

# CUET-UG Information Practices Sample Paper - 20

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.** Consider a Python Pandas Series S defined as: `S = pd.Series([12, 24, 36, 48], index=['w', 'x', 'y', 'z'])`. What will be the output of the statement `print(S[1:3])`?

- (A) w 12, x 24
- (B) x 24, y 36
- (C) x 24, y 36, z 48
- (D) y 36, z 48

**Q2.** A DataFrame df contains columns 'Sales' and 'Profit'. Which of the following commands will correctly display the 'Profit' for rows where 'Sales' is greater than 5000?

- (A) `df.loc[df['Sales'] > 5000, 'Profit']`
- (B) `df.iloc[df['Sales'] > 5000, 1]`
- (C) `df[df['Sales'] > 5000]['Profit']`
- (D) Both (A) and (C)

**Q3.** In SQL, which function 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

**Q4.** What will be the output of the following SQL query?

```
SELECT ROUND(157.48, -1) + TRUNCATE(15.9, 0);
```

- (A) 175
- (B) 172
- (C) 170
- (D) 165

**Q5.** Which property of a Pandas DataFrame is used to check if the DataFrame is empty?

- (A) `df.isempty`
- (B) `df.empty`
- (C) `df.isnull()`
- (D) `df.na()`

**Q6.** A network topology where every node is connected to exactly two other nodes, forming a single continuous pathway for signals, is known as:

- (A) Star Topology
- (B) Mesh Topology
- (C) Ring Topology
- (D) Tree Topology

**Q7.** Which SQL clause is used to filter the results produced by the GROUP BY clause?

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



- Q8.** The method used to import data from a CSV file into a Pandas DataFrame, specifically skipping the first 5 rows, is:
- (A) `pd.read_csv('data.csv', skip=5)`
  - (B) `pd.read_csv('data.csv', skiprows=5)`
  - (C) `pd.read_csv('data.csv', header=5)`
  - (D) `pd.from_csv('data.csv', start=5)`
- Q9.** What is the correct way to add a legend to a Matplotlib plot where multiple lines are plotted?
- (A) `plt.show(legend=True)`
  - (B) `plt.legend()`
  - (C) `plt.label(legend='show')`
  - (D) `plt.display_legend()`
- Q10.** The 'Right to be Forgotten' falls under which domain of Societal Impacts?
- (A) Intellectual Property Rights
  - (B) Privacy and Data Protection
  - (C) Cybercrime Prevention
  - (D) E-waste Management
- Q11.** A Relational Algebra operation that yields a relation consisting of all tuples appearing in both relations  $R$  and  $S$  is:
- (A) Union
  - (B) Difference
  - (C) Intersection
  - (D) Cartesian Product
- Q12.** Which of the following SQL functions returns the current system date and time?
- (A) `CURDATE()`



- (B) SYSDATE()
- (C) NOW()
- (D) Both (B) and (C)

**Q13.** Given a DataFrame `df` with 5 rows and 3 columns, what will `df.size` return?

- (A) 5
- (B) 3
- (C) 15
- (D) (5, 3)

**Q14.** Which of the following is an example of an Intellectual Property Rights (IPR) violation?

- (A) Phishing
- (B) Plagiarism
- (C) Identity Theft
- (D) Cyber Stalking

**Q15.** In Matplotlib, which function is used to create a histogram?

- (A) `plt.bar()`
- (B) `plt.hist()`
- (C) `plt.plot(kind='hist')`
- (D) `plt.histogram()`

**Q16.** What is the result of the SQL expression: `SELECT SUBSTR("CUET EXAMINATION 2026", 6, 11);`

- (A) "EXAMINATION"
- (B) "CUET EXAM"
- (C) " EXAMINATION"
- (D) "EXAMIN"



- Q17.** An IP address consists of \_\_\_\_\_ bits.
- (A) 32
  - (B) 48
  - (C) 64
  - (D) 128
- Q18.** Which device is primarily used to connect different networks operating on different protocols?
- (A) Switch
  - (B) Hub
  - (C) Gateway
  - (D) Repeater
- Q19.** What will be the output of `print(pd.Series([10, 20]) + pd.Series([5, 5], index=[0, 2]))`?
- (A) 0 15, 1 25, 2 5
  - (B) 0 15, 1 NaN, 2 NaN
  - (C) 0 15, 1 20, 2 5
  - (D) 0 10, 1 20, 2 5
- Q20.** Which SQL join returns all records from the left table and the matched records from the right table?
- (A) INNER JOIN
  - (B) RIGHT JOIN
  - (C) LEFT JOIN
  - (D) FULL JOIN
- Q21.** In Pandas, `df.iloc[2:4, :]` will select:
- (A) Rows at index 2 and 3



- (B) Rows at index 2, 3, and 4
- (C) Columns at index 2 and 3
- (D) Rows at index 1 and 2

**Q22.** Which command is used to export a DataFrame `df` to a file named "output.csv" without the index?

- (A) `df.to_csv("output.csv", index=False)`
- (B) `df.write_csv("output.csv", index=None)`
- (C) `pd.to_csv(df, "output.csv")`
- (D) `df.csv("output.csv", index=False)`

**Q23.** What will be the output of `SELECT LENGTH(TRIM(" IP "));`

- (A) 6
- (B) 4
- (C) 2
- (D) 0

**Q24.** In a relational database, a \_\_\_\_\_ is a minimal set of attributes that uniquely identifies a tuple.

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

**Q25.** Which Matplotlib function is used to change the label of the X-axis?

- (A) `plt.xname()`
- (B) `plt.xlabel()`
- (C) `plt.set_x()`
- (D) `plt.axis(x='label')`



- Q26.** The unique 48-bit hardware address assigned to a Network Interface Card (NIC) is the:
- (A) IP Address
  - (B) MAC Address
  - (C) Port Address
  - (D) Socket Address
- Q27.** Which of the following is NOT a valid string function in SQL?
- (A) LOWER()
  - (B) UPPER()
  - (C) TITLE()
  - (D) REVERSE()
- Q28.** What is the default value for the kind parameter in `df.plot()`?
- (A) 'bar'
  - (B) 'line'
  - (C) 'scatter'
  - (D) 'hist'
- Q29.** Which SQL aggregate function ignores NULL values?
- (A) COUNT(\*)
  - (B) SUM()
  - (C) AVG()
  - (D) Both (B) and (C)
- Q30.** Which type of Cybercrime involves sending fraudulent emails to trick individuals into revealing sensitive information?
- (A) Hacking
  - (B) Phishing



- (C) Ransomware
- (D) Denial of Service

**Q31.** In Pandas, which attribute returns a tuple representing the dimensionality of the DataFrame?

- (A) `df.ndim`
- (B) `df.shape`
- (C) `df.size`
- (D) `df.index`

**Q32.** What will be the output of `SELECT MOD(15, 4);`

- (A) 3.75
- (B) 3
- (C) 4
- (D) 1

**Q33.** A DataFrame `sales` has columns 'Month' and 'Amount'. To find the total amount per month, the correct command is:

- (A) `sales.groupby('Month').Amount.sum()`
- (B) `sales.sum().groupby('Month')`
- (C) `sales.groupby('Month').sum('Amount')`
- (D) `sales['Amount'].groupby('Month').total()`

**Q34.** Which of the following is an example of Open Source Software?

- (A) MS Windows
- (B) Adobe Photoshop
- (C) Linux
- (D) MS Office



- Q35.** The `head(n)` function in Pandas returns the first  $n$  rows. If  $n$  is not specified, how many rows are returned by default?
- (A) 1
  - (B) 5
  - (C) 10
  - (D) All rows
- Q36.** Which SQL command is used to display the structure of a table named Student?
- (A) `SHOW Student;`
  - (B) `DESC Student;`
  - (C) `STRUCT Student;`
  - (D) `VIEW Student;`
- Q37.** Which part of the digital footprint includes data you intentionally share, like social media posts?
- (A) Passive Digital Footprint
  - (B) Active Digital Footprint
  - (C) Permanent Digital Footprint
  - (D) Temporary Digital Footprint
- Q38.** In Matplotlib, to save a plot as an image file, we use:
- (A) `plt.save()`
  - (B) `plt.export()`
  - (C) `plt.savefig()`
  - (D) `plt.store()`
- Q39.** Which of the following SQL functions is used to find the number of characters in a string?
- (A) `COUNT()`



- (B) LENGTH()
- (C) LEN()
- (D) CHAR()

**Q40.** The `loc` indexer in Pandas is \_\_\_\_\_ based.

- (A) Integer
- (B) Label
- (C) Boolean
- (D) Position

**Q41.** What will be the output of `SELECT DAYOFMONTH('2026-05-15');`

- (A) 5
- (B) 15
- (C) Friday
- (D) 2026

**Q42.** To display the first 3 columns of a DataFrame `df` using `iloc`, which code is correct?

- (A) `df.iloc[:, 0:3]`
- (B) `df.iloc[0:3, :]`
- (C) `df.iloc[:, 3]`
- (D) `df.iloc[1:3, 1:3]`

**Q43.** Which type of network device works at the Physical layer of the OSI model and simply broadcasts signals to all ports?

