CUET-UG Computer Science Sample Paper-10

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. A database administrator wants to ensure that the value in the 'Salary' column of the 'Employees' table never falls below ₹ 25,000. Which SQL constraint should be applied?

(A) PRIMARY KEY

(B) UNIQUE

(C) CHECK

(D) DEFAULT

Q2. A Python program opens a file using 'open("data.bin", "rb+")'. If the programmer wants to overwrite the 10th byte of the file without deleting the existing content, which sequence of commands is most appropriate?

(A) 'f.seek(10); f.write(b'X')'

(B) 'f.seek(9); f.write(b'X')'

(C) 'f.read(10); f.write(b'X')'

(D) 'f.seek(0, 10); f.write(b'X')'

Q3. In a stack-based evaluation of the postfix expression 8 2 / 5 2 $*$ +, what is the final value pushed onto the stack?

(A) 14

(B) 10

(C) 24

(D) 18

Q4. Which of the following sorting algorithms will consistently perform $O(n^2)$ comparisons regardless of whether the input list is already sorted or completely reversed?

(A) Bubble Sort

(B) Insertion Sort

(C) Selection Sort

(D) Quick Sort

Q5. In Relational Algebra, if you want to find the names of all students who have enrolled in 'Computer Science' but NOT in 'Mathematics', which operation is required?

(A) Union

(B) Intersection

(C) Set Difference

(D) Cartesian Product

Q6. A network uses a subnet mask of 255.255.255.192. How many usable host addresses are available in each subnet?

(A) 64

(B) 62

(C) 126

(D) 30

Q7. What will be the output of the following SQL command? 'SELECT LTRIM(RTRIM(' CUET 2026 ')); '

(A) 'CUET 2026'

(B) ' CUET 2026'

(C) 'CUET 2026 '

(D) 'CUET2026'

Q8. In Python, which exception is specifically raised when a dictionary is accessed using a key that does not exist?

(A) IndexError

(B) KeyError

(C) ValueError

(D) TypeError

Q9. Which protocol is responsible for resolving a known IP address into a physical MAC address on a local network?

(A) DNS

(B) DHCP

(C) ARP

(D) ICMP

Q10. While sorting a list $[50, 20, 40, 10, 30]$ using Bubble Sort, how many swaps are performed in the very first pass?

(A) 2

(B) 3

(C) 4

(D) 1

Q11. Which of the following SQL functions is used to return the position of the first occurrence of a substring within a string?

(A) LOCATE()

(B) POSITION()

(C) INSTR()

(D) All of the above

Q12. In a circular queue of size 6, if 'front = 4' and 'rear = 1', how many elements are currently present in the queue?

(A) 3

(B) 4

(C) 2

(D) 5

Q13. Which network topology is most prone to a single point of failure if the central controller/hub fails, but allows for easy addition of new nodes?

(A) Mesh

(B) Bus

(C) Star

(D) Ring

Q14. Consider the Python code: 'L = [1, [2, 3], 4]'. What will be the result of 'L[1][1] = 5; print(L)'?

(A) '[1, [2, 5], 4]'

(B) '[1, [5, 3], 4]'

(C) '[1, 5, 4]'

(D) 'Error'

Q15. In a database transaction, which property ensures that all sub-operations are treated as a single unit, such that either all succeed or none do?

(A) Consistency

(B) Isolation

(C) Durability

(D) Atomicity

Q16. What is the result of the SQL query: 'SELECT ROUND(45.926, 1);'

(A) 45.9

(B) 46.0

(C) 45.93

(D) 45

Q17. Which field in the IPv4 header is decremented by every router that the packet passes through?

(A) Version

(B) Total Length

(C) Time to Live (TTL)

(D) Header Checksum

Q18. In Python file handling, which method is used to obtain the current position of the file pointer?

(A) 'seek()'

(B) 'where()'

(C) 'tell()'

(D) 'position()'

Q19. Convert the infix expression '(A + B) * (C - D)' to its postfix equivalent.

(A) 'AB+CD-*'

(B) 'ABCD+-*'

(C) '+AB*-CD'

(D) 'AB+CD*-'

Q20. Which SQL clause is used to eliminate duplicate rows from the results of a SELECT statement?

(A) UNIQUE

(B) DISTINCT

(C) GROUP BY

(D) CHECK

Q21. Which sorting algorithm works by building a sorted array one item at a time, and is particularly efficient for small datasets or nearly sorted lists?

(A) Quick Sort

(B) Selection Sort

(C) Insertion Sort

(D) Merge Sort

Q22. A network device that works at the Data Link Layer and filters traffic based on MAC addresses is a:

(A) Hub

(B) Repeater

(C) Switch

(D) Router

Q23. In Python, what is the output of '3 ** 1 ** 3'?

(A) 27

(B) 1

(C) 3

(D) 9

Q24. Which SQL command is used to delete a specific column from an existing table?

(A) 'DELETE COLUMN'

(B) 'REMOVE COLUMN'

(C) 'ALTER TABLE ... DROP COLUMN'

(D) 'DROP TABLE ... COLUMN'

Q25. If a linear queue is implemented using an array of size 5, and we perform 5 enqueues followed by 2 dequeues, what is the current value of the ‘rear‘ index (starting from 0)?

(A) 4

(B) 2

(C) 5

(D) 3

Q26. In Relational Algebra, which operation is used to combine all rows from two relations that have the same schema, including duplicates in the result?

(A) Union

(B) Intersection

(C) Cartesian Product

(D) None of the above (Standard Union removes duplicates)

Q27. Which protocol is used by web browsers to securely exchange information with a web server using encryption?

(A) HTTP

(B) HTTPS

(C) FTP

(D) SMTP

Q28. What will be the output of ‘print(”python”.capitalize())‘?

(A) ‘PYTHON‘

(B) ‘Python‘

(C) ‘python‘

(D) ‘PyThOn‘

Q29. In SQL, which aggregate function returns the average of the values in a column, ignoring NULL values?

(A) ‘TOTAL()‘

(B) ‘AVG()‘

(C) ‘SUM()/COUNT(*)‘

(D) ‘MEAN()‘

Q30. Which searching algorithm requires the data to be in a sorted order and has a time complexity of $O(\log n)$?

(A) Linear Search

(B) Bubble Search

(C) Binary Search

(D) Selection Search

Q31. In Python, which block is used to catch and handle specific errors during code execution?

(A) ‘catch‘

(B) ‘try‘

(C) ‘except‘

(D) ‘finally‘

Q32. Which SQL wildcard is used to match exactly one character in a LIKE clause?

(A) ‘

(B) ‘*‘

(C) ‘.

(C) ‘¿

Q33. A network topology where every device has a point-to-point link to every other device is called:

(A) Mesh

(B) Star

(C) Ring

(D) Bus

Q34. What is the time complexity of pushing an element into a stack implemented using an array?

(A) $O(n)$

(B) $O(1)$

(C) $O(\log n)$

(D) $O(n^2)$

(E) $O(n \log n)$

Q35. In SQL, which command is used to permanently save all changes made during a transaction?

(A) 'SAVE'

(B) 'ROLLBACK'

(C) 'COMMIT'

(D) 'UPDATE'

Q36. What is the result of 'len("Informatics".split("i"))' in Python?

(A) 2

(B) 3

(C) 1

(D) 11

Q37. Which layer of the OSI model is responsible for providing end-to-end communication services, including error recovery and flow control?

(A) Network Layer

(B) Transport Layer

(C) Session Layer

(D) Data Link Layer

Q38. While performing Selection Sort on [64, 25, 12, 22, 11], what is the state of the list after the first pass?

(A) '[11, 25, 12, 22, 64]'

(B) '[11, 64, 25, 12, 22]'

(C) '[12, 25, 64, 22, 11]'

(D) '[25, 12, 22, 11, 64]'

Q39. Which SQL function is used to calculate the number of days between two dates?

(A) 'DATEDIFF()'

(B) 'DIFF()'

(C) 'SUBDATE()'

(D) 'DAYCOUNT()'

Q40. In Python, what does 'pickle.dump(obj, f)' do?

