CUET-UG Computer Science Sample Paper-4

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 of the following SQL functions is used to find the remainder when one number is divided by another?

(A) DIV()

(B) REMAINDER()

(C) MOD()

(D) SUBSTR()

Q2. A candidate key that is not selected as the primary key of a relation is known as:

(A) Foreign Key

(B) Alternate Key

(C) Secondary Key

(D) Composite Key

Q3. In which network topology is every node connected to a central hub or switch?

(A) Mesh

(B) Bus

(C) Ring

(D) Star

Q4. The process of inserting an element into a stack is called _____, and it follows the _____ principle.

(A) Pop, FIFO

(B) Push, LIFO

(C) Enqueue, LIFO

(D) Dequeue, FIFO

Q5. What is the worst-case time complexity of the Binary Search algorithm?

(A) $O(n)$

(B) $O(n^2)$

(C) $O(\log n)$

(D) $O(1)$

Q6. Which block in Python exception handling is executed regardless of whether an exception occurred or not?

(A) try

(B) except

(C) else

(D) finally

Q7. Which protocol is primarily used for transferring files between a client and a server on a computer network?

(A) HTTP

(B) SMTP

(C) FTP

(D) POP3

Q8. Which SQL function would you use to extract the first three characters from a string column 'City'?

(A) LEFT(City, 3)

(B) SUBSTR(City, 1, 3)

(C) MID(City, 3)

(D) Both (A) and (B)

Q9. In Relational Algebra, the unary operation that selects specific columns from a relation while removing duplicate rows is:

(A) Selection (σ)

(B) Projection (π)

(C) Join ($\bowtie$)

(D) Union ($\cup$)

Q10. What is the unique hardware address assigned to a Network Interface Card (NIC) by the manufacturer?

(A) IP Address

(B) MAC Address

(C) Gateway Address

(D) Subnet Mask

Q11. Convert the following Infix expression to Postfix: $A + B * C$

(A) $ABC * +$

(B) $AB + C*$

(C) $*CBA+$

(D) $ABC + *$

Q12. Which sorting algorithm repeatedly steps through the list, compares adjacent elements, and swaps them if they are in the wrong order?

(A) Selection Sort

(B) Insertion Sort

(C) Bubble Sort

(D) Merge Sort

(E) Quick Sort

Q13. Which method is used to write a list of strings into a text file in Python, where each string represents a line?

(A) write()

(B) writelines()

(C) dump()

(D) writeall()

Q14. A network device that connects multiple computers and uses the MAC address to send data packets only to the specific destination port is a:

(A) Hub

(B) Switch

(C) Repeater

(D) Router

Q15. Which SQL function is used to get the current system date and time?

(A) DATE()

(B) CURDATE()

(C) NOW()

(D) SYSDATE()

Q16. Which of the following is a valid method to remove the last element from a list named 'L' in Python and return it?

(A) L.remove()

(B) L.delete()

(C) L.pop()

(D) L.discard()

Q17. Which SQL aggregate function is used to calculate the total number of non-null values in a specific column?

(A) SUM()

(B) TOTAL()

(C) COUNT()

(D) ADD()

Q18. In a database, which property ensures that once a transaction is committed, it remains so even in the event of a power failure or system crash?

(A) Atomicity

(B) Consistency

(C) Isolation

(D) Durability

Q19. Which network protocol is used to automatically assign IP addresses to devices on a network?

(A) DNS

(B) DHCP

(C) ARP

(D) ICMP

Q20. What is the postfix equivalent of the infix expression: $(A + B)/C$

(A) $AB + C/$

(B) $ABC + /$

(C) $/+ABC$

(D) $AB/C+$

Q21. Which sorting algorithm has a best-case time complexity of $O(n)$ when the input list is already sorted?

(A) Selection Sort

(B) Bubble Sort (Optimized)

(C) Binary Search

(D) Quick Sort

Q22. In Python, what is the default mode when opening a file using the open() function?

(A) 'w' (Write)

(B) 'a' (Append)

(C) 'r' (Read)

(D) 'rb' (Read Binary)

Q23. Which type of malware disguises itself as legitimate software but performs malicious actions once executed?

(A) Worm

(B) Trojan Horse

(C) Spyware

(D) Ransomware

Q24. In SQL, which clause is used to filter the results of an aggregate function?

(A) WHERE

(B) ORDER BY

(C) HAVING

(D) GROUP BY

Q25. Which network device operates at the Physical Layer (Layer 1) and simply regenerates a signal to extend the network distance?

(A) Router

(B) Bridge

(C) Repeater

(D) Gateway

Q26. In a Queue data structure, if elements A, B, and C are inserted in that order, which element will be removed first?

(A) C

(B) B

(C) A

(D) None of the above

Q27. Which Python exception is raised when a calculation exceeds the maximum limit for a numeric data type?

(A) ValueError

(B) TypeError

(C) OverflowError

(D) ZeroDivisionError

Q28. The SQL command used to modify the structure of an existing table, such as adding a new column, is:

(A) UPDATE

(B) ALTER

(C) MODIFY

(D) CHANGE

Q29. Which protocol is used for secure communication over a computer network, widely used on the Internet for banking and shopping?

(A) HTTP

(B) HTTPS

(C) FTP

(D) Telnet

Q30. What is the time complexity of the Linear Search algorithm in the worst-case scenario?

(A) $O(1)$

(B) $O(\log n)$

(C) $O(n)$

(D) $O(n^2)$

Q31. Which of the following sorting algorithms is considered 'stable' because it preserves the relative order of equal elements?

(A) Selection Sort

(B) Quick Sort

(C) Insertion Sort

(D) Heap Sort

Q32. In SQL, which constraint is used to ensure that all values in a column are different?

(A) NOT NULL

(B) CHECK

(C) UNIQUE

(D) DEFAULT

Q33. What is the size of an IPv6 address?

(A) 32 bits

(B) 64 bits

(C) 128 bits

(D) 256 bits

Q34. Which Python function is used to check the data type of a variable?

(A) datatype()

(B) type()

(C) check()

(D) kind()

Q35. Which data structure is best suited for implementing a recursive function call mechanism in a computer system?

(A) Queue

(B) Linked List

(C) Stack

(D) Tree

Q36. Which SQL command is used to delete all rows from a table without deleting the table structure itself, and is categorized as DDL?

(A) DELETE

(B) DROP

(C) TRUNCATE

(D) REMOVE

Q37. Which transmission medium uses light pulses to transmit data at very high speeds over long distances?

(A) Coaxial Cable

(B) Twisted Pair Cable

(C) Fiber Optic Cable

(D) Infrared

Q38. In Python, what will be the result of the expression: 16//3?

(A) 5.33

(B) 5

(C) 1

(D) 5.0

Q39. Which of the following is an example of an open-source operating system?

(A) Windows

(B) macOS

(C) Linux

(D) iOS

Q40. In a relational database, a column that refers to the primary key of another table is called a:

(A) Candidate Key

(B) Super Key

(C) Foreign Key

(D) Secondary Key

Q41. Which searching technique is efficient only when the list is already sorted?

(A) Linear Search

(B) Sequential Search

(C) Binary Search

(D) Hash Search

Q42. Which HTTP status code indicates that a requested page was not found on the server?

(A) 200

(B) 404

(C) 500

(D) 301

Q43. Which of the following methods is used to read all lines from a file and return them as a list of strings in Python?

(A) read()

(B) readline()

(C) readlines()

(D) listread()

Q44. Which network topology has a single backbone cable to which all nodes are connected via drop lines?

(A) Star

(B) Ring

(C) Bus

(D) Mesh

Q45. Which SQL function is used to find the average value of a numeric column?

(A) MEAN()

(B) AVG()

(C) AVERAGE()

(D) SUM()/COUNT()

Q46. A Stack follows which of the following mechanisms?

(A) FIFO

(B) LILO

(C) LIFO

(D) Random Access

Q47. Which layer of the OSI model is responsible for routing data packets across different networks?

(A) Data Link Layer

(B) Physical Layer

(C) Network Layer

(D) Transport Layer

Q48. In Python exception handling, which block is used to catch and handle the error?

(A) catch

(B) try

(C) except

(D) throw

Q49. What does SQL stand for?