- (A) Switch
- (B) Hub
- (C) Bridge
- (D) Router



- Q44.** To rename a column 'Old' to 'New' in a DataFrame df, we use:
- (A) `df.rename(columns={'Old': 'New'})`
  - (B) `df.change(column='Old', to='New')`
  - (C) `df.columns['Old'] = 'New'`
  - (D) `df.update('Old', 'New')`
- Q45.** What does the `explode()` function do in Matplotlib's pie chart?
- (A) It increases the size of the whole chart.
  - (B) It separates a slice from the main pie.
  - (C) It changes the colors of the slices.
  - (D) It adds a shadow to the chart.
- Q46.** Which SQL operation is used to combine the results of two SELECT statements, removing duplicates?
- (A) JOIN
  - (B) UNION
  - (C) COMBINE
  - (D) MERGE
- Q47.** E-waste contains hazardous substances like \_\_\_\_\_.
- (A) Lead and Mercury
  - (B) Carbon and Hydrogen
  - (C) Iron and Aluminum
  - (D) Nitrogen and Oxygen
- Q48.** In Pandas, `df.T` is used for:
- (A) Transforming data types
  - (B) Transposing the DataFrame
  - (C) Truncating the DataFrame



(D) Sorting the DataFrame

**Q49.** What will be the output of `SELECT POW(3, 2) + ROUND(4.6);`

(A) 13

(B) 14

(C) 11

(D) 15

**Q50.** Which attribute of a Pandas Series returns only the underlying values as a NumPy array?

(A) `S.index`

(B) `S.values`

(C) `S.data`

(D) `S.array`



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

The fundamental concept here is **Slicing in a Pandas Series**. A Series is a 1D labeled array. Slicing can be performed using positional indices or labels. When using integer slicing like `S[start:stop]`, Python follows the standard convention where the `start` index is inclusive and the `stop` index is exclusive. This means the element at the `stop` position is not included in the resulting subset. It is also important to note that slicing a Series returns both the index labels and their corresponding values, preserving the data's metadata.

**Solution:**1. **Initial Data Analysis:**

The Series `S` is defined with values `[12, 24, 36, 48]` and index labels `['w', 'x', 'y', 'z']`.

The internal positional mapping is:

- Position 0: Label 'w', Value 12
- Position 1: Label 'x', Value 24
- Position 2: Label 'y', Value 36
- Position 3: Label 'z', Value 48

2. **Executing the Slice:**

The statement `S[1:3]` requests data starting from position 1 up to (but not including) position 3.

- Included position: 1 (x 24)
- Included position: 2 (y 36)
- Excluded position: 3 (z 48)

3. **Output Construction:**

Pandas will return a new Series containing only positions 1 and 2.

The output format displays the index on the left and the value on the right:

```
x 24
y 36
```

4. **Conclusion:**

This positional slicing logic is consistent across Python lists, NumPy arrays, and Pandas Series, making it a critical pattern for data manipulation.

**Final Answer:** The correct option is (B).

**Answer: (B)**



Q2.

**Solution****Concept:**

This question focuses on **Data Filtering and Selection** in a Pandas DataFrame. Filtering involves selecting specific rows based on a logical condition (Boolean Indexing), while selection refers to choosing specific columns. Pandas provides two primary ways to handle this:

- **Boolean Indexing with Square Brackets:** Using `df[condition]` to filter rows and then `['column']` to select the column.
- **The .loc Accessor:** Using `df.loc[row_indexer, column_indexer]`, which is the more explicit and optimized method.

Understanding the difference between these methods and when they are interchangeable is essential for writing robust code.

**Solution:**

1. **Analyzing Option (A):**

`df.loc[df['Sales'] > 5000, 'Profit']` uses the label-based accessor. The first argument is a Boolean mask filtering for Sales above 5000, and the second argument specifies the 'Profit' column. This is the "best practice" approach.

2. **Analyzing Option (B):**

`df.iloc[df['Sales'] > 5000, 1]` is technically problematic in many Pandas versions because `.iloc` expects integer positions, not a Boolean Series, for row indexing. While some versions may allow it, it is not the standard usage.

3. **Analyzing Option (C):**

`df[df['Sales'] > 5000]['Profit']` is known as chained indexing. It first creates a filtered DataFrame and then selects the 'Profit' column. Although it can trigger warnings during data assignment, it is perfectly valid for data retrieval and display.

4. **Conclusion:**

Since both (A) and (C) are widely used and correct syntactical ways to retrieve the filtered Profit data, the most accurate answer is that both options work.

**Final Answer:** The correct option is (D).

**Answer: (D)**



Q3.

**Solution****Concept:**

The core concept involves **SQL String Search Functions**. In relational databases like MySQL, finding the location of a specific character or substring within a text field is a common requirement. Since different database engines evolved independently, several functions were created to perform this task. CUET-UG Information Practices specifically covers MySQL, which supports multiple synonyms for substring location to ensure compatibility with standard SQL and other proprietary systems. These functions return the 1-based index of the first occurrence of the substring.

**Solution:**1. **Understanding LOCATE():**

The `LOCATE(substr, str)` function returns the position of the first occurrence of the substring within the string. It is a very common function in MySQL.

2. **Understanding POSITION():**

The `POSITION(substr IN str)` function is the ANSI SQL standard version of the locate function. MySQL supports this standard syntax for portability.

3. **Understanding INSTR():**

The `INSTR(str, substr)` function is another synonym used in both MySQL and Oracle. Note that the order of arguments in `INSTR` is the reverse of `LOCATE`.

4. **Synthesis:**

All three functions—`LOCATE`, `POSITION`, and `INSTR`—return the index of the substring. If the substring is not found, all these functions return 0. Because all three are valid and functional in a standard SQL environment, they are all correct answers for the requirement.

**Final Answer:** The correct option is (D).

**Answer: (D)**



Q4.

**Solution****Concept:**

This question tests the application of **SQL Mathematical Functions**, specifically `ROUND()` and `TRUNCATE()`.

- **ROUND(number, decimals):** Rounds a number to a specified number of decimal places. If the decimal value is negative, it rounds digits to the left of the decimal point (tens, hundreds, etc.).
- **TRUNCATE(number, decimals):** Simply removes digits beyond the specified decimal place without rounding. It "chops" the number.

Calculating the result of an expression involving both functions requires a step-by-step approach to handle the differing logic of rounding versus truncation.

**Solution:**

1. **Evaluating `ROUND(157.48, -1)`:**

The second argument is  $-1$ , which indicates rounding to the nearest ten.

The units digit is 7. Since 7 is 5 or greater, we round the tens digit up.

157.48 rounded to the nearest ten becomes 160.

2. **Evaluating `TRUNCATE(15.9, 0)`:**

The second argument is 0, meaning we want zero decimal places.

Truncation ignores the value of the decimal part. Even though .9 is close to the next whole number, `TRUNCATE` simply removes it.

15.9 truncated to 0 places becomes 15.

3. **Final Calculation:**

Substitute the values back into the expression:  $160 + 15 = 175$ .

4. **Common Errors:**

A common mistake is rounding 15.9 to 16, which would result in 176, or rounding 157.48 to the nearest unit (157), leading to 172. Careful attention to the negative index in `ROUND` is vital.

**Final Answer:** The correct option is (A).

**Answer: (A)**



Q5.

**Solution****Concept:**

The concept here is **Pandas DataFrame Attributes**. Attributes are properties of an object that provide information about its state or structure. When dealing with real-world data, datasets are often generated dynamically. Before performing complex operations (like calculating an average or merging), it is a best practice to verify if the DataFrame actually contains any data. An "empty" DataFrame is one that has no rows or no columns. Pandas provides a highly efficient boolean attribute to perform this check without needing to count rows manually.

**Solution:**1. **Defining the empty attribute:**

The `df.empty` attribute returns True if the DataFrame has zero elements. This is internally checked by looking at the axes (index and columns) to see if the length of any axis is zero.

2. **Contrasting with functions:**

- `df.isnull()` is a method used to find missing (NaN) values within the data. A DataFrame could be full of missing values but still not be "empty" in terms of its structure.

- `df.isempty` is not a valid Pandas attribute; it is a common naming confusion from other programming languages like Java or MATLAB.

3. **Technical Nuance:**

Because `empty` is an attribute and not a method, you do not use parentheses when calling it. Using `df.empty()` would result in a `TypeError`.

4. **Conclusion:**

To verify if a container has no data in Pandas, `df.empty` is the correct and most efficient property to use.

**Final Answer:** The correct option is (B).

**Answer: (B)**



Q6.

**Solution****Concept:**

The concept here is **Network Topologies**. A network topology describes the physical or logical arrangement of nodes (computers, printers, etc.) and the connections between them. Different topologies offer varying levels of reliability, cost, and ease of installation. A circular arrangement where data passes through each node in a specific sequence is a fundamental pattern in networking. Understanding how data circulates through these paths is essential for diagnosing network latency and identifying points of failure in a communication system.

**Solution:**1. **Analyzing the Star Topology:**

In a Star topology, all nodes are connected to a central device like a hub or switch. If the central device fails, the entire network goes down, but if a single node fails, the rest of the network remains functional. This does not match the description of a "continuous pathway through two other nodes."

2. **Analyzing the Mesh Topology:**

In a Mesh topology, every node is potentially connected to every other node. This provides high redundancy but is very expensive and complex to wire. It does not follow a simple "two-node" connection rule for every device.

3. **Analyzing the Ring Topology:**

In a Ring topology, each node connects to exactly two neighboring nodes, one on each side. This creates a closed loop or ring. Data travels from node to node in one direction (or two in a dual-ring setup), and each node acts as a repeater to keep the signal strong. This matches the question's description perfectly.

4. **Conclusion:**

Because the question specifies that every node is connected to exactly two others to form a single continuous pathway, the Ring topology is the only arrangement that fits this geometric and logical definition.

**Final Answer:** The correct option is (C).

**Answer: (C)**



Q7.

**Solution****Concept:**

The core concept here is the **SQL HAVING Clause** and its distinction from the WHERE clause. In SQL, data is often aggregated using the GROUP BY clause to find totals, averages, or counts for specific categories. However, we often need to filter these aggregated results. Since the WHERE clause operates on individual rows *before* they are grouped, it cannot be used to filter groups based on aggregate values (like "show only departments with a total salary over 100,000"). This necessitates a specialized clause that acts on the final grouped output.

**Solution:**1. **The Role of WHERE:**

The WHERE clause is used to filter records from a table based on specific criteria before any grouping occurs. For example, WHERE Salary > 5000 filters individual employees.

