CUET-UG Information Practices Sample Paper-6

Duration: 1 Hour **Maximum Marks: 250**

Instructions

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

Q1. A network administrator needs to connect two different star topology networks located in different buildings using a device that operates at the Data Link Layer and can filter traffic based on physical addresses. Which device is most suitable?

(A) Hub

(B) Repeater

(C) Bridge

(D) Gateway

Q2. Consider a Series object 'S' created using Pandas. If the command `print(S[S > 100].index[0])` results in a KeyError, what is the most likely reason?

(A) The Series contains no values greater than 100.

(B) The Series has non-integer indices.

(C) The Series 'S' has not been initialized.

(D) The index of the first filtered element is null.

Q3. In SQL, which function will return the position of the first occurrence of a substring "data" within a string "Informatics", and what is the specific result?

(A) INSTR("Informatics", "data"); Result: 0

(B) LOCATE("data", "Informatics"); Result: NULL

(C) SUBSTR("Informatics", "data"); Result: 0

(D) POSITION(”data” IN ”Informatics”); Result: -1

Q4. A relational database table ’Project’ has a composite primary key (EmpID, ProjID). Which of the following statements regarding the integrity of this table is strictly true?

(A) EmpID can be NULL if ProjID is unique.

(B) Neither EmpID nor ProjID can contain NULL values.

(C) The combination (EmpID, ProjID) can be repeated once.

(D) ProjID must be a foreign key from another table.

Q5. Which SQL clause is mandatory when using an aggregate function in the SELECT list along with a non-aggregated column?

(A) ORDER BY

(B) HAVING

(C) GROUP BY

(D) WHERE

Q6. A network technician needs to connect 20 computers in a local area network where the failure of one cable should not affect the rest of the network, but the central connection point must be active. Which topology and device combination is best suited for this?

(A) Ring Topology with a Repeater

(B) Star Topology with a Switch

(C) Bus Topology with a Terminator

(D) Mesh Topology with a Hub

Q7. Consider a DataFrame `df` with 5 rows and 3 columns. If you execute `df.iloc[1:4, 0:2]`, what will be the dimensions (shape) of the resulting object?

(A) (4, 2)

(B) (3, 3)

(C) (3, 2)

(D) (2, 3)

Q8. Which of the following SQL commands will correctly display the names of students in uppercase whose names end with the letter 'a'?

(A) `SELECT UPPER(Name) FROM Student WHERE Name LIKE 'a%';`

(B) `SELECT UPPER(Name) FROM Student WHERE Name = '%a';`

(C) `SELECT UPPER(Name) FROM Student WHERE Name LIKE '%a';`

(D) `SELECT Name.UPPER() FROM Student WHERE Name LIKE '%a';`

Q9. In a relational database, what is the result of a 'Projection' (π) operation on a table?

(A) It selects a subset of rows based on a condition.

(B) It selects a subset of columns from the table.

(C) It combines two tables based on a common attribute.

(D) It removes duplicate rows from the entire database.

Q10. Which SQL keyword is used to filter the results produced by a `GROUP BY` clause based on a specific condition?

(A) WHERE

(B) DISTINCT

(C) HAVING

(D) LIMIT

Q11. A network administrator is configuring a corporate network. To ensure security, they want to map a unique, permanent hardware identifier of each network interface card to a dynamic logical address. Which two address types are being referred to?

(A) SSID and IP Address

(B) MAC Address and IP Address

(C) Port Number and MAC Address

(D) Gateway Address and Subnet Mask

Q12. Consider a Pandas DataFrame 'df'. Which of the following code snippets will correctly return the last two rows and the first two columns of the DataFrame?

(A) `df.iloc[-2:, :2]`

(B) `df.iloc[:2, -2:]`

(C) `df.loc[-2:, :2]`

(D) `df.iloc[2:, 2:]`

Q13. In SQL, what will be the output of the command: `SELECT ROUND(157.48, -1), TRUNCATE(157.48, -1);`?

(A) 160, 150

(B) 157, 157

(C) 160, 157

(D) 157.5, 150

Q14. A table 'Employee' has columns (EmpID, Name, DeptID, Salary). Which Relational Algebra expression will display the 'Name' of all employees whose 'Salary' is greater than 50000?

(A) $\pi_{Name}(\sigma_{Salary>50000}(Employee))$

(B) $\sigma_{Name}(\pi_{Salary>50000}(Employee))$

(C) $\pi_{Salary>50000}(\sigma_{Name}(Employee))$

(D) $\sigma_{Salary>50000}(Employee)$

Q15. When performing a Cartesian Product ($A \times B$) between table A (5 rows, 3 columns) and table B (4 rows, 2 columns), what is the degree and cardinality of the resulting table?

(A) Degree: 6, Cardinality: 9

(B) Degree: 5, Cardinality: 20

(C) Degree: 20, Cardinality: 5

(D) Degree: 6, Cardinality: 20

Q16. A Python developer is using the Pandas library to handle a dataset. They want to create a Series object from a dictionary where the dictionary keys are strings. However, they want the Series to follow a specific order provided in a list. Which parameter of the Series constructor should they use?

(A) `dtype`

(B) `index`

(C) `copy`

(D) `name`

Q17. In SQL, consider a table `Sales` with columns `SaleID, Product, Quantity`. Which of the following queries will display the average `Quantity` sold for each `Product`, but only for those products whose average quantity exceeds 50?

(A) `SELECT Product, AVG(Quantity) FROM Sales WHERE AVG(Quantity) > 50 GROUP BY Product;`

(B) `SELECT Product, AVG(Quantity) FROM Sales GROUP BY Product HAVING AVG(Quantity) > 50;`

(C) `SELECT Product, AVG(Quantity) FROM Sales GROUP BY Product WHERE Quantity > 50;`

(D) `SELECT Product, AVG(Quantity) FROM Sales HAVING AVG(Quantity) > 50;`

Q18. Which property of a Pandas DataFrame allows you to access the underlying NumPy array representation of the data?