(A) Reads an object from a file

(B) Serializes an object and writes it to a file

(C) Deletes an object from memory

(D) Creates a text representation of a binary file

Q41. The maximum number of nodes in a binary tree of height 3 (where the root is at height 0) is:

(A) 7

(B) 15

(C) 8

(D) 3

Q42. Which protocol is used to transfer electronic mail from one mail server to another?

(A) POP3

(B) IMAP

(C) SMTP

(D) HTTP

Q43. In SQL, which keyword is used with ‘ORDER BY‘ to sort results in ascending order (default)?

(A) ‘DESC‘

(B) ‘ASC‘

(C) ‘UP‘

(D) ‘SORT‘

Q44. What will be the output of ‘print(range(5)[-1])‘ in Python?

(A) 5

(B) 4

(C) 0

(D) -1

Q45. In a relational database, which key is used to uniquely identify a record in another table, creating a relationship between them?

(A) Primary Key

(B) Candidate Key

(C) Foreign Key

(D) Composite Key

Q46. Which malware type spreads by self-replicating and sending copies of itself to other computers over a network without human intervention?

(A) Virus

(B) Worm

(C) Trojan Horse

(D) Spyware

Q47. What is the output of 'SELECT SUBSTR('COMPUTER', 4, 3);'?

(A) 'PUT'

(B) 'MPU'

(C) 'TER'

(D) 'COM'

Q48. In Python, which file mode opens a file for writing and positions the pointer at the end of the file?

(A) 'w'

(B) 'r+'

(C) 'a'

(D) 'w+'

Q49. Which type of switching technology breaks data into small packets and sends them independently through the network?

(A) Circuit Switching

(B) Message Switching

(C) Packet Switching

(D) Line Switching

Q50. What is the worst-case time complexity of Quick Sort?

(A) $O(n \log n)$

(B) $O(n)$

(C) $O(n^2)$

(D) $O(1)$

Detailed Solutions

Q1.

Solution

Concept:

The ‘CHECK‘ constraint in SQL is a Data Integrity constraint used to limit the value range that can be placed in a column. It ensures that all values in a column satisfy a specific Boolean condition.

Solution:

1. To enforce a rule that salary must be above a certain value, we use the ‘CHECK‘ clause during table creation or alteration. 2. The syntax is: ‘CHECK (Salary >= 25000)‘. 3. Other options: - ‘PRIMARY KEY‘ is for unique identification of rows. - ‘UNIQUE‘ ensures all values in a column are different. - ‘DEFAULT‘ provides a value when none is specified, but does not prevent low values from being entered manually. 4. Therefore, ‘CHECK‘ is the correct mechanism for validating data against a specific range or limit.

Final Answer: The correct constraint is CHECK.

Answer: (C)

Q2.

Solution

Concept:

In Python file handling, the ‘seek(offset, whence)‘ method moves the file pointer to a specific position. The ‘offset‘ is the number of bytes, and ‘whence‘ defines the reference point (0 for beginning, 1 for current, 2 for end).

Solution:

1. Files are zero-indexed in Python. To reach the 10th byte, the pointer must be moved to index 9. 2. The command ‘f.seek(9)‘ moves the pointer to the 10th character/byte from the start of the file. 3. Once the pointer is at position 9, ‘f.write(b’X’)‘ will write ’X’ at that exact location, overwriting whatever was previously there. 4. Option A (‘seek(10)‘) would put the pointer at the 11th byte. Option C (‘read(10)‘) moves the pointer but is less direct than ‘seek‘.

Final Answer: The sequence is f.seek(9); f.write(b’X’).

Answer: (B)

Q3.

Solution

Concept:

Postfix expression evaluation uses a stack. Numbers are pushed onto the stack, and when an operator is encountered, the top two elements are popped, the operation is applied, and the result is pushed back.

Solution:

1. Scan the expression: '8 2 / 5 2 * +' 2. Push 8, Push 2. Stack: '[8, 2]' 3. Encounter '/': Pop 2 and 8. Perform '8 / 2 = 4'. Push 4. Stack: '[4]' 4. Push 5, Push 2. Stack: '[4, 5, 2]' 5. Encounter '*': Pop 2 and 5. Perform '5 * 2 = 10'. Push 10. Stack: '[4, 10]' 6. Encounter '+': Pop 10 and 4. Perform '4 + 10 = 14'. Push 14. Stack: '[14]' 7. The final result is 14.

Final Answer: The final value is 14.

Answer: (A)

Q4.

Solution

Concept:

Time complexity for sorting algorithms depends on the number of comparisons and swaps. Selection Sort is unique because its comparison logic is independent of the initial order of the elements.

Solution:

1. Selection Sort works by finding the minimum element in the unsorted part and swapping it into the correct position. 2. To find the minimum, it always compares every element in the remaining unsorted sublist. 3. It performs $\frac{n(n-1)}{2}$ comparisons regardless of whether the list is sorted, reverse-sorted, or random. 4. Unlike Bubble Sort or Insertion Sort, which can be optimized to $O(n)$ for sorted data, Selection Sort always takes $O(n^2)$ time.

Final Answer: The algorithm is Selection Sort.

Answer: (C)

Q5.

Solution

Concept:

Relational Algebra operations allow us to manipulate sets of data. To find elements that exist in one set but not another, we use the Set Difference operation.

Solution:

1. Let A be the set of students in 'Computer Science' and B be the set of students in 'Mathematics'. 2. The requirement is to find $A - B$. 3. The Set Difference operator (denoted by $-$) returns all tuples that are in the first relation but not in the second. 4. Note: For Set Difference to work, both relations must be union-compatible (same schema).

Final Answer: The operation is Set Difference.

Answer: (C)

Q6.

Solution

Concept:

Subnetting is the practice of dividing a network into two or more smaller networks. The number of usable hosts in a subnet is determined by the number of bits remaining for the host portion after the subnet mask has been applied. A standard IPv4 address has 32 bits. The formula for the total number of addresses in a subnet is 2^n, where n is the number of host bits. However, the number of "usable" host addresses is $2^n - 2$ because the first address is reserved for the Network ID and the last address is reserved for the Broadcast Address.

Solution:

1. **Identify the Subnet Mask:** The given mask is 255.255.255.192. 2. **Analyze the Last Octet:** The first three octets (255.255.255) are full, meaning they belong to the network portion. We look at the last octet: 192. 3. **Convert to Binary:** The decimal number 192 in binary is 11000000_2. 4. **Determine Host Bits:** In the last octet, the first 2 bits are '1' (network/subnet bits) and the remaining 6 bits are '0' (host bits). Therefore, $n = 6$. 5. **Calculate Total Addresses:** The total number of addresses per subnet is $2^6 = 64$. 6. **Calculate Usable Addresses:** Subtract 2 for the network and broadcast addresses: $64 - 2 = 62$. 7. Thus, in a network with a .192 mask, each subnet can support 62 unique physical devices or hosts.

Final Answer: There are 62 usable host addresses.

Answer: (B)

Q7.

Solution

Concept:

In SQL, string manipulation functions are used to clean and format data. 'LTRIM()' stands for "Left Trim," which removes leading spaces from a string. 'RTRIM()' stands for "Right Trim," which removes trailing spaces. When nested, these functions work from the inside out to "strip" the string of all surrounding whitespace.

Solution:

1. **Inner Function Execution:** The inner function is 'RTRIM(' CUET 2026 ')'. 2. This function looks at the right side of the string and removes the spaces following "2026". The result of this step is '' CUET 2026''. 3. **Outer Function Execution:** The result from step 2 is then passed to the outer function: 'LTRIM(' CUET 2026')'. 4. The 'LTRIM' function looks at the left side of the string and removes the three leading spaces before "CUET". 5. **Final Result:** After removing both the leading and trailing spaces, the remaining string is ''CUET 2026''. 6. Note that spaces *between* words (the space between CUET and 2026) are not affected by trim functions; only exterior spaces are removed. This ensures data integrity for strings that contain multiple words.

Final Answer: The output is 'CUET 2026'.

Answer: (A)

Q8.

Solution

Concept:

Python uses a robust Exception Handling mechanism to deal with runtime errors. Different errors trigger specific exception classes. A dictionary in Python consists of Key-Value pairs. When we attempt to retrieve a value using a key, Python searches the internal hash table for that specific key.

