

CUET-UG Computer Science Sample Paper-17

Duration: 1 Hour

Maximum Marks: 250

Instructions

- This paper contains a total of 50 Multiple Choice Questions.
- Each correct answer carries **+5 marks**.
- Each incorrect answer carries **-1 mark**.
- No negative marking for unattempted questions.

Q1. Which SQL aggregate function would you use to find the standard deviation of values in a column named 'Score'?

- (A) *STDDEV()*
- (B) *VAR()*
- (C) *STDEV_POP()*
- (D) *SQRT_VAR()*

Q2. In Python, what will be the result of the expression $2 * *3 * *2$?

- (A) 64
- (B) 512
- (C) 4096
- (D) 128

Q3. A network administrator wants to connect two different types of networks that use different protocols. Which device is essential for this task?

- (A) Bridge
- (B) Gateway
- (C) Switch
- (D) Repeater



- Q4.** Which of the following sorting algorithms has the best average-case time complexity of $O(n \log n)$ but can degrade to $O(n^2)$ in the worst case?
- (A) Merge Sort
 - (B) Heap Sort
 - (C) Quick Sort
 - (D) Selection Sort
- Q5.** In Relational Algebra, which operation is used to combine all rows from two relations, including duplicates?
- (A) Union
 - (B) Union All
 - (C) Intersection
 - (D) Cartesian Product
- Q6.** What is the size of an IPv6 address in bits?
- (A) 32
 - (B) 64
 - (C) 128
 - (D) 256
- Q7.** Which SQL function is used to convert a string to lowercase?
- (A) *LOWER()*
 - (B) *LCASE()*
 - (C) *Both(A)and(B)*
 - (D) *MIN()*
- Q8.** In Python file handling, which mode opens a file for both reading and writing in binary format, overwriting the file if it exists?
- (A) *wb*



- (B) $rb+$
- (C) $wb+$
- (D) $ab+$

Q9. In a stack, if elements 'A', 'B', and 'C' are pushed in order, which element will be the first one to be popped?

- (A) A
- (B) B
- (C) C
- (D) None

Q10. Which protocol is used for sending and receiving web pages over the Internet securely?

- (A) HTTP
- (B) HTTPS
- (C) FTP
- (D) SMTP

Q11. What is the time complexity of searching for an element in a balanced Binary Search Tree?

- (A) $O(n)$
- (B) $O(n^2)$
- (C) $O(\log n)$
- (D) $O(1)$

Q12. Which SQL clause is used to filter groups formed by the GROUP BY clause?

- (A) WHERE
- (B) HAVING
- (C) ORDER BY



(D) LIMIT

Q13. In Python, what is the purpose of the 'pass' statement?

- (A) To exit a loop
- (B) To skip the current iteration
- (C) As a placeholder for future code
- (D) To raise an exception

Q14. Which layer of the OSI model is responsible for routing data packets across different networks?

- (A) Data Link Layer
- (B) Network Layer
- (C) Transport Layer
- (D) Application Layer

Q15. Which sorting algorithm is often considered the most efficient for large datasets due to its $O(n \log n)$ worst-case time complexity?

- (A) Bubble Sort
- (B) Insertion Sort
- (C) Merge Sort
- (D) Selection Sort

Q16. Which of the following is a DDL (Data Definition Language) command in SQL?

- (A) UPDATE
- (B) TRUNCATE
- (C) SELECT
- (D) INSERT

Q17. In Python, what is the output of the following code snippet?

```
L = [1, 2, 3]
```



```
L.extend([4, 5])  
print(len(L))
```

- (A) 3
- (B) 4
- (C) 5
- (D) Error

Q18. Which network topology requires a multipoint connection and uses a single cable to connect all nodes?

- (A) Star
- (B) Ring
- (C) Bus
- (D) Mesh

Q19. Which data structure is typically used to implement recursion?

- (A) Queue
- (B) Stack
- (C) Linked List
- (D) Tree

Q20. Which of the following is NOT a valid exception handling keyword in Python?

- (A) try
- (B) except
- (C) catch
- (D) finally

Q21. What is the purpose of the GROUP BY clause in a SQL query?

- (A) To sort the result set
- (B) To filter rows based on a condition



- (C) To arrange identical data into groups
- (D) To join two tables

Q22. A MAC address is also known as the:

- (A) Logical address
- (B) Physical address
- (C) IP address
- (D) Network address

Q23. Which sorting algorithm works by repeatedly inserting an element from an unsorted part into its correct position in a sorted part?

- (A) Bubble Sort
- (B) Selection Sort
- (C) Insertion Sort
- (D) Quick Sort

Q24. In Python, which function is used to get the current position of the file pointer?

- (A) *seek()*
- (B) *tell()*
- (C) *where()*
- (D) *pos()*

Q25. Which protocol is used to map a numeric IP address to a physical MAC address?

- (A) DNS
- (B) ARP
- (C) DHCP
- (D) ICMP

Q26. What is the result of the SQL expression `SELECT ROUND(15.678, 2);`?



- (A) 15.6
- (B) 15.67
- (C) 15.68
- (D) 16.00

Q27. In a linear queue, if the 'rear' index is equal to the size of the array minus one, the queue is said to be:

- (A) Empty
- (B) Full
- (C) Overflowed
- (D) Underflowed

Q28. Which type of switching breaks data into small pieces before transmitting them over the network?

- (A) Circuit Switching
- (B) Packet Switching
- (C) Message Switching
- (D) Line Switching

Q29. In Python, what does the `pickle.dump()` method do?

- (A) Reads an object from a file
- (B) Serializes an object to a file
- (C) Deletes an object
- (D) Displays an object

Q30. Which attribute uniquely identifies a tuple in a relation?

- (A) Foreign Key
- (B) Primary Key
- (C) Alternate Key



(D) Candidate Key

Q31. Which SQL function is used to return the current date and time in 'YYYY-MM-DD HH:MM:SS' format?

(A) *CURDATE()*

(B) *NOW()*

(C) *SYSDATE()*

(D) *DATE()*

Q32. What is the postfix equivalent of the infix expression: $(A + B) * (C - D)$?

(A) $AB + CD - *$

(B) $ABCD + -*$

(C) $AB + CD * -$

(D) $AB * CD - +$

Q33. In Python, which exception is raised when a local or global name is not found?

(A) ImportError

(B) ValueError

(C) NameError

(D) TypeError

Q34. Which network device operates at the Physical Layer (Layer 1) of the OSI model and acts as a multiport repeater?

(A) Switch

(B) Hub

(C) Router

(D) Bridge

Q35. In Relational Algebra, the σ (sigma) operator is used for:



- (A) Projection
- (B) Selection
- (C) Join
- (D) Union
- (E)

Q36. Which sorting algorithm has a consistent time complexity of $O(n^2)$ for best, average, and worst cases?

- (A) Insertion Sort
- (B) Bubble Sort
- (C) Selection Sort
- (D) Quick Sort

Q37. What does the 'MAC' in MAC address stand for?

