are processed after all macrotasks in the queue are completed

Answer: A

Q1243. What is the Temporal Dead Zone for let and const declarations?

- A. The gap between an event being triggered and its handler being executed
- B. The period between scope entry and declaration where access throws error
- C. The time period between page load and when DOMContentLoaded event fires
- D. The delay between a network request being sent and response being received

Answer: B

Q1244. How does prototypal inheritance differ from classical inheritance?

- A. Prototypal objects inherit directly from other objects via prototype chains
- B. Prototypal is only available in strict mode while classical works anywhere
- C. Prototypal inheritance requires TypeScript while classical uses JavaScript
- D. Prototypal uses classes and constructors while classical uses object linking

Answer: A

Q1245. What is the purpose of WeakMap and how does it differ from Map?

- A. WeakMap automatically expires entries after a time while Map keeps forever
- B. WeakMap encrypts its values while Map stores all values as plain text data
- C. WeakMap holds weak references to keys allowing garbage collection of them
- D. WeakMap stores only string keys while Map can store any data type as key

Answer: C

Q1246. How does the JavaScript engine optimize code using hidden classes?

- A. It hides CSS class names from the DOM to improve rendering performance
- B. It compresses class definitions to reduce memory usage in the heap space
- C. It encrypts class methods to prevent reverse engineering of the codebase
- D. It creates internal type representations to speed up property access times

Answer: D

Q1247. What problem does the structured clone algorithm solve that JSON parse/stringify cannot?

- A. It validates data types during cloning and throws errors for mismatches
- B. It encrypts the cloned object data for secure storage in local sessions
- C. It handles circular references, typed arrays, and more complex data types
- D. It compresses the cloned data to use less memory than the JSON approach

Answer: C

Q1248. How does tail call optimization affect recursive functions in JavaScript?

- A. It limits recursion depth to a fixed number regardless of memory available
- B. It caches results of previous recursive calls to speed up the computation
- C. It automatically converts recursive functions into iterative loop versions
- D. It reuses the current stack frame for tail position calls preventing overflow

Answer: D

Q1249. What is the difference between a shallow copy and a deep copy of an object?

- A. Shallow copies only copy top-level properties while deep clones all nested
- B. Shallow copies are faster but create read-only versions of the object data
- C. Deep copies can only be created using third-party libraries not native code
- D. Shallow copies duplicate all nested levels while deep copies only copy top

Answer: A

Q1250. How does optional chaining (?.) handle deeply nested property access safely?

- A. It logs a warning to the console but continues execution with null value
- B. It automatically creates missing nested properties with default values set
- C. It returns undefined instead of throwing when accessing nullish references
- D. It throws a custom error message when a nested property is not found yet

Answer: C

Q1251. How does the JavaScript engine handle memory management and garbage collection?

- A. Developers must manually allocate and free memory using explicit commands
- B. The engine uses mark-and-sweep to automatically reclaim unreachable objects
- C. Memory is allocated once at startup and never freed during page lifetime
- D. The engine writes unused data to disk swap files for later reuse if needed

Answer: B

Q1252. What is the purpose of SharedArrayBuffer and Atomics in JavaScript?

- A. They share CSS stylesheets between multiple browser tabs simultaneously
- B. They automatically split large arrays into smaller chunks for processing
- C. They enable shared memory and atomic operations between web worker threads
- D. They synchronize database transactions across multiple server instances

Answer: C

Q1253. How does the Reflect API complement the Proxy API in JavaScript?

- A. Reflect provides default implementations for Proxy trap operations used
- B. Reflect converts Proxy objects back into normal JavaScript plain objects
- C. Reflect enables Proxy objects to intercept CSS style computations too
- D. Reflect optimizes Proxy performance by caching trap results in memory

Answer: A

Q1254. What is the difference between concurrency and parallelism in JavaScript?

- A. They are identical concepts with no practical difference in any language
- B. Concurrency manages multiple tasks interleaved while parallelism runs simultaneously
- C. Parallelism is single-threaded while concurrency requires multiple threads
- D. Concurrency is only possible in Node.js while parallelism works in browsers

Answer: B

Q1255. How do private class fields (#) differ from closure-based privacy in JavaScript?

- A. Private fields work only in strict mode while closures work in any context
- B. Private fields are only syntactic sugar with no actual encapsulation enforced
- C. Private fields are slower because they require additional proxy wrapper objects
- D. Private fields use hard privacy enforced by the engine unlike closure convention

Answer: D

Q1256. What is the purpose of the FinalizationRegistry in JavaScript?

- A. It manages the order in which CSS animations complete their final frames
- B. It ensures all pending network requests complete before page unload event
- C. It finalizes database transactions and ensures data integrity on commit
- D. It registers callbacks invoked when registered objects are garbage collected

Answer: D

Q1257. How does the import() dynamic import function differ from static import?

- A. Dynamic import returns a Promise and loads modules conditionally at runtime
- B. Dynamic import loads modules at compile time while static loads at runtime
- C. Dynamic import only works with JSON files while static works with all types
- D. Dynamic import is synchronous while static import is always asynchronous

Answer: A

Q1258. What is the Transferable interface and how does it optimize worker communication?

- A. It transfers ownership of data to a worker without copying the memory buffer
- B. It compresses messages between workers to reduce inter-thread bandwidth use
- C. It copies data between workers using a shared clipboard buffer mechanism
- D. It translates data formats between workers running different script versions

Answer: A

Q1259. How does the TC39 proposal process work for new JavaScript features?

- A. Proposals go through stages 0-4 with increasing maturity before inclusion
- B. The ECMA committee implements features first then writes specifications after
- C. Features are voted on by developers through public polls on social media
- D. Features are added by browser vendors independently without coordination

Answer: A

Q1260. What is the Iterator Helpers proposal and what problem does it solve?

- A. It helps iterate over CSS selectors to find matching DOM elements faster
- B. It provides helper functions for converting arrays into linked list types
- C. It adds chainable methods like map and filter directly to iterator objects
- D. It creates automatic pagination helpers for iterating over API result sets

Answer: C

Q1261. How does Svelte's compile-time approach differ from React's runtime approach?

- A. Svelte only works with TypeScript while React supports only plain JavaScript
- B. Svelte uses a larger runtime library while React compiles away its framework
- C. Svelte compiles components to imperative DOM code eliminating virtual DOM need
- D. Svelte requires manual DOM updates while React automates them with diffing

Answer: C

Q1262. What is the islands architecture and which frameworks implement it?

- A. A network architecture where each microservice runs on separate server nodes
- B. A pattern of interactive components in static HTML used by Astro and Fresh
- C. A database sharding pattern where each table is stored on a separate island
- D. A pattern where each page is an isolated island with no shared state at all

Answer: B

Q1263. How does React's Suspense mechanism work for data fetching and code splitting?

- A. It pauses server execution and buffers all response data until completion
- B. It suspends the browser event loop until all pending promises are resolved
- C. It shows fallback content while waiting for lazy components or async data
- D. It automatically caches all fetch responses to prevent redundant API calls

Answer: C

Q1264. What is the purpose of Angular's change detection strategy OnPush?

- A. It only checks the component when its inputs change or events fire within
- B. It enables automatic push notifications from the server to client updates
- C. It pushes all changes to a remote server for persistence before rendering
- D. It forces synchronous rendering of all child components before the parent

Answer: A

Q1265. How does Vue's Composition API improve code organization over the Options API?

- A. It eliminates the need for templates and uses only render functions for views
- B. It compiles Vue components into native mobile applications automatically
- C. It converts Options API components into Composition API at build time process
- D. It enables logic reuse through composable functions grouped by feature concern

Answer: D

Q1266. What is the fiber architecture in React and what problem does it solve?

- A. It provides fiber optic network optimizations for faster data fetching use
- B. It enables incremental rendering allowing React to pause and resume work
- C. It creates a fiber-based threading model for true parallel component rendering
- D. It converts React components into web components for cross-framework usage

Answer: B

Q1267. How does Solid.js achieve fine-grained reactivity without a virtual DOM?

- A. It delegates all DOM updates to web workers for off-main-thread rendering
- B. It requires developers to manually specify which DOM nodes need updating
- C. It uses a virtual DOM internally but markets itself as not having one
- D. It uses signals and reactive primitives that update only affected DOM nodes

Answer: D

Q1268. What are the tradeoffs of using micro-frontends with module federation?

- A. Module federation eliminates all build complexity and has no tradeoffs at all
- B. It enables independent deployment but adds complexity in shared dependency management
- C. It requires all micro-frontends to use the same exact version of every dependency
- D. Module federation only works with React and cannot be used with other frameworks

Answer: B

Q1269. How does streaming SSR improve time-to-first-byte compared to traditional SSR?

- A. It progressively sends HTML chunks as they render without waiting for all data
- B. It splits the HTML into multiple separate HTTP requests sent in parallel order
- C. It sends the complete HTML document only after all data fetching is finished
- D. It compresses the HTML response using streaming compression before any sending

Answer: A

Q1270. What is the concept of signals in modern frontend frameworks like Preact and Angular?

- A. Network signals indicating connection quality used for adaptive content loading
- B. Reactive primitives that track dependencies and update only subscribers directly
- C. Audio signals processed by the Web Audio API for multimedia web applications
- D. Database signals that notify the frontend when backend data changes on server

Answer: B

Q1271. How does database sharding improve scalability for high-traffic applications?

- A. It compresses database tables to fit more records into less storage space
- B. It creates a complete copy of the database on every single server available
- C. It caches all database queries in memory eliminating disk read operations
- D. It distributes data across multiple databases reducing load on each individual

Answer: D

Q1272. What is the saga pattern and when is it used in distributed systems?

- A. A database indexing strategy for optimizing full-text search across multiple tables
- B. A pattern for managing distributed transactions across microservices with compensations
- C. A UI design pattern for creating step-by-step wizard forms in web applications
- D. A testing pattern for writing narrative-style test cases in behavior-driven development

Answer: B

Q1273. How does event sourcing differ from traditional CRUD data storage?

- A. Event sourcing stores state changes as events while CRUD stores current state
- B. Event sourcing is synchronous while CRUD operations are always asynchronous
- C. Event sourcing deletes old data while CRUD keeps complete history of changes
- D. Event sourcing only works with NoSQL while CRUD requires relational databases

Answer: A

Q1274. What is the CQRS pattern and what problem does it solve?

- A. It compresses query results on the server before sending them to clients
- B. It combines all read and write operations into a single optimized data model
- C. It creates a queue for all server requests to process them sequentially only
- D. It separates read and write models for independent optimization and scaling

Answer: D

Q1275. How does a message queue improve backend architecture reliability?

- A. It queues CSS style calculations to prevent layout thrashing in the browser
- B. It stores database query results in a queue for faster repeated access times
- C. It decouples services by buffering messages for asynchronous processing safely
- D. It creates a queue of user interface updates for batch rendering on client

Answer: C

Q1276. What is the difference between horizontal and vertical scaling?