(A) `df.array`

(B) `df.values`

(C) `df.numpy()`

(D) `df.to`$_{a}rray()$

Q19. Under the IT Act and societal norms, what is the term used for the trails of data people leave behind while using the internet, including social media activities and website visits?

(A) Digital Evidence

(B) Cyber Trail

(C) Digital Footprint

(D) Electronic Shadow

Q20. What is the correct SQL syntax to extract the name of the month from a date column 'DOB' in a MySQL database?

(A) `MONTHNAME(DOB)`

(B) `GETMONTH(DOB)`

(C) `EXTRACTMONTH(DOB)`

(D) $\texttt{NAME}_O F_M ONTH(DOB)$

Q21. A network administrator is designing a network for a small office. They need to ensure that the network can handle high traffic volumes efficiently by reducing collisions. Which device should be used to segment the network into multiple collision domains, where each port represents its own collision domain?

(A) Hub

(B) Repeater

(C) Switch

(D) Passive Hub

Q22. Which of the following Pandas operations will result in a `ValueError` when attempting to add a new column to a DataFrame `df`?

(A) Assigning a scalar value to a new column name.

(B) Assigning a list of values whose length is equal to the number of rows in `df`.

(C) Assigning a Series with a different index but the same length as `df`.

(D) Assigning a list of values whose length is less than the number of rows in `df`.

Q23. Consider the SQL statement: `SELECT SUBSTR("Global Warming", -7, 4);`. What will be the output in a MySQL environment?

(A) Warm

(B) armi

(C) Warmin

(D) Global

Q24. Which relational algebra operator is used to remove rows from one relation that are also present in another relation, provided both relations are union-compatible?

(A) Union ($\cup$)

(B) Set Difference ($-$)

(C) Intersection ($\cap$)

(D) Cartesian Product ($\times$)

Q25. In a database of a retail store, the 'Orders' table has a 'Quantity' column. To find the total quantity sold for each 'ProductID' but only include orders where the 'Quantity' is greater than 10, which order of clauses is correct?

(A) `SELECT... FROM... GROUP BY... HAVING... WHERE...`

(B) `SELECT... FROM... WHERE... GROUP BY...`

(C) `SELECT... FROM... GROUP BY... WHERE...`

(D) `SELECT... FROM... WHERE... HAVING...`

Q26. A network administrator is connecting two different departments. Department A uses a 10BaseT network and Department B uses a 100BaseTX network. Which device is required to connect these two segments while also managing the difference in transmission speeds?

(A) Hub

(B) Repeater

(C) Switch

(D) Passive Hub

Q27. Consider a Pandas DataFrame 'df' with a MultiIndex. Which method is most efficient for swapping the levels of the MultiIndex?

(A) `df.reindex()`

(B) `df.swaplevel()`

(C) `df.transpose()`

(D) `df.pivot()`

Q28. In SQL, you need to find the total number of characters in the 'Description' column of a table, excluding any trailing spaces. Which combination of functions is correct?

(A) `LENGTH(RTRIM(Description))`

(B) `COUNT(TRIM(Description))`

(C) `LEN(LTRIM(Description))`

(D) $\texttt{CHAR}_L ENGTH(RTRIM(Description))$

Q29. A database administrator wants to ensure that a value in the 'DeptID' column of the 'Employee' table must exist in the 'DeptID' column of the 'Department' table. This is an example of which constraint?

(A) Primary Key Constraint

(B) Check Constraint

(C) Referential Integrity Constraint

(D) Unique Constraint

Q30. Which SQL join returns all rows from the left table and the matched rows from the right table, filling with NULL values where there is no match?

(A) Inner Join

(B) Right Join

(C) Left Join

(D) Full Join

Q31. A network designer is implementing a structure for a large campus. They want to use a topology that provides the highest level of redundancy by connecting every node to every other node. What is the total number of physical links required for a network with n nodes in this topology?

(A) $n(n-1)$

(B) n^2

(C) $\frac{n(n-1)}{2}$

(D) $2n-1$

Q32. You have two Pandas Series, `S1` and `S2`. `S1` has indices [0, 1, 2] and `S2` has indices [1, 2, 3]. If you perform the operation `S1 + S2`, what will be the index of the resulting Series?

(A) $[0, 1, 2, 3]$

(B) $[0, 1, 2]$

(C) $[1, 2, 3]$

(D) $[1, 2]$

Q33. A database user executes the following query: `SELECT ROUND(291.56, -2);`. What is the result, and why?

(A) 290; because it rounds to the nearest ten.

(B) 300; because it rounds to the nearest hundred.

(C) 291; because -2 is an invalid parameter.

(D) 200; because it truncates the tens and units.

Q34. Which property of a database ensures that once a transaction has been committed, it will remain so, even in the event of a power loss or system crash?

(A) Atomicity

(B) Consistency

(C) Isolation

(D) Durability

Q35. In SQL, which join is equivalent to an Inner Join where the join condition is based on all columns in the two tables that have the same name?

(A) Equi Join

(B) Natural Join

(C) Cross Join

(D) Self Join

Q36. A network administrator needs to connect two offices using a device that can read the logical IP addresses of incoming packets and determine the best path for data to travel across different networks. Which device is required?

(A) Bridge

(B) Switch

(C) Router

(D) Repeater

Q37. Given a DataFrame `df`, which of the following commands will display only the rows where the value in the 'Salary' column is between 30000 and 50000 (both inclusive)?

(A) `df[df['Salary'] == 30000 : 50000]`

(B) `df[(df['Salary'] >= 30000) & (df['Salary'] <= 50000)]`

(C) `df[(df['Salary'] >= 30000) | (df['Salary'] <= 50000)]`

(D) `df.filter('Salary' between 30000 and 50000)`

Q38. In SQL, what will be the result of the following expression: `SELECT MOD(11, 3), POWER(3, 2), INSTR('DATABASE', 'A');`?