- (A) Media Access Control
- (B) Memory Access Control
- (C) Main Access Channel
- (D) Message Authentication Code

Q38. In Python, the `seek(offset, from_where)` method's `from_where` argument value of 1 represents:

- (A) Beginning of the file
- (B) Current position of the file pointer
- (C) End of the file
- (D) Second character of the file

Q39. Which SQL keyword is used to remove all records from a table, including the space allocated for the records, but keeps the table structure?

- (A) DELETE



- (B) DROP
- (C) TRUNCATE
- (D) REMOVE

Q40. A firewall primarily provides which of the following security functions?

- (A) Data Encryption
- (B) Packet Filtering
- (C) Virus Removal
- (D) User Authentication

Q41. Which data structure works on the FIFO (First-In-First-Out) principle?

- (A) Stack
- (B) Queue
- (C) Tree
- (D) Graph

Q42. In SQL, which wildcard character represents zero, one, or multiple characters in a LIKE pattern?

- (A) _ (Underscore)
- (B) * (Asterisk)
- (C) % (Percentage)
- (D) ? (Question Mark)

Q43. What is the time complexity of the Bubble Sort algorithm in the best-case scenario (when the list is already sorted)?

- (A) $O(n)$
- (B) $O(n \log n)$
- (C) $O(n^2)$
- (D) $O(1)$



- Q44.** Which protocol is primarily used to download email messages from a mail server to a local client?
- (A) SMTP
 - (B) POP3
 - (C) HTTP
 - (D) FTP
- Q45.** In Python file handling, which method returns a list containing each line in the file as a list item?
- (A) *read()*
 - (B) *readline()*
 - (C) *readlines()*
 - (D) *readall()*
- Q46.** Which constraint in SQL ensures that all values in a column are unique and not null?
- (A) UNIQUE
 - (B) PRIMARY KEY
 - (C) FOREIGN KEY
 - (D) CHECK
- Q47.** Which topology connects every node to exactly two other nodes, forming a single continuous pathway for signals?
- (A) Star
 - (B) Bus
 - (C) Ring
 - (D) Mesh
- Q48.** In Python, which of the following is used to handle exceptions that may occur during the execution of a block of code?



- (A) *try ... except*
- (B) *if ... else*
- (C) *for ... in*
- (D) *def ... return*

Q49. Which layer of the OSI model is responsible for establishing, managing, and terminating connections between applications?

- (A) Presentation Layer
- (B) Session Layer
- (C) Transport Layer
- (D) Data Link Layer

Q50. In SQL, what is the default order of sorting for the ORDER BY clause?

- (A) Descending
- (B) Ascending
- (C) Random
- (D) No default order



Detailed Solutions**Q1.****Solution****Concept:**

In SQL, aggregate functions perform a calculation on a set of values and return a single value. While common functions include SUM() and AVG(), statistical functions like standard deviation are used to measure the amount of variation or dispersion in a set of values.

Solution:

- (a) The question asks for the standard deviation of values in a specific column.
- (b) In standard SQL, the function used for population standard deviation is STDDEV() or STDDEV_POP().
- (c) VAR() is used for variance, not standard deviation.
- (d) SQRT_VAR() is not a standard SQL aggregate function.
- (e) Therefore, STDDEV_POP() is the correct choice.

Final Answer: The function is STDDEV_POP().

Answer: (C)

Q2.**Solution****Concept:**

In Python, the exponentiation operator ‘**’ has right-to-left associativity. This means that in an expression with multiple exponentiation operators, the operation on the right is evaluated first.

Solution:

1. The given expression is $2 * 3 * 2$. 2. Due to right-to-left associativity, we first evaluate the rightmost part: $3 * 2$. 3. $3 * 2 = 9$. 4. Now, the expression becomes $2 * 9$. 5. $2^9 = 2 \times 2 \times 2 \times 2 \times 2 \times 2 \times 2 \times 2 \times 2 = 512$. 6. If it were left-to-right, it would be $(2^3)^2 = 8^2 = 64$, which is incorrect in Python.

Final Answer: The result is 512.

Answer: (B)



Q3.

Solution**Concept:**

Network interconnection devices operate at different layers of the OSI model[cite: 3179]. A device that connects two different networks using different protocols must perform protocol translation.

Solution:

1. A **Bridge** connects two similar network segments at the Data Link Layer[cite: 3395]. 2. A **Switch** connects devices within a single network using MAC addresses[cite: 688]. 3. A **Repeater** simply regenerates signals to extend distance. 4. A **Gateway** is a sophisticated device (often a router with additional software) that acts as an entry/exit point for a network and can translate data between different protocol suites (e.g., connecting a TCP/IP network to an IPX/SPX network).

Final Answer: The device is a Gateway.

Answer: (B)

Q4.

Solution**Concept:**

Time complexity measures how the runtime of an algorithm grows with input size[cite: 392]. Different sorting algorithms have different characteristics based on the pivot or partitioning method used.

Solution:

1. Merge Sort always performs at $O(n \log n)$ regardless of input order. 2. Selection Sort always performs at $O(n^2)$. 3. Quick Sort uses a "divide-and-conquer" approach with a pivot. 4. In the average case, Quick Sort is highly efficient at $O(n \log n)$. 5. However, if the pivot is poorly chosen (e.g., the smallest or largest element in an already sorted list), it degrades to its worst-case complexity of $O(n^2)$.

Final Answer: The algorithm is Quick Sort.

Answer: (C)



Q5.

Solution**Concept:**

Relational Algebra and SQL provide operations to combine data from different relations. While standard Set Theory operations (like Union) handle unique elements, specific variations exist for duplicates.

Solution:

1. In pure Relational Algebra, the **Union** (\cup) operation combines tuples from two union-compatible relations and removes duplicates. 2. In SQL, the 'UNION' operator also removes duplicates by default. 3. To combine all rows including duplicates, the 'UNION ALL' operator is used. 4. Intersection finds only common elements. 5. Cartesian Product (\times) pairs every row of the first relation with every row of the second.

Final Answer: The operation is Union All.

Answer: (B)

Q6.

Solution**Concept:**

IPv6 (Internet Protocol version 6) was developed by the Internet Engineering Task Force (IETF) to deal with the long-anticipated problem of IPv4 address exhaustion. It uses a much larger address space.

Solution:

1. An IPv4 address is 32 bits long, which allows for approximately 4.3 billion unique addresses. 2. An IPv6 address is 128 bits long, represented as eight groups of four hexadecimal digits. 3. The number of bits in IPv6 is calculated as 2^{128} , providing a vastly larger address space (approximately 3.4×10^{38} addresses). 4. Therefore, the size is 128 bits.

Final Answer: The size is 128 bits.

Answer: (C)



Q7.

Solution**Concept:**

SQL provides various string functions to manipulate text data. Case conversion functions are essential for standardizing data during retrieval or comparison.