Solution:

1. **Scenario:** Suppose we have a dictionary ‘student = ”name”: ”Amit”, ”age”: 20‘. 2. If we try to access ‘student[”grade”]‘, Python realizes that the key ”grade” does not exist in the dictionary object. 3. **Exception Triggered:** In this specific case, Python raises a ‘KeyError‘. This is distinct from an ‘IndexError‘, which is raised when you try to access an invalid index in a sequence like a list or a tuple. 4. **Other Exceptions:** - ‘ValueError‘ is raised when a function receives an argument of the correct type but inappropriate value (e.g., ‘int(”abc”)‘). - ‘TypeError‘ is raised when an operation is applied to an object of an inappropriate type (e.g., adding a string to an integer). 5. Understanding these specific exceptions allows developers to write precise ‘except‘ blocks to handle errors gracefully without crashing the program.

Final Answer: The exception raised is KeyError.

Answer: (B)

Q9.

Solution

Concept:

In a Local Area Network (LAN), devices communicate using physical MAC (Media Access Control) addresses. However, applications and the OS use IP (Internet Protocol) addresses to identify targets. The Address Resolution Protocol (ARP) acts as the bridge between the Network Layer (Layer 3) and the Data Link Layer (Layer 2).

Solution:

1. When Host A wants to send data to Host B on the same subnet, it knows Host B’s IP address but needs its MAC address to construct an Ethernet frame. 2. **ARP Request:** Host A broadcasts an ARP request packet to the entire network asking, ”Who has IP address X.X.X.X? Tell me your MAC address.” 3. **ARP Reply:** The device with that specific IP address sends an ARP reply directly back to Host A, containing its hardware MAC address. 4. **Caching:** Host A then stores this mapping in its ”ARP Cache” for future use to avoid repeated broadcasts. 5. **Comparison with other protocols:** - ‘DNS‘ resolves domain names (like google.com) to IP addresses. - ‘DHCP‘ assigns dynamic IP addresses to hosts. - ‘ICMP‘ is used for diagnostics and error reporting (like ping).

Final Answer: The protocol is ARP.

Answer: (C)

Q10.

Solution

Concept:

Bubble Sort is a simple sorting algorithm that repeatedly steps through the list, compares adjacent elements, and swaps them if they are in the wrong order. This process is repeated until the list is sorted. In each ”pass,” the largest unsorted element ”bubbles up” to its correct position at the end of the list.

Solution:

1. **Initial List:** ‘[50, 20, 40, 10, 30]‘ 2. **Pass 1 - Comparison 1:** Compare 50 and 20. Since 50 > 20, **Swap** them. List becomes: ‘[20, 50, 40, 10, 30]‘ (Swap 1) 3. **Pass 1 - Comparison 2:** Compare 50 and 40. Since 50 > 40, **Swap** them. List becomes: ‘[20, 40, 50, 10, 30]‘ (Swap 2) 4. **Pass 1 - Comparison 3:** Compare 50 and 10. Since 50 > 10, **Swap** them. List becomes: ‘[20, 40, 10, 50, 30]‘ (Swap 3) 5. **Pass 1 - Comparison 4:** Compare 50 and 30. Since 50 > 30, **Swap** them. List becomes: ‘[20, 40, 10, 30, 50]‘ (Swap 4) 6. **End of Pass 1:** The largest element (50) is now at the last position. The total number of swaps performed during this first pass is exactly 4. 7. Note: The number of comparisons in the first pass is also 4 (n-1), and in this specific case, every comparison resulted in a swap.

Final Answer: The number of swaps is 4.

Answer: (C)

Q11.

Solution

Concept:

In SQL (Structured Query Language), string manipulation functions are essential for searching and data extraction. When we need to find the location of a specific character or a string within another string, different database systems (like MySQL, PostgreSQL, or Oracle) offer various functions that essentially perform the same task: returning the 1-based index of the first occurrence of the target substring.

Solution:

1. **LOCATE(substring, string):** This is commonly used in MySQL. It returns the position of the first occurrence of a substring. If the substring is not found, it returns 0. 2. **POSITION(substring IN string):** This is an ANSI SQL standard function. It provides the same result as LOCATE but uses a slightly different syntax with the ‘IN‘ keyword. 3. **INSTR(string, substring):** This function is widely used in Oracle and MySQL. Note that the order of arguments is reversed compared to LOCATE (the string comes first, then the substring). 4. **Conclusion:** Since all three functions—‘LOCATE()‘, ‘POSITION()‘, and ‘INSTR()‘—are designed to identify the starting index of a substring within a main string, they are all valid answers depending on the specific SQL dialect being used. In a general competitive exam context, knowing that multiple functions exist for this purpose is vital.

Final Answer: All of the above functions can be used.

Answer: (D)

Q12.

Solution

Concept:

A Circular Queue is a linear data structure that follows the FIFO (First In First Out) principle but connects the last position back to the first position to make the queue ”circular.” This solves the problem of memory wastage in a standard linear queue where space at the front cannot be reused. Calculating the number of elements in a circular queue requires understanding how the ‘front‘ and ‘rear‘ pointers move.

Solution:

1. **Given Data:** ‘Size (N) = 6‘, ‘front = 4‘, and ‘rear = 1‘. 2. **Understanding the Pointers:** - ‘front‘ usually points to the element to be removed (or just before it, depending on implementation). - ‘rear‘ points to the last element added. 3. **Calculating Count:** - If ‘rear >= front‘, the formula is simply ‘rear - front + 1‘. - If ‘rear < front‘ (which is the case here), it means the queue has ”wrapped around” the end of the array. The formula becomes: ‘(Size - front) + (rear + 1)‘. 4. **Step-by-Step Calculation:** - Elements from ‘front‘ to the end: Indices 4 and 5 (2 elements). - Elements from the start to ‘rear‘: Indices 0 and 1 (2 elements). - Total elements = 2 + 2 = 4. 5. Alternatively, using the modulo formula: ‘(rear - front + Size)

Final Answer: There are 4 elements currently in the queue.

Answer: (B)

Q13.

Solution

Concept:

Network Topology refers to the physical or logical arrangement of nodes and connections in a network. The Star Topology is the most common arrangement used in modern Local Area Networks (LANs). In this setup, every individual node (computer, printer, etc.) is connected to a central connection point, such as a Hub, Switch, or Router.

Solution:

1. **Structure:** All data traffic in a Star topology travels through the central hub before reaching its destination. The hub acts as a signal repeater or a forwarder. 2. **Advantages:** Adding or removing a new device is extremely easy and does not disrupt the rest of the network. If one peripheral cable fails, only that specific node goes offline. 3. **The ”Single Point of Failure”:** The primary disadvantage is the central dependency. If the central Hub or Switch fails, the entire network becomes inoperable because no node can communicate with another. 4. **Comparison:** - In a **Bus** topology, a break in the main cable takes the network down. - In a **Mesh** topology, there is high redundancy, so there is no single point of failure, but it is very expensive and complex. - In a **Ring** topology, a failure in one node or cable typically breaks the entire loop.

Final Answer: The topology is the Star topology.

Answer: (C)

Q14.

Solution

Concept:

Python lists are mutable sequences that can contain other lists as elements (nested lists). Accessing an element in a nested list requires multiple indices: the first index selects the inner list, and the second index selects the specific item within that inner list.

Solution:

1. **Initial State:** ‘L = [1, [2, 3], 4]‘. - ‘L[0]‘ is ‘1‘. - ‘L[1]‘ is the sublist ‘[2, 3]‘. - ‘L[2]‘ is ‘4‘. 2. **Accessing the Sublist:** The command ‘L[1]‘ targets the list ‘[2, 3]‘. 3. **Modifying the Element:** Within ‘L[1]‘, the index ‘[1]‘ refers to the second element of that sublist, which is ‘3‘. 4. **The Assignment:** ‘L[1][1] = 5‘ replaces the value ‘3‘ with ‘5‘. 5. **Final Result:** The inner list ‘[2, 3]‘ now becomes ‘[2, 5]‘. When the entire list ‘L‘ is printed, it reflects this change because lists are modified ”in-place.” 6. The updated list is ‘[1, [2, 5], 4]‘. This demonstrates Python’s ability to handle complex, nested data structures efficiently through reference-based indexing.