- A. Horizontal only works for databases while vertical only works for web servers
- B. Horizontal adds more servers while vertical adds more resources to one server
- C. Horizontal adds resources to one server while vertical adds more server nodes
- D. Horizontal scaling is free while vertical scaling always requires paid licenses

Answer: B

Q1277. How does the strangler fig pattern help migrate monoliths to microservices?

- A. It immediately replaces the entire monolith with microservices in one deployment
- B. It gradually replaces monolith functionality by routing to new services incrementally
- C. It copies the monolith code into separate repositories without any code changes
- D. It strangles the monolith by reducing its allocated server resources over time

Answer: B

Q1278. What is the purpose of a dead letter queue in message processing?

- A. To create a backup of all messages sent during system maintenance downtime
- B. To queue messages from deactivated user accounts for archival purposes only
- C. To permanently delete messages that have been successfully processed already
- D. To store messages that fail processing for later inspection and retry handling

Answer: D

Q1279. How does the outbox pattern ensure reliable event publishing in microservices?

- A. It stores events in a database table atomically with data then publishes them
- B. It creates an outbox email notification for every event published by service
- C. It sends events directly to the message broker before committing the database
- D. It buffers events in browser local storage until network connectivity returns

Answer: A

Q1280. What is the bulkhead pattern in distributed systems?

- A. A pattern that isolates failures by partitioning resources into separate pools
- B. A pattern for compressing bulk API responses to reduce network bandwidth use
- C. A pattern for creating bulk database insert operations for high throughput
- D. A pattern for generating bulk SSL certificates for multiple domain endpoints

Answer: A

Q1281. How does the N+1 query problem occur and how can it be resolved?

- A. It occurs when a query returns N+1 duplicate records from the database table
- B. It occurs when loading related data triggers a query per item instead of batching
- C. It occurs when N+1 users simultaneously access the same server resource endpoint
- D. It occurs when the database connection pool has exactly N+1 connections open

Answer: B

Q1282. What is the purpose of a write-ahead log in database systems?

- A. To record changes before applying them ensuring durability and crash recovery
- B. To write CSS styles ahead of rendering to prevent layout shift on the page
- C. To prewrite API documentation before implementing the actual endpoint logic
- D. To log all user write operations for billing and usage tracking purposes only

Answer: A

Q1283. How does connection multiplexing work in HTTP/2 server implementations?

- A. It multiplies the server bandwidth by splitting traffic across many networks
- B. It duplicates every request to multiple servers for redundancy and failover
- C. It creates multiple separate TCP connections for each HTTP request sent out
- D. It sends multiple requests and responses over a single TCP connection stream

Answer: D

Q1284. What is the leader election pattern in distributed server systems?

- A. A pattern for selecting which CSS framework the development team should adopt
- B. A pattern where users vote for their preferred server through a web interface
- C. A pattern where the fastest server automatically becomes the primary database
- D. A pattern where one node is chosen to coordinate tasks among distributed nodes

Answer: D

Q1285. How does server-side request forgery (SSRF) attack work and how is it prevented?

- A. SSRF sends forged DNS responses to redirect traffic to malicious web servers
- B. SSRF exploits SQL injection to forge database records in production tables
- C. SSRF tricks a server into making requests to unintended internal resources
- D. SSRF forges client-side cookies to impersonate legitimate authenticated users

Answer: C

Q1286. What is the purpose of graceful shutdown in server applications?

- A. To finish in-progress requests and release resources before stopping the server
- B. To immediately terminate all active connections without completing any requests
- C. To reduce server CPU speed gradually to lower power consumption over time
- D. To send shutdown notifications to all connected users via email before stop

Answer: A

Q1287. How does the sidecar pattern work in server-side architecture?

- A. It deploys a helper process alongside the main service for cross-cutting tasks
- B. It creates a secondary copy of the application for A/B testing user features
- C. It adds a sidebar navigation component to every page of the web application
- D. It places the database alongside the web server on the same physical machine

Answer: A

Q1288. What is back-pressure in server systems and how should it be handled?

- A. It is the physical pressure on server hardware from high temperature levels
- B. It is a database optimization that compresses data under high storage pressure
- C. It is a mechanism to slow producers when consumers cannot keep up with load
- D. It is the CSS pressure applied to buttons to create depth visual effects

Answer: C

Q1289. How does the actor model approach concurrency differently than shared memory?

- A. Actors require operating system level support and cannot run on standard hardware
- B. Actors communicate via message passing with isolated state avoiding shared memory
- C. Actors use a single thread and cannot process multiple messages simultaneously
- D. Actors share a global memory space and use locks for synchronization of access

Answer: B

Q1290. What is the purpose of distributed tracing in microservice architectures?

- A. To track requests across multiple services for debugging and performance analysis
- B. To monitor network cable connections between physical server rack units in DC
- C. To trace CSS rendering performance across multiple browser rendering engines
- D. To trace database query execution plans across multiple table join operations

Answer: A

Q1291. How does database replication improve availability and what are its consistency challenges?

- A. Replication creates copies across servers but may cause stale reads temporarily
- B. Replication guarantees immediate consistency across all replicas without any delay
- C. Replication eliminates the need for backups since data exists on multiple servers
- D. Replication only works with NoSQL databases and not with relational SQL databases

Answer: A

Q1292. What is the write-ahead log and how does it ensure data durability?

- A. It logs all SELECT queries for audit purposes before returning results to clients
- B. It writes database documentation ahead of implementation for planning purposes
- C. It records changes to a log before applying them enabling crash recovery of data
- D. It pre-generates auto-increment IDs ahead of insert operations for performance

Answer: C

Q1293. How does database connection pool sizing affect application performance?

- A. Larger pools always improve performance with no diminishing returns whatsoever
- B. Connection pools should always match the exact number of application server threads
- C. Too few connections cause queuing while too many waste resources and add contention
- D. Pool size has no impact on performance and is purely a configuration preference

Answer: C

Q1294. What is the difference between optimistic and pessimistic locking?

- A. Optimistic is for read operations only while pessimistic is for writes exclusively
- B. Optimistic locking is faster but pessimistic locking uses less memory on the server
- C. Optimistic assumes no conflicts and checks at commit while pessimistic locks upfront
- D. Optimistic locks the resource immediately while pessimistic waits until commit time

Answer: C

Q1295. How do graph databases differ from relational databases for relationship-heavy data?

- A. Graph databases require data to be normalized into third normal form before insertion
- B. Graph databases cannot store any properties on nodes and only support edge connections
- C. Graph databases store data in tables but use graphical user interfaces for queries
- D. Graph databases model relationships as first-class citizens enabling efficient traversal

Answer: D

Q1296. What is the purpose of a materialized view and when should it be used?

- A. It stores precomputed query results physically for fast access on complex queries
- B. It visualizes database schema as a material design diagram for documentation use
- C. It generates a physical backup file of the database view definitions and metadata
- D. It creates a virtual view that recomputes on every access for fresh data always

Answer: A

Q1297. How does database partitioning differ from sharding?

- A. Partitioning distributes data across servers while sharding splits within one database
- B. Partitioning and sharding are identical concepts with no practical difference at all
- C. Partitioning is for NoSQL only while sharding works exclusively with SQL databases
- D. Partitioning splits data within one database while sharding distributes across servers

Answer: D

Q1298. What is the CAP theorem's impact on database selection for distributed systems?

- A. CAP only applies to NoSQL databases and has no relevance to SQL database systems
- B. CAP requires all three properties to be met for any distributed database to function
- C. CAP states you can only guarantee two of consistency, availability, partition tolerance
- D. CAP states distributed systems must choose between consistency and partition tolerance

Answer: C

Q1299. How does a time-series database optimize for temporal data compared to relational databases?

- A. It stores all timestamps as plain text strings for human-readable log file output
- B. It requires data to be inserted in reverse chronological order for index efficiency
- C. It uses time-based partitioning and compression optimized for append-heavy workloads
- D. It only supports queries for the most recent hour of data and archives everything else

Answer: C

Q1300. What is the change data capture pattern and how does it enable event-driven architectures?

- A. It captures network packet changes for monitoring database connection health status
- B. It captures CSS property changes and generates transition animations automatically
- C. It captures UI changes in the browser DOM and sends them to the server for storage
- D. It captures database changes as events enabling downstream systems to react to them

Answer: D

Q1301. How does a timing attack exploit web applications and how is it mitigated?

- A. It measures response time differences to infer secret data using constant-time comparison
- B. It exploits expired TLS certificates during the brief window before renewal occurs
- C. It sends requests at precisely timed intervals to bypass rate limiting protections
- D. It attacks during specific time zones when server administrators are offline

Answer: A

Q1302. What is the difference between symmetric and asymmetric encryption in web security?

- A. Symmetric uses two keys while asymmetric uses a single shared secret key only
- B. Symmetric uses one shared key while asymmetric uses a public-private key pair
- C. Symmetric is more secure than asymmetric in all possible use case scenarios today
- D. Symmetric is for encrypting files while asymmetric is for encrypting databases

Answer: B

Q1303. How does a JWT token compromise affect application security and what mitigations exist?

- A. JWT compromise only affects the single request where the token was intercepted
- B. JWT tokens cannot be compromised because they are encrypted with AES-256 always
- C. JWT compromise requires the attacker to also know the user password to exploit
- D. JWT compromise grants access until expiration mitigated by short TTL and rotation

Answer: D

Q1304. What is the purpose of a Web Application Firewall and how does it differ from a network firewall?

- A. WAF operates at the network layer while network firewall operates at application layer
- B. WAF only protects static websites while network firewall protects dynamic applications
- C. WAF replaces the need for HTTPS encryption while network firewall enhances encryption
- D. WAF protects against application layer attacks while network firewall handles lower layers

Answer: D

Q1305. How does certificate pinning enhance TLS security and what are its drawbacks?

- A. It pins SSL certificates to specific IP addresses for geographic restriction purposes
- B. It pins certificates in DNS records requiring DNSSEC for validation of every connection
- C. It binds a host to specific certificates preventing man-in-middle with CA compromise risk
- D. It permanently pins a certificate preventing any future certificate rotation or updates

Answer: C

Q1306. What is a supply chain attack in web development and how can it be mitigated?

- A. A vulnerability in the CDN supply chain that serves corrupted static asset files
- B. A compromised dependency that injects malicious code mitigated by lockfiles and audits
- C. An attack on the physical supply chain of server hardware during manufacturing
- D. An attack where the attacker buys the domain name of a deprecated npm package

Answer: B

Q1307. How does Subresource Integrity protect against CDN compromises?

- A. It checks that CDN servers have valid SSL certificates before loading any resources
- B. It encrypts all resources loaded from CDNs to prevent eavesdropping during transfer
- C. It creates a backup copy of CDN resources on the origin server for failover support
- D. It verifies resource integrity using cryptographic hashes ensuring files are untampered

Answer: D

Q1308. What is the OAuth 2.0 authorization code flow and why is PKCE important?