Solution:

1. 'LOWER()' is the standard SQL function used to convert all characters in a string to lowercase. 2. 'LCASE()' is a synonym for 'LOWER()' in many SQL dialects like MySQL. 3. Since both are commonly accepted and functional in the environments typically tested in CUET, the most inclusive answer is "Both (A) and (B)". 4. 'MIN()' is an aggregate function used to find the minimum value in a set.

Final Answer: Both (A) and (B).

Answer: (C)

Q8.

Solution**Concept:**

Python's 'open()' function uses specific mode strings to determine how a file is accessed. Binary modes are used for non-text files like images or compiled data.

Solution:

1. The 'w' in a mode indicates "writing" and starts by truncating (overwriting) the file if it already exists. 2. The 'b' indicates "binary" mode. 3. The '+' indicates that the file is open for both reading and writing (updating). 4. 'rb+' opens for reading and writing but does not overwrite the file initially; it starts at the beginning. 5. 'wb+' opens for writing and reading in binary, and it **overwrites** the existing file immediately upon opening. 6. 'ab+' is for appending and reading in binary.

Final Answer: The mode is wb+.

Answer: (C)



Q9.

Solution**Concept:**

A Stack is a linear data structure that follows the LIFO (Last-In-First-Out) principle. This means the last element added to the stack is the first one to be removed.

Solution:

1. Elements are pushed in the order: 'A', then 'B', then 'C'. 2. The stack visualization from bottom to top would be: - [C] <- Top - [B] - [A] 3. When a 'pop' operation is performed, it removes the element currently at the "Top" of the stack. 4. Since 'C' was the last element pushed, it is at the top. 5. Therefore, 'C' will be the first element to be popped.

Final Answer: The element is C.

Answer: (C)

Q10.

Solution**Concept:**

Network protocols define the rules for communication. Secure communication on the web requires encryption to protect data integrity and privacy.

Solution:

1. **HTTP** (HyperText Transfer Protocol) is the standard protocol for the web but transmits data in plain text, making it insecure. 2. **HTTPS** (HyperText Transfer Protocol Secure) is the secure version of HTTP. It uses SSL/TLS encryption to secure the communication channel. 3. **FTP** is for file transfer. 4. **SMTP** is for sending emails. 5. In the modern web, "securely" implies the use of HTTPS.

Final Answer: The protocol is HTTPS.

Answer: (B)



Q11.

Solution**Concept:**

Binary Search Tree (BST) is a hierarchical data structure where each node has at most two children. The defining property of a BST is that for any given node, all elements in the left subtree are smaller than the node, and all elements in the right subtree are larger. This structure is specifically designed to optimize searching operations by following a principle similar to the binary search algorithm used on sorted arrays.

Solution:

1. To search for a value in a BST, we start at the root. If the target value is smaller than the root, we move to the left child; if larger, we move to the right child. 2. In a **balanced** BST (such as an AVL tree or Red-Black tree), the height of the tree is maintained at approximately $\log_2(n)$, where n is the total number of nodes. 3. Since each comparison allows us to discard half of the remaining nodes (or one entire subtree), the number of steps required to find an element (or determine its absence) is proportional to the height of the tree. 4. Therefore, the time complexity for searching in a balanced BST is $O(\log n)$. 5. It is important to note that if the tree is **unbalanced** (skewed), the complexity could degrade to $O(n)$, but the question specifically mentions a balanced tree.

Final Answer: The time complexity is $O(\log n)$.

Answer: (C)

Q12.

Solution**Concept:**

The SQL 'GROUP BY' clause is used to arrange raw data into identical groups based on one or more columns. Once groups are formed, aggregate functions (like 'COUNT', 'SUM', 'AVG') are often applied. However, the standard 'WHERE' clause cannot be used to filter these aggregated results because 'WHERE' acts on individual rows before grouping occurs.

Solution:

1. The 'HAVING' clause was specifically introduced in SQL to provide a way to filter the results of a 'GROUP BY' operation. 2. The sequence of execution in a SQL query is: 'FROM' → 'WHERE' (row filtering) → 'GROUP BY' (grouping) → 'HAVING' (group filtering) → 'SELECT' → 'ORDER BY'. 3. For example, if you want to find departments where the average salary is greater than ₹ 50,000, you would use: 'GROUP BY Department HAVING AVG(Salary) > 50000'. 4. 'WHERE' is used for non-aggregated columns, 'ORDER BY' is for sorting, and 'LIMIT' is for restricting the number of rows returned. 5. Therefore, the 'HAVING' clause is the correct tool for filtering groups.

Final Answer: The clause is HAVING.

Answer: (B)



Q13.

Solution**Concept:**

In Python, the syntax requires that code blocks (like those following 'if', 'for', 'def', or 'class') cannot be empty. However, there are many scenarios in software development where a programmer might want to define the structure of a function or a loop without implementing the logic immediately.

Solution:

1. The 'pass' statement is a null operation; nothing happens when it executes. 2. It serves as a **syntactic placeholder**. If you leave a block empty, Python will raise an 'IndentationError'. Adding 'pass' prevents this error. 3. It differs from 'break', which exits a loop, and 'continue', which jumps to the next iteration of a loop. 4. It is also different from a comment; while comments are ignored by the interpreter, 'pass' is a valid statement that is processed but results in no action. 5. This is particularly useful during the "Top-Down" design of a program where you define functions but fill in the code later.

Final Answer: It serves as a placeholder for future code.

Answer: (C)

Q14.

Solution**Concept:**

The OSI (Open Systems Interconnection) model is a conceptual framework that standardizes the functions of a telecommunication or computing system into seven abstraction layers. Each layer serves the layer above it and is served by the layer below it.

Solution:

1. The **Data Link Layer** (Layer 2) is responsible for node-to-node delivery and uses MAC addresses within a local network. 2. The **Network Layer** (Layer 3) is responsible for host-to-host delivery. Its primary job is **Routing**, which involves determining the best physical path for data to travel from the source to the destination across multiple networks. 3. Devices like Routers operate at this layer and use Logical Addresses (IP addresses) to make routing decisions. 4. The **Transport Layer** (Layer 4) handles end-to-end communication and error recovery, while the **Application Layer** (Layer 7) provides the interface for end-user software. 5. Consequently, routing is the signature function of the Network Layer.

Final Answer: The layer is the Network Layer.

Answer: (B)



Q15.

Solution**Concept:**

Efficiency in sorting algorithms is measured by time complexity. While simple algorithms like Bubble or Insertion sort are easy to implement, their $O(n^2)$ complexity makes them unsuitable for large-scale data. Advanced algorithms use a "Divide and Conquer" strategy to achieve better performance.

Solution:

1. **Bubble Sort** and **Insertion Sort** have a worst-case complexity of $O(n^2)$, which grows quadratically with data size. 2. **Merge Sort** works by recursively splitting the list into halves until each sublist contains a single element, and then merging them back in a sorted manner. 3. The key advantage of Merge Sort is its **stability** and its guaranteed time complexity of $O(n \log n)$ in all cases (best, average, and worst). 4. While **Quick Sort** is often faster in practice for average cases, its worst-case is $O(n^2)$. 5. Because Merge Sort maintains $O(n \log n)$ even in the worst-case scenario, it is often preferred for very large datasets where performance consistency is critical.