Final Answer: The output is [1, [2, 5], 4].

Answer: (A)

Q15.

Solution

Concept:

In Database Management Systems (DBMS), the ACID properties (Atomicity, Consistency, Isolation, Durability) ensure that database transactions are processed reliably. These properties are the cornerstone of data integrity, especially in environments where multiple users access the data simultaneously.

Solution:

1. **Atomicity:** This property follows the ”All or Nothing” rule. A transaction often involves multiple steps (e.g., in a bank transfer, deducting money from account A and adding it to account B). Atomicity ensures that if any single step fails, the entire transaction is aborted, and the database is rolled back to its previous state. 2. **Consistency:** Ensures the database moves from one valid state to another, maintaining all predefined rules and constraints. 3. **Isolation:** Ensures that the concurrent execution of transactions results in a system state that would be obtained if transactions were executed serially (one after another). 4. **Durability:** Guarantees that once a transaction has been committed, it will remain so, even in the event of a power failure or system crash. 5. **Conclusion:** Since the question asks about treating multiple sub-operations as a single unit that either completely succeeds or completely fails, **Atomicity** is the defining property.

Final Answer: The property is Atomicity.

Answer: (D)

Q16.

Solution

Concept:

The ‘ROUND()‘ function in SQL is a mathematical function used to round a numeric value to a specified number of decimal places. The syntax is generally ‘ROUND(number, decimals)‘. If the ‘decimals‘ argument is positive, it rounds to the specified decimal position; if it is zero, it rounds to the nearest integer.

Solution:

1. **Input Analysis:** The query is ‘SELECT ROUND(45.926, 1);‘. 2. **Precision Level:** The second argument is ‘1‘, which means we need to round the number to one decimal place (the tenths position). 3. **Rounding Rule:** We look at the digit to the right of the target decimal place. The target digit is ‘9‘ (tenths), and the digit to its right is ‘2‘ (hundredths). 4. **Logic Application:** Since the digit ‘2‘ is less than 5, we do not increment the target digit. We simply keep it as ‘9‘ and discard the subsequent digits. 5. **Final Computation:** 45.926 rounded to 1 decimal place becomes 45.9. 6. Note: If the query had been ‘ROUND(45.926, 2)‘, the result would have been 45.93 because the third decimal ‘6‘ is greater than 5, causing the ‘2‘ to round up.

Final Answer: The result is 45.9.

Answer: (A)

Q17.

Solution

Concept:

The IPv4 (Internet Protocol version 4) header contains several fields that manage the routing and delivery of packets. One critical field for network stability is the ”Time to Live” (TTL). This field acts as a safety mechanism to prevent packets from circulating indefinitely in the network in case of a routing loop.

Solution:

1. **Function of TTL:** When a packet is created, the sender sets an initial TTL value (e.g., 64 or 128). 2. **Router Processing:** Every time the packet arrives at a router (a ”hop”), the router processes the header. Before forwarding the packet to the next destination, the router decrements the TTL value by exactly 1. 3. **Discarding Packets:** If a router receives a packet and, after decrementing, the TTL reaches 0, the router discards the packet and sends an ICMP ”Time Exceeded” message back to the sender. 4. **Role of other fields:** - **Version:** Indicates the IP version (4 or 6) and does not change. - **Total Length:** Defines the size of the entire packet. - **Header Checksum:** This field is actually recalculated by each router because the TTL changes, but it is the TTL itself that is the primary field being decremented.

Final Answer: The field is Time to Live (TTL).

Answer: (C)

Q18.

Solution

Concept:

When working with file objects in Python, the system maintains a ”file pointer” or ”cursor” that indicates the current byte position where the next read or write operation will occur. To manage file I/O effectively, especially in random access files (like binary files), programmers must be able to track this position.

Solution:

(a) **The `tell()` Method:** The built-in method `file_object.tell()` returns the current byte position.

(b) **Contrast with `seek()`:** Used to change the position.

(c) **Usage Example:** After reading 10 characters, `f.tell()` returns 10.

(d) **Data types:** In binary mode, it strictly tracks bytes.

(e) Other options like `where()` or `position()` do not exist.

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

Answer: (C)

Q19.

Solution

Concept:

Infix notation is the standard mathematical notation where operators are placed between operands. Postfix notation (also known as Reverse Polish Notation) places operators after their operands, which eliminates the need for parentheses and simplifies computer-based evaluation using a stack.

Solution:

1. **Input Expression:** ‘(A + B) * (C - D)‘ 2. **Order of Operations:** According to BODMAS/PEMDAS rules, expressions inside parentheses are evaluated first. 3. **Step 1 (First Parenthesis):** Convert ‘(A + B)‘. The operands are A and B, and the operator is +. In postfix, this becomes ‘AB+‘. 4. **Step 2 (Second Parenthesis):** Convert ‘(C - D)‘. The operands are C and D, and the operator is -. In postfix, this becomes ‘CD-‘. 5. **Step 3 (Final Operator):** Now treat ‘(AB+)‘ and ‘(CD-)‘ as two separate operands. The operator connecting them is ‘*‘. 6. **Combining:** Place the operator after the two operands: ‘[Operand1] [Operand2] [*]‘. 7. Result: ‘AB+CD-*‘. 8. This matches the logic used by compilers to parse mathematical expressions efficiently.

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

Answer: (A)

Q20.

Solution

Concept:

Databases often contain redundant data or multiple rows with the same values across specific columns. When generating reports, it is often necessary to see only the unique values. SQL provides a specific keyword to filter out these repeating entries from the result set of a ‘SELECT‘ query.

Solution:

1. **The ‘DISTINCT‘ Keyword:** The ‘DISTINCT‘ clause is used immediately after the ‘SELECT‘ keyword to instruct the database engine to evaluate the resulting rows and remove any exact duplicates. 2. **Syntax Example:** ‘SELECT DISTINCT City FROM Students;‘ would return each city name only once, even if hundreds of students live in the same city. 3. **Comparison with other clauses:** - ‘UNIQUE‘ is a constraint applied during table creation to prevent duplicates from being entered, but it is not used in a ‘SELECT‘ list. - ‘GROUP BY‘ is used for aggregating data (like counting or summing), and while it can produce unique lists, its primary purpose is calculation. - ‘CHECK‘ is a validation constraint for data ranges. 4. ‘DISTINCT‘ operates on the entire row selected. If you select multiple columns, the database only removes rows where the combination of all selected columns is identical.

Final Answer: The clause is DISTINCT.

Answer: (B)

Q21.

Solution

Concept:

Insertion Sort is an iterative sorting algorithm that builds a sorted sublist at the beginning of the array. It works by taking one element from the unsorted portion and ”inserting” it into its correct relative position within the already sorted portion. It is highly efficient for small datasets and is ”stable,” meaning it preserves the relative order of equal elements.

Solution:

1. **How it Works:** Imagine you are sorting a deck of cards in your hands. You pick one card at a time and slide it into the correct spot among the cards you are already holding. This is the exact logic of Insertion Sort. 2. **Algorithmic Steps:** - The first element is considered sorted by default. - For the next element, the algorithm compares it with the elements in the sorted sublist (to its left). - If the new element is smaller than the element being compared, the larger element is shifted one position to the right to make space. - This shifting continues until the correct ”gap” is found, where the new element is then placed. 3. **Efficiency:** For a list that is already sorted, Insertion Sort only performs $O(n)$ comparisons and zero swaps, making it much faster than Selection Sort in such cases. However, in the worst-case (reverse sorted), it takes $O(n^2)$ time. 4. **Comparison:** Unlike Selection Sort, which finds the minimum and swaps, Insertion Sort shifts elements. Unlike Bubble Sort, it doesn't just swap adjacent items repeatedly; it finds the specific destination for the current item.

Final Answer: The algorithm is Insertion Sort.

Answer: (C)

Q22.

Solution

Concept:

The OSI (Open Systems Interconnection) model and the TCP/IP model define different layers for network communication. Networking hardware is often categorized by the ”layer” at which it operates. The Data Link Layer (Layer 2) is responsible for node-to-node data transfer—handling the physical addressing and error detection of frames within a single local network.