(A) 2, 9, 2

(B) 3, 6, 1

(C) 2, 8, 3

(D) 1, 9, 2

Q39. In Relational Algebra, which operation is used to combine all rows from two relations R and S, provided they have the same number of attributes and matching domains?

(A) Set Difference

(B) Intersection

(C) Union

(D) Cartesian Product

Q40. Which SQL aggregate function should be used to find the total number of unique values in the 'Department' column, ignoring NULLs?

(A) `COUNT(*)`

(B) `SUM(DISTINCT Department)`

(C) `COUNT(DISTINCT Department)`

(D) `TOTAL(DISTINCT Department)`

Q41. A network engineer is troubleshooting a large enterprise network. They notice that the broadcast traffic from one department is affecting the performance of another department. Which device or technology should be implemented to logically segment the network into different broadcast domains?

(A) Hub

(B) VLAN on a Switch

(C) Bridge

(D) Repeater

Q42. Consider two DataFrames, `df1` and `df2`, both having a column named 'ID'. Which Pandas function is most appropriate to combine them based on the common 'ID' column, similar to a SQL JOIN?

(A) `pd.concat([df1, df2])`

(B) `df1.append(df2)`

(C) `pd.merge(df1, df2, on='ID')`

(D) `df1.combine(df2)`

Q43. What will be the output of the following SQL query if the 'DOB' column contains the value '2005-08-15'?

```
SELECT DAYOFMONTH(DOB), MONTH(DOB), YEAR(DOB);
```

(A) 15, Aug, 2005

(B) 15, 08, 2005

(C) 15, 8, 2005

(D) 08, 15, 2005

Q44. In a relational database schema, a 'Domain' is defined as:

(A) The number of attributes in a relation.

(B) The pool of atomic values from which an attribute draws its actual values.

(C) The total number of rows in a table.

(D) A unique identifier for a tuple.

Q45. A database contains a table `Orders` with columns `OrderID`, `CustomerID`, `OrderDate`, `Amount`. Which query will show the total `Amount` for each `CustomerID`, sorted from highest total to lowest?

(A)
```
SELECT CustomerID, SUM(Amount) FROM Orders GROUP BY CustomerID
ORDER BY SUM(Amount) DESC;
```

(B)
```
SELECT CustomerID, SUM(Amount) FROM Orders GROUP BY CustomerID
ORDER BY Amount ASC;
```

(C) `SELECT CustomerID, SUM(Amount) FROM Orders ORDER BY SUM(Amount) DESC;`

(D) `SELECT CustomerID, TOTAL(Amount) FROM Orders GROUP BY CustomerID SORT BY Amount;`

Q46. A network administrator is concerned about E-waste and wants to implement a policy for the sustainable disposal of old office computers. Which of the following is considered the most environmentally friendly method under the 'Societal Impacts' of IT?

(A) Discarding them in local landfills.

(B) Selling components to uncertified local scrap dealers.

(C) Refurbishing and donating to schools or recycling through authorized vendors.

(D) Burning the plastic components to recover copper wiring.

Q47. A data scientist has a DataFrame `df` with a 'Date' column. They want to set this 'Date' column as the index of the DataFrame permanently. Which command should they use?

(A) `df.index = df['Date']`

(B) $\texttt{df.set}_{i}ndex('Date', inplace = True)$

(B) `df.reindex('Date')`

(C) $\texttt{df.add}_{i}ndex('Date')$

Q48. In SQL, what is the difference between `CHAR(10)` and `VARCHAR(10)` data types?

(A) `CHAR` is for numbers, `VARCHAR` is for text.

(B) `CHAR` is variable length, `VARCHAR` is fixed length.

(C) `CHAR` is fixed length (pads with spaces), `VARCHAR` is variable length.

(D) There is no difference; they are used interchangeably.

Q49. A user wants to create a bar chart in Matplotlib to compare the 'Sales' of five

different 'Regions'. Which function and customization parameter should be used to change the color of the bars to 'green' and add a label for the y-axis?

(A) `plt.bar(regions, sales, color='green'); plt.ylabel('Sales')`

(B) `plt.plot(regions, sales, type='bar'); plt.y('Sales')`

(C) `plt.hist(regions, sales, c='g'); plt.label`$_y('Sales')$

(C) `plt.barplot(regions, sales, color='green'); plt.yaxis('Sales')`

Q50. Which of the following is a violation of Intellectual Property Rights (IPR) where someone uses another person's creative work without permission or credit?

(A) Phishing

(B) Plagiarism

(C) Hacking

(D) Spamming

Detailed Solutions

Q1.

Solution

Concept:

Network devices operate at different layers of the OSI model. A Bridge is a Layer 2 (Data Link Layer) device that connects two different network segments. It maintains a MAC address table to filter traffic, ensuring that data is only forwarded to the specific segment where the destination MAC address resides, thereby reducing network congestion.

Solution:

1. A Hub is a Layer 1 device that broadcasts data to all ports, lacking filtering capabilities. 2. A Repeater is a Layer 1 device used to amplify signals but does not filter based on addresses. 3. A Bridge operates at the Data Link Layer and filters traffic using MAC (physical) addresses, making it ideal for connecting two buildings while managing traffic efficiently. 4. A Gateway operates at higher layers and is typically used to connect networks with different protocols.

Final Answer: The device that filters traffic based on physical addresses at the Data Link Layer is a Bridge.

Answer: (C)

Q2.

Solution

Concept:

In Pandas, boolean indexing allows us to filter data. When we use `S[S > 100]`, it returns a new Series containing only the elements that satisfy the condition. If no elements meet the criteria, the result is an empty Series. Accessing `.index[0]` on an empty Series results in an IndexError or KeyError because the index list is empty.

Solution:

1. The expression `S > 100` creates a boolean mask. 2. `S[S > 100]` applies this mask to the Series. 3. If no values in 'S' are greater than 100, the filtered Series is empty. 4. An empty Series has an empty index (`.index`). 5. Attempting to access the first element of an empty index (`[0]`) causes the error because there is no "first element" to retrieve.