Final Answer: The algorithm is Merge Sort.

Answer: (C)

Q16.

Solution**Concept:**

SQL commands are categorized into functional groups. Data Definition Language (DDL) consists of commands used to define or modify the database schema (structure), such as creating tables, altering their structure, or deleting the structure entirely. Unlike Data Manipulation Language (DML), DDL changes are generally auto-committed and affect the structure rather than the data content itself.

Solution:

1. 'UPDATE', 'SELECT', and 'INSERT' are part of Data Manipulation Language (DML) or Data Query Language (DQL). They are used to retrieve, add, or modify the actual data rows stored within the tables. 2. 'TRUNCATE' is a unique command. While it deletes data, it is classified as a DDL command because it operates by deallocating the data pages used by the table. It is essentially a faster, non-logged way to clear a table compared to 'DELETE'. 3. Other common DDL commands include 'CREATE', 'ALTER', and 'DROP'. 4. Therefore, 'TRUNCATE' is the correct DDL command among the options provided.

Final Answer: The command is TRUNCATE.

Answer: (B)



Q17.

Solution**Concept:**

In Python, lists are mutable sequences. There are multiple ways to add elements to a list, and understanding the difference between 'append()' and 'extend()' is crucial. While 'append()' adds its argument as a single element, 'extend()' iterates over its argument (which must be an iterable) and adds each element to the list individually.

Solution:

1. Initially, we have a list 'L = [1, 2, 3]'. The length of this list is 3. 2. The code then calls 'L.extend([4, 5])'. 3. The 'extend' method takes the list '[4, 5]', unpacks its elements, and appends them one by one to the end of 'L'. 4. After this operation, the list 'L' becomes '[1, 2, 3, 4, 5]'. 5. Now, calling 'len(L)' will count the total number of items currently in the list. 6. There are 5 items in the list. 7. Note: If we had used 'L.append([4, 5])', the list would have become '[1, 2, 3, [4, 5]]' with a length of 4.

Final Answer: The length is 5.

Answer: (C)

Q18.

Solution**Concept:**

Network topology refers to the physical or logical arrangement of nodes and connections in a network. Different topologies use different methods for signal distribution and media access. A "multipoint" connection implies that multiple devices share the same medium simultaneously.

Solution:

1. In a **Star** topology, each node has a dedicated point-to-point connection to a central hub/switch. 2. In a **Ring** topology, each device is connected to exactly two neighbors, forming a closed loop (point-to-point connections). 3. In a **Mesh** topology, nodes have multiple point-to-point connections to other nodes. 4. A **Bus** topology uses a single backbone cable (the "bus") as a shared communication medium. All nodes are connected to this single cable via "drops" and "taps." Because all devices use the same wire, it is a multipoint connection. 5. Signals transmitted by any node travel the length of the bus and are available to all other nodes.

Final Answer: The topology is Bus.

Answer: (C)



Q19.

Solution**Concept:**

Recursion is a programming technique where a function calls itself to solve smaller instances of the same problem. To keep track of these nested function calls, the computer needs a way to store the state of each call (including local variables and return addresses) so it can resume correctly once a sub-call finishes.

Solution:

1. When a function is called, its state is "pushed" onto a memory region called the **Call Stack**.
2. If that function calls itself again, the new state is pushed on top of the previous one.
3. This follows the LIFO (Last-In-First-Out) principle: the most recent function call must complete and be "popped" before the previous one can continue.
4. Data structures like Queues (FIFO) or Linked Lists (general storage) do not inherently provide the automatic "return to previous state" mechanism required for nested execution.
5. Therefore, the **Stack** is the fundamental data structure used by compilers and interpreters to manage recursion.

Final Answer: The data structure is Stack.

Answer: (B)

Q20.

Solution**Concept:**

Exception handling in Python allows a program to deal with unexpected errors (exceptions) gracefully without crashing. Python uses a specific set of keywords to define blocks that monitor, catch, and clean up after errors.

Solution:

- (a) The `try` block contains the code that might raise an error.
- (b) The `except` block contains the code that handles the error if one occurs.
- (c) The `finally` block contains code that runs regardless of whether an error occurred.
- (d) `else` runs if no exception occurred.
- (e) The keyword `catch` is used in languages like C++, Java, and C#.
- (f) In Python, the equivalent keyword is `except`.

Final Answer: The invalid keyword is `catch`.

Answer: (C)



Q21.

Solution**Concept:**

In SQL, the 'GROUP BY' clause is a powerful tool used to collect data across multiple records and group the results by one or more columns. It is fundamentally different from a standard sort or filter because it changes the level of granularity of the output. Instead of seeing individual rows, the database engine collapses rows with the same values in the grouped columns into a single summary row.

Solution:

1. The primary purpose of 'GROUP BY' is to arrange identical data into groups. This is almost always done in conjunction with aggregate functions such as 'SUM()', 'AVG()', 'COUNT()', 'MAX()', or 'MIN()'. 2. For example, if you have a table of "Sales" and you want to see the total revenue per "Region", you would use 'GROUP BY Region'. The database identifies all unique regions and sums the sales for each. 3. It does not perform sorting; that is the job of the 'ORDER BY' clause. 4. It does not filter raw rows; that is the job of the 'WHERE' clause. 5. It does not perform table joins; that is handled by the 'JOIN' or 'FROM' clauses. 6. In a CUET-level context, it is critical to remember that any column in the 'SELECT' list that is not part of an aggregate function **must** be included in the 'GROUP BY' clause.

Final Answer: The purpose is to arrange identical data into groups.

Answer: (C)

Q22.

Solution**Concept:**

Networking requires unique identification of devices at different layers of the OSI model. There are two primary types of addresses: Logical addresses (like IP addresses) which can change based on the network location, and Physical addresses (like MAC addresses) which are hardcoded into the hardware.

Solution:

1. A MAC (Media Access Control) address is a unique identifier assigned to a Network Interface Controller (NIC). 2. It is often referred to as the **Physical Address** because it is burned into the hardware during the manufacturing process. 3. While an IP address (Logical Address) helps in routing data across different networks (Global identification), the MAC address is used to deliver data to the specific device within a local segment (Local identification). 4. MAC addresses are 48 bits (6 bytes) long and are typically represented in hexadecimal format (e.g., 00:1A:2B:3C:4D:5E). 5. In the Data Link Layer (Layer 2) of the OSI model, switches use the physical address to forward frames to the correct port.

Final Answer: A MAC address is also known as the Physical address.

Answer: (B)



Q23.

Solution**Concept:**

Sorting is the process of arranging elements in a specific order (ascending or descending). Insertion Sort is a simple, intuitive algorithm that builds the final sorted array one item at a time. It is much less efficient on large lists than more advanced algorithms like Quick Sort or Merge Sort, but it performs well for small datasets or lists that are already partially sorted.