Solution:

1. **The Switch:** A Switch is an intelligent networking device that operates at the Data Link Layer. It maintains a ”MAC Address Table” (also known as a CAM table). 2. **Filtering and Forwarding:** When a frame arrives, the switch looks at the destination MAC address. Instead of broadcasting the data to every port (like a Hub does), the switch forwards the frame only to the specific port where that MAC address is located. 3. **Collision Domains:** Switches provide dedicated bandwidth to each port, effectively creating separate collision domains, which significantly improves network performance compared to a Hub. 4. **Other Devices:** - **Hub/Repeater:** Operates at the Physical Layer (Layer 1). It does not understand MAC addresses and simply repeats signals. - **Router:** Operates at the Network Layer (Layer 3) and uses IP addresses to route data between different networks. 5. Therefore, a device that filters traffic based on MAC addresses in a LAN is a Switch.

Final Answer: The device is a Switch.

Answer: (C)

Q23.

Solution

Concept:

Python evaluates expressions based on operator precedence and associativity. The exponentiation operator (‘**‘) is unique in Python because it has ”Right-to-Left” associativity. Most other operators, like addition or multiplication, are evaluated from left to right.

Solution:

1. **The Expression:** ‘3 ** 1 ** 3‘. 2. **Associativity Rule:** Because ‘**‘ is right-associative, Python evaluates the expression from the rightmost power first. This is equivalent to writing ‘3 ** (1 ** 3)‘. 3. **Step 1:** Calculate the rightmost part: ‘1 ** 3‘. 1 raised to any power is 1. So, ‘1 ** 3 = 1‘. 4. **Step 2:** Substitute the result back into the main expression: ‘3 ** 1‘. 5. **Step 3:** Calculate the final power: ‘3 ** 1 = 3‘. 6. **Common Mistake:** If a student evaluates from left to right, they would do ‘(3 ** 1) ** 3‘, which is ‘3 ** 3 = 27‘. However, following the standard Python rules of precedence and associativity, this is incorrect. 7. This distinction is critical in competitive programming to avoid logic errors in mathematical computations.

Final Answer: The output is 3.

Answer: (C)

Q24.

Solution

Concept:

Data Definition Language (DDL) commands in SQL are used to define or modify the structure of database objects like tables. While `DROP TABLE` removes an entire table, the `ALTER TABLE` command is used when you want to modify specific parts of a table's structure without deleting the existing data.

Solution:

(a) **The Task:** To remove an unnecessary column from a table, we must "alter" the existing structure.

(b) **Syntax:** `ALTER TABLE table_name DROP COLUMN column_name;`

(c) **Key Components:**

- `ALTER TABLE`: Targets the table.
- `DROP COLUMN`: Removes the column.

(d) **Important Note:** This operation is usually permanent.

(e) **Contrast:** `DELETE` removes rows, not columns.

Final Answer: The command is `ALTER TABLE ... DROP COLUMN`.

Answer: (C)

Q25.

Solution

Concept:

A Linear Queue is a data structure that follows the First-In-First-Out (FIFO) principle. It uses two pointers: ‘front‘ (to track where the next removal happens) and ‘rear‘ (to track where the next insertion happens). In a standard array-based implementation, these indices start at -1 or 0 and increment as operations occur.

Solution:

1. **Initial State:** Array size = 5. Indices are 0, 1, 2, 3, 4. Assuming starting index -1 for an empty queue. 2. **Action 1 (5 Enqueues):** - 1st enqueue: ‘rear = 0‘ - 2nd enqueue: ‘rear = 1‘ - 3rd enqueue: ‘rear = 2‘ - 4th enqueue: ‘rear = 3‘ - 5th enqueue: ‘rear = 4‘ - At this point, the queue is full. ‘rear‘ is at the last index (4). 3. **Action 2 (2 Dequeues):** - When a dequeue happens, the ‘front‘ pointer moves forward (from -1 to 0, then 0 to 1). - Crucially, the ‘rear‘ pointer **does not move** during a dequeue operation in a linear queue. It only moves during enqueues. 4. **Conclusion:** After 5 enqueues, ‘rear‘ reached 4. The 2 dequeues affected only the ‘front‘ pointer. Therefore, the value of ‘rear‘ remains 4. 5. Note: In a circular queue, the behavior might differ, but in a standard linear queue, the rear only advances.

Final Answer: The current value of the rear index is 4.

Answer: (A)

Q26.

Solution

Concept:

Relational Algebra is a procedural query language used to model the data stored in relations and define operations on them. The **Union** ($\cup$) operation is a fundamental set-theoretic operation. For two relations to be compatible for a Union operation, they must be "Union Compatible," meaning they must have the same number of attributes with identical domains (data types).

Solution:

1. **Standard Union Behavior:** In mathematical set theory and standard Relational Algebra, a Union operation combines all tuples from relation R and relation S. By definition, sets do not contain duplicate elements. Therefore, if a row exists in both R and S, the standard Union operation will only include it once in the result set. 2. **The "Union All" Concept:** To include duplicates (i.e., every single row from both tables regardless of whether they are identical), modern SQL uses the 'UNION ALL' command. However, in strict Relational Algebra, the operator that preserves all rows including duplicates is sometimes referred to as a multiset union. 3. **Question Context:** The question asks for an operation that combines rows and *includes duplicates*. In standard relational theory, a Union removes duplicates. If the options provided suggest that a standard Union does this, it is often a point of distinction between SQL behavior and formal Algebra. 4. **Conclusion:** Since the standard Union operation is defined to eliminate duplicates to maintain set integrity, a specific variant or a "None of the above" choice is required if we are strictly following the rule that standard Union results in a unique set.

Final Answer: None of the above (as standard Relational Union removes duplicates).

Answer: (D)

Q27.

Solution

Concept:

Network security and data integrity are maintained through protocols that encrypt the communication between a client (browser) and a server. **HTTPS** (HyperText Transfer Protocol Secure) is the secure version of HTTP. It uses cryptographic protocols to provide end-to-end encryption.

Solution:

1. **How it Works:** HTTPS uses **TLS** (Transport Layer Security) or previously **SSL** (Secure Sockets Layer) to encrypt normal HTTP requests and responses. This ensures that even if a hacker intercepts the data packets (a Man-in-the-Middle attack), they cannot read the content because it is encrypted. 2. **Authentication:** HTTPS also provides authentication, ensuring that the user is communicating with the intended website and not an imposter. This is verified through Digital Certificates issued by Certificate Authorities (CAs). 3. **Visual Cues:** In a web browser, a secure connection is usually indicated by a padlock icon in the address bar and a URL starting with ‘https://‘. 4. **Comparison:** - **HTTP:** Sends data in ”Plain Text.” Anyone on the network can see passwords or credit card info. - **FTP:** Used for file transfers, not standard web browsing. - **SMTP:** Used for sending emails, not for loading web pages.

Final Answer: The protocol is HTTPS.

Answer: (B)

Q28.

Solution

Concept:

Python provides several built-in string methods for formatting. The ‘.capitalize()‘ method is specifically designed to handle sentence-style casing. It is important to distinguish it from other methods like ‘.upper()‘, ‘.lower()‘, or ‘.title()‘.

Solution:

1. **The ‘.capitalize()‘ Logic:** This method performs two specific actions on a string: - It converts the very first character of the string to **Uppercase**. - It converts all remaining characters in the string to **Lowercase**. 2. **Trace:** Input string is ‘”python”‘. - First character ’p’ becomes ’P’. - Remaining characters ’ython’ are already lowercase, so they stay ’ython’. - Result: ‘”Python”‘. 3. **Contrast with other methods:** - ‘”python”.upper()‘ would result in ‘”PYTHON”‘. - ‘”python”.title()‘ would also result in ‘”Python”‘ in this case, but if the string was ‘”hello world”‘, ‘title()‘ would produce ‘”Hello World”‘ while ‘capitalize()‘ would produce ‘”Hello world”‘. 4. This method is particularly useful when you want to ensure the first word of a user-inputted string starts with a capital letter regardless of how it was typed.

Final Answer: The output is ’Python’.

Answer: (B)

Q29.

Solution

Concept:

Aggregate functions in SQL perform a calculation on a set of values and return a single value. These functions are highly optimized for handling large datasets. A key characteristic of most aggregate functions (except 'COUNT(*)') is that they automatically ignore 'NULL' values during calculation.