(A) Structured Query Language

(B) Standard Query List

(C) Simple Query Language

(D) Sequential Query Logic

Q50. Which sorting algorithm works by selecting the smallest element from the unsorted part and swapping it with the first element of the unsorted part?

(A) Bubble Sort

(B) Insertion Sort

(C) Selection Sort

(D) Merge Sort

Detailed Solutions

Q1.

Solution

Concept:

SQL (Structured Query Language) provides various mathematical functions to perform arithmetic operations on data. To find the remainder of a division, we use the modulo operation.

Solution:

1. In SQL, the MOD() function returns the remainder of one number divided by another. 2. The syntax is typically MOD(dividend, divisor) or using the % operator in some dialects. 3. DIV() is often used for integer division, while REMAINDER() is specific to certain databases like Oracle but MOD() is the standard mapping for academic exams like CUET. 4. SUBSTR() is a string function, not a mathematical one.

Final Answer: The correct function is MOD().

Answer: (C)

Q2.

Solution

Concept:

In a Relational Database Management System (RDBMS), keys are used to uniquely identify records. A table can have multiple keys capable of being the primary key.

Solution:

1. Candidate Keys are all those keys/attributes that can uniquely identify a row in a table. 2. Out of all the Candidate Keys, one is selected by the database designer to be the Primary Key. 3. The remaining Candidate Keys that were not chosen to be the Primary Key are officially termed Alternate Keys. 4. Foreign Keys are used to create links between tables, and Composite Keys are keys consisting of more than one attribute.

Final Answer: An unselected candidate key is an Alternate Key.

Answer: (B)

Q3.

Solution

Concept:

Network topology refers to the physical or logical layout of a network, defining how different nodes (computers, printers, or servers) are connected to one another. Each topology varies in terms of cost, cable length, ease of installation, and fault tolerance. In modern Local Area Networks (LANs), the architecture focuses on centralized management to ensure that data traffic can be efficiently directed without causing collisions or performance bottlenecks across the entire system.

Solution:

1. In a **Star Topology**, all peripheral nodes are individually connected to a central connection point, which is typically a **Hub** or a **Switch**. This central device acts as a server or a "traffic cop," managing the transmission of data between nodes. 2. Because each node has its own dedicated cable connecting it to the center, a failure in one cable only affects that specific node, leaving the rest of the network operational. This provides high reliability compared to a Bus topology. 3. The choice of the central device is crucial: a **Hub** broadcasts data to every port, whereas a **Switch** (the modern standard) is intelligent enough to send data only to the specific destination port by reading the MAC address. 4. While this topology requires more cabling than a Bus or Ring layout, it is much easier to troubleshoot and expand. If the central hub itself fails, however, the entire network becomes incapacitated, which is the primary disadvantage of this structure. 5. In contrast, a Mesh topology involves multiple interconnections between nodes, a Bus uses a single backbone cable, and a Ring connects nodes in a continuous closed loop where data travels in one direction.

Final Answer: The topology described is the Star Topology.

Answer: (D)

Q4.

Solution

Concept:

A Stack is a linear data structure that serves as a collection of elements with two principal operations: adding an item and removing an item. Unlike an array, where you can access any index at any time, a stack restricts access to only one end, known as the ”top.” This behavioral constraint makes stacks ideal for specific computational tasks, such as managing function calls in memory (the call stack), reversing strings, or evaluating mathematical expressions in compilers.

Solution:

1. The specific operation of adding an element to the top of a stack is formally known as a **Push** operation. When an item is pushed, it becomes the new ”top” of the stack. 2. Conversely, removing an element from the stack is known as a **Pop** operation. This operation always removes the element that was most recently added. 3. Because of this restricted entry and exit point, stacks follow the **LIFO** principle, which stands for **Last-In, First-Out**. This means the most recent object placed on the stack is the first one to be taken off. 4. To visualize this, consider a stack of physical trays in a cafeteria: the tray placed last on the top of the pile is the first one a customer picks up. 5. In the context of the given options: **Enqueue** and **Dequeue** are terms used specifically for **Queues**, which follow the **FIFO** (First-In, First-Out) principle. Therefore, the combination of ”Push” and ”LIFO” is the only technically accurate description for a stack’s behavior.

Final Answer: The process is Push and the principle is LIFO.

Answer: (B)

Q5.

Solution

Concept:

Time complexity is a representation of the computational complexity that describes the amount of time an algorithm takes to run as a function of the length of the input. For searching algorithms, efficiency is measured by how many comparisons the algorithm must perform to find a target value in the worst possible scenario (the "Worst-Case"). Binary Search is a highly efficient "Divide and Conquer" algorithm, but it carries a strict prerequisite: the dataset must be sorted before the search begins.

Solution:

1. Binary Search works by repeatedly halving the search space. It starts by comparing the target value to the middle element of the sorted list. 2. If the target is smaller than the middle element, the algorithm discards the upper half of the list and repeats the process on the lower half. If it is larger, it discards the lower half and focuses on the upper half. 3. Every single comparison reduces the number of elements to be searched by exactly half. 4. Mathematically, if we have n elements, we want to know how many times we can divide n by 2 before we are left with only one element. This is represented by the equation $2^k = n$, where k is the number of steps. Solving for k, we get $k = \log_2 n$. 5. Therefore, the maximum number of comparisons needed is proportional to the logarithm of the input size. Even if the list grows exponentially (e.g., from 1,000 to 1,000,000 items), the number of steps only increases slightly (from 10 to 20). 6. In contrast, a Linear Search has a worst-case complexity of $O(n)$ because it may have to check every single element in the list.

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

Answer: (C)

Q6.

Solution

Concept:

Python provides a robust mechanism for error handling using the try-except-finally blocks. This ensures that the program can handle unexpected errors gracefully without crashing and can perform cleanup actions.

Solution:

1. The **try** block contains the code that might raise an exception. 2. The **except** block handles the exception if one occurs within the try block. 3. The **else** block runs only if no exceptions were raised in the try block. 4. The **finally** block is the "cleanup" block. It is guaranteed to execute whether an exception was raised or not, and even if the program encountered a return statement. This makes it ideal for closing files or releasing network resources.

Final Answer: The block that always executes is finally.

Answer: (D)

Q7.

Solution

Concept:

Network protocols are sets of rules that govern how data is transmitted across a network. Different protocols are designed for different types of data exchange, such as web browsing, email, or file transfer.

Solution:

1. HTTP (HyperText Transfer Protocol) is used for transferring web pages and related content. 2. SMTP (Simple Mail Transfer Protocol) and POP3 (Post Office Protocol) are specifically used for sending and receiving emails. 3. FTP (File Transfer Protocol) is a standard communication protocol used specifically for the transfer of computer files from a server to a client, or vice versa, on a computer network. 4. FTP allows users to upload, download, and manage files on a remote server efficiently.

Final Answer: The protocol used for transferring files is FTP.

Answer: (C)

Q8.

Solution

Concept:

In SQL, string manipulation functions allow users to extract portions of text from a column. Depending on the SQL dialect (MySQL, PostgreSQL, etc.), different functions may achieve the same result.

Solution:

1. The LEFT(string, n) function returns the specified number of characters from the left side of the string. So, LEFT(City, 3) extracts the first 3 characters. 2. The SUBSTR(string, start, length) function (also known as SUBSTRING) extracts a substring. By setting the start position to 1 and the length to 3, we get the first 3 characters. 3. MID() is another synonym used in some databases like MySQL to perform the same task as SUBSTR. 4. In competitive exams like CUET, knowledge of multiple syntax variations is often tested. Both LEFT() and SUBSTR() are standard ways to achieve this.

Final Answer: Both (A) and (B) are correct methods.

Answer: (D)

Q9.

Solution

Concept:

Relational Algebra is a formal language used to model the data stored in relational databases and the operations performed on them. Unary operations are those that operate on a single relation.

Solution:

1. Selection (σ) is used to filter rows (tuples) based on a specific condition. 2. Projection (π) is used to select specific columns (attributes) from a table. A key characteristic of the Projection operation is that it automatically removes duplicate rows from the result set to maintain the mathematical definition of a relation as a set. 3. Join and Union are binary operations because they require two relations to operate. 4. Since the question specifies selecting columns and removing duplicates, it refers to Projection.

Final Answer: The operation is Projection (π).

Answer: (B)

Q10.

Solution

Concept:

Network identification involves two primary types of addresses: Logical addresses (IP) and Physical addresses (MAC).