2. **The Role of GROUP BY:**

The GROUP BY clause collects data across multiple records and groups the results by one or more columns. It is used in conjunction with aggregate functions like SUM(), AVG(), and COUNT().

3. **The Role of HAVING:**

The HAVING clause was added to SQL because the WHERE keyword could not be used with aggregate functions. HAVING filters the groups created by GROUP BY. For instance, GROUP BY Dept HAVING COUNT(ID) > 5 filters for departments with more than 5 members.

4. **Conclusion:**

Because the question asks for the clause that filters results *produced by* the GROUP BY clause, HAVING is the correct technical choice.

**Final Answer:** The correct option is (C).

**Answer:** (C)



Q8.

**Solution****Concept:**

The concept here is **Importing Data using Pandas**. The `read_csv()` function is one of the most powerful and frequently used tools in the Pandas library. It allows for a high degree of customization when reading external data files. Real-world CSV files often contain headers, metadata, or comments at the top of the file that are not part of the actual dataset. To handle this, Pandas provides parameters to instruct the parser to ignore a specific number of lines at the beginning of the file.

**Solution:**1. **Understanding the skiprows parameter:**

The `skiprows` parameter is specifically designed for this task. It can take an integer, which tells Pandas to skip that many lines from the start of the file. For example, `skiprows=5` tells the function to ignore the first five rows and start reading from the sixth row.

2. **Analyzing incorrect options:**

- `skip=5`: This is not a valid parameter in the `read_csv` function.
- `header=5`: This parameter tells Pandas that the column names are located on the 6th row (index 5), but it may still attempt to process earlier rows or behave differently than a simple skip.
- `from_csv`: This is an older, deprecated method. Modern Pandas code strictly uses `read_csv`.

3. **Practical Application:**

Using `skiprows` ensures that the DataFrame starts with the intended data structure, preventing errors that would occur if non-data text (like a file description or a report title) were read into a numeric column.

4. **Conclusion:**

The most direct and standard way to skip a set number of lines at the beginning of a file during the import process is the `skiprows` parameter.

**Final Answer:** The correct option is (B).

**Answer: (B)**



Q9.

**Solution****Concept:**

The concept involved here is **Matplotlib Visualization Customization**. When creating a plot with multiple datasets (lines, bars, or scatter points), it is critical for the viewer to know which color or style represents which variable. A legend is a small box on the chart that provides this identifying information. In Matplotlib, labels are typically defined during the plotting stage, but they do not automatically appear on the final figure. A specific command must be invoked to render the legend box based on those labels.

**Solution:**1. **Assigning Labels:**

When calling a plotting function like `plt.plot(x, y, label='Data 1')`, you provide a label string. This string is stored internally by Matplotlib.

2. **Rendering the Legend:**

The `plt.legend()` function scans the current axes for any labeled plots and creates a legend box using those labels. Without calling this function, the labels exist in the data structure but remain invisible on the actual chart.

3. **Analyzing incorrect options:**

- `plt.show(legend=True)`: The `show()` function is used to display the figure and does not take a legend argument.

- `plt.label()`: This is not a standard Matplotlib function for creating legends; it is a common confusion with `xlabel` or `ylabel`.

- `plt.display_legend()`: This function does not exist in the Matplotlib library.

4. **Conclusion:**

The standard, built-in method to display the legend in a Matplotlib plot is the `legend()` function.

**Final Answer:** The correct option is (B).

**Answer: (B)**

Q10.

**Solution****Concept:**

The concept here is **Societal Impacts and Data Privacy**. In the digital age, information about individuals can persist indefinitely on the internet. The "Right to be Forgotten" (also known as the right to erasure) is a legal concept developed to give individuals the power to request the removal of personal data from search engines and other online platforms when the data is no longer relevant, is inaccurate, or violates their privacy interests. This is a major pillar of modern data protection laws like the GDPR (General Data Protection Regulation).

**Solution:**

1. **Distinguishing between domains:**

- **Intellectual Property Rights (IPR):** This deals with ownership of creations like books, inventions, and software (Copyrights, Patents). It is not about personal identity.
- **Cybercrime Prevention:** This deals with illegal activities like hacking, phishing, and spreading malware.
- **E-waste Management:** This deals with the environmental impact of disposing of electronic devices.

2. **Focusing on Privacy:**

The Right to be Forgotten is fundamentally about **Privacy and Data Protection**. It addresses how personal information is stored, processed, and displayed to the public. It balances the public's right to know with an individual's right to keep their past private.

3. **Conclusion:**

Because the right focuses on the control of personal information and the mitigation of privacy risks associated with permanent digital records, it falls squarely within the category of Privacy and Data Protection.

**Final Answer:** The correct option is (B).

**Answer: (B)**



Q11.

**Solution****Concept:**

The concept here is Relational Algebra Operations. Relational algebra is a formal language used to describe how data is manipulated in a relational database. It consists of a set of operations that take one or two relations as input and produce a new relation as output. Among the set operations, we have Union, Intersection, and Difference. These operations are used to combine or compare rows from two different tables. To find commonalities between two datasets, we look for tuples that are present in both the first set and the second set simultaneously.

**Solution:**

## 1. Understanding the Union Operation:

The Union operation combines all tuples from two relations, removing any duplicates. It essentially represents the logic of set A or set B. This does not fit the requirement of finding only common elements.

## 2. Understanding the Difference Operation:

The set difference operation ( $R \text{ minus } S$ ) finds tuples that are present in the first relation  $R$  but are not present in the second relation  $S$ . It is used to find unique elements in one set.

## 3. Understanding the Intersection Operation:

The Intersection operation yields a relation consisting of all tuples that appear in both relations  $R$  and  $S$ . In logical terms, this represents set A and set B. For this operation to be valid, the two relations must be union-compatible, meaning they must have the same number of attributes and the attributes must have matching data types.

## 4. Understanding the Cartesian Product:

The Cartesian Product ( $R \text{ cross } S$ ) combines every tuple of the first relation with every tuple of the second relation. It results in a much larger relation and is not used for finding common records.

## 5. Conclusion:

Since the question asks for the operation that identifies tuples appearing in both relations, Intersection is the mathematically and logically correct term.

**Final Answer:** The correct option is (C).

**Answer: (C)**



Q12.

**Solution****Concept:**

The concept here is SQL Date and Time Functions. In database management, tracking when a record was created or updated is a standard requirement. SQL provides several built-in functions to retrieve the current date and time from the system's clock. Depending on the level of precision required—whether just the date or both the date and the time—different functions are employed. Furthermore, different SQL dialects may provide synonymous functions to achieve the same result to maintain compatibility with legacy systems.

**Solution:**

## 1. Analysis of CURDATE():

The CURDATE() function in MySQL returns the current date only, in the format YYYY-MM-DD. It does not include the time component. While useful for daily reporting, it does not fulfill the requirement of providing both date and time.

## 2. Analysis of SYSDATE():

The SYSDATE() function returns the current date and time as a value in YYYY-MM-DD HH:MM:SS format. A unique characteristic of SYSDATE() is that it returns the exact time at which the function executes, which might differ slightly from the start of the query if the query takes a long time to run.

## 3. Analysis of NOW():

The NOW() function also returns the current date and time in the YYYY-MM-DD HH:MM:SS format. Unlike SYSDATE(), the value returned by NOW() is established at the moment the query begins execution and remains constant for all rows processed by that query.

## 4. Conclusion:

Both SYSDATE() and NOW() are valid functions for retrieving the system's current date and time. Therefore, the most accurate answer encompassing both standard options is (D).

**Final Answer:** The correct option is (D).

**Answer: (D)**



Q13.

**Solution****Concept:**

The concept here is Pandas DataFrame Attributes for Dimensionality. When working with the Pandas library in Python, attributes provide metadata about the DataFrame object. Understanding the difference between size, shape, and ndim is a common point of confusion for beginners but is essential for data auditing. The size attribute specifically measures the total volume of data points stored in the structure, regardless of how many rows or columns they are organized into.

**Solution:**

1. Defining the shape attribute:

The `df.shape` attribute returns a tuple representing the dimensions of the DataFrame. For a DataFrame with 5 rows and 3 columns, `df.shape` would return (5, 3).

2. Defining the ndim attribute:

The `df.ndim` attribute returns an integer representing the number of axes or dimensions. For any standard DataFrame, the result of `ndim` is always 2 (rows and columns).

3. Calculating the size attribute:

The `df.size` attribute calculates the total number of elements in the DataFrame. This is mathematically determined by multiplying the number of rows by the number of columns. Calculation: 5 rows multiplied by 3 columns equals 15 total elements.

4. Practical Interpretation:

If you think of a DataFrame as a grid or a spreadsheet, the size tells you the total number of individual cells available for storing data. In this scenario, with a 5 by 3 grid, you have 15 cells.

5. Conclusion:

Based on the mathematical definition of the size attribute in Pandas, the result for a 5 by 3 DataFrame is 15.

**Final Answer:** The correct option is (C).

**Answer:** (C)



Q14.

**Solution****Concept:**

The concept here is Intellectual Property Rights (IPR) and Ethics. Intellectual Property refers to creations of the mind, such as inventions, literary and artistic works, designs, and symbols used in commerce. IPR is protected by law, which enables people to earn recognition or financial benefit from what they invent or create. Violating these rights involves using or claiming someone else's creative work without permission or proper attribution. This is distinct from general cybercrimes that target data or systems directly.

**Solution:**

## 1. Evaluating Phishing and Identity Theft:

Phishing and Identity Theft are forms of cyber fraud. Phishing involves tricking users into revealing sensitive info, while Identity Theft is the use of someone's personal details for gain. These are crimes against individuals or organizations but are not specifically IPR violations.

## 2. Evaluating Cyber Stalking:

Cyber stalking is a form of online harassment and is categorized as a crime against a person's safety and mental well-being. It does not involve the theft of creative ideas or property.