Solution:

1. **The 'AVG()' Function:** This function calculates the arithmetic mean of a numeric column. 2. **Handling NULLs:** If a column has values '[10, 20, NULL, 30]', the 'AVG()' function will calculate $(10 + 20 + 30)/3$. It treats the row with the 'NULL' as if it does not exist for both the sum and the divisor. 3. **Manual Calculation:** If you were to use 'SUM(column) / COUNT(*)', you might get a different (and often incorrect) result because 'COUNT(*)' counts all rows including NULLs, whereas 'AVG()' only counts rows with actual data. 4. **Other Options:** - 'TOTAL()' is used in some dialects like SQLite but is not standard. - 'MEAN()' is not a standard SQL aggregate function; it is more common in Python libraries like Pandas or NumPy. - 'SUM()/COUNT(*)' is a manual way to find an average but is prone to errors if NULLs are present.

Final Answer: The function is AVG().

Answer: (B)

Q30.

Solution

Concept:

Searching is the process of finding a specific element within a collection of data. Binary Search is an "interval-based" search algorithm that is significantly more efficient than Linear Search, but it comes with a strict prerequisite: the data must be sorted.

Solution:

1. **How it Works:** Binary Search works on the "Divide and Conquer" principle. It compares the target value to the middle element of the array. 2. **The Logic:** - If the target is equal to the middle element, the search is complete. - If the target is smaller, the search continues in the left half. - If the target is larger, the search continues in the right half. 3. **Efficiency:** Because the search space is cut in half with every single comparison, the time complexity is $O(\log n)$. For example, in a list of 1,024 elements, Linear Search might take 1,024 comparisons, while Binary Search will take at most 10 ($\log_2 1024 = 10$). 4. **Requirement:** If the list is not sorted, the middle-element comparison would be meaningless, as there would be no way to know which half to discard. Therefore, sorting is mandatory.

Final Answer: The algorithm is Binary Search.

Answer: (C)

Q31.

Solution

Concept:

Exception handling in Python is a critical programming construct used to manage runtime errors gracefully, ensuring the program does not crash unexpectedly. The structure primarily uses four keywords: ‘try‘, ‘except‘, ‘finally‘, and ‘else‘. While the ‘try‘ block contains the code that might fail, the ‘except‘ block is where the actual ”handling” occurs.

Solution:

1. **The ‘except‘ Block:** This block is designed to ”catch” exceptions that are ”raised” in the preceding ‘try‘ block. If an error occurs, Python stops the execution of the ‘try‘ block and jumps immediately to the matching ‘except‘ block. 2. **Specific vs. General:** Programmers can specify the type of error to catch, such as ‘except ValueError:‘ or ‘except ZeroDivisionError:‘. This allows for precise error messaging and recovery strategies tailored to the specific problem. 3. **Execution Flow:** If no exception occurs in the ‘try‘ block, the ‘except‘ block is completely skipped. 4. **Other Components:** - ‘try‘: The section where you place ”risky” code. - ‘finally‘: A block that runs no matter what (even if the program crashes or returns), often used for cleaning up resources like closing files or database connections. 5. **Conclusion:** Because the question asks specifically for the block used to **catch and handle** errors, the ‘except‘ block is the correct technical answer.

Final Answer: The block used is except.

Answer: (C)

Q32.

Solution

Concept:

The `LIKE` operator in SQL is used within a `WHERE` clause to search for a specified pattern. To define these patterns, SQL uses special characters known as "wildcards." The two most common wildcards in standard SQL (ANSI) are the percent sign (%) and the underscore (_).

Solution:

(a) **The Underscore (_) Wildcard:** This character is a placeholder that represents **exactly one** single character. For example, the pattern `'M_y'` would match "May" or "Moy" but not "My" or "Maddy."

(b) **The Percent (%) Wildcard:** This character represents zero, one, or multiple characters. For instance, `'S%'` matches "Sky," "Sun," or "Saturn."

(c) **Usage in Logic:** If you need to find all names that are exactly four letters long and start with 'A', you would use the criteria `LIKE 'A___'`. Each underscore forces the database to ensure a character exists at that specific position.

(d) **Platform Variations:** While _ is standard for SQL, some applications like MS Access might use a question mark (?) instead.

Final Answer: The wildcard is the underscore (_).

Answer: (C)

Q33.

Solution

Concept:

A network topology defines how different nodes in a network are interconnected. In a **Mesh Topology**, the concept of redundancy is taken to its maximum. Every node has a dedicated point-to-point link to every other node in the network. This creates a highly resilient system where data can take multiple paths to reach its destination.

Solution:

1. **Connectivity Formula:** In a "Full Mesh" network with n nodes, the total number of physical links is calculated using the formula $n(n-1)/2$. For example, a network with 5 computers would require 10 separate cables. 2. **Advantages:** The primary benefit is fault tolerance. If one link fails, the network traffic can be rerouted through other nodes. Additionally, it offers high privacy and security because data travels on dedicated lines. 3. **Disadvantages:** The main drawback is the high cost of installation and the complexity of cabling as the number of nodes increases. It also requires each device to have many I/O ports. 4. **Contrast:** - **Star:** Nodes connect to a central hub. - **Bus:** Nodes connect to a single backbone cable. - **Ring:** Nodes connect in a closed loop. 5. Given the "point-to-point link to every other device" description, it perfectly describes a Mesh topology.

Final Answer: The topology is Mesh.

Answer: (A)

Q34.

Solution

Concept:

Time complexity is a measure of the amount of time an algorithm takes to run as a function of the length of the input. In the context of a Stack (a LIFO data structure), the most common operations are 'push' (insertion) and 'pop' (removal). When implemented using an array, these operations are highly optimized.

Solution:

1. **Mechanism:** In an array-based stack, we maintain a variable called 'top' that stores the index of the highest element. 2. **Push Logic:** To perform a 'push' operation: - Increment the 'top' index ($top = top + 1$). - Assign the value to that index ($Array[top] = value$). 3. **Constant Time:** Both of these steps are basic arithmetic and memory assignment operations. They do not depend on how many elements are already in the stack (n). Whether the stack contains 5 elements or 5 million elements, the time required to add one more is exactly the same. 4. **Big O Notation:** In computer science, an operation that takes a constant amount of time regardless of input size is denoted as $O(1)$. 5. **Exceptions:** If the array is full and needs to be "resized" (dynamic array), a single push might take $O(n)$, but the "amortized" or standard time complexity for a standard stack push is considered $O(1)$.

Final Answer: The time complexity is O(1).

Answer: (B)

Q35.

Solution

Concept:

A transaction in a database is a sequence of operations performed as a single logical unit of work. To maintain the ACID properties, a database must have a way to finalize these changes so they become a permanent part of the data store. This is handled by Transaction Control Language (TCL) commands.

Solution:

1. **The ‘COMMIT‘ Command:** This command is used to save all changes made during the current transaction to the database permanently. Once a ‘COMMIT‘ is executed, the changes cannot be undone using a rollback. 2. **The ‘ROLLBACK‘ Command:** In contrast, if an error occurs or the user decides not to proceed, ‘ROLLBACK‘ is used to undo all changes made since the last commit, returning the database to its previous consistent state. 3. **Visibility:** Before a ‘COMMIT‘ is issued, the changes made by a user are typically ”provisional”—they are visible to that user but not to others accessing the database simultaneously (depending on isolation levels). 4. **Context:** In many SQL environments like MySQL (with autocommit off) or Oracle, you must explicitly type ‘COMMIT;‘ to ensure your ‘INSERT‘, ‘UPDATE‘, or ‘DELETE‘ operations are actually written to the disk. 5. **Conclusion:** Since the question asks for the command to **permanently save** changes, ‘COMMIT‘ is the correct term.

Final Answer: The command is COMMIT.

Answer: (C)

Q36.

Solution

Concept:

The ‘split()‘ method in Python is a built-in string function used to break a string into a list based on a specified delimiter (separator). If the delimiter is found within the string, Python removes it and treats the text around it as individual elements of a list. If the delimiter is not found, it returns a list containing the original string as the only element.