Solution:

1. The algorithm works by conceptually splitting the list into a "sorted" part and an "unsorted" part. Initially, the sorted part contains only the first element. 2. In each step, the algorithm takes the next element from the unsorted part and compares it with the elements in the sorted part. 3. It "shifts" the sorted elements to the right to create a "gap" or "hole" where the new element fits in its correct relative order. 4. This mechanism of "taking an element and inserting it into the correct position" is the defining characteristic of **Insertion Sort**. 5. This is distinct from Selection Sort (which looks for the minimum) or Bubble Sort (which swaps adjacent pairs).

Final Answer: The algorithm is Insertion Sort.

Answer: (C)

Q24.

Solution**Concept:**

In Python file handling, the file pointer acts like a cursor that tracks where the next read or write operation will occur. To effectively manipulate files, especially binary files or large text files, developers need to be able to both move this pointer and query its current location.

Solution:

1. Python provides two primary methods for pointer manipulation: 'seek()' and 'tell()'. 2. The 'seek(offset, whence)' method is used to **move** the pointer to a specific location. 3. The 'tell()' method takes no arguments and **returns** the current position of the file pointer as an integer representing the number of bytes from the beginning of the file. 4. For example, if you open a file and read 10 characters, calling 'f.tell()' will return 10. 5. This is extremely useful when you need to record a specific location in a file to return to it later after performing other operations. 6. Option (C) and (D) are not valid Python file object methods.

Final Answer: The function is tell().

Answer: (B)



Q25.

Solution**Concept:**

For communication to occur successfully between two hosts on an Ethernet network, the sending host must know both the IP address (for logical routing) and the MAC address (for physical delivery on the local link) of the destination. However, users and applications usually only know the IP address.

Solution:

1. **ARP** (Address Resolution Protocol) is the protocol used to discover the hardware address (MAC) of a node derived from its known IP address. 2. When a host wants to send data, it checks its "ARP Cache". If the MAC address is not found, it sends an "ARP Request" broadcast to the entire local network asking, "Who has this IP?". 3. The device with that IP sends an "ARP Reply" containing its MAC address. 4. **DNS** (Domain Name System) maps domain names (google.com) to IP addresses. 5. **DHCP** (Dynamic Host Configuration Protocol) automatically assigns IP addresses to devices. 6. **ICMP** (Internet Control Message Protocol) is used for error reporting (like Ping). 7. Therefore, ARP is the correct protocol for mapping IP to MAC.

Final Answer: The protocol is ARP.

Answer: (B)

Q26.

Solution**Concept:**

The 'ROUND()' function in SQL is a mathematical function used to reduce a numeric value to a specific number of decimal places. This is essential in financial applications where currency must be limited to two decimal points or in scientific data where significant figures are required. The function follows standard rounding rules: if the digit following the rounding position is 5 or greater, the rounding digit is incremented.

Solution:

1. The provided SQL expression is 'SELECT ROUND(15.678, 2);'. 2. Here, the first argument '15.678' is the number to be rounded, and the second argument '2' specifies the decimal places. 3. To round to 2 decimal places, we look at the third decimal digit, which is '8'. 4. Since '8' is greater than or equal to '5', we must round up the second decimal digit. 5. The second decimal digit is '7', so incrementing it by 1 makes it '8'. 6. Therefore, the resulting value is '15.68'. 7. If the expression had been 'ROUND(15.678, 1)', the result would have been '15.7'. 8. If it had been 'ROUND(15.678, 0)', the result would have been '16'. 9. The database ensures that the mathematical precision is maintained while adhering to the specified formatting constraints.

Final Answer: The result is 15.68.

Answer: (C)



Q27.

Solution**Concept:**

A Queue is a linear data structure that operates on the First-In-First-Out (FIFO) principle. When implemented using a fixed-size array, two pointers—'front' and 'rear'—are typically used to track the positions where elements are removed and added, respectively. Managing these pointers correctly is vital to identifying the boundary conditions of the data structure.

Solution:

1. In a standard linear queue implementation using an array of size 'N', the valid indices range from '0' to 'N-1'. 2. The 'rear' pointer tracks the index of the last element inserted into the queue. 3. When the queue is empty, 'rear' is usually initialized to '-1'. 4. As elements are added (enqueued), the 'rear' pointer increments. 5. If the 'rear' pointer reaches the final possible index, which is 'size - 1', it indicates that there is no more space at the back of the array to accommodate new elements. 6. In this state, the queue is considered ****Full****. 7. Attempting to add another element when the queue is full results in a condition called ****Overflow****. 8. It is important to distinguish this from a "Circular Queue," where the 'rear' could wrap around to the beginning if space is available at the front. However, in a simple linear queue, 'rear == size - 1' signifies the queue is full.

Final Answer: The queue is Full.

Answer: (B)



Q28.

Solution**Concept:**

Switching is a method used to establish a path for data transfer across a large network. The two most common types are Circuit Switching (historically used in telephone networks) and Packet Switching (used in the modern Internet). Each method has a different approach to how the transmission medium is utilized and how data is formatted for travel.

Solution:

1. In **Circuit Switching**, a dedicated physical path is established between the sender and receiver for the duration of the communication. Data flows as a continuous stream. 2. In **Packet Switching**, data is not sent as a continuous stream. Instead, the message is divided into smaller, manageable units called **packets**. 3. Each packet contains a portion of the data along with a header that includes source and destination addresses, sequence numbers, and error-checking information. 4. These packets travel independently through the network, potentially taking different routes, and are reassembled in the correct order at the destination. 5. This method is highly efficient because it allows multiple users to share the same network resources simultaneously without requiring a dedicated line. 6. **Message Switching** involves sending the entire message as one unit from node to node (Store-and-Forward), but it does not break it into small pieces like Packet Switching does.

Final Answer: The type of switching is Packet Switching.

Answer: (B)

Q29.

Solution**Concept:**

Serialization is the process of converting a complex data structure or object state into a format that can be stored (like a file or memory buffer) or transmitted (across a network) and reconstructed later. In Python, the 'pickle' module is the standard library used for this purpose, allowing users to save almost any Python object.

Solution:

1. The 'pickle' module provides two main functions for serialization: 'dump()' and 'dumps()'. 2. The 'pickle.dump(object, file)' method is specifically used to convert a Python object (like a list, dictionary, or custom class instance) into a byte stream and write that stream directly into an open file object. 3. This process is called **Pickling** or **Serialization**. 4. The counterpart to this is 'pickle.load(file)', which reads the byte stream from the file and reconstructs the original Python object, a process known as **Unpickling** or **Deserialization**. 5. Using 'pickle' is highly convenient for preserving the state of a program's data between different sessions. 6. It is distinct from text-based storage (like writing strings) because it preserves the actual Python data types and structures without manual conversion.

Final Answer: It serializes an object to a file.

Answer: (B)



Q30.

Solution**Concept:**

In the Relational Model, a table (relation) consists of rows (tuples) and columns (attributes). To maintain data integrity and allow for efficient data retrieval, every row in a table must be unique. Constraints are used to enforce these rules.