## 3. Evaluating Plagiarism:

Plagiarism is the act of presenting someone else's work, ideas, or expressions as your own without consent or credit. It is a direct violation of the creator's right to be recognized for their original work. In academic and professional settings, plagiarism is a significant IPR-related ethical and sometimes legal violation.

## 4. Conclusion:

Because plagiarism specifically involves the unauthorized use and misattribution of creative and intellectual output, it is the correct example of an IPR violation among the given choices.

**Final Answer:** The correct option is (B).

**Answer: (B)**



Q15.

**Solution****Concept:**

The concept here is Data Visualization with Matplotlib. Matplotlib is the most widely used library for 2D plotting in the Python ecosystem. Different types of charts are used to represent different types of data. A histogram is a specific type of bar chart that represents the distribution of numerical data by grouping data points into ranges (known as bins). Unlike a standard bar chart that compares discrete categories, a histogram is used to visualize frequency distributions.

**Solution:**1. Analysis of `plt.bar()`:

The `bar()` function is used to create a vertical bar chart. It requires two main arguments: the categories and their corresponding heights. It is used for comparing independent categories, not for showing the distribution of a single continuous variable.

2. Analysis of `plt.hist()`:

The `hist()` function is the dedicated function for creating histograms in Matplotlib. It automatically calculates the bins and the frequency of data points falling into each bin. It is the most direct and efficient way to create this specific type of visualization.

## 3. Analysis of other methods:

While some libraries like Pandas allow you to call `.plot(kind='hist')` on a Series, in the core Matplotlib library (`plt`), the function is specifically `hist()`. Functions like `plt.histogram()` do not exist as plotting commands; they might exist in NumPy for data calculation, but not for rendering a chart.

## 4. Conclusion:

For the purpose of plotting a histogram using the Matplotlib pyplot module, the `hist()` function is the correct and standard tool.

**Final Answer:** The correct option is (B).

**Answer: (B)**



Q16.

**Solution****Concept:**

The concept involved here is SQL String Extraction using the SUBSTR() function. In SQL, SUBSTR() (or SUBSTRING()) is used to extract a specific portion of a string. The function typically takes three arguments: the source string, the starting position (which is 1-indexed in SQL), and the length of the substring to be extracted. Understanding how SQL handles indices is vital, as it differs from languages like Python where indexing starts at 0. If the length argument is provided, the function counts that many characters starting from the start position.

**Solution:**

1. Identifying the Source String:

The input string is "CUET EXAMINATION 2026".

2. Determining the Start Position:

The second argument is 6.

Let us map the characters:

C(1), U(2), E(3), T(4), [Space](5), E(6).

The extraction begins at the 6th character, which is the letter 'E'.

3. Applying the Length Argument:

The third argument is 11. This means we need to count 11 characters starting from the 'E' at position 6.

Counting: E(1), X(2), A(3), M(4), I(5), N(6), A(7), T(8), I(9), O(10), N(11).

The sequence of characters extracted is "EXAMINATION".

4. Conclusion:

By carefully mapping the 1-based indexing of SQL and counting the exact length specified, the resulting substring is the word "EXAMINATION".

**Final Answer:** The correct option is (A).

**Answer:** (A)



Q17.

**Solution****Concept:**

The concept here is Network Addressing, specifically the structure of an Internet Protocol (IP) address. An IP address is a unique numerical label assigned to each device connected to a computer network that uses the Internet Protocol for communication. There are two versions of IP addresses in use: IPv4 and IPv6. However, in standard networking contexts and most academic curricula unless specified otherwise, "IP address" typically refers to IPv4. Understanding the bit-level structure is essential for sub-netting and logical network organization.

**Solution:**

## 1. Understanding IPv4 Structure:

An IPv4 address is composed of 4 octets (groups of 8 bits). Since there are 4 octets, the total number of bits is 4 multiplied by 8, which equals 32 bits. These are usually expressed in dotted-decimal notation, such as 192.168.1.1.

## 2. Contrasting with MAC Addresses:

A Media Access Control (MAC) address is a physical address and consists of 48 bits. It is important not to confuse the hardware address length with the logical IP address length.

## 3. Contrasting with IPv6:

IPv6 was developed to deal with the exhaustion of IPv4 addresses and uses 128 bits. However, standard multiple-choice questions in this subject tier identify the 32-bit address as the primary definition for a standard IP address.

## 4. Conclusion:

Based on the standard architecture of the IPv4 protocol, the address length is 32 bits.

**Final Answer:** The correct option is (A).

**Answer:** (A)



Q18.

**Solution****Concept:**

The concept here is Network Interconnectivity Devices. In a computer network, different devices are used to manage data traffic. While some devices like hubs and switches connect devices within the same network (LAN), other devices are required to bridge the gap between entirely different networks. These networks might use different communication protocols, different architectures, or different data formats. A specialized device is required to perform protocol translation and ensure data can flow from one system to another seamlessly.

**Solution:**

## 1. Evaluating Hubs and Switches:

Hubs and switches operate at the physical and data link layers respectively. They are used to connect devices within a single local area network. They do not have the intelligence to translate between different network protocols.

## 2. Evaluating Repeaters:

A repeater simply amplifies or regenerates a signal to allow it to travel longer distances. It does not look at the data or the protocol; it only deals with the electrical or optical signal.

## 3. Evaluating the Gateway:

A Gateway is a network node that connects two networks operating on different protocols. It acts as an entry and exit point for the network. Unlike a router, which primarily directs traffic between similar networks, a gateway can perform complex protocol conversion, making it the correct choice for connecting disparate systems.

## 4. Conclusion:

Because the task involves connecting different networks with different protocols, the Gateway is the specific device designed for this high-level translation and connectivity.

**Final Answer:** The correct option is (C).

**Answer: (C)**



Q19.

**Solution****Concept:**

The concept here is Vectorized Operations and Index Alignment in Pandas. One of the most powerful features of Pandas is its ability to automatically align data based on index labels during arithmetic operations. When you add two Series, Pandas looks for matching labels in both objects. If a label exists in both, it adds the values. If a label exists in only one of the objects, the result for that label will be NaN (Not a Number) because it cannot perform addition with a missing value.

**Solution:**

1. Mapping Series 1:

```
pd.Series([10, 20])
```

Index 0 → Value 10

Index 1 → Value 20

2. Mapping Series 2:

```
pd.Series([5, 5], index=[0, 2])
```

Index 0 → Value 5

Index 2 → Value 5

3. Performing the Addition:

- For Index 0: Both have values (10 and 5).  $10 + 5 = 15$ .

- For Index 1: Series 1 has 20, but Series 2 has no Index 1. Result = NaN.

- For Index 2: Series 2 has 5, but Series 1 has no Index 2. Result = NaN.

4. Resulting Output:

The new Series will have indices 0, 1, and 2. The values will be 15, NaN, and NaN respectively.

5. Conclusion:

This behavior highlights how Pandas prioritizes data integrity through label alignment rather than simple positional addition.

**Final Answer:** The correct option is (B).

**Answer: (B)**



Q20.

**Solution****Concept:**

The concept here is SQL Joins. Joins are used to combine rows from two or more tables based on a related column between them. In real-world databases, data is normalized across multiple tables, and joins are the primary mechanism for retrieving a complete view of that data. There are several types of joins, each defining how the database should handle rows that do not have a match in the other table. A join that prioritizes the data in the "left" (first mentioned) table is a standard requirement in data reporting.

**Solution:**

## 1. Analysis of INNER JOIN:

An Inner Join only returns records where there is a match in both tables. If a row in the left table does not have a corresponding value in the right table, it is excluded from the results.

## 2. Analysis of RIGHT JOIN:

A Right Join returns all records from the right table and the matched records from the left table. If there is no match, the left side will contain NULL values.

## 3. Analysis of LEFT JOIN:

A Left (Outer) Join returns all records from the left table and the matched records from the right table. Even if there is no match in the right table for a specific row in the left table, that row will still appear in the result set, with NULL values in the columns belonging to the right table. This exactly matches the scenario described in the question.

## 4. Conclusion:

By definition, the join that ensures all records from the first mentioned table are kept while pulling matching data from the second is the Left Join.

**Final Answer:** The correct option is (C).

**Answer: (C)**



Q21.

**Solution****Concept:**

The concept here is Positional Indexing using the `.iloc` accessor in Pandas. The `.iloc` indexer is strictly integer-location based for selection by position. It follows the standard Python slicing syntax where the start index is inclusive and the stop index is exclusive. When working with DataFrames, the `.iloc` indexer accepts two arguments separated by a comma: the first represents the row selection and the second represents the column selection. If a colon (`:`) is used as an argument, it indicates that all elements along that axis should be selected.

**Solution:**

1. Analyzing the row selection:

The command is `df.iloc[2:4, :]`. The first part `[2:4]` refers to the rows.

Following the inclusive-start and exclusive-stop rule, this selects index 2 and index 3. Index 4 is excluded.

Therefore, it selects the 3rd and 4th rows of the DataFrame (since indexing starts at 0).

2. Analyzing the column selection:

The second part of the command is a colon (`:`).

In Pandas slicing, a colon without specified boundaries means "all".

This implies that for the selected rows (2 and 3), all columns in the DataFrame will be included in the output.

3. Combining the logic:

The operation selects rows at integer positions 2 and 3 and includes every column available.

4. Conclusion:

Based on the positional slicing rules of the `.iloc` accessor, the selection corresponds to rows at index 2 and 3. This is a fundamental skill for data subsetting and cleaning where specific records need to be isolated by their physical order rather than their labels.

**Final Answer:** The correct option is (A).

**Answer: (A)**



Q22.

**Solution****Concept:**

The concept here is Data Exporting in Pandas. After processing and analyzing data in a DataFrame, it is often necessary to save the results back into a file format like CSV (Comma Separated Values) for sharing or long-term storage. The `to_csv()` method is the primary tool for this. A common requirement when exporting is to control whether the DataFrame's index (the row labels) should be saved as the first column in the file. By default, Pandas includes the index, but in many cases, especially when the index is just a set of default integers, it is better to exclude it to keep the file clean.