Solution:

1. **The Input:** The string is ‘”Informatics”‘ and the method is ‘.split(”i”)‘. 2. **Case Sensitivity:** Python is a case-sensitive language. In the string ‘”Informatics”‘, the first letter is an uppercase ‘’I’‘ and the letter near the end is a lowercase ‘’i’‘. 3. **Execution:** - Python looks for lowercase ‘’i’‘. It finds one occurrence near the end (in the ”...tic...” part). - The string is split at that lowercase ‘’i’‘. - The part before the ‘’i’‘ is ‘”Informat”‘. - The part after the ‘’i’‘ is ‘”cs”‘. 4. **The Resulting List:** The list generated is ‘[’Informat’, ’cs’]‘. 5. **Length Calculation:** The ‘len()‘ function then counts the number of elements in this list. Since there are two strings in the list, the length is 2. 6. **Note:** If the split was on lowercase ‘’n’‘, the result would still be a list of length 2. If it was on uppercase ‘’I’‘, the result would be ‘[”, ’nformatics’]‘, which also has a length of 2.

Final Answer: The result is 2.

Answer: (A)

Q37.

Solution

Concept:

The Transport Layer (Layer 4) of the OSI model acts as the critical liaison between the upper application-oriented layers and the lower network-oriented layers. Its primary responsibility is to ensure that data is delivered error-free, in sequence, and without losses or duplications.

Solution:

1. **End-to-End Communication:** Unlike the Network Layer, which only gets packets from one router to the next, the Transport Layer handles the communication from the source process to the destination process. 2. **Segmentation:** It breaks large data chunks from the session layer into smaller units called ”segments” and reassembles them at the receiving end. 3. **Error Recovery:** It uses mechanisms like Acknowledgments (ACK) and Retransmission to ensure that if a segment is lost during transit, it is sent again. 4. **Flow Control:** It manages the data transmission rate to prevent a fast sender from overwhelming a slow receiver by using techniques like ”sliding windows.” 5. **Protocols:** The most common protocols at this layer are TCP (Transmission Control Protocol), which is connection-oriented and reliable, and UDP (User Datagram Protocol), which is connectionless and faster but less reliable.

Final Answer: The layer is the Transport Layer.

Answer: (B)

Q38.

Solution

Concept:

Selection Sort is a comparison-based sorting algorithm. It works by dividing the input list into a sorted part (at the left) and an unsorted part. In each ”pass,” the algorithm searches the entire unsorted part to find the absolute minimum element and then swaps it with the leftmost element of the unsorted part.

Solution:

1. **Initial List:** ‘[64, 25, 12, 22, 11]‘ 2. **Objective of Pass 1:** Find the minimum value in the entire list (index 0 to 4). 3. **Finding Minimum:** - Compare 64 and 25: 25 is smaller. - Compare 25 and 12: 12 is smaller. - Compare 12 and 22: 12 is smaller. - Compare 12 and 11: 11 is smaller. - The minimum value found is **11**. 4. **The Swap:** Swap the minimum value (11) with the value at the first position (64). 5. **State after Swap:** The 11 moves to index 0, and 64 moves to index 4. 6. **Resulting List:** ‘[11, 25, 12, 22, 64]‘. 7. **Conclusion:** After the first pass, the smallest element is ”selected” and placed in its final sorted position. The rest of the elements remain in their relative order except for the one swapped with the minimum.

Final Answer: The state is [11, 25, 12, 22, 64].

Answer: (A)

Q39.

Solution

Concept:

Working with dates is a common requirement in database management. SQL provides specialized date and time functions to perform arithmetic on date objects.

Solution:

(a) **The `DATEDIFF()` Function:** In most SQL dialects, `DATEDIFF(end_date, start_date)` is used.

(b) **Return Value:** Returns number of days.

(c) **Syntax variation:**

- MySQL: `DATEDIFF('2023-12-31', '2023-12-01')`
- SQL Server: `DATEDIFF(day, '2023-12-01', '2023-12-31')`

(d) **Other Options:** `SUBDATE()` subtracts intervals.

(e) Used for age calculation, durations, etc.

Final Answer: The function is `DATEDIFF()`.

Answer: (A)

Q40.

Solution

Concept:

Object Serialization is the process of converting a data structure or object into a format that can be stored or transmitted. In Python, the `pickle` module is used for this purpose.

Solution:

(a) **Pickling (Serialization):** `pickle.dump(obj, file_object)` converts a Python object into a byte stream and writes it to a file.

(b) **Why use it?** It allows storing complex data structures easily.

(c) **Unpickling:** `pickle.load(file_object)` recreates the object.

(d) **Binary Mode:** Use `'wb'` and `'rb'`.

(e) **Security Note:** Only load trusted data.

Final Answer: It serializes an object and writes it to a file.

Answer: (B)

Q41.

Solution

Concept:

A binary tree is a hierarchical data structure where each node has at most two children, referred to as the left child and the right child. The height of a tree is defined as the number of edges on the longest path from the root to a leaf. Calculating the maximum capacity of a tree involves understanding geometric progression based on the levels of the tree.

Solution:

1. **Defining Levels:** If the height (h) is 3 and the root is at height 0: - Level 0: Root node ($2^0 = 1$ node) - Level 1: Children of the root ($2^1 = 2$ nodes) - Level 2: Grandchildren ($2^2 = 4$ nodes) - Level 3: Great-grandchildren ($2^3 = 8$ nodes) 2. **Summing the Nodes:** To find the maximum total number of nodes in a perfect binary tree, we sum the nodes at each level: 1+2+4+8. 3. **Formula Application:** The general formula for the maximum number of nodes in a binary tree of height h is $2^{h+1} - 1$. 4. **Calculation:** Substituting $h = 3$ into the formula:

$$2^{3+1} - 1 = 2^4 - 1 = 16 - 1 = 15$$

5. **Contextual Note:** It is important to check the definition of height provided in the question. Some textbooks define height as the number of nodes on the longest path (where root is height 1). However, since this question specifies "root is at height 0," the standard edge-based calculation results in 15 nodes.

Final Answer: The maximum number of nodes is 15.

Answer: (B)

Q42.

Solution

Concept:

Email communication relies on a suite of protocols within the TCP/IP stack. While protocols like POP3 and IMAP are used by clients to retrieve or ”pull” messages from a server, the process of sending or ”pushing” mail from a client to a server, or between servers themselves, requires a different protocol.

Solution:

1. **SMTP (Simple Mail Transfer Protocol):** This is the primary protocol used to send electronic mail over the internet. It operates at the Application Layer and typically uses TCP port 25 or 587. 2. **How it Works:** When you hit ”send” on an email, your client uses SMTP to upload the message to your Mail Transfer Agent (MTA). That server then uses SMTP again to relay the message across the internet to the recipient's mail server. 3. **Comparison with Retrieval Protocols:** - **POP3 (Post Office Protocol v3):** Downloads the email from the server to the local device and usually deletes it from the server. - **IMAP (Internet Message Access Protocol):** Synchronizes the email between the server and multiple devices, allowing you to view the same mail on a phone and laptop. 4. **Conclusion:** Because the question specifically asks for the protocol used to **transfer** mail from one server to another, SMTP is the only correct choice. It is the ”delivery” engine of the email world.

Final Answer: The protocol is SMTP.

Answer: (C)

Q43.

Solution

Concept:

The ‘ORDER BY‘ clause in SQL is used to sort the result set of a query in either ascending or descending order. Sorting is a fundamental operation for data analysis, allowing users to see the highest sales, alphabetical lists of names, or chronological events.

Solution:

1. **Default Behavior:** In SQL, if you use the ‘ORDER BY‘ clause without specifying a direction, the database engine defaults to an ascending sort (A to Z, or smallest number to largest). 2. **The ‘ASC‘ Keyword:** To be explicit, programmers use the keyword ‘ASC‘. For example: ‘SELECT * FROM Products ORDER BY Price ASC;‘. 3. **The ‘DESC‘ Keyword:** Conversely, if you want the results in reverse order (highest to lowest), you must use the ‘DESC‘ keyword. 4. **Multiple Columns:** You can sort by multiple columns by separating them with commas. For instance, ‘ORDER BY Department ASC, Salary DESC‘ would sort departments alphabetically and then show the highest salaries within each department first. 5. **Logic:** Sorting does not change the data in the table; it only changes how the data is presented in the output buffer of that specific query. It is a Data Query Language (DQL) feature.