Solution:

1. A **Primary Key** is a specific attribute or a set of attributes that uniquely identifies each tuple in a relation. 2. By definition, a Primary Key must contain unique values, and it cannot contain 'NULL' values. This ensures that every record can be precisely located and distinguished from others. 3. A **Candidate Key** is any attribute (or set) that *could* serve as a primary key because it is unique. One of these is chosen to be the Primary Key. 4. An **Alternate Key** is a candidate key that was not chosen to be the primary key. 5. A **Foreign Key** is an attribute in one table that refers to the primary key of another table, used to establish a relationship between the two. 6. While multiple keys might provide uniqueness, the Primary Key is the formal, unique identifier assigned by the database designer to represent the row.

Final Answer: The attribute is the Primary Key.

Answer: (B)

Q31.

Solution**Concept:**

SQL provides several built-in functions to handle date and time values. In a database, tracking the exact moment a record was created or modified is a fundamental requirement. While some functions return only the date or only the time, certain functions provide a combined timestamp.

Solution:

1. The 'CURDATE()' function returns only the current date in 'YYYY-MM-DD' format, without any time component. 2. The 'NOW()' function is the most commonly used function in MySQL and similar dialects to retrieve the current date and time together. It returns the value in 'YYYY-MM-DD HH:MM:SS' format. 3. 'SYSDATE()' is very similar to 'NOW()', but there is a subtle difference: 'NOW()' returns the time at which the query *started* executing, whereas 'SYSDATE()' returns the actual time at which the function *itself* executes. For most standard exam purposes, 'NOW()' is the primary answer for current date and time. 4. 'DATE()' is typically used to extract the date part from a datetime expression. 5. Given the specific format 'YYYY-MM-DD HH:MM:SS', 'NOW()' is the standard and correct choice.

Final Answer: The function is NOW().

Answer: (B)



Q32.

Solution**Concept:**

In computer science, expressions are written in different notations: Infix (operators between operands), Postfix (operators after operands), and Prefix (operators before operands). Postfix notation (also known as Reverse Polish Notation) is preferred by compilers because it eliminates the need for parentheses and respects operator precedence without ambiguity.

Solution:

1. The original infix expression is: $(A + B) * (C - D)$. 2. We follow the order of operations, starting with the expressions inside the parentheses. 3. Convert the first parenthetical expression: $(A + B)$ becomes $AB+$. 4. Convert the second parenthetical expression: $(C - D)$ becomes $CD-$. 5. Now the expression looks like: $[AB+] * [CD-]$. 6. Treat each bracketed term as a single operand. The operator is $*$. 7. Move the operator to the end of the two operands: $AB + CD - *$. 8. Removing the spaces, we get $AB + CD - *$. 9. This matches option (A).

Final Answer: The postfix expression is $AB+CD-*$.

Answer: (A)

Q33.

Solution**Concept:**

Python's runtime environment raises specific exceptions when it encounters errors during execution. Understanding these exceptions helps in debugging and writing robust code. A "name" in Python refers to variables, functions, modules, or classes.

Solution:

1. A 'NameError' is raised when the interpreter tries to execute code that references a name that has not been defined in the current local or global scope. 2. For example, if you type 'print(x)' but have never assigned a value to 'x', Python raises a 'NameError'. 3. 'ImportError' occurs when an 'import' statement fails to find the module or a name within the module. 4. 'ValueError' occurs when a function receives an argument of the right type but an inappropriate value (e.g., 'int("abc")'). 5. 'TypeError' occurs when an operation or function is applied to an object of an inappropriate type (e.g., adding a string to an integer). 6. Since the question specifically mentions a name not being found, 'NameError' is the correct exception.

Final Answer: The exception is NameError.

Answer: (C)



Q34.

Solution**Concept:**

Networking hardware is classified based on the layer of the OSI model at which it functions. The Physical Layer (Layer 1) deals with the actual electrical, optical, or radio signals. Devices at this layer do not "read" the data packets; they simply handle bits.

Solution:

1. A **Hub** is a non-intelligent device that operates at the Physical Layer. 2. When a hub receives a signal on one of its ports, it simply regenerates the signal and broadcasts (repeats) it to all other ports. Because it has multiple ports and performs the same function as a repeater, it is called a **multiport repeater**. 3. A **Switch** operates at the Data Link Layer (Layer 2) and forwards data only to the specific destination port based on MAC addresses. 4. A **Router** operates at the Network Layer (Layer 3). 5. A **Bridge** is a Layer 2 device used to connect two network segments. 6. Thus, the Hub is the only device in the list that fits the definition of a Layer 1 multiport repeater.

Final Answer: The device is a Hub.

Answer: (B)

Q35.

Solution**Concept:**

Relational Algebra is a formal language used to model the operations performed on relational databases. It uses Greek symbols to represent different operations. Understanding these symbols is fundamental to understanding how SQL queries are processed by a database engine.

Solution:

1. The σ (sigma) operator represents the **Selection** operation. It is used to select a subset of tuples (rows) from a relation that satisfy a specific condition. It acts as a horizontal filter (similar to the 'WHERE' clause in SQL). 2. The π (pi) operator represents **Projection**, which is used to select specific columns (attributes) from a relation, acting as a vertical filter. 3. The \bowtie symbol represents a **Join** operation. 4. The \cup symbol represents **Union**. 5. Therefore, sigma (σ) is strictly for Selection. For example, $\sigma_{age>20}(Students)$ would return all rows from the Students table where the age is greater than 20.

Final Answer: The operator is used for Selection.

Answer: (B)



Q36.

Solution**Concept:**

Sorting algorithms vary significantly in their performance characteristics. While some algorithms like Insertion Sort or Bubble Sort can achieve $O(n)$ time complexity in the best-case scenario (when the list is already sorted), others maintain a fixed number of operations regardless of the initial order of the data. Selection Sort is a prime example of an algorithm that does not adapt to the existing order of elements.

Solution:

1. **Selection Sort** operates by repeatedly finding the minimum element from the unsorted portion of the list and swapping it with the first element of that unsorted portion. 2. The algorithm uses two nested loops. The outer loop runs $n - 1$ times to place each element in its correct position. The inner loop always scans the remaining unsorted elements to find the absolute minimum. 3. Because the inner loop always performs the same number of comparisons to find the minimum, regardless of whether the elements were already in order or not, the total number of comparisons is always $\frac{n(n-1)}{2}$. 4. Mathematically, this results in a time complexity of $O(n^2)$ for the best, average, and worst cases. 5. In contrast, **Bubble Sort** and **Insertion Sort** can be optimized to finish in $O(n)$ if the list is sorted, and **Quick Sort** varies between $O(n \log n)$ and $O(n^2)$. 6. Thus, Selection Sort is unique for its consistent $O(n^2)$ behavior across all scenarios.

Final Answer: The algorithm is Selection Sort.

Answer: (C)

Q37.