**Solution:****1. Understanding the `to_csv()` method:**

This is an instance method, meaning it is called directly on the DataFrame object (`df.to_csv()`). It requires the filename or path as the first argument.

**2. Managing the Index:**

To prevent Pandas from writing the row index to the CSV file, we use the boolean parameter `index`. Setting `index=False` ensures that only the data columns are written to the file.

**3. Evaluating the options:**

- **Option (A)** `df.to_csv("output.csv", index=False)` is the correct and standard syntax.
- **Option (B)** uses `write_csv` and `index=None`, which are not standard Pandas commands.
- **Option (C)** attempts to use `pd.to_csv(df, ...)`, which is incorrect because `to_csv` is a method of the DataFrame object, not the top-level Pandas module.
- **Option (D)** uses `df.csv`, which is not a valid method.

**4. Conclusion:**

The standard and most efficient way to export data without including the index labels is using `df.to_csv` with the `index` parameter set to `False`.

**Final Answer:** The correct option is (A).

**Answer:** (A)



Q23.

**Solution****Concept:**

The concept here is SQL String Cleaning and Length Calculation. In database management, data often comes with "noise" such as leading or trailing spaces. These spaces can cause issues during comparison or display. SQL provides the TRIM() function to remove these unnecessary spaces from both ends of a string. Additionally, the LENGTH() function is used to count the total number of characters remaining in the string. Understanding how these functions nest within each other is a key aspect of query writing.

**Solution:**

1. Analyzing the inner function:

The inner part of the expression is TRIM(" IP ").

The original string " IP " has two spaces at the beginning, the letters I and P, and two spaces at the end.

The TRIM function removes all leading and trailing spaces.

Result = "IP".

2. Analyzing the outer function:

The outer part is SELECT LENGTH("IP").

The LENGTH() function counts every character in the provided string.

In the string "IP", there are exactly two characters: 'I' and 'P'.

3. Calculating the final result:

The length of the trimmed string "IP" is 2.

4. Conclusion:

Even though the initial string had a total length of 6 (2 spaces + 2 letters + 2 spaces), the application of TRIM reduced it to the core content before the LENGTH function was applied. This nested logic is essential for data validation and cleaning in SQL.

**Final Answer:** The correct option is (C).

**Answer: (C)**



Q24.

**Solution****Concept:**

The concept here is Database Keys and Table Constraints. In a relational database, keys are used to establish and identify relationships between tables and to ensure that each record in a table is unique. A table can have several potential keys that could uniquely identify a row. These are known as Candidate Keys. From these candidates, the database designer selects one to be the Primary Key. Understanding the hierarchy and definition of these keys is fundamental to database normalization and integrity.

**Solution:**

## 1. Defining the Candidate Key:

A Candidate Key is a column or a set of columns that can uniquely identify any record in a table without any redundant data. It must contain unique values and cannot be NULL. A table can have multiple candidate keys.

## 2. Defining the Primary Key:

The Primary Key is the specific Candidate Key chosen by the database administrator to be the main unique identifier for the table. While it is the "minimal set of attributes" that identifies a tuple, the term "Candidate Key" more broadly fits the definition of any set that satisfies this requirement before one is officially selected as Primary.

## 3. Evaluating other keys:

- A Foreign Key is used to link two tables together.
- An Alternate Key is a Candidate Key that was not chosen as the Primary Key.

## 4. Conclusion:

The term that specifically refers to a minimal set of attributes capable of uniquely identifying a tuple (making it a "candidate" for the primary role) is the Candidate Key. This ensures that no subset of the key is sufficient to provide uniqueness.

**Final Answer:** The correct option is (B).

**Answer: (B)**



Q25.

**Solution****Concept:**

The concept here is Matplotlib Axis Labeling. When creating any data visualization, labeling the axes is one of the most critical steps to ensure the chart is interpretable. Without labels, a viewer cannot know what the units or categories on the horizontal (X) and vertical (Y) axes represent. Matplotlib, through its `pyplot` module, provides straightforward functions to add text to these axes. These functions are highly customizable, allowing for changes in font size, color, and weight, though their basic usage is very simple.

**Solution:****1. Identifying the correct function:**

The function specifically designed to add a label to the X-axis is `plt.xlabel()`. It takes a string as its primary argument, which it then displays centered below the X-axis.

**2. Analyzing incorrect options:**

- `plt.xname()`: This is not a valid Matplotlib function.
- `plt.set_x()`: This might be confused with methods used in object-oriented Matplotlib (`ax.set_xlabel`), but as a standalone `pyplot` function, it is incorrect.
- `plt.axis()`: This function is used to get or set some axis properties like limits (`xmin`, `xmax`), but not the text label itself.

**3. Practical usage:**

A typical command would be `plt.xlabel("Time in Seconds")`. This provides immediate context to the numerical data displayed along the bottom of the graph.

**4. Conclusion:**

For the standard task of adding descriptive text to the horizontal axis in a Matplotlib plot, the `xlabel()` function is the correct and only standard tool provided by the library.

**Final Answer:** The correct option is (B).

**Answer: (B)**



Q26.

**Solution****Concept:**

The concept here is Network Hardware Identification, specifically the MAC Address. Every device that connects to a network requires a physical piece of hardware known as a Network Interface Card (NIC). To ensure that data packets are delivered to the correct physical device on a local network, each NIC is assigned a unique, permanent identification number by the manufacturer. This is known as the Media Access Control (MAC) address. Unlike IP addresses, which are logical and can change depending on the network a device joins, the MAC address is "burned" into the hardware.

**Solution:**

## 1. Structure of a MAC Address:

A MAC address consists of 48 bits, usually represented as six groups of two hexadecimal digits (e.g., 00:1A:2B:3C:4D:5E). The first 24 bits identify the manufacturer (OUI), and the last 24 bits are a unique serial number assigned by that manufacturer.

## 2. Contrasting with IP Address:

An IP address is a logical address used for routing data across different networks (like the Internet). IPv4 addresses are 32 bits, and IPv6 addresses are 128 bits. They are assigned by software or network administrators, not hardware manufacturers.

## 3. Role of the NIC:

The Network Interface Card acts as the middleman between the computer and the network medium. The MAC address resides at the Data Link Layer (Layer 2) of the OSI model.

## 4. Conclusion:

Because the question specifies a 48-bit unique hardware address assigned to the NIC, the MAC Address is the correct technical term. It serves as the "fingerprint" of the networking hardware.

**Final Answer:** The correct option is (B).

**Answer: (B)**



Q27.

**Solution****Concept:**

The concept here is SQL String Functions. SQL provides a wide array of functions to manipulate text data stored in tables. These functions allow developers to format output, search for patterns, or transform data for better readability. While many functions like those for changing case (LOWER/UPPER) or reversing strings are standardized across many database systems like MySQL, certain function names are specific to certain programming languages and are not found in the SQL standard. Distinguishing between language-specific methods and SQL-specific functions is a key skill.

**Solution:**

## 1. Evaluating LOWER() and UPPER():

These are standard SQL functions used to convert a string into all lowercase or all uppercase characters, respectively. They are essential for performing case-insensitive searches and standardizing data display.

## 2. Evaluating REVERSE():

The REVERSE() function is a built-in SQL function that takes a string and returns it with the character order inverted. It is commonly used in various data processing tasks.

## 3. Evaluating TITLE():

In Python, there is a string method called .title() which capitalizes the first letter of every word. However, in standard SQL (including MySQL), there is no built-in function named TITLE(). To achieve "Title Case" in SQL, one would typically need a complex combination of SUBSTR, UPPER, and CONCAT functions or a custom stored function.

## 4. Conclusion:

Since LOWER, UPPER, and REVERSE are all valid built-in SQL functions, but TITLE is a Python-specific method not found in standard SQL, (C) is the correct answer for an invalid function.

**Final Answer:** The correct option is (C).

**Answer: (C)**



Q28.

**Solution****Concept:**

The concept here is Matplotlib Data Visualization Defaults. When using the Pandas `.plot()` method or the high-level Matplotlib interface, there are certain default behaviors designed to make plotting as simple as possible. If a user provides a dataset and calls the plot function without specifying the "kind" of chart (such as bar, pie, or scatter), the library assumes a default visualization style. Understanding these defaults helps developers write cleaner code by only specifying parameters when they need to deviate from the standard behavior.

**Solution:**

1. Understanding the `plot()` function:

The `df.plot()` function in Pandas is a wrapper around Matplotlib. It takes a parameter named "kind" which determines the visual representation of the data.

2. Identifying the default:

When the "kind" parameter is omitted, the default value is 'line'. Line charts are the standard choice for visualizing trends and continuous data points, which is why they were chosen as the baseline for the library.

3. Analyzing other chart types:

- 'bar' must be explicitly called for categorical comparisons.
- 'scatter' is used for identifying correlations between two variables and requires X and Y parameters.
- 'hist' is used for frequency distributions.

4. Conclusion:

For any DataFrame or Series, calling `.plot()` without further arguments will result in a line chart. Therefore, 'line' is the default value for the kind parameter.

**Final Answer:** The correct option is (B).

**Answer: (B)**



Q29.

**Solution****Concept:**

The concept here is SQL Aggregate Functions and NULL Handling. Aggregate functions perform a calculation on a set of values and return a single value. A critical aspect of database management is understanding how "NULL" (missing or unknown data) affects these calculations. In SQL, most aggregate functions are designed to skip over NULL values to prevent them from corrupting the mathematical result (e.g., an average shouldn't be zero just because one record is missing data). However, there is a specific nuance regarding the COUNT function that every database professional must know.