- A. OAuth uses basic auth and PKCE adds an extra password layer for security measure
- B. Auth code flow exchanges code for tokens and PKCE prevents authorization code interception
- C. OAuth only works with social login providers and PKCE extends it to email passwords
- D. Auth code flow is deprecated and PKCE is the entirely new protocol that replaced it

Answer: B

Q1309. How does a zero-day vulnerability differ from a known vulnerability in web security?

- A. Zero-day is a vulnerability discovered on the first day the software was released
- B. Zero-day is an unknown exploit with no patch while known vulnerabilities have fixes
- C. Zero-day means the vulnerability has zero impact while known ones have high impact
- D. Zero-day vulnerabilities can only affect server software and not client applications

Answer: B

Q1310. What is the purpose of security response headers like Permissions-Policy?

- A. To set file system permissions for static assets served by the web server only
- B. To control which browser features and APIs a document is allowed to use safely
- C. To define which team members have permission to deploy code to production servers
- D. To manage database user permissions through HTTP headers sent with each request

Answer: B

Q1311. How does the PRPL pattern optimize web application loading performance?

- A. It Parses HTML, Resolves dependencies, Processes CSS, and Loads JavaScript
- B. It Prioritizes above-fold, Reduces payload, Prefetches links, Lowers latency
- C. It Pre-renders, Re-renders, Pre-loads, and Lazy-loads all page components
- D. It Pushes critical resources, Renders route, Pre-caches, and Lazy-loads rest

Answer: D

Q1312. What is the impact of long tasks on web performance and how are they identified?

- A. Long tasks are HTTP requests taking over five seconds to receive a server response
- B. Long tasks are CSS animations exceeding one second that cause janky scrolling
- C. Tasks that write large amounts of data to localStorage causing storage quota errors
- D. Tasks exceeding 50ms block the main thread and are identified via Long Tasks API

Answer: D

Q1313. How does resource prioritization work in the browser and how can developers influence it?

- A. Browsers assign priority based on resource type and developers use fetchpriority hints
- B. Browsers load all resources with equal priority and developers cannot change this
- C. Resource priority is determined solely by the order resources appear in the HTML file
- D. Developers must manually specify download order for every single resource on page

Answer: A

Q1314. What is the purpose of the Performance Observer API?

- A. To observe network performance and automatically switch to lower quality resources
- B. To asynchronously observe performance entries like marks, measures, and paint timing
- C. To observe and report CSS rendering performance for animation smoothness checks
- D. To monitor server CPU and memory performance from the client browser JavaScript

Answer: B

Q1315. How does Brotli compression compare to Gzip for web content delivery?

- A. Brotli only works with HTML files while Gzip supports all content types served
- B. Brotli and Gzip produce identical compression ratios for all web content types
- C. Brotli achieves higher compression ratios than Gzip especially for static content
- D. Brotli requires client-side decompression libraries while Gzip is built into browsers

Answer: C

Q1316. What is layout thrashing and how can it be prevented?

- A. Layout thrashing occurs when too many CSS classes are applied to a single element
- B. It occurs when multiple CSS files conflict with each other causing rendering delays
- C. It occurs when repeated reads and writes to the DOM force synchronous layout recalculations
- D. Layout thrashing happens when the server sends too many layout instructions in response

Answer: C

Q1317. How does the speculation rules API improve navigation performance?

- A. It allows prerendering of pages the user is likely to navigate to using rules set
- B. It guesses user input and pre-fills form fields to reduce time to completion
- C. It predicts server load and automatically scales resources before traffic spikes
- D. It speculates which CSS styles will be needed and pre-compiles them for speed

Answer: A

Q1318. What is the difference between First Contentful Paint and Largest Contentful Paint?

- A. FCP is a server metric while LCP is a client metric and they measure different things
- B. FCP measures when first content renders while LCP measures largest content element render
- C. FCP measures first byte received while LCP measures last byte received from server
- D. FCP only measures text rendering while LCP only measures image rendering on the page

Answer: B

Q1319. How does edge computing improve web application performance?

- A. It runs computation at network edge locations closer to users reducing latency seen
- B. It prioritizes content at the top edge of the page for above-the-fold rendering
- C. It sharpens image edges using advanced algorithms for better visual quality display
- D. It processes data at the edge of the database table for faster query result times

Answer: A

Q1320. What is the Interaction to Next Paint metric and why did it replace First Input Delay?

- A. INP measures overall responsiveness across all interactions not just the first input event
- B. INP measures network latency to the next server while FID measured to the current server
- C. INP only measures touch interactions while FID measured all types of user input events
- D. INP measures the time between two consecutive paint operations on the canvas element

Answer: A

Q1321. How does visual regression testing work and what tools support it?

- A. It manually reviews visual design changes in a design review meeting with stakeholders
- B. It tests CSS regression properties by comparing computed styles programmatically
- C. It runs CSS animations in reverse to test visual consistency during undo operations
- D. It captures and compares screenshots pixel-by-pixel using tools like Percy or Chromatic

Answer: D

Q1322. What is property-based testing and how does it differ from example-based testing?

- A. Property-based only tests object properties while example-based tests entire components
- B. Property-based uses specific examples while example-based uses random test input data
- C. Property-based requires no assertions while example-based requires explicit assertions
- D. Property-based generates random inputs to verify invariants rather than specific examples

Answer: D

Q1323. How does contract testing ensure compatibility between microservices?

- A. It verifies that service interactions conform to agreed-upon API contracts defined
- B. It verifies that all microservices use the same programming language and framework
- C. It ensures all services are deployed in the same data center for low latency access
- D. It tests the legal contracts associated with third-party service license agreements

Answer: A

Q1324. What is mutation testing and how does it evaluate test suite quality?

- A. It introduces small code changes to verify that tests catch defects effectively now
- B. It tests how the application handles mutated or corrupted database records safely
- C. It mutates the test code itself to verify it can be easily maintained over time
- D. It tests the application with different user input mutations for edge case coverage

Answer: A

Q1325. How does chaos engineering apply to web application testing?

- A. It intentionally introduces failures to test system resilience and recovery ability
- B. It generates random database schemas to test ORM compatibility across configurations
- C. It tests the application with randomly generated CSS styles for visual robustness
- D. It creates chaotic user interface layouts to test responsive design breakpoints

Answer: A

Q1326. What is the difference between synthetic and real user monitoring?

- A. Synthetic measures server metrics while RUM measures only client-side rendering speed
- B. Synthetic uses scripted tests in controlled environments while RUM tracks actual users
- C. Synthetic only works in development while RUM only works in staging environments used
- D. Synthetic monitors real users while RUM uses automated scripts for testing performance

Answer: B

Q1327. How does test parallelization improve CI/CD pipeline performance?

- A. It parallelizes code compilation and testing in the same process for efficiency gains
- B. It runs all tests sequentially on multiple servers for consistent result comparison
- C. It creates parallel copies of the database for each test to run in complete isolation
- D. It runs independent tests simultaneously across workers reducing total execution time

Answer: D

Q1328. What is the role of test fixtures and factories in maintaining test data?

- A. Fixtures are physical hardware devices used for hardware-in-the-loop testing only
- B. Fixtures store CSS styling for test result report pages displayed in the browser
- C. Fixtures and factories create consistent reusable test data reducing setup duplication
- D. Factories generate random production data for populating staging environments daily

Answer: C

Q1329. How does fuzz testing help discover security vulnerabilities?

- A. It applies fuzzy logic to determine which test cases should be prioritized first
- B. It feeds random or malformed input to find crashes and unexpected behavior patterns
- C. It tests the application with fuzzy or blurred images to check image processing
- D. It tests fuzzy string matching algorithms for accuracy in search functionality

Answer: B

Q1330. What is the testing trophy pattern and how does it differ from the testing pyramid?

- A. The trophy inverts the pyramid putting the most end-to-end tests at the base level
- B. The trophy eliminates unit tests entirely and focuses only on end-to-end test coverage
- C. The trophy has more integration tests emphasizing confidence per test written effort
- D. The trophy adds a new layer of performance tests that the pyramid does not include

Answer: C

Q1331. How does GitOps differ from traditional CI/CD approaches?

- A. GitOps eliminates the need for any version control system in the deployment process
- B. GitOps only works with GitHub and cannot be used with GitLab or Bitbucket platforms
- C. GitOps requires manual approval for every deployment while CI/CD is fully automated
- D. GitOps uses Git as the single source of truth for declarative infrastructure and apps

Answer: D

Q1332. What is the purpose of a service mesh in container orchestration?

- A. It manages service-to-service communication with observability, security, and routing
- B. It generates a mesh network of CDN nodes for distributed static asset delivery use
- C. It creates a visual mesh overlay of all CSS grid layouts on the deployed web page
- D. It meshes multiple databases together into a single unified query interface for apps

Answer: A

Q1333. How does immutable infrastructure improve deployment reliability?

- A. It replaces servers entirely instead of updating them ensuring consistent clean state
- B. It freezes all application features and prevents any new deployments from occurring
- C. It makes database records immutable so they cannot be updated or deleted by anyone
- D. It prevents any developer from modifying the infrastructure code in the repository

Answer: A

Q1334. What is the purpose of a deployment rollback strategy?

- A. To create a backup of the deployment pipeline configuration for disaster recovery
- B. To gradually roll forward by deploying incrementally to larger user group segments
- C. To roll back CSS changes to a previous design version when stakeholders disagree
- D. To quickly revert to a previous working version when a deployment causes issues

Answer: D

Q1335. How does Kubernetes handle self-healing of deployed applications?

- A. It automatically rewrites application code to fix bugs detected in production use
- B. It heals network connections between services by automatically rerouting all traffic
- C. It monitors pod health and restarts or replaces failed pods to maintain desired state
- D. It repairs corrupted database tables by restoring from the most recent backup file

Answer: C

Q1336. What is the difference between horizontal and vertical pod autoscaling in Kubernetes?

- A. Horizontal only scales CPU while vertical only scales memory allocation for each pod
- B. Horizontal adjusts resources while vertical adds replicas for load distribution work
- C. Horizontal adds more pod replicas while vertical adjusts pod resource limits upward
- D. Horizontal requires manual triggers while vertical scales automatically based on load

Answer: C

Q1337. How does a progressive delivery strategy improve deployment safety?

- A. It progressively increases the server CPU speed during deployment for faster builds
- B. It delivers CSS and JavaScript progressively using chunked transfer encoding only
- C. It deploys all changes at once but progressively enables monitoring after deployment
- D. It gradually exposes changes to users with feature flags, canaries, and observability

Answer: D

Q1338. What is the purpose of chaos engineering in production deployments?

- A. To generate chaotic test data that stress-tests database storage capacity under load
- B. To randomly assign deployment tasks to team members for knowledge sharing purposes
- C. To intentionally inject failures to verify system resilience and recovery mechanisms
- D. To create chaotic deployment schedules that test team response time capabilities

Answer: C

Q1339. How does multi-region deployment improve application availability?