Solution:

1. An IP Address is a logical address assigned by the network software and can change depending on the network the device is connected to. 2. A MAC (Media Access Control) Address is a unique physical identifier assigned to a Network Interface Card (NIC) by the manufacturer at the time of production. 3. It is a 48-bit (6-byte) address usually represented in hexadecimal format (e.g., 00:1A:2B:3C:4D:5E). 4. Unlike IP addresses, the MAC address is ”burned” into the hardware and remains constant regardless of the network environment.

Final Answer: The unique hardware address is the MAC Address.

Answer: (B)

Q11.

Solution

Concept:

The conversion from Infix notation (where operators are between operands) to Postfix notation (where operators follow operands) is a fundamental application of the Stack data structure. This process is governed by operator precedence and associativity rules. In the standard hierarchy, exponentiation has the highest precedence, followed by multiplication and division (which share equal priority), and finally addition and subtraction. When an expression contains operators of the same precedence, they are typically evaluated from left to right (left-associative).

Solution:

1. To convert $A + B * C$, we begin by scanning the expression from left to right. 2. The first symbol is the operand A, which is immediately placed into the postfix output. Output: A. 3. The next symbol is the operator +. Since the stack is empty, we push + onto the stack. Stack: $[+]$. 4. The next symbol is the operand B, which is added to the output. Output: AB. 5. The next symbol is the operator $*$. We compare its precedence with the operator at the top of the stack (+). Since multiplication $(*)$ has a higher precedence than addition $(+)$, we do not pop the stack; instead, we push $*$ onto the stack. Stack: $[+, *]$. 6. The final symbol is the operand C, which is added to the output. Output: ABC. 7. After scanning the entire expression, we pop all remaining operators from the stack one by one and append them to the output. 8. First, $*$ is popped. Output: $ABC*$. 9. Next, + is popped. Output: $ABC * +$.

Final Answer: The resulting postfix expression is $ABC * +$.

Answer: (A)

Q12.

Solution

Concept:

Sorting is the process of arranging data in a specific order (ascending or descending). Bubble Sort is one of the simplest sorting algorithms. It is an internal, comparison-based sort named for the way smaller or larger elements ”bubble” to the top of the list through successive swaps. While it is highly intuitive and easy to implement, it is inefficient for large datasets compared to algorithms like Quick Sort or Merge Sort, having an average and worst-case time complexity of $O(n^2)$.

Solution:

1. Bubble Sort works by starting at the beginning of the list and comparing the first two elements. 2. If the first element is greater than the second (for ascending order), they are swapped. 3. This comparison and potential swap move to the next pair (2nd and 3rd), then the 3rd and 4th, and so on, until the end of the list is reached. 4. After the first complete pass through the list, the largest element is guaranteed to have ”bubbled up” to its correct final position at the end of the list. 5. The process repeats for the remaining $n - 1$ elements, then $n - 2$, until no more swaps are required or all passes are completed. 6. Other algorithms work differently: Selection Sort finds the minimum element and places it at the start; Insertion Sort builds a sorted array one element at a time; Merge Sort uses a divide-and-conquer approach. 7. The specific mechanism of ”comparing adjacent elements and swapping” is the defining characteristic of Bubble Sort.

Final Answer: The algorithm described is Bubble Sort.

Answer: (C)

Q13.

Solution

Concept:

File handling in Python allows developers to create, read, update, and delete files. When dealing with text files, Python provides built-in methods through the file object. Writing data to a file can be done in multiple ways depending on whether the data is a single string or a collection of strings. Understanding the difference between these methods is crucial for efficient data persistence and ensuring that formatting (like line breaks) is handled correctly by the programmer.

Solution:

1. The **write()** method takes a single string as an argument and writes it to the file. To write multiple lines with write(), one would need to use a loop or manually concatenate strings with newline characters (\n). 2. The **writelines()** method is specifically designed to take an iterable (like a list or a tuple of strings) and write its contents to the file sequentially. 3. It is important to note that writelines() does not automatically add newline characters at the end of each string in the list; the programmer must include \n within the strings if they want separate lines in the text file. 4. The **dump()** function is part of the 'pickle' or 'json' modules and is used for object serialization, not for standard text file writing. 5. Therefore, for a list of strings intended for a text file, writelines() is the most direct and appropriate method.

Final Answer: The correct method is writelines().

Answer: (B)

Q14.

Solution

Concept:

Network hardware devices serve as the backbone of data communication. Each device operates at a different layer of the OSI (Open Systems Interconnection) model and possesses varying levels of "intelligence" regarding how it handles data packets. A Hub is a physical layer device that broadcasts data to all ports, while a Switch is a data link layer device (Layer 2) that can identify specific hardware on a network.

Solution:

1. When a data packet enters a Switch, the switch looks at the destination MAC address in the packet header. 2. The Switch maintains a MAC address table (also called a CAM table) that maps MAC addresses to the specific physical ports where those devices are connected. 3. Instead of broadcasting the data to every connected device (which creates unnecessary traffic and security risks, as a Hub does), the Switch establishes a temporary point-to-point connection between the sender and the receiver. 4. This process significantly improves network efficiency and bandwidth utilization because multiple pairs of devices can communicate simultaneously without collisions. 5. Routers operate at the Network Layer (Layer 3) and use IP addresses to route data between different networks, whereas the question specifically focuses on MAC addresses within a network.

Final Answer: The device described is a Switch.

Answer: (B)

Q15.

Solution

Concept:

Database management systems often need to record the exact moment a transaction occurs or a record is updated. SQL provides built-in temporal functions to fetch the current date and time from the operating system where the database server is running. While different SQL versions (MySQL, Oracle, SQL Server) have slight variations in function names, certain functions are widely recognized in standard academic curricula.

Solution:

1. The **DATE()** function is generally used to extract the date part from a date-time expression. 2. The **CURDATE()** function returns only the current date (YYYY-MM-DD) without the time component. 3. The **NOW()** function is a very common function in MySQL and other systems that returns the current date and time as a single value (YYYY-MM-DD HH:MM:SS). 4. The **SYSDATE()** function is also used to return the current date and time. In many systems, NOW() returns the time at which the statement began to execute, while SYSDATE() returns the exact time at which the function itself executes. 5. In the context of most Computer Science examinations and standard SQL practice, both NOW() and SYSDATE() are valid for obtaining both date and time, but NOW() is the most frequently cited answer in general MCQ contexts for "current system date and time."

Final Answer: The function used is NOW().

Answer: (C)

Q16.

Solution

Concept:

Python lists are dynamic arrays that provide several built-in methods for data manipulation. Removing elements is a frequent operation, and Python offers different ways to do so depending on whether you know the value, the index, or simply want to treat the list like a stack. Understanding the return value of these methods is essential because some modify the list in place without returning the removed item, while others return the item to the caller for further processing.

Solution:

1. The **remove()** method is used to delete the first occurrence of a specific value. It does not return the value; it returns None. If the value is not found, it raises a ValueError. 2. The **delete** keyword (used as del L[index]) removes an item at a specific position but does not return it. 3. The **pop()** method is specifically designed to remove and return an element from the list. By default, if no index is specified (i.e., L.pop()), it removes and returns the very last element in the list. This makes the list behave like a LIFO (Last-In-First-Out) stack. 4. The **discard()** method is not a list method; it is used with sets to remove an element without raising an error if the element is missing. 5. In academic examinations, pop() is highlighted as the primary method for stack-based operations where the removed data needs to be captured in a variable.

Final Answer: The correct method is L.pop().

Answer: (C)

Q17.

Solution

Concept:

Aggregate functions in SQL perform a calculation on a set of values and return a single value. These are widely used for generating reports and summarizing data. A critical aspect of aggregate functions that is frequently tested in exams is how they handle NULL values. In a relational database, NULL represents missing or unknown data, and most aggregate functions are designed to ignore these entries to ensure the mathematical result is representative of the actual data present.

Solution:

1. The **SUM()** function calculates the total numerical addition of values in a column, ignoring NULLs. 2. The **COUNT()** function has two primary uses: COUNT(*) and COUNT(column_name). 3. COUNT(*) counts every row in the table, regardless of whether the columns contain NULL values or not. 4. **COUNT(column_name)**, however, specifically counts only the non-null entries within that specific column. This is the precise functionality required by the question. 5. TOTAL() and ADD() are not standard SQL aggregate functions; SUM() and COUNT() are the standard keywords. 6. When an examiner asks for the "number of values," they are looking for a count, whereas "total" usually implies a sum of numerical values. Therefore, COUNT() is the logically correct choice for finding the number of entries.