**Solution:**

## 1. Analysis of SUM() and AVG():

Both the SUM() function and the AVG() function automatically ignore NULL values. If you are calculating the average salary of 10 employees and 2 have NULL salaries, SQL will calculate the sum of the 8 known salaries and divide by 8, not 10.

## 2. Analysis of COUNT(\*):

The COUNT(\*) function is unique. It counts the total number of rows in a table, regardless of whether any specific columns contain NULL values. It essentially counts the existence of the record itself.

## 3. Analysis of COUNT(column\_name):

Unlike COUNT(\*), the version of count that specifies a column name \*will\* ignore NULL values in that specific column.

## 4. Conclusion:

Since both SUM() and AVG() are designed to ignore NULLs to maintain mathematical accuracy of the non-missing data, both (B) and (C) are correct. Thus, (D) is the most comprehensive answer.

**Final Answer:** The correct option is (D).

**Answer: (D)**



Q30.

**Solution****Concept:**

The concept here is Cybercrime and Online Security. As society becomes more reliant on digital communication, criminals use various psychological and technical tactics to exploit users. Phishing is a specific form of social engineering where attackers masquerade as a trusted entity (like a bank, a government agency, or a popular service) to steal sensitive information such as usernames, passwords, and credit card details. This is usually done through deceptive emails, text messages, or websites that look nearly identical to the legitimate versions.

**Solution:**

## 1. Identifying Phishing:

The hallmark of phishing is "baiting" the user. An email might claim there is an "unauthorized login attempt" and provide a link to a fake login page. Once the user enters their credentials, the attacker captures them. This matches the description of sending fraudulent emails to trick individuals.

## 2. Contrasting with Hacking:

Hacking generally refers to the unauthorized access to a computer system or network by exploiting technical vulnerabilities. While phishing can be a tool used by hackers, the specific act of tricking a person via email is defined as phishing.

## 3. Contrasting with Ransomware:

Ransomware is a type of malware that encrypts a victim's files, with the attacker demanding a ransom to restore access.

## 4. Contrasting with Denial of Service (DoS):

A DoS attack aims to make a service or network unavailable by overwhelming it with a flood of traffic.

## 5. Conclusion:

Because the question specifically describes the act of using fraudulent communication to deceive individuals into revealing information, Phishing is the correct classification for this cybercrime.

**Final Answer:** The correct option is (B).

**Answer: (B)**



Q31.

**Solution****Concept:**

The concept here is Pandas DataFrame Metadata and the shape attribute. In the Pandas library, after creating or loading a DataFrame, it is vital to understand its structure. Attributes like shape, size, and ndim provide this structural information. While ndim tells you the number of dimensions and size tells you the total number of elements, the shape attribute provides the specific count of rows and columns. This is returned in the form of a Python tuple, which is an immutable sequence type.

**Solution:**

1. Defining the shape attribute:

The `df.shape` attribute returns a tuple representing the dimensionality of the DataFrame. For example, if a DataFrame has 10 rows and 5 columns, `df.shape` will return (10, 5).

2. Analyzing incorrect options:

- `df.ndim`: This returns an integer representing the number of dimensions. For a DataFrame, this is always 2 (rows and columns). It does not provide the specific count of those rows or columns.

- `df.size`: This returns an integer representing the total number of elements (rows multiplied by columns).

- `df.index`: This returns the index (row labels) of the DataFrame, which is an Index object, not a tuple of dimensions.

3. Practical Application:

Data scientists frequently use `df.shape[0]` to get the number of rows and `df.shape[1]` to get the number of columns. This is essential for loop iterations, data validation, and ensuring that two DataFrames are compatible for operations like concatenation or merging.

4. Conclusion:

Because the question asks for the attribute that returns a tuple representing the dimensionality (rows, columns), the shape attribute is the correct technical choice.

**Final Answer:** The correct option is (B).

**Answer: (B)**



Q32.

**Solution****Concept:**

The concept here is the SQL MOD() function. In SQL, the MOD() function is used to perform the modulo operation, which finds the remainder after the division of one number by another. It is a mathematical function that is highly useful in various scenarios, such as determining if a number is even or odd (MOD(n, 2)) or performing calculations that involve cycles or patterns. The function typically takes two arguments: the dividend and the divisor.

**Solution:**

1. Understanding the syntax:

The syntax is MOD(N, M), which returns the remainder of N divided by M. This can also be written using the modulo operator (

2. Performing the calculation:

In the query SELECT MOD(15, 4), the number 15 is the dividend and 4 is the divisor.

- 15 divided by 4 equals 3 with a remainder.

- 4 multiplied by 3 equals 12.

- 15 minus 12 equals 3.

3. Analyzing the result:

The remainder of the division of 15 by 4 is 3. It is important to distinguish this from the result of the division itself (which would be 3.75). The MOD() function specifically extracts the integer remainder that is left over after the maximum number of whole divisors has been subtracted.

4. Conclusion:

By applying the standard mathematical definition of the modulo operation, the remainder of 15 divided by 4 is 3.

**Final Answer:** The correct option is (B).

**Answer: (B)**



Q33.

**Solution****Concept:**

The concept here is the Pandas GroupBy operation. Grouping data is a fundamental part of data analysis, allowing you to split a dataset into groups based on some criteria, apply a function (like sum, mean, or count), and combine the results. This is often referred to as the Split-Apply-Combine pattern. In Pandas, the `groupby()` method is used to create a GroupBy object, and then an aggregation function is called on a specific column or the entire group.

**Solution:**

1. Understanding the `groupby()` syntax:

To group by a specific column, you use `df.groupby('ColumnName')`. This creates a GroupBy object where rows with the same value in 'ColumnName' are clustered together.

2. Applying the aggregation:

After grouping, you need to specify which column you want to aggregate. If you want the total 'Amount' for each 'Month', you select the 'Amount' column and call the `.sum()` method.

3. Evaluating the options:

- Option (A) `sales.groupby('Month').Amount.sum()` is a standard and efficient way to perform this operation. It groups by Month, selects the Amount column, and then sums the values within each group.

- Option (C) `sales.groupby('Month').sum('Amount')` is also a valid syntax in many Pandas versions where you call `sum` on the group and specify the column as a parameter, though (A) is more common in instructional materials.

4. Conclusion:

The most idiomatic and standard way to find the total amount per month is to group by the categorical column (Month) and then sum the numeric column (Amount). Option (A) represents this logic perfectly.

**Final Answer:** The correct option is (A).

**Answer:** (A)



Q34.

**Solution****Concept:**

The concept here is Open Source Software (OSS). Software is generally categorized into two types: Proprietary and Open Source. Proprietary software is owned by a company, and its source code is kept secret. Users must usually pay for a license to use it and are not allowed to modify or redistribute it. Open Source Software, on the other hand, is software whose source code is made available to the public. Anyone can inspect, modify, and distribute the code, often for free. This promotes collaboration and transparency in software development.

**Solution:**

## 1. Analyzing MS Windows and MS Office:

Both MS Windows (an operating system) and MS Office (a productivity suite) are developed by Microsoft. They are proprietary software. The source code is not public, and you must purchase a license to use them legally.

## 2. Analyzing Adobe Photoshop:

Adobe Photoshop is a professional image editing software developed by Adobe. It is also proprietary. It is a commercial product with restricted access to its internal code.

## 3. Analyzing Linux:

Linux is the most famous example of open-source software. It is a kernel (and often refers to the entire operating system distribution) that was originally created by Linus Torvalds and is now maintained by a global community of developers. Its source code is freely available under licenses like the GNU GPL.

## 4. Conclusion:

Because Linux is the only option where the source code is public and freely modifiable by the community, it is the correct example of Open Source Software.

**Final Answer:** The correct option is (C).

**Answer:** (C)



Q35.

**Solution****Concept:**

The concept here is the Pandas `head()` function. When working with large datasets, it is often impractical to display the entire DataFrame. The `head()` function is a convenient tool used to preview the first few rows of a DataFrame or Series. This allows the user to quickly check column names, data types, and the general structure of the data. The function accepts an optional integer argument  $n$  that specifies the number of rows to return.

**Solution:**

1. Understanding the function parameters:

The signature is `df.head(n=5)`. The parameter  $n$  has a default value assigned by the Pandas developers.

2. Identifying the default value:

If the user calls `df.head()` without providing an integer, the function defaults to returning the first 5 rows. This was chosen as a sensible default to give a quick overview without cluttering the display.

3. Contrasting with `tail()`:

Similarly, the `tail()` function returns the last  $n$  rows of a DataFrame, and it also defaults to 5 rows if no argument is provided.

4. Practical Application:

If a user wanted to see the first 10 rows, they would explicitly write `df.head(10)`. If they only wanted to see the very first row, they would write `df.head(1)`. But for the majority of quick checks, the default of 5 is the most frequently encountered result.

5. Conclusion:

Based on the standard library definition in Pandas, the default number of rows returned by the `head()` function is 5.

**Final Answer:** The correct option is (B).

**Answer: (B)**



Q36.

**Solution****Concept:**

The concept here is SQL Metadata and Table Inspection. In a Relational Database Management System (RDBMS), it is essential for a developer to be able to view the structure of a table they have created. This includes seeing the column names, the data types assigned to each column (like INT, VARCHAR, or DATE), whether a column can accept NULL values, and identifying which columns serve as keys. Every SQL dialect provides a specific command to "describe" or "explain" this internal metadata to the user.

**Solution:**

1. Understanding the DESCRIBE command:

In MySQL, which is the primary focus of the CUET Information Practices syllabus, the command DESCRIBE (or its shorthand DESC) is used to display the table structure. When you execute DESC Student;, the database returns a result set with headers like Field, Type, Null, Key, Default, and Extra.

2. Analyzing incorrect options:

- SHOW Student;: This is incorrect. The SHOW command is used for broader objects, such as SHOW TABLES; or SHOW DATABASES;, but it does not describe the specific columns of a single table.