Solution**Concept:**

Network communication relies on multiple layers of addressing. The lowest level of addressing for hardware devices on a local area network (LAN) is the MAC address. This address is essentially the "fingerprint" of the network interface card (NIC) and is used for frame delivery within the Data Link Layer.

Solution:

1. The acronym **MAC** stands for **Media Access Control**. 2. It refers to the sub-layer of the Data Link Layer (Layer 2) in the OSI model that provides the hardware addressing and channel access control mechanisms. 3. A MAC address is a 48-bit (6-byte) identifier. The first 3 bytes are the Organizationally Unique Identifier (OUI) which identifies the manufacturer, and the last 3 bytes are the device's specific serial number. 4. "Media" refers to the physical transmission medium (like copper wire, fiber, or air), and "Access Control" refers to the protocols that manage how multiple devices share that medium without causing collisions. 5. Other options like "Memory Access Control" or "Main Access Channel" are incorrect technical terms in the context of networking hardware identifiers.

Final Answer: MAC stands for Media Access Control.

Answer: (A)



Q38.

Solution**Concept:**

The `seek()` method in Python is used to change the position of the file handle (pointer) to a specific byte offset. This allows random access within a file. The method signature is:

```
file_object.seek(offset, whence)
```

Solution:

- (a) The `whence` argument defines the reference point from which the offset is calculated.
- (b) Standard values are:
 - **0**: Beginning of the file
 - **1**: Current file pointer position
 - **2**: End of the file
- (c) Therefore, `seek(10, 1)` moves the pointer 10 bytes forward from the current position.
- (d) In Python 3, relative seeking using 1 is mainly supported in binary mode.

Final Answer: The value 1 represents the current position of the file pointer.

Answer: (B)

Q39.

Solution**Concept:**

Managing large datasets in SQL often requires clearing out all data from a table. There are three primary commands that can remove data: 'DELETE', 'DROP', and 'TRUNCATE'. While they might seem similar, they have very different impacts on the database structure, performance, and transaction logs.

Solution:

1. The 'DELETE' command removes rows one by one and logs each deletion. It allows for a 'WHERE' clause to filter specific rows.
2. The 'DROP' command removes the entire table structure along with all the data, indexes, and permissions. The table ceases to exist in the database.
3. The 'TRUNCATE' command is a Data Definition Language (DDL) operation. It removes all rows from a table by deallocating the data pages used to store the table data.
4. Crucially, 'TRUNCATE' **preserves the table structure** (the columns, constraints, and indexes remain intact for future use).
5. It is much faster than 'DELETE' for large tables because it does not generate individual row-level logs.
6. Therefore, for removing all records while keeping the structure and resetting identity seeds, 'TRUNCATE' is the correct keyword.

Final Answer: The keyword is TRUNCATE.

Answer: (C)



Q40.

Solution**Concept:**

Network security involves protecting data and resources from unauthorized access and threats. A firewall is a fundamental security component that acts as a barrier between a trusted internal network and an untrusted external network (like the Internet). It monitors and controls incoming and outgoing network traffic based on predetermined security rules.

Solution:

1. The primary mechanism a firewall uses is **Packet Filtering**. 2. It examines the header of each data packet (containing source IP, destination IP, port numbers, and protocol type) and compares it against a set of Access Control Lists (ACLs). 3. If the packet meets the criteria defined in the rules, it is allowed to pass; otherwise, it is dropped or rejected. 4. While advanced modern firewalls (Next-Generation Firewalls) can perform encryption or virus scanning, the "primary" and fundamental function of a standard firewall as defined in the OSI model is filtering traffic. 5. It does not typically "remove" viruses from a system (that is the job of Antivirus software), nor is its primary purpose data encryption (that is the job of VPNs or SSL/TLS).

Final Answer: The primary function is Packet Filtering.

Answer: (B)

Q41.

Solution**Concept:**

A Queue is a linear data structure that manages data based on the **First-In-First-Out (FIFO)** principle. This means the first element added to the queue will be the first one to be removed. It is analogous to a real-world line of people waiting for a service; the person who arrives first is served first. This structure is essential in computer science for tasks such as printer spooling, CPU task scheduling, and handling asynchronous data transfers.

Solution:

1. In a **Stack**, the principle followed is **LIFO** (Last-In-First-Out), where the most recently added item is the first to be removed. 2. In a **Queue**, two primary operations are performed: 'Enqueue' (adding an element to the rear) and 'Dequeue' (removing an element from the front). 3. Because removal always happens from the front, the element that has been in the queue the longest (the "First" one "In") is the one that is removed (the "First" one "Out"). 4. **Trees** and **Graphs** are non-linear data structures and do not strictly follow the FIFO principle in their basic definitions, although specific traversal methods like Breadth-First Search (BFS) utilize queues. 5. Therefore, the Queue is the definitive FIFO data structure.

Final Answer: The data structure is Queue.

Answer: (B)



Q42.

Solution**Concept:**

Pattern matching in SQL is performed using the 'LIKE' operator, which allows for flexible searching within string columns. To facilitate this, SQL provides two primary "wildcard" characters that act as placeholders for unknown characters in a search string. Understanding the difference between these wildcards is fundamental for writing precise queries.

Solution:

1. The **Underscore ()** wildcard represents exactly **one** single character. It is used when you know the specific length of the string but not certain characters. 2. The **Percentage (%)** wildcard represents an unspecified number of characters. 3. For example, 'LIKE 'A%. 4. The **Asterisk (*)** and **Question Mark (?)** are common wildcards in file systems (like Windows or Linux shells), but they are not the standard wildcards used in the SQL 'LIKE' clause. 5. Consequently, the Percentage sign is the correct character for representing an unspecified number of characters.

Final Answer: The wildcard character is %.

Answer: (C)

Q43.

Solution**Concept:**

The efficiency of Bubble Sort, like many sorting algorithms, is highly dependent on the initial state of the input data. Bubble Sort works by repeatedly stepping through the list, comparing adjacent elements, and swapping them if they are in the wrong order. This process repeats until the entire list is sorted. However, a standard implementation can be optimized to recognize if a list is already sorted.

Solution:

1. In a standard, non-optimized Bubble Sort, the algorithm always runs two nested loops, resulting in $O(n^2)$ regardless of the data. 2. However, in an **optimized** version (which is the standard expectation in academic exams), a flag is used to track if any swaps occurred during a pass. 3. If the list is **already sorted** (the Best-Case scenario), the algorithm will complete its first pass through the n elements, realize that zero swaps were made, and immediately terminate. 4. This first pass involves $n - 1$ comparisons. 5. In Big O notation, this linear performance is expressed as $O(n)$. 6. If the list is not sorted, the complexity remains $O(n^2)$ for average and worst cases.

Final Answer: The best-case time complexity is $O(n)$.

Answer: (A)



Q44.

Solution**Concept:**

Email communication involves different protocols for different stages of the process. Sending an email requires one protocol, while retrieving or downloading that email from a server to a local device (like Outlook or a mobile app) requires another. The two primary protocols for retrieval are POP3 and IMAP.