Final Answer: The error occurs because the Series contains no values greater than 100, leading to an empty filtered object.

Answer: (A)

Q3.

Solution

Concept:

SQL provides various string functions to locate substrings. The `INSTR(string, substring)` function returns the position of the first occurrence of a substring within a string. In MySQL and standard SQL, positions are 1-indexed. If the substring is not found, the function returns 0.

Solution:

1. The target string is ”Informatics”. 2. The substring to search for is ”data”. 3. Scanning ”Informatics”, the sequence ”d-a-t-a” does not exist within it. 4. Since the substring is missing, the `INSTR` function returns 0. 5. Other functions like `LOCATE` or `POSITION` would also indicate a lack of match, but `INSTR` specifically returning 0 is the standard behavior for a ”not found” result.

Final Answer: The function returns 0 because ”data” is not present in ”Informatics”.

Answer: (A)

Q4.

Solution

Concept:

The Entity Integrity Constraint states that no primary key attribute can contain a NULL value. This is because a primary key is used to uniquely identify each row in a table. In the case of a composite primary key (a key made of multiple columns), this rule applies to every column that constitutes the key.

Solution:

1. A composite primary key is formed by the combination of (EmpID, ProjID). 2. For the combination to be unique and reliable for identification, every part of that combination must be known. 3. Therefore, EmpID cannot be NULL, and ProjID cannot be NULL. 4. If either were NULL, the database could not guarantee the uniqueness of the record or identify it properly.

Final Answer: Neither EmpID nor ProjID can contain NULL values as they are parts of the primary key.

Answer: (B)

Q5.

Solution

Concept:

When a SQL query includes both an aggregate function (like `SUM`, `COUNT`, `AVG`) and a regular, non-aggregated column in the `SELECT` statement, the database needs to know how to group the calculations. This is handled by the `GROUP BY` clause.

Solution:

1. Consider a query: `SELECT Department, COUNT(EmpID) FROM Employee`. 2. Without a grouping instruction, the database does not know whether to count all employees or count them per department. 3. To fix this, you must include `GROUP BY Department`. 4. Standard SQL rules dictate that any column in the `SELECT` list that is not part of an aggregate function must be present in the `GROUP BY` clause.

Final Answer: The GROUP BY clause is mandatory for organizing results when mixing aggregate and non-aggregate columns.

Answer: (C)

Q6.

Solution

Concept:

In a Star topology, all nodes are connected to a central connection point (like a switch or hub). If a single cable connecting a computer to the central node fails, only that computer goes offline; the rest of the network remains functional. A Switch is preferred over a Hub because it directs data only to the intended recipient, improving efficiency.

Solution:

1. In a Bus topology, a single cable failure can bring down the entire network. 2. In a Ring topology, a break in the cable usually disrupts the whole loop. 3. Star topology provides the required fault isolation for individual cables. 4. Using a Switch (Layer 2) is more efficient than a Hub (Layer 1) for managing 20 computers.

Final Answer: Star Topology with a Switch is the best combination for cable failure resilience and central management.

Answer: (B)

Q7.

Solution

Concept:

The `.iloc` indexer in Pandas follows the standard Python slicing rule: `start:stop`, where the `start` index is inclusive and the `stop` index is exclusive. The number of elements retrieved is calculated as `stop - start`.

Solution:

1. For rows: `1:4` means indices 1, 2, and 3. Calculation: $4 - 1 = 3$ rows. 2. For columns: `0:2` means indices 0 and 1. Calculation: $2 - 0 = 2$ columns. 3. The resulting object will therefore have 3 rows and 2 columns. 4. The shape attribute of this resulting DataFrame will be `(3, 2)`.

Final Answer: The dimensions of the resulting object will be (3, 2).

Answer: (C)

Q8.

Solution

Concept:

To manipulate and filter string data in SQL, we use the `UPPER()` function for case conversion and the `LIKE` operator with wildcards for pattern matching. The `%` wildcard represents zero or more characters.

Solution:

1. `UPPER(Name)` converts the values in the Name column to capital letters. 2. To find names ending with 'a', the pattern must be `'%a'` (anything followed by 'a'). 3. `'a%'` would search for names starting with 'a'. 4. Option C combines the correct function `UPPER()` with the correct pattern matching syntax `LIKE '%a'`.

Final Answer: The correct command is SELECT UPPER(Name) FROM Student WHERE Name LIKE '%a';.

Answer: (C)

Q9.

Solution

Concept:

Relational Algebra consists of several operations. 'Selection' (σ) filters rows, while 'Projection' (π) filters columns. Projection allows a user to view only specific attributes of interest from a relation, effectively reducing the "vertical" size of the table.

Solution:

1. Selection (σ) is used to pick rows (tuples) that satisfy a condition (e.g., Age > 20). 2. Projection (π) is used to pick specific columns (attributes) like Name and RollNo. 3. It does not combine tables (that is a Join) nor does it primarily focus on row duplicates across the whole database. 4. Therefore, the result of a Projection is a subset of columns.

Final Answer: A Projection operation selects a subset of columns from the table.

Answer: (B)

Q10.

Solution

Concept:

The `WHERE` clause is used to filter individual rows before any grouping occurs. To filter groups created by the `GROUP BY` clause (for example, filtering based on the result of an aggregate function like `SUM` or `COUNT`), the `HAVING` clause must be used.

Solution:

1. `WHERE` cannot be used with aggregate functions directly. 2. If we want to find departments where the `COUNT(*)` is greater than 5, we must group them first. 3. After `GROUP BY Department`, we apply `HAVING COUNT(*) > 5`. 4. `HAVING` acts as a filter on the summarized data.

Final Answer: The HAVING keyword is used to filter results after the GROUP BY clause.

Answer: (C)

Q11.

Solution

Concept:

Every Network Interface Card (NIC) has a unique, permanent physical address assigned by the manufacturer called the MAC (Media Access Control) Address. For communication over a network, a logical, dynamic address called the IP (Internet Protocol) Address is assigned to the device.