Final Answer: The aggregate function used is COUNT().

Answer: (C)

Q18.

Solution

Concept:

The ACID properties (Atomicity, Consistency, Isolation, Durability) are a set of principles that guarantee database transactions are processed reliably. In the context of database management systems (DBMS), a transaction is a single logical unit of work. Without ACID properties, a system crash during a bank transfer, for example, could result in money being debited from one account without being credited to another.

Solution:

1. **Atomicity** ensures that a transaction is ”all or nothing”—either the entire transaction succeeds or none of it is applied. 2. **Consistency** ensures that a transaction brings the database from one valid state to another, maintaining all predefined rules and constraints. 3. **Isolation** ensures that concurrent execution of transactions leaves the database in the same state as if they were executed sequentially. 4. **Durability** is the specific property concerned with the permanence of the data. It guarantees that once a user receives a confirmation that a transaction has been committed, the changes are saved to non-volatile memory (like a hard drive). Even if the system crashes or loses power immediately after the commit, the data will be recovered and intact upon restart. 5. This is achieved through transaction logs that record changes before they are permanently written to the data files.

Final Answer: The property that ensures permanence is Durability.

Answer: (D)

Q19.

Solution

Concept:

Network administration involves managing how devices identify each other. Manually assigning an IP address to every computer, smartphone, and IoT device in a large organization would be prone to errors and incredibly time-consuming. To solve this, networking protocols were developed to automate the management of logical addresses within a local area network (LAN).

Solution:

1. **DNS** (Domain Name System) translates human-readable domain names (like www.google.com) into IP addresses but does not assign IPs to devices. 2. **DHCP** (Dynamic Host Configuration Protocol) is a network management protocol used on UDP/IP networks. A DHCP server dynamically assigns an IP address and other network configuration parameters (like subnet mask and default gateway) to each device on a network so they can communicate with other IP networks. 3. When a device connects to the network, it broadcasts a "DHCP Discover" packet. The server responds with an "Offer," and once accepted, the device is "leased" an IP address for a specific duration. 4. **ARP** (Address Resolution Protocol) is used to find a MAC address when only the IP address is known. 5. **ICMP** (Internet Control Message Protocol) is used for diagnostic and error-reporting purposes, such as the 'ping' command.

Final Answer: The protocol used for automatic IP assignment is DHCP.

Answer: (B)

Q20.

Solution

Concept:

Postfix notation (also known as Reverse Polish Notation) is used by compilers to evaluate mathematical expressions without the need for parentheses to define the order of operations. The transformation from Infix to Postfix follows a strict hierarchy of operations: Parentheses have the highest priority, followed by exponents, then multiplication/division, and finally addition/subtraction. In postfix, the operator always follows its operands.

Solution:

1. The given expression is $(A + B)/C$. 2. According to the rules of precedence, we must evaluate the expression inside the parentheses first, even though addition generally has lower priority than division. 3. Within the parentheses $(A + B)$, the operands are A and B, and the operator is $+$. In postfix, this becomes $AB+$. 4. We can now treat $(AB+)$ as a single operand. Let's call it X. 5. The remaining expression is X/C. 6. In postfix, the operator $/$ moves to the end. This becomes $XC/$. 7. Substituting X back with its postfix value $(AB+)$, the full expression becomes $AB + C/$. 8. This format allows a computer using a stack to push A, push B, see the $+$ and add them, then push C, see the $/$ and divide the previous sum by C. 9. Note that without the parentheses, $A+B/C$ would have resulted in $ABC/+$.

Final Answer: The postfix equivalent is $AB + C/$.

Answer: (A)

Q21.

Solution

Concept:

Time complexity is a computational measure that describes the amount of time an algorithm takes to run as a function of the length of the input. Most sorting algorithms, such as Selection Sort or standard Quick Sort, have a worst-case and average-case complexity of $O(n^2)$ or $O(n \log n)$. However, certain algorithms can be ”optimized” to recognize when a list is already in the correct order. This is particularly relevant for the Bubble Sort algorithm, where the primary operation is comparing adjacent elements and swapping them if they are out of order.

Solution:

1. In a standard Bubble Sort, the algorithm performs $n - 1$ passes regardless of the initial order of the elements. 2. An **Optimized Bubble Sort** introduces a ”flag” (a boolean variable) to monitor whether any swaps occurred during a specific pass. 3. If the algorithm completes a full pass through the list without performing a single swap, it concludes that the list is already sorted and terminates immediately. 4. When an already sorted list is provided as input, the optimized algorithm will make exactly one pass ($n - 1$ comparisons) to verify the order. Since it only does one linear scan, the time complexity is $O(n)$. 5. In contrast, Selection Sort always takes $O(n^2)$ because it must scan the remaining unsorted part to find the minimum element every time. Quick Sort typically takes $O(n \log n)$ but can degrade to $O(n^2)$ depending on pivot selection. Binary Search is a searching algorithm, not a sorting one.

Final Answer: Optimized Bubble Sort has a best-case complexity of $O(n)$.

Answer: (B)

Q22.

Solution

Concept:

File I/O (Input/Output) in Python is handled through the built-in **open()** function. This function creates a file object and serves as a gateway to read from or write to a file stored on the disk. When calling this function, the programmer can specify the ”mode,” which dictates the permissions and the starting position of the file pointer. If the mode is not explicitly provided by the user, Python defaults to a safe setting to prevent accidental data loss.

Solution:

1. The **'w'** (Write) mode opens a file for writing. If the file exists, it truncates (erases) the content. If it doesn't exist, it creates a new one. 2. The **'a'** (Append) mode opens a file for writing but keeps existing content, adding new data to the end of the file. 3. The **'r'** (Read) mode is the **default mode** in Python. It opens the file for reading only. If the specified file does not exist, Python raises a FileNotFoundError. 4. The **'rb'** mode is used for reading binary files (like images or compiled code), which is not the default for standard text files. 5. Using 'r' as the default is a design choice by Python to ensure that simply opening a file (perhaps by mistake or omission of arguments) does not lead to the deletion of important data, which would happen if 'w' were the default.

Final Answer: The default mode is 'r' (Read).

Answer: (C)

Q23.

Solution

Concept:

Malware (Malicious Software) is an umbrella term for various types of intrusive or hostile software. In computer security, understanding the behavior and delivery method of malware is essential for defense. Malicious programs are categorized based on how they replicate, how they hide, and the specific damage they cause. One of the most common threats relies on social engineering—tricking the user into believing the software is helpful or necessary.

Solution:

1. A **Worm** is standalone malware that replicates itself to spread to other computers, often using a network without any human intervention. 2. A **Trojan Horse** (or simply Trojan) is named after the Ancient Greek story. It is a type of malware that **disguises itself** as a legitimate, useful, or harmless program (like a free game, a system update, or an antivirus tool). Once the user executes the program, it performs its hidden malicious payload, such as stealing data or creating a ”backdoor” for hackers. 3. **Spyware** focuses on gathering information about a person or organization without their knowledge. 4. **Ransomware** encrypts a user's files and demands payment (ransom) to regain access. 5. The distinguishing feature mentioned in the question—disguising as legitimate software—is the classic definition of a Trojan Horse.

Final Answer: The malware described is a Trojan Horse.

Answer: (B)

Q24.

Solution

Concept:

In SQL, filtering data is a primary task. Most users are familiar with the **WHERE** clause, which filters individual rows based on a specific condition before any grouping or aggregation occurs. However, when we use aggregate functions like SUM(), COUNT(), or AVG(), we often need to filter the ”results” of those calculations. Because of the SQL order of execution, the WHERE clause cannot be used with aggregate functions directly.

Solution:

1. The **GROUP BY** clause is used to arrange identical data into groups. For example, grouping sales by ”Region.” 2. The **HAVING** clause was specifically added to SQL to allow users to filter the data produced by a GROUP BY clause based on an aggregate condition. 3. For instance, if you want to see only those regions where the total sales (SUM) exceed ₹ 10,000, you must use: *GROUP BY Region HAVING SUM(Sales) > 10000*. 4. **WHERE** filters the source data; **HAVING** filters the summarized data. 5. **ORDER BY** is used for sorting the final result set in ascending or descending order and does not perform any filtering. 6. Therefore, the HAVING clause is the only logical choice for filtering based on aggregate results.