- STRUCT Student;: While "structure" is what we are looking for, STRUCT is not a valid SQL keyword for this operation.

- VIEW Student;: The VIEW command is used in the context of creating or managing virtual tables (Views), not for inspecting the schema of a base table.

3. Practical Application:

Before writing an INSERT statement, a programmer will often run DESC TableName; to ensure they are providing values in the correct order and format required by the database schema.

4. Conclusion:

The standard and most common command for inspecting the schema and attributes of a table in SQL is DESC (or DESCRIBE).

**Final Answer:** The correct option is (B).

**Answer: (B)**



Q37.

**Solution****Concept:**

The concept here is Digital Footprints and Online Privacy. A digital footprint is the trail of data you create while using the internet. It includes the websites you visit, emails you send, and information you submit to online services. Digital footprints are generally categorized into two types: Active and Passive. Understanding the difference between these two is vital for digital citizenship and protecting one's online reputation, as some data is shared knowingly while other data is collected without the user's explicit awareness.

**Solution:**

## 1. Defining the Active Digital Footprint:

An Active Digital Footprint is created when a user deliberately and intentionally shares information about themselves. Examples include posting on social media (Facebook, Instagram, X), sending emails, uploading videos to YouTube, or filling out online forms. The user is fully aware that they are contributing data to the digital environment.

## 2. Defining the Passive Digital Footprint:

A Passive Digital Footprint is created when data is collected about a user without their active participation or sometimes even without their knowledge. Examples include web servers logging your IP address, websites tracking your browsing history via cookies, or apps accessing your location data in the background.

## 3. Evaluating the question:

The question specifically mentions "data you intentionally share, like social media posts." This aligns perfectly with the definition of an Active Digital Footprint.

## 4. Conclusion:

Because the action of posting to social media is a conscious and intentional act of data sharing, it falls under the category of an Active Digital Footprint.

**Final Answer:** The correct option is (B).

**Answer: (B)**



Q38.

**Solution****Concept:**

The concept here is Matplotlib File Exporting. Matplotlib is not only useful for displaying charts on a screen during an interactive session but also for generating high-quality image files for reports, websites, or publications. To move from a transient plot in memory to a permanent file on a disk, a specific function is required. This function allows the user to specify the file name and the format (such as .png, .jpg, .pdf, or .svg). It also provides options for controlling the resolution (DPI) and the transparency of the output.

**Solution:**

1. Identifying the correct function:

The standard function in Matplotlib's pyplot module for saving a figure is `savefig()`. Its basic syntax is `plt.savefig('filename.png')`.

2. Analyzing incorrect options:

- `plt.save()`: While it sounds logical, this is not the name of the function in Matplotlib. It is a common mistake for those used to other libraries.
- `plt.export()`: This is not a valid Matplotlib function.
- `plt.store()`: This is not a valid Matplotlib function.

3. Practical Nuance:

It is important to call `plt.savefig()` *before* calling `plt.show()`. This is because `plt.show()` typically clears the current figure after displaying it, which could result in a blank image if `savefig` is called afterward.

4. Conclusion:

The dedicated and standard command used to write the current plot to a physical file is `savefig()`.

**Final Answer:** The correct option is (C).

**Answer: (C)**



Q39.

**Solution****Concept:**

The concept here is SQL String Length Functions. When dealing with text data in a database, knowing the number of characters in a string is useful for data validation, formatting, and analysis. For example, you might want to ensure a password meets a minimum length or truncate a description that is too long. While different SQL versions have slightly different naming conventions (like LEN in SQL Server), the CUET-UG syllabus focuses on MySQL standards.

**Solution:**

1. Understanding the LENGTH() function:

In MySQL, the LENGTH() function returns the length of a string measured in bytes. For standard English characters, this corresponds directly to the number of characters. Another function, CHAR\_LENGTH(), specifically counts characters regardless of multi-byte encoding.

2. Analyzing incorrect options:

- COUNT(): This is an aggregate function used to count the number of rows in a result set, not the number of characters in a single string.

- LEN(): This is the function name used in Microsoft SQL Server, but it is not the standard function name in MySQL.

- CHAR(): This function is used to return the character associated with a specific ASCII/Unicode code (e.g., CHAR(65) returns 'A'), it does not measure length.

3. Practical Usage:

A typical query would be SELECT Name, LENGTH(Name) FROM Student;. This would return each student's name alongside the count of characters in that name.

4. Conclusion:

For the purpose of determining the character count of a string in a MySQL environment, LENGTH() is the correct and most widely used function.

**Final Answer:** The correct option is (B).

**Answer: (B)**



Q40.

**Solution****Concept:**

The concept here is Label-based Selection in Pandas using `.loc`. The Pandas library provides two main ways to access data: `.loc` and `.iloc`. Understanding the distinction between these two is one of the most important milestones in learning Python for data handling. While `.iloc` uses integer-based positional indexing (similar to how you access elements in a Python list), `.loc` is designed to use the labels of the rows and columns. This makes the code much more readable and robust, as it doesn't depend on the physical order of the data.

**Solution:**

## 1. Defining Label-based Indexing:

When you use `df.loc[row_label, column_label]`, you are telling Pandas exactly which index names to look for. For example, if your index is `['Alpha', 'Beta', 'Gamma']`, you would use `df.loc['Alpha']` to get the first row.

2. Contrast with `.iloc`:

The `.iloc` indexer is "Position-based." It does not care about the names of the indices; it only cares about where they are (0, 1, 2...).

3. Behavior of `.loc` with Slicing:

A unique feature of `.loc` is that when you slice with labels (e.g., `.loc['Alpha':'Gamma']`), both the start and the end labels are **\*\*inclusive\*\***. This is different from standard Python/`iloc` slicing where the end is exclusive.

## 4. Conclusion:

Because the `.loc` accessor is specifically designed to work with the user-defined labels (names) of the indices and columns, it is classified as "Label-based."

**Final Answer:** The correct option is (B).

**Answer: (B)**



Q41.

**Solution****Concept:**

The concept here is SQL Date Extraction Functions. In database systems, date and time data is often stored in a combined format (YYYY-MM-DD). To perform meaningful analysis, we often need to extract a specific component of that date, such as the year, the month, or the day of the month. MySQL provides dedicated functions for this purpose, such as YEAR(), MONTH(), and DAYOFMONTH() (which is synonymous with DAY()). Understanding these functions is vital for generating reports like monthly sales summaries or identifying events occurring on specific days.

**Solution:**

1. Analyzing the input:

The function call is `SELECT DAYOFMONTH('2026-05-15');`. The input is a string representing a date in the standard ISO format.

2. Understanding the function logic:

The DAYOFMONTH() function extracts the numeric day part from the given date. It returns a value ranging from 1 to 31, depending on the month. - Year: 2026 - Month: 05 (May) - Day: 15

3. Evaluating the options:

- Option (A) 5: This is the result you would get if you used the MONTH() function.
- Option (B) 15: This is the correct numerical day extracted from the date.
- Option (C) Friday: This is what the DAYNAME() function would return, providing the textual day of the week.
- Option (D) 2026: This is the result you would get if you used the YEAR() function.

4. Conclusion:

Since the query specifically requests the "day of the month" component, the function returns the integer 15.

**Final Answer:** The correct option is (B).

**Answer: (B)**



Q42.

**Solution****Concept:**

The concept here is Positional Indexing and Slicing with `.iloc` in Pandas. Slicing a DataFrame is a two-dimensional operation: you must define which rows you want and which columns you want. The syntax for `.iloc` is `df.iloc[row_slice, column_slice]`. Because `.iloc` is integer-positional, it follows the Python rule where the starting index is included and the ending index is excluded. This allows for precise extraction of data subsets based on their physical location in the table.

**Solution:**

1. Understanding the requirement:

The goal is to display the "first 3 columns" of a DataFrame named `df`.

2. Analyzing row selection:

Since the question does not specify a row filter, we want to include all rows. In Pandas slicing, the colon symbol (`:`) represents "all" along an axis. Therefore, the first argument in `.iloc` should be `:`.

3. Analyzing column selection:

We need the first 3 columns. In zero-based indexing, these are columns at positions 0, 1, and 2. To capture these using a slice, we must start at 0 and stop at 3 (because the stop value is exclusive). The column slice is `0:3`.

4. Evaluating the options:

- Option (A) `df.iloc[:, 0:3]`: This correctly selects all rows (`:`) and columns 0, 1, and 2 (`0:3`).
- Option (B) `df.iloc[0:3, :]`: This does the opposite; it selects the first 3 rows and all columns.
- Option (C) `df.iloc[:, 3]`: This would select only the 4th column, not a range of the first three.
- Option (D) `df.iloc[1:3, 1:3]`: This selects a sub-grid starting from the second row and second column.

5. Conclusion:

The correct syntactical representation for selecting all rows and the first three columns using integer-based positions is `df.iloc[:, 0:3]`.

**Final Answer:** The correct option is (A).

**Answer: (A)**



Q43.

**Solution****Concept:**

The concept here is Network Hardware and the OSI Physical Layer. Networking devices vary in their "intelligence"—that is, their ability to process and direct data based on its content. A Hub is the most basic networking device. It operates at the Physical Layer (Layer 1) of the OSI model. When a data packet arrives at one port of a hub, it does not check the destination MAC address. Instead, it simply copies the electrical signal and broadcasts it out to all other ports. This is inefficient and can lead to data collisions, but it is a fundamental concept in legacy networking.

**Solution:**

## 1. Analyzing the Hub:

A hub is a "dumb" device. It does not maintain an address table. Its only job is to repeat the signal it receives to every other device connected to it. This matches the description of "simply broadcasts signals to all ports."

## 2. Analyzing the Switch:

A switch is more intelligent. It operates at the Data Link Layer (Layer 2) and uses MAC addresses to send data only to the specific port where the destination device is connected. This significantly reduces network traffic.