Solution:

1. MAC Address is a 48-bit hardware identifier (e.g., 00-0C-29-44-53-A1). 2. IP Address is a logical address (e.g., 192.168.1.1) that can change based on the network. 3. SSIDs are used for wireless network names, and Port numbers identify specific services/applications. 4. Therefore, the mapping involves MAC and IP addresses.

Final Answer: The unique hardware identifier is the MAC address and the dynamic logical address is the IP address.

Answer: (B)

Q12.

Solution

Concept: In Pandas `.iloc`, negative indexing can be used to count from the end of the DataFrame. The syntax $df.iloc[row_selection, column_selection]$ uses slices where `-n:` indicates "from the $n-th$ last element to the end".

Solution:

(a) To get the last two rows, we use the row slice `-2:`.

(b) To get the first two columns, we use the column slice `:2` (which implies indices 0 to 1).

(c) Combining these into `.iloc`, we get `df.iloc[-2:, :2]`.

(d) `.loc` is label-based and would not typically use negative integer slices unless the labels themselves were integers.

Final Answer: The correct snippet is `df.iloc[-2:, :2]`.

Answer: (A)

Q13.

Solution

Concept:

The `ROUND(n, d)` function rounds n to d decimal places. If d is negative, it rounds to the left of the decimal point (tens, hundreds, etc.). The `TRUNCATE(n, d)` function simply chops off the numbers to the specified place without rounding.

Solution:

1. `ROUND(157.48, -1)`: We round at the tens place. Since the unit digit is 7 (which is ≥ 5), the tens digit increases. Result: 160. 2. `TRUNCATE(157.48, -1)`: We move to the tens place and replace everything to its right with zero without checking for rounding. Result: 150. 3. Comparing the two, we get 160 and 150.

Final Answer: The results are 160 and 150 respectively.

Answer: (A)

Q14.

Solution

Concept:

In Relational Algebra, operations are nested. The inner operation is performed first. We use Selection (σ) to filter the rows based on a condition and Projection (π) to select the specific columns to be displayed.

Solution:

1. First, we need to find the specific rows where Salary > 50000. This is $\sigma_{Salary>50000}(Employee)$. 2. From that resulting set of rows, we only want to see the 'Name' column. This requires applying π_{Name} to the result of the selection. 3. The nested expression is $\pi_{Name}(\sigma_{Salary>50000}(Employee))$. 4. Option B is incorrect because you cannot select a condition from a projection if that projection has already discarded the 'Salary' column.

Final Answer: The correct expression is $\pi_{Name}(\sigma_{Salary>50000}(Employee))$.

Answer: (A)

Q15.

Solution

Concept:

The Cartesian Product ($A \times B$) combines every row of table A with every row of table B. - **Degree:** The number of columns (Sum of degrees of both tables). - **Cardinality:** The number of rows (Product of cardinalities of both tables).

Solution:

1. Table A Degree (columns) = 3; Table B Degree = 2. Resulting Degree = $3 + 2 = 5$. 2. Table A Cardinality (rows) = 5; Table B Cardinality = 4. Resulting Cardinality = $5 \times 4 = 20$. 3. Therefore, the resulting table has a degree of 5 and a cardinality of 20.

Final Answer: Degree: 5, Cardinality: 20.

Answer: (B)

Q16.

Solution

Concept:

When creating a Pandas Series from a dictionary, the keys automatically become the index. However, if a specific order or a subset of keys is required, the `index` parameter can be explicitly passed to the constructor. If the `index` parameter contains labels not present in the dictionary, Pandas will assign `NaN` to those positions.

Solution:

1. The Series constructor syntax is `pd.Series(data, index=index, ...)`. 2. `dtype` defines the data type of the values. 3. `index` allows the user to specify the sequence of labels for the Series. 4. If a dictionary is provided as `data`, Pandas aligns the dictionary values with the labels provided in the `index` list.

Final Answer: The index parameter is used to specify a custom order or selection of labels.

Answer: (B)

Q17.

Solution

Concept:

To filter groups based on an aggregate result (like `AVG`, `SUM`, `COUNT`), SQL requires the `HAVING` clause. The `WHERE` clause cannot filter groups because it is evaluated before the `GROUP BY` operation occurs.

Solution:

1. We must first group the data by `Product` to calculate averages per product. 2. The aggregate calculation `AVG(Quantity)` is performed for each group. 3. To filter these groups, we use `HAVING AVG(Quantity) > 50`. 4. Option A is invalid because `AVG()` cannot be used in a `WHERE` clause.

Final Answer: The correct syntax uses GROUP BY followed by HAVING to filter the aggregate results.

Answer: (B)

Q18.

Solution

Concept:

Pandas is built on top of NumPy. Sometimes developers need to access the raw underlying array data without row indices or column labels.

Solution:

(a) `df.values` returns the DataFrame data as a NumPy `ndarray`.

(b) In newer versions of Pandas, `df.to_numpy()` is also commonly used.

(c) However, `df.values` remains a standard attribute for accessing the underlying NumPy representation.

(d) `df.array` returns a Pandas extension array, not a standard NumPy array.

Final Answer: The `.values` attribute provides the NumPy representation of the DataFrame data.

Answer: (B)

Q19.

Solution

Concept:

Everything a user does online contributes to their permanent online record. This includes ”passive” data collection (like IP addresses logged by websites) and ”active” data sharing (like posting on social media).

Solution:

1. Digital Footprint is the specific term for the record of one's interactions and presence in the digital environment. 2. It can be categorized into Active (data intentionally shared) and Passive (data collected without the user's active participation). 3. It is a critical concept in digital ethics and cybersecurity as it defines a user's online reputation.

Final Answer: The term for the trail of data left behind online is Digital Footprint.

Answer: (C)

Q20.