Final Answer: The clause used to filter aggregate results is HAVING.

Answer: (C)

Q25.

Solution

Concept:

A Repeater is a network device used to regenerate or replicate a signal. As electronic signals travel along a cable, they degrade and become weaker due to resistance and interference, a phenomenon known as attenuation. If a cable is too long, the signal may become unreadable by the receiving device. Repeaters are used to overcome this physical limitation by receiving the weakened signal and retransmitting it at its original strength.

Solution:

1. The OSI (Open Systems Interconnection) model consists of seven layers. The **Physical Layer (Layer 1)** is the lowest layer, dealing with the actual physical connection and the transmission of raw bits over a medium. 2. A **Repeater** operates strictly at Layer 1. It does not look at MAC addresses, IP addresses, or the data content. It simply receives bits and repeats them. 3. A **Router** operates at Layer 3 (Network Layer) and makes decisions based on IP addresses. 4. A **Bridge** operates at Layer 2 (Data Link Layer) and uses MAC addresses to filter traffic. 5. A **Gateway** can operate at any level but is generally associated with higher layers as it translates between different protocol suites. 6. Since the question specifies ”simply regenerates a signal” and ”Layer 1,” the device is a Repeater.

Final Answer: The device is a Repeater.

Answer: (C)

Q26.

Solution

Concept:

A Queue is a fundamental linear data structure characterized by the First-In-First-Out (FIFO) principle. This means that the first element added to the collection will be the first one to be removed. It mimics real-world scenarios, such as a line of people waiting for a service, where the person who arrives first is served first. In a queue, insertions occur at one end called the ”Rear” (enqueue), and deletions occur at the opposite end called the ”Front” (dequeue). This separation of entry and exit points is what distinguishes it from a Stack.

Solution:

1. According to the problem description, three elements are inserted into the queue in a specific sequence: First A, then B, and finally C. 2. In the internal representation of the queue, A occupies the ”Front” position because it was the first to arrive. B is in the middle, and C occupies the ”Rear” position as the most recent addition. 3. When a ”remove” or ”dequeue” operation is invoked, the rules of the FIFO data structure dictate that the element currently at the Front must be processed and removed. 4. Since A was the very first element to enter the system, it is the one currently at the Front of the line. 5. Consequently, no matter how many elements (B, C, or more) are added after A, they must wait until A is removed before they can move to the Front. Therefore, A will be the first to be removed. 6. This structure is essential in computer science for tasks like printer spooling, CPU task scheduling, and handling asynchronous data transfers where order of arrival must be strictly preserved.

Final Answer: The element removed first will be A.

Answer: (C)

Q27.

Solution

Concept:

Python uses exceptions to signal that an error has occurred during the execution of a script. While syntax errors are caught during the parsing stage, exceptions are ”runtime errors” that occur when the logic of the code hits a boundary condition the computer cannot handle. Python has a rich hierarchy of built-in exceptions, allowing developers to catch specific types of errors rather than using a generic ”catch-all” approach. This specificity is crucial for debugging and creating resilient software that can recover from predictable mathematical or logical failures.

Solution:

1. **ValueError** is raised when a function receives an argument that has the right type but an inappropriate value (e.g., trying to find the square root of a negative number). 2. **TypeError** occurs when an operation or function is applied to an object of an inappropriate type (e.g., adding a string to an integer). 3. **OverflowError** is the specific exception raised when the result of an arithmetic calculation is too large to be represented by the computer's numerical data type. While Python's integers can theoretically grow to any size limited by available memory, floating-point numbers (floats) have a fixed bit-size. If a float operation (like 2.0^{10000}) exceeds the maximum limit, an OverflowError is triggered. 4. **ZeroDivisionError** is a specific type of arithmetic error raised when the second argument of a division or modulo operation is zero. 5. Given that the question focuses on a calculation exceeding the maximum numeric limit, OverflowError is the correct technical designation for this event.

Final Answer: The exception raised is OverflowError.

Answer: (C)

Q28.

Solution

Concept:

SQL commands are broadly categorized into several sub-languages, the two most common being DDL (Data Definition Language) and DML (Data Manipulation Language). DML is used to manage the data within the rows (using commands like INSERT, UPDATE, and DELETE). However, when we need to change the blueprint or the container of that data—the table structure itself—we must use DDL commands. Understanding the distinction between modifying the ”content” of a row versus the ”definition” of a column is a cornerstone of database administration.

Solution:

1. The **UPDATE** command is a DML operation. It is used to change the actual values stored inside existing rows of a table. It cannot add a new column or change a data type. 2. The **ALTER** command is a DDL operation. It is used to modify the structure of an existing database object, such as a table, view, or database. 3. Specifically, to add a new column to a table, the syntax is *ALTER TABLE table_name ADD column_name datatype;*. 4. **MODIFY** and **CHANGE** are often keywords used ”within” an ALTER statement (e.g., *ALTER TABLE ... MODIFY ...*), but they are not the primary commands themselves. 5. In academic exams like CUET, the distinction is clear: if you are changing the schema (the ”skeleton”) of the table, you are ”altering” it. If you are changing the data (the ”flesh”), you are ”updating” it. Since the question asks about adding a column, ALTER is the correct answer.

Final Answer: The command used is ALTER.

Answer: (B)

Q29.

Solution

Concept:

The modern internet relies heavily on data privacy and integrity. Standard web traffic originally used HTTP (HyperText Transfer Protocol), which sends data in ”plain text.” This means that anyone sitting between your computer and the server (like a hacker on a public Wi-Fi) could potentially read sensitive information, such as passwords or credit card numbers. To combat this, security protocols were integrated into web browsing to ensure that even if data is intercepted, it remains unreadable to unauthorized parties.

Solution:

1. **HTTP** (HyperText Transfer Protocol) is the foundation of data communication for the World Wide Web, but it lacks encryption, making it insecure for financial transactions. 2. **HTTPS** stands for HyperText Transfer Protocol **Secure**. It is essentially HTTP layered on top of the SSL/TLS protocol. 3. This protocol provides three key layers of protection: **Encryption** (encoding the data so it cannot be read by eavesdroppers), **Data Integrity** (ensuring data cannot be modified during transfer without being detected), and **Authentication** (proving that the user is communicating with the intended website and not an imposter). 4. **FTP** (File Transfer Protocol) is used for file movement and is not a web-browsing security protocol. 5. **Telnet** is an old protocol used for remote access to computers, which is notoriously insecure because it sends all data, including passwords, in plain text. 6. For secure activities like banking, shopping, and entering personal details, HTTPS is the global standard denoted by the ”padlock” icon in browser address bars.

Final Answer: The secure communication protocol is HTTPS.

Answer: (B)

Q30.

Solution

Concept:

Searching is the process of finding the location of a specific element (the ”key”) within a collection of data. Linear Search (or Sequential Search) is the most basic search algorithm. It does not require the data to be sorted, which is its main advantage. However, because it makes no assumptions about the arrangement of the data, it must be prepared for the possibility that the item it is looking for is at the very end of the list or not in the list at all.

Solution:

1. In a Linear Search, the algorithm starts at the first element (index 0) and compares it with the target key. 2. If they do not match, it moves to the second element, then the third, and continues sequentially through the list. 3. The **worst-case scenario** occurs when: a) The target element is at the very last position in a list of n elements. b) The target element is not present in the list at all. 4. In both of these cases, the algorithm must perform exactly n comparisons to reach a conclusion. 5. Therefore, the time taken grows ”linearly” with the number of elements in the input. In Big O notation, this is expressed as $O(n)$. 6. While other algorithms like Binary Search are faster ($O(\log n)$), they require the list to be sorted. Linear Search remains $O(n)$ regardless of the order, making it inefficient for very large datasets but simple for small, unsorted ones.

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

Answer: (C)

Q31.

Solution

Concept:

In the study of sorting algorithms, ”stability” is a critical property that refers to the ability of an algorithm to maintain the relative order of records with equal keys (values). If a list contains two identical elements, say 5 (red) and 5 (blue) in that order, a stable sorting algorithm ensures that in the sorted output, 5 (red) still appears before 5 (blue). This property is essential when sorting complex data structures where multiple fields are involved, such as sorting a list of students first by name and then by their exam marks.