Final Answer: The keyword is ASC.

Answer: (B)

Q44.

Solution

Concept:

In Python, the ‘range()‘ function generates a sequence of numbers. Understanding how to access elements in this sequence using indexing, especially negative indexing, is a key skill. Negative indexing allows you to access elements relative to the end of the sequence.

Solution:

1. **The Sequence:** ‘range(5)‘ generates a sequence of integers starting from 0 up to (but not including) 5. The elements are: ‘0, 1, 2, 3, 4‘. 2. **Indexing:** - Index ‘0‘ refers to ‘0‘. - Index ‘1‘ refers to ‘1‘. - Index ‘4‘ refers to ‘4‘. 3. **Negative Indexing:** Python allows the use of negative integers to count backwards from the end. - Index ‘-1‘ refers to the **last** element in the sequence. - Index ‘-2‘ refers to the second to last element, and so on. 4. **Execution:** In the sequence ‘[0, 1, 2, 3, 4]‘, the last element is ‘4‘. 5. **Evaluation:** Therefore, ‘range(5)[-1]‘ evaluates to ‘4‘. This is a common and efficient way to retrieve the final value of a range without needing to calculate the length of the sequence manually.

Final Answer: The output is 4.

Answer: (B)

Q45.

Solution

Concept:

The relational database model relies on relationships between tables to minimize data redundancy. These relationships are enforced through the use of keys. While a Primary Key identifies a record within its own table, a different type of key is needed to link that record to another table.

Solution:

1. **Foreign Key Definition:** A Foreign Key is a column (or a set of columns) in one table that refers to the Primary Key in another table. It acts as a cross-reference between tables. 2. **Enforcing Integrity:** The primary purpose of a Foreign Key is to maintain ”Referential Integrity.” This ensures that you cannot add a record to the child table if the corresponding ”parent” record does not exist in the primary table. 3. **Example:** Consider a ‘Students‘ table and a ‘Courses‘ table. If the ‘Students‘ table has a column ‘CourseID‘, that column is a Foreign Key. It ensures that a student can only be enrolled in a ‘CourseID‘ that actually exists in the ‘Courses‘ table. 4. **Comparison:** - **Candidate Key:** Any column that has the potential to be a primary key. - **Composite Key:** A primary key made up of more than one column. 5. **Conclusion:** The specific mechanism for creating a relationship by identifying a record in another table is the Foreign Key.

Final Answer: The key is the Foreign Key.

Answer: (C)

Q46.

Solution

Concept:

Malware (Malicious Software) is categorized based on how it spreads and its primary behavior. Two common types of self-spreading malware are Viruses and Worms. While they are often confused, their method of propagation is fundamentally different. A ”Worm” is a standalone program that replicates itself to spread to other computers.

Solution:

1. **Characteristics of a Worm:** A network worm is a sophisticated piece of malware that uses computer networks (like the internet or a LAN) to send copies of itself to other nodes. 2. **Independence:** Unlike a virus, which requires a ”host” file or a human to execute an infected program to spread, a worm is independent. Once it is on a system, it can find other vulnerable computers on the same network and ”crawl” over to them automatically. 3. **Network Impact:** Because worms replicate so rapidly without human intervention, they often consume vast amounts of network bandwidth and system memory, causing significant slowdowns or complete network crashes even if they don't carry a destructive ”payload” (like deleting files). 4. **Other Malware Types:** - **Virus:** Attaches to a file; needs human action to spread. - **Trojan Horse:** Disguises itself as useful software but contains hidden malicious code. - **Spyware:** Secretly monitors user activity and steals sensitive information. 5. **Conclusion:** The key phrase ”without human intervention” and ”over a network” definitively identifies the malware as a Worm.

Final Answer: The malware type is a Worm.

Answer: (B)

Q47.

Solution

Concept:

The ‘SUBSTR()‘ (or ‘SUBSTRING()‘) function in SQL is used to extract a specific portion of a string. It requires three parameters: the source string, the starting position, and the number of characters to extract. In SQL, string indexing is 1-based, meaning the first character is at position 1.

Solution:

1. **The Input String:** ‘’COMPUTER’‘ 2. **The Starting Position:** ‘4‘. - Position 1: C - Position 2: O - Position 3: M - Position 4: **P** 3. **The Length:** ‘3‘. This tells the database to take three characters starting from the ’P’. 4. **Extraction:** - 1st character: P - 2nd character: U - 3rd character: T 5. **Result:** The extracted substring is ‘’PUT’‘. 6. **Common Confusion:** In Python, ‘s[4:3]‘ would return an empty string because the second number is the end index. In SQL, the second number is the ”count” or ”length” of characters to be retrieved, which is a fundamental difference in syntax logic.

Final Answer: The output is ’PUT’.

Answer: (A)

Q48.

Solution

Concept:

Python file modes determine how a file is accessed and where the file pointer is placed upon opening. Choosing the correct mode is essential to prevent accidental data loss. The 'append' mode is specifically designed for adding data to the end of an existing file without disturbing the current content.

Solution:

1. **Mode 'a' (Append):** When a file is opened with ''a'', Python checks if the file exists. If it does, it opens the file and immediately moves the file pointer to the very end. Any subsequent 'write()' operations will add data starting from that point. 2. **Mode 'w' (Write):** If you use ''w'', Python will "truncate" or delete the entire content of the file as soon as it opens, which is dangerous if you want to keep existing data. 3. **Mode 'r+' (Read/Write):** This opens the file for both reading and writing but places the pointer at the **beginning**. If you write immediately, you will overwrite the existing text at the start of the file. 4. **Mode 'w+' (Write/Read):** Similar to 'w', it clears the file first, then allows reading and writing. 5. **Conclusion:** For the specific requirement of writing data at the end of the file, the Append mode (''a'') is the standard and correct choice.

Final Answer: The file mode is 'a'.

Answer: (C)

Q49.

Solution

Concept:

Switching techniques determine how data travels across a network from source to destination. In modern data networks and the internet, the most efficient method is one that does not require a dedicated physical path for the entire duration of the communication.

Solution:

1. **Packet Switching:** In this technology, the data (like a large file or email) is broken down into small units called "packets." Each packet contains the actual data plus a header with the destination address and sequence number. 2. **Independent Routing:** Packets are sent into the network and can take different physical paths depending on network congestion. They are reassembled into the original message at the destination. 3. **Comparison with Circuit Switching:** - **Circuit Switching** (like old landline phones) creates a dedicated physical connection between two points. No one else can use that path until the call ends, which is inefficient for data. 4. **Advantages of Packet Switching:** It maximizes the use of network resources because multiple users can share the same cables simultaneously. If one link fails, packets are simply rerouted. 5. **Conclusion:** The definition of breaking data into packets and sending them independently is the core characteristic of Packet Switching.

Final Answer: The technology is Packet Switching.

Answer: (C)

Q50.

Solution

Concept:

Quick Sort is a highly efficient ”Divide and Conquer” sorting algorithm. It works by selecting a ”pivot” element and partitioning the array into two sub-arrays: elements smaller than the pivot and elements larger than the pivot. While its average performance is excellent, its efficiency can degrade significantly based on the choice of the pivot.

Solution:

1. **The Best/Average Case:** If the pivot consistently divides the list into two nearly equal halves, the complexity is $O(n \log n)$. This is typical for random data. 2. **The Worst Case Scenario:** This occurs when the pivot is the smallest or largest element every single time. This usually happens when the input list is **already sorted** or **reverse sorted**, and the first or last element is chosen as the pivot. 3. **Result of Poor Partitioning:** In this case, instead of dividing the problem in half, the algorithm only reduces the problem size by one element (1 and $n - 1$). This leads to n nested levels of recursion. 4. **Complexity Calculation:** Since there are n levels and each level takes $O(n)$ work for partitioning, the total time complexity becomes $n \times n$. 5. **Big O:** The worst-case time complexity is expressed as $O(n^2)$. Most modern implementations use a ”Randomized Pivot” to avoid this specific performance trap.

Final Answer: The worst-case time complexity is $O(n^2)$.

Answer: (C)

Answer Key

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