- A. It deploys across geographic regions providing failover and reduced latency globally
- B. It runs multiple instances of the database on the same server for read performance
- C. It deploys the application to multiple folders within a single server for redundancy
- D. It creates multiple copies of the source code repository in different version controls

Answer: A

Q1340. What is the twelve-factor app methodology and how does it influence deployment?

- A. It mandates twelve team members minimum for any cloud deployment project management
- B. It requires twelve separate deployment stages before code can reach production servers
- C. It limits applications to twelve files maximum for simplified deployment management
- D. It defines twelve principles for building scalable, maintainable cloud-native applications

Answer: D

Q1341. How does GraphQL handle the over-fetching and under-fetching problems of REST?

- A. GraphQL automatically paginates responses to prevent over-fetching of large datasets
- B. GraphQL requires multiple round trips but each returns exactly one field per request
- C. GraphQL lets clients specify exactly which fields they need in a single query request
- D. GraphQL caches all possible response combinations to eliminate unnecessary fetching

Answer: C

Q1342. What are the tradeoffs between REST, GraphQL, and gRPC for API design?

- A. gRPC only works between frontend and backend while REST is only for backend services
- B. REST is always the best choice and GraphQL and gRPC offer no advantages over it
- C. REST is simple and cacheable, GraphQL is flexible for clients, gRPC is fast for services
- D. GraphQL is the fastest protocol while REST and gRPC are slower but more user-friendly

Answer: C

Q1343. How does the BFF (Backend for Frontend) pattern optimize API consumption?

- A. It creates dedicated backend services tailored to specific frontend client requirements
- B. It creates a beautiful frontend framework that replaces the need for API calls entirely
- C. It bundles all frontend files and serves them from the backend for faster initial loads
- D. It buffers frontend failures and retries API calls automatically without user awareness

Answer: A

Q1344. What is the purpose of API circuit breakers and how do they improve reliability?

- A. They stop calling failing APIs after threshold breaches preventing cascade failures
- B. They create circular API routes for redundant request handling and load balancing
- C. They break API responses into smaller circuits for parallel processing on the client
- D. They physically disconnect API servers from the network during maintenance windows

Answer: A

Q1345. How does event-driven API design differ from request-response patterns?

- A. Event-driven APIs push events to subscribers asynchronously instead of synchronous calls
- B. Event-driven APIs only support read operations and cannot handle data write operations
- C. Event-driven APIs require clients to poll the server continuously for new events data
- D. Event-driven APIs are limited to WebSocket transport and cannot use HTTP protocol at all

Answer: A

Q1346. What is the purpose of API schema validation and how does it prevent issues?

- A. It validates CSS schemas applied to API response rendering in the browser viewport
- B. It validates request and response structures against defined schemas catching errors early
- C. It validates server configuration files to ensure API routing is correctly configured
- D. It validates database schemas before running API queries to prevent SQL syntax errors

Answer: B

Q1347. How does API versioning through content negotiation work?

- A. Clients specify the desired API version in the URL path like /api/v1/resource endpoint
- B. Clients specify the desired version through Accept headers rather than URL path changes
- C. Servers randomly assign API versions to clients based on their geographic location data
- D. Clients download different API versions and choose which one to install on their device

Answer: B

Q1348. What is the purpose of the saga pattern in managing distributed API transactions?

- A. It tracks the complete saga or history of all API calls for audit compliance purposes
- B. It creates epic narrative documentation for complex API integration workflows only
- C. It generates sequential API call stories for automated testing of integration scenarios
- D. It coordinates long-running distributed transactions with compensating actions for failures

Answer: D

Q1349. How does server-sent events differ from WebSockets for real-time API communication?

- A. SSE requires a dedicated server while WebSockets work on any standard HTTP web server
- B. SSE only supports binary data while WebSockets only support plain text message formats
- C. SSE is bidirectional over HTTP while WebSockets only support server-to-client messages
- D. SSE provides one-way server-to-client streaming over HTTP while WebSockets are bidirectional

Answer: D

Q1350. What is the purpose of API observability and what does it encompass?

- A. It observes API source code changes and generates release notes documentation only
- B. It encompasses logging, metrics, and tracing for understanding API behavior and health
- C. It observes competitor API features and suggests improvements for your own API design
- D. It monitors the visual appearance of API documentation pages for design consistency

Answer: B

Q1351. How does partial hydration improve performance compared to full hydration?

- A. Partial hydration only hydrates interactive components reducing JavaScript sent to client
- B. Partial hydration pre-renders CSS while full hydration renders it at runtime on client
- C. Partial hydration downloads less HTML while full hydration downloads the complete page
- D. Partial hydration splits the database queries while full hydration runs them all at once

Answer: A

Q1352. What is the purpose of React Server Components in the modern web ecosystem?

- A. They are server-side templates that generate static HTML without any React features used
- B. They create server-to-server communication channels for microservice architectures only
- C. They render on the server reducing bundle size while keeping client interactivity where needed
- D. They replace all client components entirely and run everything on the server only

Answer: C

Q1353. How does the module federation pattern enable micro-frontend architectures?

- A. It creates federated databases where each micro-frontend has its own data storage silo
- B. It federates CSS modules across multiple applications for shared styling consistency
- C. It allows applications to share and load modules from other builds at runtime dynamically
- D. It merges all micro-frontend codebases into a single monolithic bundle for deployment

Answer: C

Q1354. What is the purpose of observability in modern web applications?

- A. To observe CSS rendering performance and generate visual reports for stakeholders
- B. To observe user behavior through screen recording for UX research purposes only
- C. To monitor competitor websites and observe their feature changes for benchmarking
- D. To provide insight into system behavior through logs, metrics, and distributed tracing

Answer: D

Q1355. How does the streaming architecture pattern benefit modern web applications?

- A. It streams video content to users using specialized video encoding protocols only
- B. It processes and delivers data incrementally reducing time to first meaningful content
- C. It creates a stream of database backups continuously for disaster recovery purposes
- D. It streams CSS styles progressively so the page styling appears gradually over time

Answer: B

Q1356. What is the difference between build-time and runtime in modern web frameworks?

- A. Build-time only applies to CSS while runtime only applies to JavaScript code execution
- B. Build-time requires Node.js while runtime requires a completely different language runtime
- C. Build-time is slower but more secure while runtime is faster but has security vulnerabilities
- D. Build-time generates optimized output ahead of deployment while runtime executes on request

Answer: D

Q1357. How does the AI-assisted development trend impact modern web practices?

- A. AI development tools are limited to backend development and cannot assist with frontend
- B. AI completely replaces the need for human developers in all web development tasks today
- C. AI tools assist with code generation, review, and debugging while humans provide oversight
- D. AI only works for generating CSS styles and cannot help with JavaScript or HTML coding

Answer: C

Q1358. What is the concept of edge-first architecture in modern web development?

- A. Designing buildings with edge-shaped architecture for modern tech company headquarters
- B. Developing applications that only work on Microsoft Edge browser for optimal performance
- C. Designing applications to run primarily at network edge for minimal latency and global scale
- D. Prioritizing the visual edges and borders of UI elements for a modern design aesthetic

Answer: C

Q1359. How does the concept of composability influence modern web architecture?

- A. It composes CSS animations by chaining multiple keyframe definitions in sequence only
- B. It composes music and sound effects for web applications using the Web Audio API
- C. It creates composite database indexes by combining multiple columns for query speed
- D. It builds systems from interchangeable independent modules that can be freely combined

Answer: D

Q1360. What are the tradeoffs of adopting a serverless architecture for web applications?

- A. Serverless only works for static websites and cannot handle any dynamic content at all
- B. Serverless eliminates all costs since you never pay for idle server time or any resources
- C. Serverless has no tradeoffs and is always the best choice for every web application type
- D. Serverless reduces ops overhead but introduces cold starts and vendor lock-in challenges

Answer: D

Q1361. What is the main trade-off when choosing between a monolithic architecture and a microservices architecture for a web application?

- A. Monoliths are always slower than microservices
- B. Microservices add operational complexity but allow independent scaling and deployment of services
- C. Monoliths cannot use databases
- D. Microservices require a single programming language

Answer: B

Q1362. How does WebAssembly complement JavaScript in modern web development?

- A. It completely replaces JavaScript in all browsers
- B. It allows near-native performance for compute-intensive tasks while JavaScript handles DOM interaction
- C. It is only used for styling web pages
- D. It is a server-side-only technology

Answer: B

Q1363. What problem does the micro-frontend architecture solve in large-scale web development?

- A. It reduces CSS file size
- B. It allows independent teams to develop, deploy, and maintain separate parts of a frontend application
- C. It eliminates the need for a backend
- D. It only works with Angular applications

Answer: B

Q1364. What is the significance of the critical rendering path in web performance?

- A. It determines how the server compiles JavaScript
- B. It is the sequence of steps the browser takes to convert HTML, CSS, and JS into rendered pixels on screen
- C. It only affects mobile devices
- D. It refers to the deployment pipeline

Answer: B

Q1365. What is the purpose of edge computing in the context of modern web applications?

- A. To increase the physical size of data centers
- B. To run computations closer to the end user, reducing latency and improving responsiveness
- C. To replace all client-side JavaScript
- D. To eliminate the need for caching

Answer: B

Q1366. How does the concept of progressive enhancement differ from graceful degradation?

- A. They are the same strategy
- B. Progressive enhancement builds from a baseline experience upward, while graceful degradation starts with full features and falls back
- C. Progressive enhancement only applies to CSS
- D. Graceful degradation does not work in modern browsers

Answer: B

Q1367. What is the role of a build tool like Webpack or Vite in modern web development?

- A. It hosts websites on the internet
- B. It bundles, transforms, and optimizes source files for production deployment
- C. It is used exclusively for database management
- D. It replaces the need for version control

Answer: B

Q1368. Why is accessibility (a11y) considered a fundamental aspect of professional web development?

- A. It only improves search engine rankings
- B. It ensures web content is usable by people with disabilities and often improves usability for all users
- C. It is only required by European law
- D. It slows down website performance

Answer: B

Q1369. What is the difference between Server-Sent Events and WebSockets for real-time communication?

- A. They are identical protocols
- B. SSE is unidirectional from server to client over HTTP, while WebSockets provide full-duplex bidirectional communication
- C. SSE is faster than WebSockets in all cases
- D. WebSockets cannot send text data

Answer: B

Q1370. What is the impact of third-party scripts on web application performance and security?

- A. Third-party scripts have no impact on performance
- B. They can increase page load time, introduce security vulnerabilities, and affect Core Web Vitals
- C. They always improve website speed
- D. They are automatically sandboxed by browsers

Answer: B

Q1371. What is the split-brain problem in distributed web systems?

- A. A CSS rendering issue
- B. When network partitions cause parts of a distributed system to operate independently with conflicting states
- C. A JavaScript memory leak pattern
- D. A browser compatibility issue

Answer: B

Q1372. How does the QUIC protocol improve web performance compared to TCP?

- A. QUIC is slower but more secure than TCP
- B. QUIC eliminates head-of-line blocking and reduces connection establishment latency by combining transport and encryption handshakes
- C. QUIC only works on local networks
- D. QUIC replaces HTTP entirely

Answer: B

Q1373. What is the purpose of the saga pattern in distributed architectures?