## 3. Analyzing Bridges and Routers:

- A Bridge connects two network segments but still filters traffic based on MAC addresses.
- A Router operates at the Network Layer (Layer 3) and uses IP addresses to route data between different networks.

## 4. Conclusion:

Because the question describes a device that lacks filtering capabilities and merely broadcasts signals at the physical level, the Hub is the only correct answer.

**Final Answer:** The correct option is (B).

**Answer: (B)**



Q44.

**Solution****Concept:**

The concept here is Renaming DataFrame Columns in Pandas. Data scientists often need to clean or standardize column names to make them more descriptive or easier to type (e.g., removing spaces or converting to lowercase). The `rename()` method in Pandas is a versatile tool for this. It allows you to change specific labels without having to reassign the entire columns list. It uses a dictionary-like mapping where the keys are the "old" names and the values are the "new" names.

**Solution:**

1. Understanding the `rename()` method:

The correct syntax is `df.rename(columns='OldName': 'NewName')`. The "columns" parameter explicitly tells Pandas that you are mapping column labels.

2. Analyzing the dictionary mapping:

Inside the dictionary, you provide 'Old': 'New'. Pandas searches for the key 'Old' in the column index and replaces it with the value 'New'. If 'Old' is not found, no error is thrown, but no change occurs.

3. Evaluating the options:

- Option (A) is the standard, documented way to rename columns. It can also be done "in-place" by adding the argument `inplace=True`.
- Option (B) uses a non-existent `change()` method.
- Option (C) attempts to assign a single value to a column selection, which usually doesn't work for renaming the column header itself in this manner.
- Option (D) uses `update()`, which is used for merging data from one DataFrame into another, not for renaming labels.

4. Conclusion:

To specifically change the name of a column from 'Old' to 'New' while keeping the rest of the DataFrame intact, the `rename()` method with the `columns` parameter is the standard solution.

**Final Answer:** The correct option is (A).

**Answer: (A)**



Q45.

**Solution****Concept:**

The concept here is Matplotlib Pie Chart Customization. A pie chart is a circular statistical graphic divided into slices to illustrate numerical proportions. In a standard pie chart, all slices are joined together at the center. However, for emphasis, a user might want to "pull out" one or more slices to make them stand out from the rest of the chart. In Matplotlib, this is achieved using a parameter that defines the fraction of the radius by which each slice is offset from the center.

**Solution:**

1. Understanding the pie() function parameters:

The function `plt.pie()` takes several arguments, including labels, colors, and `autopct`. The specific parameter for separating slices is called `explode`.

2. Analyzing the `explode` parameter:

The `explode` parameter takes a list or array of values, one for each slice in the pie. A value of 0 means the slice stays in its original position. A value like 0.1 means the slice is moved outward by 10% of the radius. This visual separation is used to highlight specific data points, such as the largest share or an outlier.

3. Evaluating the options:

- Option (A) is incorrect as size is handled by the data values themselves.
- Option (B) correctly describes the "exploding" effect where a slice is separated from the main pie.
- Option (C) and (D) are handled by parameters like 'colors' and 'shadow' respectively.

4. Conclusion:

The term "explode" in the context of a Matplotlib pie chart is a visual metaphor for separating or highlighting a specific segment by moving it away from the center.

**Final Answer:** The correct option is (B).

**Answer: (B)**



Q46.

**Solution****Concept:**

The concept here is SQL Pattern Matching using the LIKE operator. In database queries, we often need to find records that match a specific text pattern rather than an exact string. SQL provides the LIKE operator along with two wildcard characters to facilitate this. The percent symbol (%) represents zero, one, or multiple characters, while the underscore symbol (\_) represents exactly one character. To find words that start with a specific letter and have a specific length, we combine the starting character with a precise number of underscores.

**Solution:****1. Understanding the requirement:**

We need to find names that start with the letter 'S' and have exactly 4 characters in total.

**2. Constructing the pattern:**

- The first character must be 'S'.
- To make the total length 4, we need 3 more characters after the 'S'.
- Since each underscore (\_) represents exactly one character, we must use three underscores.
- The resulting pattern is 'S\_\_\_'.

**3. Evaluating the options:**

- **Option (A)** 'S%': This matches any name starting with S, regardless of length (e.g., "S", "Sam", "Stephen").
- **Option (B)** 'S\_\_\_': This matches 'S' followed by exactly 3 characters, totaling 4 characters. This is the correct choice.
- **Option (C)** '\_\_\_\_S': This matches a 5-character name ending in S.
- **Option (D)** '%S%': This matches any name that contains the letter S anywhere.

**4. Conclusion:**

Using positional wildcards (underscores) is the standard method in SQL for enforcing specific string lengths during a search operation.

**Final Answer:** The correct option is (B).

**Answer: (B)**



Q47.

**Solution****Concept:**

The concept here is Data Alignment and Arithmetic in Pandas. A core strength of Pandas is its ability to align data based on labels (indices) during operations. When you perform an operation between two Series or DataFrames, Pandas does not simply add them positionally. Instead, it looks for matching index labels. If a label exists in both objects, the operation is performed. If a label is missing from one of the objects, Pandas cannot assume a value (like 0) and instead returns NaN (Not a Number) to indicate that the result is unknown or missing.

**Solution:**

1. Analyzing the data:

- Series A: index [0, 1], values [10, 20]

- Series B: index [1, 2], values [30, 40]

2. Matching the indices:

- Index 0: Present in A (10), absent in B. Result = NaN.

- Index 1: Present in A (20), present in B (30). Result =  $20 + 30 = 50$ .

- Index 2: Absent in A, present in B (40). Result = NaN.

3. Resulting Series:

The final Series will contain indices 0, 1, and 2. The values will be NaN, 50.0, and NaN. (Note: Pandas often converts the result to float because NaN is a floating-point value).

4. Conclusion:

This behavior prevents silent errors in data analysis where a missing category might otherwise be treated as zero, potentially skewing statistical results.

**Final Answer:** The correct option is (B).

**Answer: (B)**



Q48.

**Solution****Concept:**

The concept here is Network Security Devices, specifically the Firewall. In a world of increasing cyber threats, protecting a private network from unauthorized external access is critical. A firewall is a network security system that monitors and controls incoming and outgoing network traffic based on predetermined security rules. It acts as a barrier between a trusted internal network and untrusted external networks, such as the Internet. Firewalls can be implemented in both hardware and software, and they are the first line of defense in network security.

**Solution:**

## 1. Identifying the Firewall:

A firewall inspects data packets to see if they meet specific criteria (like originating from a safe IP address or using an allowed port). If the packet is deemed "unsafe," the firewall blocks it. This matches the description of "preventing unauthorized access to or from a private network."

## 2. Contrasting with other devices:

- A Router directs traffic between networks but does not inherently block traffic for security unless configured with an Access Control List (ACL).
- A Gateway connects different protocols but is not primarily a security device.
- An Antivirus is software that detects and removes malicious files (malware) from a computer, but it does not manage network-level traffic flow.

## 3. Conclusion:

Because the question focuses on the management of network access and the prevention of unauthorized intrusion at the network perimeter, the Firewall is the correct technical solution.

**Final Answer:** The correct option is (C).

**Answer:** (C)



Q49.

**Solution****Concept:**

The concept here is SQL Sorting with the ORDER BY clause. Once data is retrieved from a database, it is often necessary to present it in a specific order (e.g., alphabetical order for names or chronological order for dates). The ORDER BY clause is used at the end of a SELECT statement to sort the result set. By default, SQL sorts in ascending order (A to Z, or 0 to 9). However, this behavior can be explicitly modified using keywords to specify the direction of the sort.

**Solution:**

1. Understanding the syntax:

The clause is written as ORDER BY column\_name [ASC | DESC].

2. Identifying the default:

If no keyword is specified (e.g., ORDER BY Marks;), SQL assumes ASC, which stands for Ascending.

3. Identifying the Descending keyword:

To sort from highest to lowest or Z to A, the keyword DESC must be used. For example, ORDER BY Marks DESC; will show the highest marks at the top of the list.

4. Analyzing incorrect options:

- Option (A) DESCENDING: While descriptive, the actual SQL keyword is abbreviated to DESC.
- Option (C) DOWN: This is not a valid SQL keyword for sorting.
- Option (D) REVERSE: This is a function for reversing characters in a string, not for sorting rows in a result set.

5. Conclusion:

The standard SQL keyword used to sort data in descending order within an ORDER BY clause is DESC.

**Final Answer:** The correct option is (B).

**Answer: (B)**



Q50.

**Solution****Concept:**

The concept here is E-waste and Environmental Ethics. E-waste (electronic waste) refers to discarded electrical or electronic devices. Used electronics which are destined for refurbishment, reuse, resale, or recycling are also considered e-waste. Improper disposal of electronic waste leads to environmental pollution and health hazards due to the toxic materials like lead, mercury, and cadmium found in these components. Sustainable e-waste management involves reducing consumption, repairing devices when possible, and ensuring that recycling is done through authorized and safe channels.

**Solution:**

## 1. Evaluating "Donating to charity":

This is a highly effective way to handle e-waste. It extends the lifecycle of the product and delays its entry into the waste stream.

## 2. Evaluating "Recycling through authorized centers":

Authorized centers use specialized processes to extract valuable materials and safely dispose of toxic ones, preventing environmental contamination.

## 3. Evaluating "Throwing in a regular dustbin":

This is the incorrect and harmful way to dispose of e-waste. When electronics end up in regular landfills, toxic chemicals can leak into the soil and groundwater. This violates environmental safety guidelines.

## 4. Evaluating "Upcycling":

Finding a new, creative use for an old device is a valid form of waste reduction.

## 5. Conclusion:

Since the question asks for the method that is **NOT** a proper way to handle e-waste, throwing it in a regular dustbin is the correct answer.

**Final Answer:** The correct option is (C).

**Answer:** (C)



**Answer Key**

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