Solution:

1. **SMTP** (Simple Mail Transfer Protocol) is used exclusively for **sending** emails from a client to a server or between servers. 2. **POP3** (Post Office Protocol version 3) is designed to **download** email messages from a mail server to a local client. Historically, POP3 would move the email from the server to the local device and then delete it from the server, though modern settings can leave a copy. 3. **IMAP** (Internet Message Access Protocol) is another retrieval protocol, but it synchronizes the email across devices rather than just downloading it. 4. **HTTP** is for web browsing, and **FTP** is for general file transfers. 5. Given the specific context of downloading messages to a local client as a primary function, POP3 is the traditional and correct answer.

Final Answer: The protocol is POP3.

Answer: (B)

Q45.

Solution**Concept:**

Python provides several methods for reading data from files. Some methods read the entire file, while others operate line-by-line using newline characters (`\n`).

Solution:

- (a) The `read()` method reads the complete file as a single string.
- (b) The `readline()` method reads only one line at a time.
- (c) The `readlines()` method reads all lines and returns them as a list of strings.
- (d) Example:

```
['Line 1\n', 'Line 2\n', 'Line 3\n']
```

- (e) Each element in the list corresponds to one line in the file.
- (f) `readall()` is not a standard Python file method.

Final Answer: The method is `readlines()`.

Answer: (C)



Q46.

Solution**Concept:**

In a Relational Database Management System (RDBMS), constraints are rules applied to table columns to ensure the accuracy, reliability, and integrity of the data. Two of the most critical requirements for an identifier are that it must be unique (no two records share the same value) and it must be mandatory (the value cannot be empty). While separate constraints exist for these, SQL provides a specific constraint that combines both behaviors into a single definition.

Solution:

1. The 'UNIQUE' constraint ensures all values in a column are different, but it generally allows for one or more 'NULL' values depending on the SQL dialect. 2. The 'NOT NULL' constraint ensures that a column cannot have an empty value, but it does not prevent duplicate entries. 3. The 'PRIMARY KEY' constraint is a combination of both: it automatically enforces that the column contains only **unique** values and that **no null values** are permitted. 4. Every table should ideally have one Primary Key to uniquely identify its rows. 5. 'FOREIGN KEY' is used for cross-referencing tables, and 'CHECK' is used to limit the range of values (e.g., Age > 18). 6. Because the question specifies both uniqueness and the "not null" requirement, Primary Key is the correct answer.

Final Answer: The constraint is PRIMARY KEY.

Answer: (B)

Q47.

Solution**Concept:**

Network topologies define how devices are physically or logically connected to one another. The Ring topology is a specific configuration where data travels in a circular path. This setup is often used in specialized local area networks (LANs) like Token Ring networks, where a digital "token" is passed around the circle to manage data transmission and prevent collisions.

Solution:

1. In a **Ring** topology, each node (computer or device) is connected to exactly two other nodes—one "upstream" and one "downstream." 2. This creates a closed loop or continuous pathway for signals. Data typically travels in one direction (unidirectional), though dual-ring configurations exist for redundancy. 3. In a **Star** topology, every node connects to one central hub (not two neighbors). 4. In a **Bus** topology, all nodes connect to one single main cable. 5. In a **Mesh** topology, nodes may connect to many others, often more than two, to provide multiple paths for data. 6. The description of a "single continuous pathway" with each node having exactly two connections is the fundamental definition of a Ring network.

Final Answer: The topology is Ring.

Answer: (C)



Q48.

Solution**Concept:**

Exception handling is a critical part of modern programming that prevents a program from "crashing" when it encounters an error. Instead of the execution stopping abruptly with an error message, the programmer can "catch" the error and provide an alternative set of instructions. Python's approach to this is known as the "try-except" block, which is similar to "try-catch" in other languages.

Solution:

1. The 'try' block is the guarded section of code where you place instructions that might fail (e.g., dividing by zero or opening a missing file). 2. The 'except' block immediately follows the 'try' block. If an error occurs inside the 'try' block, Python stops executing the try code and jumps directly to the 'except' block. 3. 'if ... else' is used for conditional logic based on data values, not for handling runtime system errors or crashes. 4. 'for ... in' is a loop structure used for iteration. 5. 'def ... return' is used for defining and exiting functions. 6. Therefore, the 'try ... except' structure is the only one designed specifically for error and exception management.

Final Answer: The structure is try ... except.

Answer: (A)

Q49.

Solution**Concept:**

The Session Layer is the fifth layer of the OSI (Open Systems Interconnection) model. While lower layers (like Transport and Network) handle the movement of data bits and packets, the Session Layer is concerned with the "dialogue" between two applications. It acts as the coordinator of the communication session, ensuring that the exchange is organized and synchronized.

Solution:

1. The **Session Layer** (Layer 5) is responsible for the "three-way" management of a connection: establishment (starting), management (keeping it alive and synchronized), and termination (gracefully closing). 2. It also provides "check-pointing" or "synchronization points." If a large file transfer is interrupted, the session layer can resume from the last checkpoint rather than starting from the beginning. 3. The **Presentation Layer** (Layer 6) handles encryption and formatting. 4. The **Transport Layer** (Layer 4) handles end-to-end error recovery and flow control. 5. The **Data Link Layer** (Layer 2) handles framing and physical addressing. 6. The specific responsibility for "managing and terminating connections between applications" lies solely with the Session Layer.

Final Answer: The layer is the Session Layer.

Answer: (B)



Q50.

Solution**Concept:**

Sorting is a common requirement in SQL queries, handled by the 'ORDER BY' clause. When a user requests data to be sorted by a specific column, they may want it in increasing order (A-Z, 1-10) or decreasing order (Z-A, 10-1). Database systems are designed with a "default" behavior to make queries shorter and more readable when the standard order is desired.

Solution:

1. The 'ORDER BY' clause can take two modifiers: 'ASC' for Ascending and 'DESC' for Descending. 2. If a user writes 'SELECT * FROM Students ORDER BY Name;' without specifying the direction, the SQL engine automatically applies **Ascending** order. 3. Ascending order means alphabetical order for text (A first), numerical order for numbers (smallest first), and chronological order for dates (earliest first). 4. If a user wants the opposite, they must explicitly state 'DESC'. 5. There is no "random" default, and every SQL implementation (MySQL, PostgreSQL, Oracle, SQL Server) follows the Ascending default rule.

Final Answer: The default order is Ascending.

Answer: (B)



Answer Key

Q	Ans	Q	Ans	Q	Ans	Q	Ans	Q	Ans
1	C	2	B	3	B	4	C	5	B
6	C	7	C	8	C	9	C	10	B
11	C	12	B	13	C	14	B	15	C
16	B	17	C	18	C	19	B	20	C
21	C	22	B	23	C	24	B	25	B
26	C	27	B	28	B	29	B	30	B
31	B	32	A	33	C	34	B	35	B
36	C	37	A	38	B	39	C	40	B
41	B	42	C	43	A	44	B	45	C
46	B	47	C	48	A	49	B	50	B