Solution

Concept:

MySQL provides built-in date and time functions to manipulate temporal data. To retrieve the numeric month (1-12), `MONTH()` is used, but to retrieve the full name (e.g., 'January'), a specific function is required.

Solution:

1. `MONTHNAME(date)` is the standard MySQL function that returns the full name of the month for a given date. 2. `MONTH(date)` would only return the integer index of the month. 3. Other options like `EXTRACTMONTH` or `GETMONTH` are not valid built-in MySQL function names for this specific purpose.

Final Answer: The correct function to get the month name is MONTHNAME(DOB).

Answer: (A)

Q21.

Solution

Concept:

A Collision Domain is a network segment where data packets can collide with one another when sent on a shared medium. Hubs create a single large collision domain. Switches, however, operate at the Data Link Layer and use micro-segmentation, meaning every port on a switch is its own separate collision domain.

Solution:

1. A Hub broadcasts incoming signals to all ports; if two devices transmit simultaneously, a collision occurs. 2. A Switch maintains a MAC address table. It identifies the destination and creates a temporary dedicated path between the sender and receiver. 3. Because each port on a switch is isolated, collisions are prevented between devices on different ports. 4. Therefore, to segment a network into multiple collision domains and improve efficiency, a Switch is the required device.

Final Answer: The Switch is the device where each port represents its own collision domain.

Answer: (C)

Q22.

Solution

Concept:

When adding a new column to a Pandas DataFrame using a list or a NumPy array, the length of the new data must exactly match the number of rows (length) of the existing DataFrame. If it does not, Pandas cannot align the data and raises a `ValueError`.

Solution:

1. Assigning a scalar (e.g., `df['new'] = 5`) works via broadcasting—Pandas fills every row with 5. 2. Assigning a Series of the same length works, even if indices are different (missing indices become `NaN`). 3. Assigning a list requires an exact length match. If the DataFrame has 10 rows and the list has 8 elements, there is no logical way for Pandas to "guess" where the missing data goes. 4. This mismatch triggers a `ValueError: Length of values does not match length of index`.

Final Answer: Assigning a list of values whose length is different from the number of rows in the DataFrame results in a ValueError.

Answer: (D)

Q23.

Solution

Concept:

The `SUBSTR(string, start, length)` function extracts a part of a string. In MySQL, a negative `start` value means counting backward from the end of the string. The `length` parameter defines how many characters to extract moving forward from that starting point.

Solution:

1. The string is "Global Warming". 2. Counting 7 characters from the end: - 'g' is -1, 'n' is -2, 'i' is -3, 'm' is -4, 'r' is -5, 'a' is -6, 'W' is -7. 3. The starting character is 'W'. 4. From 'W', we extract 4 characters forward: 'W', 'a', 'r', 'm'. 5. The resulting substring is "Warm".

Final Answer: The output of the statement is "Warm".

Answer: (A)

Q24.

Solution

Concept:

The Set Difference operator ($-$) is used in relational algebra to find tuples that are present in the first relation but not in the second. For this operation to be valid, both relations must be union-compatible, meaning they must have the same number of attributes and corresponding attributes must have the same domain.

Solution:

1. Union ($\cup$) combines all rows from both. 2. Intersection ($\cap$) finds rows common to both. 3. Set Difference ($R - S$) identifies rows unique to R. It effectively "removes" from R any rows that also exist in S. 4. This is the standard operator for finding "A but not B".

Final Answer: The Set Difference ($-$) operator is used to remove rows present in another relation.

Answer: (B)

Q25.

Solution

Concept:

In SQL, the order of execution is crucial. The `WHERE` clause filters individual rows *before* they are grouped. The `GROUP BY` clause then aggregates the remaining rows. The `HAVING` clause filters the resulting groups.

Solution:

1. Requirement 1: Only include orders where Quantity > 10. This is a row-level filter, so it must use `WHERE`. 2. Requirement 2: Find total quantity for each ProductID. This requires `GROUP BY ProductID` and `SUM(Quantity)`. 3. The standard syntax structure is: `SELECT... FROM... WHERE... GROUP BY...`. 4. Note that if we wanted to filter the *totals* (e.g., total > 100), we would then add `HAVING` after the `GROUP BY`.

Final Answer: The correct sequence is WHERE (to filter rows) followed by GROUP BY (to aggregate).

Answer: (B)

Q26.

Solution

Concept:

A Switch is an ”intelligent” networking device that operates at the Data Link Layer. Unlike hubs or repeaters, switches have buffers and processing capabilities that allow them to connect segments operating at different speeds (e.g., 10 Mbps and 100 Mbps). The switch stores the incoming frame and then transmits it at the speed of the destination segment.

Solution:

1. A Hub or Repeater simply forwards bits; they cannot handle different speeds and would cause a mismatch/failure. 2. A Switch can receive data at 10 Mbps and forward it to a 100 Mbps port because it processes frames individually. 3. This process is often called ”speed buffering” or ”autosensing.” 4. Therefore, the Switch is the correct device for heterogeneous speed segments.

Final Answer: The Switch is required to manage different transmission speeds between network segments.

Answer: (C)

Q27.

Solution

Concept:

Pandas MultiIndex (hierarchical indexing) allows multiple levels of indices on an axis. Sometimes the hierarchy must be rearranged for better grouping or analysis. The `swaplevel()` method exchanges two index levels.

Solution:

(a) `swaplevel(i, j)` swaps two specified index levels.

(b) The levels may be identified using names or integer positions.

(c) `reorder_levels()` is more general and can rearrange multiple levels.

(d) `transpose()` swaps rows and columns, not index levels.

(e) `pivot()` reshapes tabular data but does not swap MultiIndex levels.

Final Answer: The `swaplevel()` method is used for swapping levels of a MultiIndex.

Answer: (B)

Q28.

Solution

Concept:

To find the character count of a string, MySQL uses `LENGTH()` (which returns the length in bytes) or `CHAR_LENGTH()` (which returns the length in characters). To remove trailing spaces, the `RTRIM()` (Right Trim) function is used.