Solution:

1. **Selection Sort** is unstable because it swaps non-adjacent elements over long distances, which can easily change the relative positions of equal keys. 2. **Quick Sort** is generally unstable in its most efficient in-place implementations for the same reason—the partitioning logic involves swapping elements across the pivot. 3. **Insertion Sort** is a **stable** algorithm. It works by taking one element at a time and ”inserting” it into its correct position within the already-sorted prefix. During the comparison phase, if an element is found to be equal to the current key, the insertion sort stops moving the key and places it immediately after the existing equal element, thus preserving their original relative order. 4. **Heap Sort** is also unstable due to the way the heap structure is built and elements are swapped during the ”heapify” process. 5. In competitive programming and database management, choosing a stable sort like Insertion Sort (for small lists) or Merge Sort (for large lists) is vital when the integrity of the original sequence matters for duplicate values.

Final Answer: The stable sorting algorithm is Insertion Sort.

Answer: (C)

Q32.

Solution

Concept:

Constraints in a relational database management system (RDBMS) are rules applied to table columns to enforce data integrity. They ensure that the data entered into the database remains accurate, reliable, and consistent with the intended business logic. Constraints can be applied at the column level or the table level. If a user attempts to perform an action that violates a constraint (such as entering a duplicate ID in a primary key column), the RDBMS will reject the transaction and return an error.

Solution:

1. The **NOT NULL** constraint ensures that a column cannot have a NULL (empty) value. It does not prevent duplicate values; it only mandates that some data must be present. 2. The **CHECK** constraint allows you to specify a condition that every value in the column must satisfy, such as $Age > 18$. 3. The **UNIQUE** constraint is used to ensure that every single value in a column is distinct from every other value in that same column. While a Primary Key also enforces uniqueness, a table can have multiple UNIQUE constraints, whereas it can only have one Primary Key. Furthermore, UNIQUE constraints usually allow one NULL value, whereas Primary Keys do not. 4. The **DEFAULT** constraint provides a predefined value for a column if no value is specified during an INSERT operation. 5. For the specific requirement of ensuring ”all values are different,” the UNIQUE constraint is the standard implementation in SQL.

Final Answer: The correct constraint is UNIQUE.

Answer: (C)

Q33.

Solution

Concept:

An IP (Internet Protocol) address is a numerical label assigned to each device connected to a computer network. As the number of devices on the internet grew exponentially with the rise of smartphones and IoT, the original addressing system, IPv4, ran out of available addresses. To solve this, the Internet Engineering Task Force (IETF) developed a new version, IPv6, which uses a much larger address space. This transition is fundamental to the continued growth of the global internet.

Solution:

1. **IPv4** uses a 32-bit address scheme, which allows for approximately 4.3 billion unique addresses. These are typically written in "dotted-decimal" notation (e.g., 192.168.1.1). 2. **IPv6** was designed to provide a virtually inexhaustible supply of addresses. It uses a **128-bit** address space. 3. This 128-bit size allows for 2^{128} addresses, which is roughly 3.4×10^{38} addresses—enough to assign trillions of addresses to every single person on Earth. 4. IPv6 addresses are represented in hexadecimal notation and divided into eight groups of four hexadecimal digits, separated by colons (e.g., 2001:0db8:85a3:0000:0000:8a2e:0370:7334). 5. The increase from 32 bits to 128 bits is not just a fourfold increase in the number of addresses, but an exponential increase that fundamentally changes how global networking and subnetting are handled.

Final Answer: The size of an IPv6 address is 128 bits.

Answer: (C)

Q34.

Solution

Concept:

Python is a dynamically typed language, meaning that you do not need to declare the type of a variable when you create it. The Python interpreter determines the type based on the value assigned to the variable. However, during the development of complex programs or when debugging, it is often necessary to verify exactly what kind of data an object holds. Python provides built-in reflection capabilities to inspect objects at runtime, which is a hallmark of its flexibility as a high-level language.

Solution:

1. Python does not use ”datatype()” or ”check()” as keywords for type identification; these are common misconceptions from other programming languages. 2. The **type()** function is the standard built-in function used for this purpose. When you pass an object or variable to type(), it returns a ”type object” representing the class of that object (e.g., <class ’int’>, <class ’str’>, or <class ’list’>). 3. For example, *print(type(10))* will output *<class ’int’>*. 4. Another related function is **isinstance()**, which is often preferred in production code because it supports inheritance checks (checking if an object is an instance of a class or a subclass thereof). 5. However, for a simple check of ”what is this variable,” type() is the most direct and common answer in academic curricula. It allows programmers to perform conditional logic based on data types, ensuring that operations like addition are only performed on compatible numeric types.

Final Answer: The function used is type().

Answer: (B)

Q35.

Solution

Concept:

Recursion is a programming technique where a function calls itself to solve smaller instances of the same problem. For recursion to work, the computer must keep track of the current state of each active function call—including its local variables and the point where execution should resume—while the ”inner” call is being processed. This management of active subroutines is handled by the operating system and the compiler through a specific architectural component.

Solution:

1. When a function is called, the system creates an ”activation record” or ”stack frame.” 2. Because function calls follow a nested pattern—where the last function called is the first one to finish—the **Stack** data structure is the perfect fit. This is the Last-In-First-Out (LIFO) principle in action. 3. Every time a recursive call occurs, a new frame is ”pushed” onto the System Stack. 4. When the base case of the recursion is reached and a function returns, its frame is ”popped” from the stack, and the control returns to the frame immediately below it. 5. A Queue would not work because it is FIFO; the first function called would be expected to finish first, which is the opposite of how nested calls work. 6. A Linked List or Tree could theoretically store the data, but the hardware-level implementation of call management is optimized specifically as a Stack (often called the ”Call Stack”).

Final Answer: The Stack is best suited for recursive call mechanisms.

Answer: (C)

Q36.

Solution

Concept:

Database commands are categorized based on their functionality into sub-languages like DDL (Data Definition Language) and DML (Data Manipulation Language). DDL commands are used to define or modify the structure (schema) of the database, such as creating, altering, or removing tables. DML commands are used to manage the data within those structures. A common point of confusion for students is the difference between removing specific rows based on a condition and clearing an entire table's data while keeping the "container" intact.

Solution:

1. The **DELETE** command is a DML operation. It removes rows one by one and records each deletion in the transaction log. It is slower for large datasets but allows for a WHERE clause to filter which rows to remove. 2. The **DROP** command is a DDL operation that removes the entire table, including its data, indexes, triggers, and the structure itself. Once dropped, the table no longer exists in the database. 3. The **TRUNCATE** command is a DDL operation. It is used to delete all rows from a table in a single operation. Unlike DELETE, it does not scan the table row-by-row; it simply deallocates the data pages used by the table. Crucially, the table structure, its columns, constraints, and indexes remain perfectly intact. 4. TRUNCATE is faster than DELETE and uses fewer system resources, but it cannot be used with a WHERE clause. 5. In academic exams like CUET, TRUNCATE is specifically identified as the DDL way to "reset" a table without "dropping" the table definition from the system.

Final Answer: The command used is TRUNCATE.

Answer: (C)

Q37.

Solution

Concept:

Transmission media, also known as communication channels, are the physical paths through which data travels from a sender to a receiver in a network. These are broadly divided into Guided (wired) and Unguided (wireless) media. Each medium has specific physical properties that determine its bandwidth, susceptibility to electromagnetic interference (EMI), and maximum transmission distance. As data demands have increased, networking technology has shifted toward media that can support terabit speeds.

Solution:

1. **Coaxial Cable** uses a central copper conductor surrounded by insulation and a braided shield. It transmits data using electrical signals and is commonly used for cable television and older Ethernet standards. 2. **Twisted Pair Cable** (like Cat6) consists of pairs of insulated copper wires twisted together to reduce EMI. It also uses electrical pulses but is limited in distance (usually 100 meters). 3. **Fiber Optic Cable** is a revolutionary medium made of thin strands of glass or plastic (the core) surrounded by a cladding layer. It transmits data as **pulses of light** using the principle of Total Internal Reflection. 4. Because light travels much faster than electrical signals and is completely immune to electromagnetic interference, fiber optics can carry vast amounts of data over hundreds of kilometers with minimal signal loss (attenuation). 5. **Infrared** is an unguided (wireless) medium that uses light but requires a direct line-of-sight and is used for very short-range communication, like remote controls. 6. Therefore, for high-speed, long-distance light-based transmission, Fiber Optic is the only correct choice.