- A. To optimize CSS rendering
- B. To manage distributed transactions by breaking them into a sequence of local transactions with compensating actions
- C. To compress network traffic
- D. To handle DNS resolution

Answer: B

Q1374. What is the difference between strong consistency and eventual consistency in distributed systems?

- A. They are the same model
- B. Strong consistency guarantees all nodes see the latest data immediately while eventual consistency allows temporary divergence
- C. Strong consistency is always faster
- D. Eventual consistency prevents data replication

Answer: B

Q1375. What is the purpose of connection pooling in web architecture?

- A. To create new database connections for every request
- B. To maintain a pool of reusable connections to reduce the overhead of establishing new connections
- C. To encrypt all network traffic
- D. To manage CSS class names

Answer: B

Q1376. How does the event-driven architecture pattern benefit web applications?

- A. It makes all operations synchronous
- B. It decouples services through asynchronous event publishing and consumption, improving scalability and resilience
- C. It eliminates the need for databases
- D. It only works with a single server

Answer: B

Q1377. What is the role of a service mesh in modern web architecture?

- A. To create HTML templates
- B. To manage service-to-service communication, providing observability, security, and traffic control
- C. To replace the frontend framework
- D. To store static files

Answer: B

Q1378. What problem does the Command Query Responsibility Segregation (CQRS) pattern address?

- A. CSS specificity conflicts
- B. It separates read and write models to optimize each independently for different performance and scaling requirements
- C. It manages DNS records
- D. It compresses HTTP responses

Answer: B

Q1379. What is the purpose of request idempotency in web architecture?

- A. To make all requests faster
- B. To ensure that making the same request multiple times produces the same result without unintended side effects
- C. To encrypt request payloads
- D. To compress response bodies

Answer: B

Q1380. How does the sidecar proxy pattern work in distributed web architectures?

- A. It replaces the main application entirely
- B. It deploys a helper process alongside each service instance to handle cross-cutting concerns like networking and observability
- C. It is a CSS layout technique
- D. It manages browser cookies

Answer: B

Q1381. How does the <portal> element differ from an <iframe> for embedding content?

- A. They are identical elements
- B. Portals allow seamless navigation to the embedded content and can be activated to become the top-level page
- C. Portals cannot load external URLs
- D. Iframes are deprecated in HTML5

Answer: B

Q1382. What is the purpose of the Declarative Shadow DOM in HTML?

- A. To add CSS animations
- B. To allow server-side rendering of shadow DOM content using HTML-only syntax without JavaScript
- C. To manage form validation
- D. To create responsive images

Answer: B

Q1383. How does the blocking attribute on script and style elements affect rendering?

- A. It has no effect on rendering
- B. It explicitly indicates which rendering operations should be blocked until the resource is processed
- C. It deletes the element from the DOM
- D. It compresses the resource

Answer: B

Q1384. What is the purpose of the rel='modulepreload' link type?

- A. To load CSS modules
- B. To preload JavaScript modules and their dependencies for faster execution when needed
- C. To create module-scoped styles
- D. To define database modules

Answer: B

Q1385. How does the is attribute enable customized built-in elements?

- A. It sets the element's ID
- B. It extends a built-in HTML element with custom behavior by associating it with a custom element class
- C. It changes the element's display property
- D. It removes default browser styling

Answer: B

Q1386. What is the purpose of the ElementInternals API in custom elements?

- A. To style elements with CSS
- B. To allow custom elements to participate in forms, manage accessibility states, and interact with browser internals
- C. To create database connections
- D. To handle routing

Answer: B

Q1387. How does the importmap script type improve module loading in HTML?

- A. It imports CSS modules
- B. It allows developers to control module specifier resolution, mapping bare specifiers to URLs without a build step
- C. It creates image maps
- D. It generates sitemaps

Answer: B

Q1388. What is the purpose of the Sanitizer API in the context of HTML?

- A. To remove all JavaScript from a page
- B. To safely parse and sanitize untrusted HTML strings before inserting them into the DOM, preventing XSS attacks
- C. To compress HTML files
- D. To validate form inputs

Answer: B

Q1389. What role does the Navigation API play in modern single-page applications?

- A. It manages CSS navigation menus
- B. It provides a standardized way to intercept and handle navigation events, replacing History API workarounds
- C. It creates breadcrumb components
- D. It manages database migrations

Answer: B

Q1390. How does the Speculation Rules API improve page navigation performance?

- A. It guesses which CSS to apply
- B. It allows developers to specify rules for prefetching or prerendering future navigations based on link patterns
- C. It creates animated page transitions
- D. It manages service workers

Answer: B

Q1391. How does the CSS paint API allow custom rendering of element backgrounds?

- A. It uses JavaScript to generate images dynamically as CSS backgrounds through registered paint worklets
- B. It only changes text colors
- C. It is a server-side rendering technique
- D. It replaces all CSS with JavaScript

Answer: A

Q1392. What is the difference between logical and physical properties in CSS?

- A. They are identical
- B. Logical properties like inline-start and block-end adapt to writing direction and text flow, while physical properties like left and top do not
- C. Logical properties only work in Firefox
- D. Physical properties are deprecated

Answer: B

Q1393. How does the CSS typed object model (Typed OM) improve performance over traditional CSSOM access?

- A. It has no performance benefit
- B. It provides typed JavaScript objects for CSS values, reducing string parsing overhead and enabling more efficient style manipulation
- C. It replaces CSS entirely
- D. It only works with inline styles

Answer: B

Q1394. What is the purpose of the CSS @scope rule?

- A. To create global styles
- B. To limit the reach of CSS selectors to a specific subtree of the DOM, preventing style leakage
- C. To define animation keyframes
- D. To import external fonts

Answer: B

Q1395. How does the CSS view transitions API work for page transitions?

- A. It uses iframes for transitions
- B. It captures screenshots of old and new states, then animates between them using CSS animations on pseudo-elements
- C. It requires WebGL
- D. It only works with React

Answer: B

Q1396. What is the purpose of CSS anchor positioning?

- A. To create navigation links
- B. To position elements relative to other elements without JavaScript, enabling native tooltip and popover positioning
- C. To set the font weight
- D. To define grid areas

Answer: B

Q1397. How do CSS nesting rules change the way styles are authored?

- A. They have no effect on authoring
- B. They allow selectors to be nested inside parent rules natively, similar to preprocessor nesting, reducing repetition
- C. They replace all existing CSS syntax
- D. They only work with class selectors

Answer: B

Q1398. What is the CSS color-mix() function used for?

- A. To add borders to elements
- B. To blend two colors together in a specified color space, producing a new color
- C. To mix fonts together
- D. To generate random colors

Answer: B

Q1399. How does the CSS scroll-driven animations API differ from scroll-based JavaScript animations?

- A. It is slower than JavaScript
- B. It runs animations on the compositor thread linked to scroll progress, avoiding main-thread JavaScript execution
- C. It only supports opacity changes
- D. It requires a third-party library

Answer: B

Q1400. What is the purpose of the CSS light-dark() function?

- A. To toggle animations
- B. To return one of two colors depending on the computed color-scheme, simplifying dark mode implementation
- C. To adjust screen brightness
- D. To create gradients

Answer: B

Q1401. How does the CSS Containment Level 2 specification improve rendering performance?

- A. It has no performance impact
- B. It allows developers to declare containment boundaries that isolate an element's layout, paint, or style from the rest of the document
- C. It only works with tables
- D. It replaces the box model

Answer: B

Q1402. What is the CSS timeline-scope property used for in scroll-driven animations?

- A. To set animation duration
- B. To extend a scroll timeline's scope to ancestor elements, allowing animations to be driven by scrolling of a descendant
- C. To create CSS variables
- D. To define grid tracks

Answer: B

Q1403. How does the CSS starting-style at-rule solve transition-on-entry problems?

- A. It sets default browser styles
- B. It defines initial styles for elements being inserted into the DOM, enabling transitions from display:none to visible states
- C. It creates startup animations
- D. It initializes CSS variables

Answer: B

Q1404. What is the purpose of the CSS individual transform properties (translate, rotate, scale)?

- A. They replace the transform property entirely
- B. They allow independent control over individual transformations, enabling separate animation timing without affecting other transforms
- C. They only work in 3D space
- D. They are identical to the transform property

Answer: B

Q1405. How does the CSS math function round() improve responsive design?

- A. It rounds element corners
- B. It rounds a value to a specified interval, enabling snap-to-grid-like sizing without JavaScript
- C. It rounds animation durations
- D. It rounds font sizes to integers only

Answer: B

Q1406. What is the relationship between the CSS paint() function and the PaintWorklet API?

- A. They are unrelated features
- B. paint() references a registered PaintWorklet to generate custom images rendered directly by the browser's paint engine
- C. paint() is a JavaScript function
- D. PaintWorklet replaces CSS entirely

Answer: B

Q1407. How does the CSS masonry layout proposal extend CSS Grid?

- A. It replaces CSS Grid entirely
- B. It enables items to pack tightly along one axis while maintaining grid alignment on the other, like a Pinterest-style layout
- C. It only creates circular layouts
- D. It is a flexbox feature

Answer: B

Q1408. What is the purpose of the CSS initial-letter property?

- A. To set the default font
- B. To create decorative drop caps by sizing the first letter of a block to span multiple lines
- C. To initialize CSS variables
- D. To set the first column width

Answer: B

Q1409. How does the @layer rule affect the CSS cascade?

- A. It has no effect on the cascade
- B. It creates explicit cascade layers that control the priority of rule sets, allowing authors to manage specificity conflicts intentionally
- C. It only applies to media queries
- D. It replaces the specificity system

Answer: B

Q1410. What problem does the CSS text-wrap: balance property solve?

- A. It aligns text to the left
- B. It distributes text more evenly across lines in headings, preventing orphans and uneven line lengths
- C. It wraps text around images
- D. It sets text to uppercase

Answer: B

Q1411. How does the JavaScript garbage collector determine which objects to free?

- A. It deletes all objects after a set time
- B. It uses a mark-and-sweep algorithm to identify objects no longer reachable from the root and frees their memory
- C. It only frees variables declared with var
- D. It requires manual memory management

Answer: B

Q1412. What is the difference between a lexical scope and a dynamic scope?

- A. They are the same concept
- B. Lexical scope is determined by where code is written in the source while dynamic scope is determined by the call stack at runtime
- C. Lexical scope is slower
- D. Dynamic scope is used in JavaScript

Answer: B

Q1413. How do labeled statements work with break and continue in JavaScript?

- A. Labels are only for comments
- B. Labels name a statement, allowing break or continue to reference it from within nested loops to control the outer loop
- C. Labels define CSS classes
- D. Labels create new variables

Answer: B

Q1414. What is the purpose of the Proxy and Reflect objects working together?