Solution:

1. `RTRIM(Description)` removes all spaces from the right side of the string. 2. `LENGTH()` then calculates the size of the resulting string. 3. `LTRIM()` only removes leading (left) spaces. 4. `TRIM()` would remove both, but the question specifically targets trailing spaces. 5. Therefore, `LENGTH(RTRIM(Description))` is the standard way to achieve this.

Final Answer: The correct combination is LENGTH(RTRIM(Description)).

Answer: (A)

Q29.

Solution

Concept:

Referential Integrity ensures that the relationship between two tables remains consistent. This is implemented via a Foreign Key. It prevents users from entering a value in a child table (Employee) that does not exist in the parent table (Department).

Solution:

1. A Primary Key uniquely identifies a row in its own table. 2. A Check constraint ensures a column value meets a specific boolean condition (e.g., Salary > 0). 3. Referential Integrity (Foreign Key) links 'DeptID' in 'Employee' to 'DeptID' in 'Department'. 4. This ensures that every employee belongs to a valid, existing department.

Final Answer: This is an example of a Referential Integrity Constraint.

Answer: (C)

Q30.

Solution

Concept:

A Join is used to combine rows from two or more tables based on a related column. An Outer Join is used when we want to include rows even if there is no match in the joining table.

Solution:

1. An Inner Join only returns rows where there is a match in both tables. 2. A Left Join (or Left Outer Join) returns all rows from the left table. If there is no match in the right table, the result set will contain NULL for the right table's columns. 3. A Right Join does the opposite (all rows from the right table). 4. A Full Join returns all rows from both tables, using NULLs where matches are missing.

Final Answer: The Left Join returns all rows from the left table and matched rows from the right.

Answer: (C)

Q31.

Solution

Concept:

A Full Mesh topology offers maximum reliability because every node is directly connected to every other node in the network. This ensures that even if multiple links fail, alternative paths always exist. To calculate the number of links, we use the combination formula ${}^{n}C_{2}$.

Solution:

1. In a network of n nodes, each node must connect to $(n-1)$ other nodes. 2. If we simply multiply $n \times (n-1)$, we count each link twice (once for each end). 3. Therefore, the number of unique physical links is $\frac{n(n-1)}{2}$. 4. For example, if $n = 4$, links $= \frac{4\times3}{2} = 6$.

Final Answer: The total number of links in a full mesh topology is $\frac{n(n-1)}{2}$.

Answer: (C)

Q32.

Solution

Concept:

When performing arithmetic operations on two Pandas Series, Pandas performs "Data Alignment." It aligns the data based on the index labels. The resulting index is the "Union" of the indices of the two objects.

Solution:

1. `S1` index: {0, 1, 2} 2. `S2` index: {1, 2, 3} 3. The union of these two sets of labels is {0, 1, 2, 3}. 4. For labels 1 and 2, the values will be added. For labels 0 and 3, which exist in only one Series, the result will be `NaN` (Not a Number).

Final Answer: The resulting index will be [0, 1, 2, 3].

Answer: (B)

Q33.

Solution

Concept:

The `ROUND(number, decimals)` function in SQL allows for negative decimal values. A value of -1 rounds to the nearest ten, -2 rounds to the nearest hundred, and so on.

Solution:

1. We are rounding 291.56 to the -2 position (hundreds place). 2. We look at the digit to the right of the hundreds place, which is 9 (the tens digit). 3. Since 9 is greater than or equal to 5, we round up the hundreds digit. 4. The 2 in the hundreds place becomes 3, and the digits to the right become zero. 5. The result is 300.

Final Answer: The result is 300 as it rounds to the nearest hundred.

Answer: (B)

Q34.

Solution

Concept:

ACID properties ensure reliable database transactions. - **Atomicity:** All or nothing. - **Consistency:** Valid state to valid state. - **Isolation:** Concurrent transactions don't interfere. - **Durability:** Committed data is permanent.

Solution:

1. Durability guarantees that once a user receives a confirmation that a transaction is successful, the changes are recorded in non-volatile memory (like a hard disk). 2. Even if the system crashes immediately after the commit, the database recovery logs will ensure the data is present when the system restarts.

Final Answer: Durability is the property that ensures committed transactions survive system failures.

Answer: (D)

Q35.

Solution

Concept:

A Natural Join is a type of Join that automatically matches columns between two tables based on their names and data types. It then performs an Equi Join on those columns and returns only one copy of each matched column.

Solution:

1. If Table A and Table B both have a column named `ID`, a Natural Join will automatically execute `WHERE A.ID = B.ID`. 2. Unlike a standard Inner Join, you do not need to specify the `ON` or `USING` clause. 3. It simplifies queries but requires strict naming conventions in the database schema.

Final Answer: The Natural Join is the join based on all columns with the same name.

Answer: (B)

Q36.

Solution

Concept:

A Router is a Layer 3 (Network Layer) device. Unlike switches or bridges that use MAC addresses to navigate within a local network, routers use logical IP addresses to direct traffic between different networks (inter-networking). Routers maintain routing tables to calculate the most efficient path for data packets.

Solution:

1. A Bridge and Switch operate at the Data Link Layer (Layer 2) and use physical addresses. 2. A Repeater operates at the Physical Layer (Layer 1) to amplify signals. 3. A Router understands IP addresses and connects different subnets or completely different network types (like a LAN to the Internet). 4. Because the requirement mentions reading IP addresses and determining the best path, a Router is the only correct choice.

Final Answer: The Router is the device that uses logical IP addresses to determine the best path across networks.

Answer: (C)

Q37.

Solution

Concept:

Filtering a Pandas DataFrame based on multiple conditions requires the use of boolean operators. In Pandas, the standard Python keywords `and` and `or` cannot be used for element-wise comparisons on Series; instead, bitwise operators `&` (AND) and `|` (OR) are used. Each condition must be enclosed in parentheses.