Final Answer: The medium is Fiber Optic Cable.

Answer: (C)

Q38.

Solution

Concept:

Python provides different types of division operators to give programmers control over the resulting data type and precision. While many languages use a single symbol for division, Python distinguishes between ”true division” and ”floor division.” Understanding this distinction is vital for competitive exams, as it tests the student’s knowledge of Python’s specific arithmetic logic and how it handles floating-point versus integer results.

Solution:

1. The single slash **/** is the **True Division** operator. In Python 3, it always returns a floating-point number, even if the numbers are perfectly divisible. For example, 16/3 would result in approximately 5.333333333333333. 2. The double slash **//** is the **Floor Division** (or Integer Division) operator. It divides the first number by the second and then ”rounds down” to the nearest whole integer (the floor). 3. Mathematically, 16 divided by 3 is 5.333.... The floor of 5.333 is 5. 4. Note that if one of the operands is a float (e.g., 16.0//3), the result would be 5.0. However, since both 16 and 3 are integers here, the result is a pure integer 5. 5. The modulo operator **%** would return the remainder, which in this case would be 1 (since $16 = 3 \times 5 + 1$). 6. Because the expression uses the // operator, the fractional part is discarded entirely, leaving only the integer quotient.

Final Answer: The result of the expression is 5.

Answer: (B)

Q39.

Solution

Concept:

Operating systems (OS) can be categorized based on their licensing and development models. A "Closed Source" or proprietary OS is owned by a company that keeps the source code secret and charges users for a license. An "Open Source" OS, however, is developed under a license that allows anyone to view, modify, and distribute the source code freely. This collaborative model has led to some of the most stable and secure computing environments in the world.

Solution:

1. **Windows** is a proprietary operating system developed by Microsoft. Users pay for a license, and the underlying code is not available to the public. 2. **macOS** and **iOS** are proprietary systems developed by Apple. They are built on a Unix-based core (Darwin) which is partially open, but the final operating systems themselves are closed-source and tied to Apple hardware. 3. **Linux** is the most famous example of an open-source operating system kernel. It was started by Linus Torvalds in 1991. 4. Because the source code is open, thousands of developers worldwide contribute to its development. This has resulted in various "distributions" (distros) like Ubuntu, Fedora, Debian, and CentOS, tailored for everything from personal laptops to massive web servers and supercomputers. 5. Open-source software promotes transparency, community-driven innovation, and lower costs for organizations, which is why Linux dominates the server market and cloud infrastructure today.

Final Answer: The open-source operating system is Linux.

Answer: (C)

Q40.

Solution

Concept:

The power of a relational database lies in its ability to link data across different tables, preventing data redundancy and ensuring consistency. This is achieved through the use of ”Keys.” While a Primary Key uniquely identifies a record within its own table, a mechanism is needed to establish a relationship between a record in one table and a related record in another. This cross-referencing is the foundation of relational modeling.

Solution:

1. A **Candidate Key** is an attribute (or set of attributes) that can uniquely identify a row in a table. One of these is chosen as the Primary Key. 2. A **Super Key** is any set of attributes that uniquely identifies a row, which could include extra attributes not strictly necessary for unique identification. 3. A **Foreign Key** is an attribute in one table (the ”child” or ”referencing” table) that points to the **Primary Key** of another table (the ”parent” or ”referenced” table). 4. For example, in a database for a school, a ’Student’ table might have a Primary Key called ’StudentID’. An ’Enrollment’ table would have a column called ’StudentID’ that links back to the ’Student’ table. In the ’Enrollment’ table, ’StudentID’ is the Foreign Key. 5. Foreign Keys enforce ”Referential Integrity,” ensuring that you cannot have an enrollment record for a student who does not exist in the master ’Student’ table. This prevents ”orphaned” records and keeps the database logical and clean.

Final Answer: The column is called a Foreign Key.

Answer: (C)

Q41.

Solution

Concept:

Searching algorithms are categorized by their strategy and their requirements regarding the data's organization. The most basic search, Linear Search, requires no specific order but is inefficient for large datasets. To achieve higher performance, we use algorithms that exploit the order of elements. If a list is sorted, we can use a "divide and conquer" strategy that significantly reduces the number of comparisons needed to find a specific target.

Solution:

1. **Linear Search** (and its synonym, **Sequential Search**) starts from the beginning and checks every element. It does not benefit from sorting; its complexity remains $O(n)$ regardless of the order. 2. **Binary Search** is the most common algorithm that **requires a sorted list** to function. It works by repeatedly dividing the search interval in half. If the value of the search key is less than the item in the middle of the interval, the search continues in the lower half; otherwise, it continues in the upper half. 3. This halving mechanism gives Binary Search a logarithmic time complexity, $O(\log n)$. For a list of 1,000 items, Linear Search might take 1,000 steps, while Binary Search takes only about 10. 4. If you attempt to use Binary Search on an unsorted list, the "middle" comparison will yield a false direction, and the algorithm will fail to find the element even if it exists. 5. **Hash Search** uses a hash table and usually offers $O(1)$ complexity, but it relies on a specific data structure mapping rather than the simple sorted order of a list. Therefore, in the context of list-based searching, Binary Search is the definitive answer.

Final Answer: The technique is Binary Search.

Answer: (C)

Q42.

Solution

Concept:

The Hypertext Transfer Protocol (HTTP) uses standard response codes to inform a client (like a web browser) about the result of a request made to a server. These three-digit codes are grouped into five classes: 1xx (Informational), 2xx (Success), 3xx (Redirection), 4xx (Client Error), and 5xx (Server Error). Understanding these codes is essential for web developers and network administrators to diagnose connectivity and resource availability issues.

Solution:

1. **200 (OK)**: This is the standard response for a successful HTTP request. It means the server found the page and transmitted it to the browser correctly. 2. **301 (Moved Permanently)**: This code indicates that the requested resource has been assigned a new permanent URI (Uniform Resource Identifier) and any future references to this resource should use one of the returned URIs. 3. **404 (Not Found)**: This is perhaps the most famous error code on the internet. It is a ”Client Error” code, meaning the server was able to communicate with the client, but the server could not find the specific resource (web page, image, or file) that the user requested. This usually happens when a link is broken or a URL is mistyped. 4. **500 (Internal Server Error)**: This is a generic ”Server Error” message, indicating that the server encountered an unexpected condition that prevented it from fulfilling the request, even though the resource itself might exist. 5. For the specific case where a page simply ”cannot be found,” 404 is the globally recognized standard.

Final Answer: The status code is 404.

Answer: (B)

Q43.

Solution

Concept:

When processing data from files in Python, programmers must choose the most efficient way to bring that data into memory. Python's file object provides several methods to extract text. The choice depends on whether the file is very large (where reading line-by-line is better to save RAM) or whether the program needs to manipulate all lines as a flexible collection, such as a list, for sorting or searching purposes.

Solution:

1. The **read()** method reads the entire content of the file into a single, giant string. This includes all newline characters (\n) and spaces. 2. The **readline()** method reads only one line at a time from the file. Each time it is called, it moves the file pointer to the next line. 3. The **readlines()** method reads the entire file but, unlike read(), it automatically splits the content based on newline characters. It returns a **list of strings**, where each element in the list represents a single line from the file (including the newline character at the end of each string). 4. For example, if a file has three lines, readlines() will return: ['Line 1\n', 'Line 2\n', 'Line 3\n']. 5. This is extremely useful for programmers who want to iterate over lines using a 'for' loop or access a specific line using an index like *lines[2]*. 6. There is no standard method called "listread()" in Python's built-in file handling library.

Final Answer: The correct method is readlines().

Answer: (C)

Q44.

Solution

Concept:

Network topology describes how the various nodes, links, and nodes in a network are arranged. The Bus Topology was one of the earliest and simplest ways to connect multiple computers in a Local Area Network (LAN). It is a passive topology, meaning the computers on the network only ”listen” for data being sent and do not move the data from one computer to the next (unlike a Ring topology).