- A. They are for styling elements
- B. Proxy intercepts operations on objects while Reflect provides default behavior methods that traps can delegate to
- C. They manage CSS animations
- D. They handle file uploads

Answer: B

Q1415. How does the structuredClone() function handle circular references compared to JSON.parse(JSON.stringify())?

- A. Neither handles circular references
- B. structuredClone() correctly preserves circular references while JSON methods throw an error on circular structures
- C. JSON handles circular references better
- D. structuredClone() removes circular references

Answer: B

Q1416. What is the purpose of the Symbol.toPrimitive well-known symbol?

- A. To create CSS symbols
- B. To define how an object is converted to a primitive value, controlling behavior in type coercion contexts
- C. To create HTML symbols
- D. To generate random numbers

Answer: B

Q1417. How does the JavaScript engine optimize property access using inline caches?

- A. It caches CSS properties
- B. It remembers the shape of objects at specific code locations to skip repeated property lookup steps on subsequent accesses
- C. It stores all properties in a database
- D. It uses localStorage for caching

Answer: B

Q1418. What is the difference between Object.create(null) and a regular object literal?

- A. They are identical
- B. Object.create(null) creates an object with no prototype chain, avoiding inherited properties like toString and hasOwnProperty
- C. Object.create(null) is slower
- D. Regular object literals have no prototype

Answer: B

Q1419. How does the for...in loop differ from Object.keys() for iterating over properties?

- A. They are identical
- B. for...in iterates over all enumerable properties including inherited ones while Object.keys() only returns own enumerable properties
- C. for...in is faster
- D. Object.keys() includes inherited properties

Answer: B

Q1420. What is the purpose of the @@iterator (Symbol.iterator) protocol in JavaScript?

- A. To iterate over CSS rules
- B. To define a standard way for objects to produce a sequence of values, making them usable with for...of and spread syntax
- C. To create database iterators
- D. To manage HTTP request sequences

Answer: B

Q1421. How does the WeakRef class work and what is its relationship to garbage collection?

- A. It prevents garbage collection
- B. It holds a weak reference to an object that does not prevent it from being garbage collected, allowing observation without retention
- C. It creates strong references
- D. It manages CSS references

Answer: B

Q1422. What is the purpose of the Atomics object and how does it work with SharedArrayBuffer?

- A. It manages CSS atoms
- B. It provides atomic operations on shared memory locations to prevent race conditions in concurrent multi-threaded JavaScript
- C. It creates atomic CSS classes
- D. It manages DOM atoms

Answer: B

Q1423. How does the using declaration (Explicit Resource Management) improve resource cleanup?

- A. It imports CSS modules
- B. It ensures resources like file handles and database connections are automatically cleaned up via Symbol.dispose when they go out of scope
- C. It creates reusable components
- D. It manages memory allocation

Answer: B

Q1424. What is the difference between Error.cause and traditional error wrapping?

- A. They are identical approaches
- B. Error.cause provides a standardized way to chain errors by attaching the original error as a cause property
- C. Error.cause is deprecated
- D. Traditional wrapping is always better

Answer: B

Q1425. How does the Temporal API improve date and time handling compared to the Date object?

- A. It has no improvements over Date
- B. It provides immutable date-time types, proper time zone support, and eliminates ambiguities present in the Date API
- C. It only works in Node.js
- D. It replaces setTimeout

Answer: B

Q1426. What is the purpose of the Decorator pattern in the TC39 proposal and how does it modify class definitions?

- A. It adds CSS decorations
- B. It allows functions to wrap class elements to add behavior, modify definitions, or add metadata at definition time
- C. It creates decorative HTML elements
- D. It manages font decorations

Answer: B

Q1427. How does the Record and Tuple proposal introduce immutable data structures to JavaScript?

- A. Records are identical to objects
- B. Records and Tuples are deeply immutable versions of objects and arrays that are compared by value rather than by reference
- C. They only work with strings
- D. They replace Map and Set

Answer: B

Q1428. What problem does the Pipeline Operator proposal solve in JavaScript?

- A. It manages network pipes
- B. It improves readability of nested function calls by allowing values to be piped through a series of functions in a left-to-right order
- C. It creates CSS pipes
- D. It handles data streaming

Answer: B

Q1429. How does the ShadowRealm proposal enable secure code isolation in JavaScript?

- A. It creates CSS shadow effects
- B. It provides a mechanism to evaluate code in an isolated global environment with its own set of built-in objects
- C. It manages Shadow DOM
- D. It creates encrypted variables

Answer: B

Q1430. What is the purpose of the AsyncContext proposal in JavaScript?

- A. It manages CSS context
- B. It provides a way to propagate context values through asynchronous operations without explicit parameter passing
- C. It creates database contexts
- D. It manages HTML contexts

Answer: B

Q1431. How does React's compiler (React Forget) optimize re-renders without manual memoization?

- A. It removes all re-renders
- B. It automatically inserts memoization at the compiler level, eliminating the need for useMemo, useCallback, and React.memo
- C. It only optimizes CSS
- D. It replaces the virtual DOM

Answer: B

Q1432. What is the concept of resumability in Qwik and how does it differ from hydration?

- A. They are the same process
- B. Resumability serializes the application state and event handlers to HTML, allowing the client to resume without re-executing component code
- C. Resumability is slower
- D. Hydration is more efficient

Answer: B

Q1433. How does fine-grained reactivity in Solid.js avoid the overhead of virtual DOM diffing?

- A. It uses a faster virtual DOM
- B. It tracks reactive dependencies at the individual expression level, updating only the specific DOM nodes that depend on changed data
- C. It skips all DOM updates
- D. It uses Web Workers for rendering

Answer: B

Q1434. What is the purpose of the React use() hook and how does it change data fetching patterns?

- A. It replaces all other hooks
- B. It allows reading the value of a resource like a Promise or Context during rendering, enabling Suspense-based data fetching
- C. It only fetches CSS resources
- D. It replaces the fetch API

Answer: B

Q1435. How does Angular's signal-based reactivity change its change detection model?

- A. It has no effect on change detection
- B. Signals enable fine-grained reactivity that automatically tracks dependencies and updates only affected components, reducing unnecessary checks
- C. Signals replace all observables
- D. Signals only work with forms

Answer: B

Q1436. What is the purpose of React Server Actions and how do they simplify data mutations?

- A. They manage CSS server-side
- B. They allow defining server-side functions that can be called directly from client components, handling form submissions without manual API routes
- C. They replace all server routes
- D. They only work with databases

Answer: B

Q1437. How does the islands architecture in Astro optimize page performance?

- A. It renders everything on the client
- B. It renders static HTML by default and only hydrates interactive components as isolated islands, minimizing client-side JavaScript
- C. It uses iframes for each component
- D. It eliminates all JavaScript

Answer: B

Q1438. What is the difference between optimistic and pessimistic UI updates in frontend applications?

- A. They are the same strategy
- B. Optimistic updates immediately reflect changes in the UI before server confirmation while pessimistic updates wait for server response
- C. Optimistic updates are always correct
- D. Pessimistic updates are faster

Answer: B

Q1439. How does streaming SSR improve the user experience compared to traditional SSR?

- A. It has no improvement over traditional SSR
- B. It sends HTML in chunks as they become ready, allowing the browser to start rendering before the entire page is generated
- C. It only streams CSS
- D. It requires WebSocket connections

Answer: B

Q1440. What is the purpose of module federation in micro-frontend architectures?

- A. It manages CSS modules
- B. It allows separately built and deployed applications to share code and components at runtime without rebuilding
- C. It merges all JavaScript into one file
- D. It only works with Angular modules

Answer: B

Q1441. How does the repository pattern improve backend architecture?

- A. It replaces all databases
- B. It abstracts data access logic behind an interface, decoupling business logic from specific database implementations
- C. It manages CSS repositories
- D. It creates HTML templates

Answer: B

Q1442. What is the purpose of database connection pooling in high-traffic applications?

- A. To create more databases
- B. To maintain a pool of reusable connections that avoid the overhead of establishing new connections for each request
- C. To pool CSS styles
- D. To merge database schemas

Answer: B

Q1443. How does the event sourcing pattern differ from traditional CRUD storage?

- A. They are identical approaches
- B. Event sourcing stores every state change as an immutable event, allowing complete reconstruction of state, while CRUD overwrites current state
- C. Event sourcing is slower
- D. CRUD stores more historical data

Answer: B

Q1444. What is the purpose of the circuit breaker pattern in backend services?

- A. To manage electrical circuits
- B. To prevent cascading failures by stopping requests to a failing service and returning fallback responses
- C. To break CSS layouts
- D. To interrupt database connections

Answer: B

Q1445. How does idempotency in API design prevent duplicate operations?

- A. It makes all requests faster
- B. It ensures that making the same request multiple times has the same effect as making it once, preventing duplicate side effects
- C. It only applies to GET requests
- D. It replaces authentication

Answer: B

Q1446. What is the purpose of the outbox pattern in microservices?

- A. To manage email outboxes
- B. To reliably publish events by writing them to a database table and processing them asynchronously, ensuring atomicity
- C. To store outgoing CSS styles
- D. To manage outbound network connections

Answer: B

Q1447. How does the CQRS pattern benefit read-heavy backend applications?

- A. It has no benefit for reads
- B. It separates read and write models, allowing the read model to be optimized with denormalized views and separate scaling
- C. It only works with NoSQL databases
- D. It replaces all caching

Answer: B

Q1448. What is the purpose of distributed locking in backend systems?

- A. To lock CSS stylesheets
- B. To coordinate exclusive access to shared resources across multiple service instances to prevent race conditions
- C. To lock database schemas
- D. To prevent code modifications

Answer: B

Q1449. How does the strangler fig pattern facilitate migration from monolith to microservices?

- A. It removes the monolith immediately
- B. It gradually replaces monolith functionality with microservices by routing specific features to new services while keeping the monolith running
- C. It only works with fig trees
- D. It creates a copy of the monolith

Answer: B

Q1450. What is the purpose of a dead letter queue in message-driven backend systems?

- A. To send farewell emails
- B. To store messages that failed processing after maximum retry attempts for later investigation and reprocessing
- C. To delete old messages
- D. To compress message queues

Answer: B

Q1451. How does the cluster module in Node.js improve server performance?

- A. It clusters CSS styles
- B. It forks multiple worker processes to utilize all CPU cores, distributing incoming connections across workers
- C. It creates database clusters
- D. It clusters HTML templates

Answer: B

Q1452. What is the purpose of database query optimization using execution plans?

- A. To plan CSS layouts
- B. To analyze how the database engine processes a query, identifying bottlenecks like full table scans that can be resolved with indexes
- C. To plan server deployments
- D. To optimize HTML rendering

Answer: B

Q1453. How does connection draining ensure zero-downtime deployments?

- A. It drains CSS connections
- B. It allows existing requests to complete while stopping new connections from reaching the old server instance during deployment
- C. It drains database connections permanently
- D. It removes all network connections

Answer: B

Q1454. What is the purpose of request coalescing in server-side caching?