Solution:

1. In a **Bus Topology**, all devices are connected to a single continuous cable called the **backbone** or ”bus.” 2. Connection to this backbone is made using ”drop lines” and ”taps.” A drop line is a connection running from the device to the main cable, and a tap is a connector that splices into the main cable. 3. At both ends of the backbone cable, components called ”terminators” are placed to absorb the signal when it reaches the end, preventing it from bouncing back and causing interference (signal reflection). 4. The main advantage of this topology is that it is easy to install and requires less cabling than a Star or Mesh topology. 5. However, its primary weakness is that if the backbone cable is damaged or cut, the entire network fails. Furthermore, only one device can transmit data at a time; if multiple devices transmit, a collision occurs. 6. This description of a ”single backbone with drop lines” is the classic architectural definition of a Bus Topology.

Final Answer: The network topology is Bus Topology.

Answer: (C)

Q45.

Solution

Concept:

Aggregate functions in SQL allow users to perform mathematical calculations on a range of values within a column to return a single summarized result. These are vital for data analysis and business intelligence. While different database systems (like MySQL, PostgreSQL, or Oracle) might have custom extensions, the standard SQL language defines a set of core aggregate functions that are consistent across almost all platforms.

Solution:

1. To find the sum of values, we use **SUM()**. To find the highest or lowest values, we use **MAX()** and **MIN()**. To count the number of entries, we use **COUNT()**. 2. For calculating the arithmetic mean (the average), standard SQL uses the **AVG()** function. 3. Although "Mean" and "Average" are synonymous in mathematics, SQL syntax does not recognize *MEAN()* or *AVERAGE()* as standard keywords. Using them would result in a syntax error. 4. The **AVG()** function automatically ignores NULL values in the column, ensuring that the average is calculated only on actual data points. 5. While a user could manually calculate an average by using *SUM(column)/COUNT(column)*, the AVG() function is the built-in, optimized, and standard way to achieve this in a single command. It is readable, efficient, and avoids potential errors related to dividing by zero if the count is not checked.

Final Answer: The SQL function is AVG().

Answer: (B)

Q46.

Solution

Concept:

A Stack is a fundamental linear data structure that operates on a restricted access principle. Unlike an array where any element can be accessed via an index, or a queue where elements are added at one end and removed from another, a stack only allows operations at a single point known as the ”Top.” This structural constraint dictates the order in which data is processed, making it an essential tool for system-level tasks such as expression evaluation, syntax parsing, and memory management.

Solution:

1. The mechanism followed by a stack is **LIFO**, which stands for **Last-In-First-Out**. 2. This means that the most recent element added (pushed) to the stack is the very first one to be removed (popped). Imagine a physical stack of dinner plates: you place a new plate on the top, and when you need one, you take the one from the top that you just placed. 3. **FIFO** (First-In-First-Out) is the mechanism used by Queues, where the first element to arrive is the first to be served. 4. **LILO** (Last-In-Last-Out) is technically synonymous with FIFO but is rarely used in standard terminology. 5. **Random Access** is a property of arrays or linked lists where you can jump to any specific element immediately. Stacks strictly forbid this; you must pop the top elements to reach the ones underneath. 6. The LIFO nature of stacks is what allows computers to handle ”undo” operations in text editors—the last action you performed is the first one ”undone.”

Final Answer: A Stack follows the LIFO mechanism.

Answer: (C)

Q47.

Solution

Concept:

The Open Systems Interconnection (OSI) model is a conceptual framework used to understand and standardize the functions of a telecommunication or computing system. It partitions the complex process of data communication into seven logical layers. Each layer serves the layer above it and is served by the layer below it. By isolating functions into layers, network engineers can develop hardware and software that are compatible across different manufacturers and technologies.

Solution:

1. The **Physical Layer (Layer 1)** deals with bit-level transmission over physical media. 2. The **Data Link Layer (Layer 2)** handles hop-to-hop delivery and uses MAC addresses to move data within a single local network. 3. The **Network Layer (Layer 3)** is the ”intelligence” center for global connectivity. Its primary responsibility is **Routing**—the process of selecting the best path for data to travel from a source on one network to a destination on an entirely different network. 4. The Network Layer uses logical addressing (IP Addresses) to identify devices and utilizes devices called Routers to manage the traffic between subnets. 5. The **Transport Layer (Layer 4)** is responsible for end-to-end communication, error recovery, and flow control (using protocols like TCP and UDP). 6. Because the question specifically mentions ”routing data packets across different networks,” the Network Layer is the only correct functional match.

Final Answer: The Network Layer is responsible for routing.

Answer: (C)

Q48.

Solution

Concept:

Error handling is a critical part of writing production-grade Python code. When a program encounters a situation it cannot handle (like dividing by zero or opening a non-existent file), it ”raises an exception.” If this exception is not ”caught,” the program terminates abruptly with an error message. To prevent this, Python provides a structured block of keywords that allows the programmer to try a risky operation and define a specific backup plan if that operation fails.

Solution:

1. In Python, the error-handling structure starts with the **try** block. This block contains the code that is monitored for errors. 2. If an error occurs within the try block, the flow of execution immediately jumps to the **except** block. 3. The **except** block is where the actual ”handling” happens. The programmer writes code here to resolve the error, log a message, or offer the user an alternative path. For example: *except FileNotFoundError: print(”File not found! Please check the path.”)* 4. Other languages like Java or C++ use the keyword **catch** for this purpose, which is a common source of confusion for students. However, in Python, ”catch” is not a valid keyword; you must use **except**. 5. The **throw** keyword is used in other languages to manually trigger an error; in Python, the equivalent keyword is **raise**. 6. Therefore, while ”try” starts the process, it is the ”except” block that is specifically designed to catch and manage the resulting error.

Final Answer: The block used to catch and handle the error is except.

Answer: (C)

Q49.

Solution

Concept:

Relational databases store data in tables, but to interact with that data—to retrieve, update, or delete it—we need a specialized language. In the 1970s, researchers at IBM developed a language that allowed users to describe ”what” data they wanted rather than ”how” the computer should find it. This declarative nature made it the industry standard for database management, leading to the diverse ecosystem of database engines we see today, from small mobile databases to massive enterprise systems.

Solution:

1. **SQL** stands for **Structured Query Language**. 2. **Structured**: This refers to the fact that the language is designed for ”structured” data—data that fits into a predefined format of rows and columns (a schema). 3. **Query**: The primary purpose of the language is to ”query” the database, which means asking questions or requesting specific subsets of data using the SELECT statement. 4. **Language**: It is a domain-specific programming language. Unlike general-purpose languages like Python or C++, it is specialized for data manipulation and definition within an RDBMS. 5. Other variations like ”Standard Query List” or ”Sequential Query Logic” are incorrect. While SQL is standardized by ISO and ANSI, the ”S” stands specifically for ”Structured.” 6. Understanding this acronym is fundamental for any student of computer science, as SQL is the backbone of almost all web applications and data analysis tools used in the modern world.

Final Answer: SQL stands for Structured Query Language.

Answer: (A)

Q50.

Solution

Concept:

Selection Sort is a comparison-based sorting algorithm known for its simplicity and the fact that it performs a minimal number of swaps compared to other algorithms. It is an in-place sort, meaning it does not require extra storage space. While it has a time complexity of $O(n^2)$, making it inefficient for very large lists, it is often taught to introduce students to the concept of dividing a list into a ”sorted” and an ”unsorted” sublist.

Solution:

1. The algorithm works by logically dividing the input list into two parts: a sorted sublist at the beginning (initially empty) and an unsorted sublist containing the rest of the items. 2. In each iteration, the algorithm scans the entire **unsorted part** to find the absolute minimum (smallest) element. 3. Once the smallest element is identified, it is **swapped** with the element at the first position of the unsorted part. 4. After the swap, this element is now considered part of the ”sorted” sublist, and the boundary of the unsorted part moves one step to the right. 5. This process repeats until only one element remains in the unsorted part, at which point the entire list is sorted. 6. Unlike **Bubble Sort**, which swaps adjacent elements, or **Insertion Sort**, which shifts elements to make room for a new one, Selection Sort only makes one swap per pass. This makes it the correct answer for the mechanism described in the question.

Final Answer: The algorithm is Selection Sort.

Answer: (C)

Answer Key

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