- A. To merge CSS requests
- B. To combine multiple identical concurrent requests into a single backend call, serving all requesters with the same response
- C. To coalesce HTML files
- D. To merge database tables

Answer: B

Q1455. How does server-side request forgery (SSRF) exploit server-side applications?

- A. It forges CSS styles
- B. It tricks the server into making requests to internal resources or services by manipulating user-controlled input used in server requests
- C. It forges HTML elements
- D. It duplicates JavaScript code

Answer: B

Q1456. What is the purpose of structured logging in server-side applications?

- A. To structure CSS files
- B. To output log entries as structured data (JSON) with consistent fields, enabling automated parsing, searching, and alerting
- C. To structure HTML templates
- D. To organize database tables

Answer: B

Q1457. How does graceful shutdown differ from abrupt termination in server applications?

- A. They are identical
- B. Graceful shutdown stops accepting new requests, finishes in-flight requests, closes connections, and releases resources before exiting
- C. Graceful shutdown is faster
- D. Abrupt termination is safer

Answer: B

Q1458. What is the purpose of the back-pressure mechanism in streaming server applications?

- A. To compress network traffic
- B. To signal data producers to slow down when consumers cannot process data fast enough, preventing memory exhaustion
- C. To increase server pressure
- D. To manage CSS pressure

Answer: B

Q1459. How does the worker thread model in Node.js differ from the cluster module for CPU-intensive tasks?

- A. They are identical
- B. Worker threads share memory within a single process for CPU tasks while cluster creates separate processes that share the server port
- C. Worker threads are slower
- D. Cluster uses less memory

Answer: B

Q1460. What is the purpose of rate limiting with sliding window algorithms?

- A. To limit CSS animations per second
- B. To track request counts over a continuously moving time window, providing smoother rate limiting than fixed windows
- C. To limit database queries
- D. To control HTML rendering speed

Answer: B

Q1461. How does Multi-Version Concurrency Control (MVCC) enable concurrent database access?

- A. It locks all tables
- B. It maintains multiple versions of data so readers see a consistent snapshot while writers create new versions without blocking reads
- C. It duplicates the database
- D. It uses CSS versioning

Answer: B

Q1462. What is the purpose of a write-ahead log (WAL) in database crash recovery?

- A. To log CSS changes
- B. To record all changes to a log file before applying them to the database, ensuring recovery by replaying the log after a crash
- C. To write HTML ahead of rendering
- D. To log JavaScript errors

Answer: B

Q1463. How does range partitioning differ from hash partitioning in database systems?

- A. They are identical strategies
- B. Range partitioning splits data by value ranges enabling range queries while hash partitioning distributes data evenly using a hash function
- C. Range partitioning is always faster
- D. Hash partitioning only works with strings

Answer: B

Q1464. What is the N+1 query problem and how do eager loading and batching solve it?

- A. It is a CSS specificity issue
- B. It occurs when fetching related data triggers N additional queries; eager loading and batching reduce this to a constant number of queries
- C. It only happens with MongoDB
- D. It is a JavaScript memory issue

Answer: B

Q1465. What is the purpose of a materialized view and how does it differ from a regular view?

- A. They are the same thing
- B. A materialized view stores precomputed query results physically, providing faster reads at the cost of requiring periodic refresh
- C. Materialized views are slower
- D. Regular views store data physically

Answer: B

Q1466. How does the change data capture (CDC) pattern enable event-driven architectures?

- A. It captures CSS changes
- B. It monitors database transaction logs for changes and publishes them as events, enabling downstream systems to react to data modifications
- C. It captures HTML changes
- D. It captures JavaScript variable changes

Answer: B

Q1467. What is the purpose of query plan caching in database systems?

- A. To cache CSS query results
- B. To store optimized execution plans for previously seen queries, avoiding expensive re-optimization on repeated executions
- C. To cache HTML pages
- D. To cache JavaScript functions

Answer: B

Q1468. How does a graph database model data differently from a relational database?

- A. They model data identically
- B. Graph databases use nodes, edges, and properties to represent and store relationships directly, enabling efficient traversal of connected data
- C. Graph databases cannot store relationships
- D. Relational databases are always better for relationships

Answer: B

Q1469. What is the purpose of connection pool monitoring in production database deployments?

- A. To monitor CSS pool usage
- B. To track pool utilization, wait times, and connection lifecycle to detect issues like pool exhaustion, leaks, or misconfigured sizing
- C. To monitor HTML elements
- D. To track JavaScript memory pools

Answer: B

Q1470. How does an LSM tree-based storage engine differ from a B-tree-based engine?

- A. They are identical
- B. LSM trees buffer writes in memory and flush sorted runs to disk, optimizing write throughput, while B-trees update pages in place, optimizing reads
- C. LSM trees are always slower
- D. B-trees cannot handle concurrent writes

Answer: B

Q1471. How does a Content Security Policy nonce-based approach prevent XSS compared to hash-based CSP?

- A. They are identical approaches
- B. Nonces use a unique random token per request added to allowed scripts, while hashes allow only scripts matching precomputed hash values
- C. Nonces are less secure
- D. Hashes require more server resources

Answer: B

Q1472. What is the purpose of the Permissions-Policy header and how does it reduce attack surface?

- A. It manages CSS permissions
- B. It controls which browser features (camera, microphone, geolocation) can be used by the page and its iframes, reducing exploitable surface
- C. It manages user login permissions
- D. It sets database access permissions

Answer: B

Q1473. How does a timing attack exploit web applications and what is the defense?

- A. It attacks CSS timing functions
- B. It measures response time differences to extract information like valid usernames; constant-time comparison functions are the defense
- C. It times JavaScript execution
- D. It attacks database query timing

Answer: B

Q1474. What is the purpose of certificate transparency logs?

- A. To log CSS certificate changes
- B. To publicly record all issued TLS certificates, enabling domain owners to detect unauthorized certificates issued for their domains
- C. To track HTML changes
- D. To log JavaScript errors

Answer: B

Q1475. How does the PKCE extension improve the OAuth 2.0 authorization code flow for public clients?

- A. It adds CSS styling to OAuth pages
- B. It prevents authorization code interception attacks by binding the code to a cryptographic verifier known only to the client
- C. It compresses OAuth tokens
- D. It speeds up token refresh

Answer: B

Q1476. What is prototype pollution and how can it compromise a Node.js application?

- A. It pollutes CSS prototypes
- B. It modifies Object.prototype through malicious input, potentially adding properties to all objects and bypassing security checks
- C. It pollutes HTML prototypes
- D. It pollutes database prototypes

Answer: B

Q1477. What is the purpose of a Web Application Firewall (WAF) and how does it differ from a network firewall?

- A. They are identical
- B. A WAF inspects HTTP traffic at the application layer for attack patterns like XSS and SQLi while a network firewall filters traffic at the network layer
- C. WAFs are less secure
- D. Network firewalls inspect HTTP content

Answer: B

Q1478. How does token binding prevent token theft in bearer token authentication systems?

- A. It binds CSS tokens to elements
- B. It cryptographically binds tokens to the TLS connection, making stolen tokens unusable from a different TLS session
- C. It binds JavaScript variables
- D. It binds database tokens to queries

Answer: B

Q1479. What is the purpose of Software Composition Analysis (SCA) in web security?

- A. To analyze CSS compositions
- B. To scan project dependencies for known vulnerabilities and license compliance issues in third-party code
- C. To compose HTML elements
- D. To analyze JavaScript compositions

Answer: B

Q1480. How does a credential stuffing attack differ from a brute force attack?

- A. They are identical attacks
- B. Credential stuffing uses previously leaked username-password pairs from data breaches while brute force systematically tries all possible combinations
- C. Brute force uses leaked credentials
- D. Credential stuffing is always faster

Answer: B

Q1481. How does the Interaction to Next Paint (INP) metric improve upon First Input Delay (FID)?

- A. It is the same metric
- B. INP measures responsiveness across all interactions throughout the page lifecycle, not just the first input, providing a more comprehensive view
- C. INP only measures first click
- D. FID is more comprehensive

Answer: B

Q1482. How does the scheduler.yield() API improve long task handling in browsers?

- A. It yields CSS rendering
- B. It allows JavaScript to voluntarily yield control back to the browser during long tasks, preventing input delay while maintaining task priority
- C. It yields database connections
- D. It yields server resources

Answer: B

Q1483. What is the purpose of the PerformanceObserver API for monitoring web performance?

- A. To observe CSS changes
- B. To asynchronously observe and collect performance metrics like LCP, FID, and CLS without polling, enabling real user monitoring
- C. To observe database performance
- D. To monitor server health

Answer: B

Q1484. How does Brotli compression compare to Gzip for reducing transfer sizes?

- A. They achieve identical compression
- B. Brotli typically achieves 15-25% better compression ratios than Gzip for text resources, though at higher compression CPU cost
- C. Gzip always compresses better
- D. Brotli only works with images

Answer: B

Q1485. What is the purpose of the fetchpriority attribute and how does it affect resource loading?

- A. It sets CSS priority
- B. It lets developers hint at the relative priority of resources (high, low, auto), influencing browser fetch scheduling for critical resources
- C. It prioritizes database queries
- D. It manages server priorities

Answer: B

Q1486. How does the speculation rules API enable instant page navigation?

- A. It speculates about CSS styles
- B. It allows prerendering entire pages in the background based on declarative rules, enabling near-instant navigations when the user clicks
- C. It predicts server loads
- D. It speculates database queries

Answer: B

Q1487. What is the impact of main thread long tasks on web performance and how are they identified?

- A. Long tasks have no impact
- B. Tasks exceeding 50ms block the main thread, delaying input processing and rendering; they are identified using the Long Tasks API or Performance profiler
- C. Long tasks improve performance
- D. Only tasks over 1 second matter

Answer: B

Q1488. How does edge computing with edge functions improve web application latency?

- A. It moves computation farther from users
- B. It runs server-side logic at CDN edge locations near users, reducing round-trip time for dynamic content generation
- C. It only serves static files
- D. It replaces all backend servers

Answer: B

Q1489. What is the purpose of priority hints in resource loading and how do they interact with browser heuristics?

- A. They replace browser loading behavior
- B. They supplement browser heuristics by letting developers explicitly adjust the fetch priority of resources like images, scripts, and fetch calls
- C. They have no effect on loading
- D. They only work with CSS files

Answer: B

Q1490. How does the content-visibility CSS property improve rendering performance for long pages?

- A. It hides content permanently
- B. It allows the browser to skip layout and paint for off-screen elements, dramatically reducing initial rendering work for long pages
- C. It only affects text visibility
- D. It replaces lazy loading

Answer: B

Q1491. How does property-based testing complement example-based testing?

- A. It replaces example-based testing
- B. It generates random inputs to verify that properties hold for all cases, discovering edge cases that manually chosen examples miss
- C. It only tests CSS properties
- D. It is slower and less useful

Answer: B

Q1492. What is the testing trophy pattern and how does it prioritize test types?

- A. It is a CSS trophy animation
- B. It emphasizes integration tests as the largest layer, with fewer unit tests, even fewer E2E tests, and static analysis at the base
- C. It is identical to the testing pyramid
- D. It prioritizes unit tests above all

Answer: B

Q1493. How does mutation testing evaluate the quality of a test suite?

- A. It mutates CSS properties
- B. It introduces small code changes (mutations) and checks if tests detect them; surviving mutations indicate gaps in test coverage
- C. It mutates database records
- D. It creates new test files

Answer: B

Q1494. What is contract testing and how does it prevent integration failures in microservices?

- A. Testing legal contracts
- B. It verifies that services adhere to agreed-upon API contracts, catching incompatible changes before deployment without full integration tests
- C. It tests CSS contracts
- D. It manages database contracts

Answer: B

Q1495. How does test sharding across multiple CI workers reduce pipeline duration?

- A. It has no effect on duration
- B. It distributes test files across parallel CI runners, with each runner executing a subset, reducing total wall-clock time proportionally
- C. It only works with unit tests
- D. It requires identical test runners

Answer: B

Q1496. What is the purpose of chaos engineering and how does it apply to web applications?

- A. Creating chaotic CSS layouts
- B. Deliberately introducing failures into production to verify system resilience and identify weaknesses before they cause outages
- C. Creating random test data
- D. Randomizing deployment order

Answer: B

Q1497. How does visual regression testing work at scale?

- A. It compares CSS file sizes
- B. It captures and pixel-compares screenshots of UI components across changes, using AI to filter false positives from legitimate changes
- C. It visually inspects database records
- D. It manually compares page designs

Answer: B

Q1498. What is the purpose of API fuzz testing and how does it discover vulnerabilities?

- A. To make APIs fuzzy
- B. To send malformed, unexpected, or random data to API endpoints to discover crashes, memory leaks, and security vulnerabilities
- C. To add blur effects to APIs
- D. To fuzz CSS values

Answer: B

Q1499. How do test containers improve integration testing reliability?

- A. They contain CSS test styles
- B. They spin up real database and service instances in Docker containers for tests, ensuring consistent environments and avoiding shared state
- C. They contain HTML elements
- D. They compress test files

Answer: B

Q1500. What is the purpose of the testing library philosophy of testing user behavior over implementation details?

- A. It tests CSS behavior
- B. It encourages testing what users see and interact with rather than internal component state, making tests more resilient to refactoring
- C. It only tests API behavior
- D. It avoids testing entirely

Answer: B

Q1501. How does GitOps use Git as the single source of truth for infrastructure and deployments?

- A. It stores CSS in Git
- B. It stores desired infrastructure and application state declaratively in Git, with automated controllers continuously reconciling actual state with the repository
- C. It only manages Git repositories
- D. It replaces CI/CD entirely

Answer: B

Q1502. What is the purpose of a service mesh in container orchestration deployments?

- A. To mesh CSS styles together
- B. To manage service-to-service communication with features like mutual TLS, traffic splitting, retries, and distributed tracing without modifying application code
- C. To create mesh network layouts
- D. To mesh database connections

Answer: B

Q1503. How does immutable infrastructure differ from mutable infrastructure and why is it preferred?

- A. They are identical approaches
- B. Immutable infrastructure replaces entire server instances instead of updating them in place, ensuring consistency and eliminating configuration drift
- C. Immutable infrastructure is more expensive
- D. Mutable infrastructure is always more reliable

Answer: B

Q1504. What is the purpose of canary analysis in progressive delivery?

- A. To analyze bird sounds
- B. To automatically compare metrics between the canary release and the stable version, rolling back if degradation is detected
- C. To analyze CSS canary values
- D. To manage canary database records

Answer: B

Q1505. How does Kubernetes handle rolling updates and what role do readiness probes play?

- A. It updates all pods simultaneously
- B. It gradually replaces old pods with new ones, using readiness probes to ensure new pods are serving traffic before terminating old ones
- C. It only works with Docker
- D. Readiness probes are optional and unused

Answer: B

Q1506. What is the purpose of a deployment feature flag system?

- A. To flag CSS features
- B. To decouple deployment from release by toggling features on or off without redeploying, enabling gradual rollouts and instant rollbacks
- C. To flag database features
- D. To create HTML feature flags

Answer: B

Q1507. How does Infrastructure as Code (IaC) with tools like Terraform improve deployment reliability?

- A. It codes CSS infrastructure
- B. It defines infrastructure declaratively in version-controlled files, enabling reproducible provisioning, change tracking, and peer review
- C. It only works with AWS
- D. It replaces all deployment tools

Answer: B

Q1508. What is the purpose of deployment runbooks and how do they improve incident response?

- A. To run CSS books
- B. To provide step-by-step procedures for deployments and incident recovery, reducing mean time to recovery through documented playbooks
- C. To manage database runbooks
- D. To create HTML runbooks

Answer: B

Q1509. How does multi-region deployment improve application resilience?

- A. It deploys CSS to multiple regions
- B. It runs application instances across geographically separate regions so that regional failures do not cause total outages
- C. It only improves database speed
- D. It requires identical hardware everywhere

Answer: B

Q1510. What is the purpose of observability in production deployments and how does it differ from monitoring?

- A. They are identical concepts
- B. Observability provides the ability to understand internal system state from external outputs (logs, metrics, traces) while monitoring checks predefined conditions
- C. Observability replaces monitoring
- D. Monitoring is more comprehensive

Answer: B

Q1511. How does GraphQL's DataLoader pattern solve the N+1 query problem?

- A. It loads CSS data
- B. It batches multiple individual data requests within a single tick into one bulk query, eliminating redundant database calls
- C. It only loads JavaScript modules
- D. It replaces GraphQL resolvers

Answer: B

Q1512. What is the purpose of API schema validation using JSON Schema or Zod?

- A. To validate CSS schemas
- B. To define and enforce the shape, types, and constraints of API request and response data at runtime, catching errors early
- C. To validate HTML schemas
- D. To validate database schemas only

Answer: B

Q1513. How does content negotiation work in REST APIs and what role does the Accept header play?

- A. It negotiates CSS content
- B. The client sends an Accept header specifying desired response formats (JSON, XML, etc.) and the server returns data in the best matching format
- C. It only works with HTML
- D. It negotiates database connections

Answer: B

Q1514. What is the purpose of API gateway patterns like Backend for Frontend (BFF)?

- A. To style frontend gateways
- B. BFF creates dedicated API gateways for each frontend type, aggregating and tailoring backend service calls to specific client needs
- C. It replaces all backend services
- D. It only works with GraphQL

Answer: B

Q1515. How does gRPC achieve better performance than REST for service-to-service communication?

- A. It uses CSS for styling
- B. It uses HTTP/2, Protocol Buffers for binary serialization, and supports streaming, resulting in smaller payloads and lower latency than JSON over HTTP
- C. It is always slower than REST
- D. It replaces HTTP entirely

Answer: B

Q1516. What is the purpose of idempotency keys in API design and how do they work?

- A. To identify CSS keys
- B. They allow clients to attach a unique key to requests so the server can detect retries and return the original response without re-executing the operation
- C. To create database keys
- D. To generate JavaScript keys

Answer: B

Q1517. How does event-driven API design with event sourcing differ from traditional request-response patterns?

- A. They are identical patterns
- B. Event-driven APIs publish domain events that consumers subscribe to asynchronously, enabling loose coupling and eventual consistency instead of synchronous calls
- C. Event-driven APIs are always slower
- D. They only work with GraphQL

Answer: B

Q1518. What is the purpose of API observability and what metrics should be tracked?

- A. To observe CSS changes
- B. To monitor API health through metrics like request rate, error rate, latency percentiles, and throughput, enabling proactive issue detection
- C. To observe HTML changes
- D. To observe database queries

Answer: B

Q1519. How does the API circuit breaker pattern improve system resilience?

- A. It breaks CSS circuits
- B. It monitors failure rates and opens the circuit to reject requests to a failing service, preventing resource exhaustion and enabling recovery
- C. It breaks database circuits
- D. It disconnects API clients

Answer: B

Q1520. What is the purpose of API contract testing with tools like Pact?

- A. To test CSS contracts
- B. To verify that API providers and consumers agree on request/response formats by testing both sides against a shared contract independently
- C. To test HTML contracts
- D. To test database contracts

Answer: B

Q1521. How does the concept of partial hydration improve performance in modern frameworks?

- A. It partially loads CSS
- B. It hydrates only interactive components on a page while leaving static content as plain HTML, reducing JavaScript payload
- C. It partially loads databases
- D. It replaces full page hydration entirely

Answer: B

Q1522. What is the purpose of the edge-first architecture pattern in modern web applications?

- A. It adds CSS edge effects
- B. It runs application logic at CDN edge locations by default, falling back to origin servers only when necessary, minimizing latency globally
- C. It only processes edge cases
- D. It replaces all server-side logic

Answer: B

Q1523. How does the signals pattern in modern frameworks differ from React's state model?

- A. They are identical patterns
- B. Signals provide fine-grained reactivity by tracking dependencies at the value level, updating only subscribers without re-rendering entire component trees
- C. Signals are slower
- D. React's model is more granular

Answer: B

Q1524. What is the purpose of the Baseline web platform compatibility initiative?

- A. To create CSS baselines
- B. To define a clear set of web features that are interoperable across major browsers, helping developers know what they can safely use
- C. To baseline database performance
- D. To measure JavaScript execution baselines

Answer: B

Q1525. How does the composability pattern influence modern web architecture?

- A. It composes CSS animations
- B. It enables building applications by combining independent, interchangeable modules and services rather than monolithic implementations
- C. It only applies to React components
- D. It replaces all existing patterns

Answer: B

Q1526. What is the purpose of the View Transitions API for multi-page applications?

- A. To create CSS views
- B. It enables smooth animated transitions between pages in traditional multi-page applications, previously only achievable in SPAs
- C. It only works in SPAs
- D. It replaces CSS animations

Answer: B

Q1527. How does AI-assisted development change modern web development workflows?

- A. It replaces all developers
- B. AI tools assist with code generation, debugging, review, and documentation while developers provide direction, review outputs, and handle complex decisions
- C. AI writes all CSS
- D. AI manages all databases

Answer: B

Q1528. What is the purpose of the Interop project for web standards?

- A. To create interoperable CSS
- B. It is a cross-browser collaboration where major browser vendors commit to improving compatibility on specific web features each year
- C. It only affects JavaScript
- D. It manages database interoperability

Answer: B

Q1529. How does the concept of zero-JavaScript pages improve performance in modern frameworks like Astro?

- A. It removes all JavaScript from the internet
- B. Pages ship no client-side JavaScript by default, rendering everything as static HTML and only adding JavaScript for explicitly interactive components
- C. It prevents JavaScript from running
- D. It only works for simple pages

Answer: B

Q1530. What are the key tradeoffs of adopting a serverless architecture for web applications?

- A. There are no tradeoffs
- B. Serverless offers automatic scaling and reduced ops but introduces cold starts, vendor lock-in, execution limits, and challenges with stateful workloads
- C. Serverless is always better
- D. Serverless cannot scale

Answer: B