Solution:

1. Condition 1: `df['Salary'] >= 30000` 2. Condition 2: `df['Salary'] <= 50000` 3. Since both must be true (inclusive range), we connect them with `&`. 4. Using `|` would return all rows where Salary is either $\geq$ 30000 or $\leq$ 50000, which would effectively return almost all numeric data. 5. Slicing with `:` is only for indices or position, not for value-based filtering.

Final Answer: The correct syntax is df[(df['Salary'] >= 30000) & (df['Salary'] <= 50000)].

Answer: (B)

Q38.

Solution

Concept:

This query tests three mathematical and string functions: - `MOD(a, b)`: Returns the remainder of a/b. - `POWER(a, b)`: Returns a^b. - `INSTR(str, sub)`: Returns the 1-indexed position of the first occurrence of the substring.

Solution:

1. `MOD(11, 3)`: $11 \div 3 = 3$ with a remainder of 2. Result: 2. 2. `POWER(3, 2)`: $3^2 = 9$. Result: 9. 3. `INSTR('DATABASE', 'A')`: In the word 'DATABASE', 'D' is at 1 and 'A' is at 2. Result: 2. 4. Combining these, we get: 2, 9, 2.

Final Answer: The output is 2, 9, 2.

Answer: (A)

Q39.

Solution

Concept:

The Union (∪) operation in Relational Algebra combines all tuples from two relations. For this to be valid, the relations must be Union-Compatible: they must have the same arity (number of columns) and the domains of corresponding attributes must match.

Solution:

1. Intersection finds only common rows. 2. Set Difference finds rows in one but not the other. 3. Cartesian Product combines columns from both and creates every possible row pair. 4. Union effectively ”stacks” the tables, including all records from both while removing duplicates (in set theory terms).

Final Answer: The Union operation combines all rows from two union-compatible relations.

Answer: (C)

Q40.

Solution

Concept:

Aggregate functions perform a calculation on a set of values. The `COUNT()` function can be modified with the `DISTINCT` keyword to only count non-duplicate values. By default, aggregate functions ignore NULL values (except `COUNT(*)`).

Solution:

1. `COUNT(*)` counts every row, including NULLs and duplicates. 2. `SUM()` is used for adding numbers, not counting unique labels. 3. `COUNT(DISTINCT Department)` first identifies the unique labels in the column, removes the NULLs, and then counts the remaining entries. 4. For example, if departments are [Sales, Sales, HR, NULL], the distinct set is {Sales, HR}, and the count is 2.

Final Answer: The correct function is COUNT(DISTINCT Department).

Answer: (C)

Q41.

Solution

Concept:

A Broadcast Domain is a portion of a network in which any device can transmit data directly to another device without going through a router. By default, all ports on a switch belong to the same broadcast domain. To restrict broadcast traffic, Virtual Local Area Networks (VLANs) are used to logically segment a single physical switch into multiple isolated networks.

Solution:

1. Hubs and Repeaters extend the collision and broadcast domains. 2. A standard Switch breaks up collision domains but maintains a single broadcast domain. 3. Implementing a VLAN (Virtual LAN) allows an administrator to group ports logically. Broadcasts sent by a device in VLAN 10 will not be received by devices in VLAN 20. 4. This improves security and network performance by reducing unnecessary traffic.

Final Answer: VLAN on a Switch is the technology used to segment a network into different broadcast domains.

Answer: (B)

Q42.

Solution

Concept:

Pandas provides several ways to combine DataFrames. `pd.concat()` is mainly used for stacking DataFrames vertically or horizontally along an axis. For database-style joins where data is aligned based on common keys (columns), `pd.merge()` is the primary and most flexible function.

Solution:

1. `pd.concat()` simply glues objects together. 2. `append()` is a simplified version of concat for adding rows (now deprecated in newer versions). 3. `pd.merge()` allows for Inner, Left, Right, and Outer joins, specifically using the `on` parameter to define the join key. 4. Since the requirement is to combine based on the 'ID' column, `pd.merge()` is the correct choice.

Final Answer: The pd.merge() function is used to combine DataFrames based on a common column.

Answer: (C)

Q43.

Solution

Concept:

SQL date functions extract specific integer components from a date string formatted as 'YYYY-MM-DD'. - `DAYOFMONTH()` returns the day (1-31). - `MONTH()` returns the numeric month (1-12). - `YEAR()` returns the four-digit year.

Solution:

1. For the date '2005-08-15': 2. `DAYOFMONTH('2005-08-15')` is 15. 3. `MONTH('2005-08-15')` is 8. 4. `YEAR('2005-08-15')` is 2005. 5. Note: `MONTH()` returns the integer without a leading zero unless formatted otherwise. Option B (08) is incorrect for the standard return type.

Final Answer: The output is 15, 8, 2005.

Answer: (C)

Q44.

Solution

Concept:

In relational database theory, a Domain represents the set of allowable values for one or more attributes. It defines the data type, range, and format that the values must adhere to.

Solution:

1. The number of attributes is called the 'Degree'. 2. The number of rows is called 'Cardinality'. 3. A unique identifier is a 'Primary Key'. 4. The Domain acts as a constraint, ensuring that an attribute like 'Age' only contains valid integers (e.g., 0 to 120), rather than text or negative numbers.

Final Answer: A Domain is the pool of atomic values from which an attribute draws its actual values.

Answer: (B)

Q45.

Solution

Concept:

To present summarized data in a specific order, SQL uses the `GROUP BY` clause for aggregation and the `ORDER BY` clause for sorting. To sort from highest to lowest, the `DESC` (descending) keyword must be used.

Solution:

1. We need to group by `CustomerID` to calculate the sum per customer. 2. `SUM(Amount)` provides the total for each group. 3. `ORDER BY SUM(Amount) DESC` takes those calculated totals and sorts the entire result set in descending order. 4. Option B sorts in ascending order (lowest to highest). Option C is invalid because you cannot use `SUM` without a `GROUP BY` if you are also selecting `CustomerID`.

Final Answer: The correct query uses GROUP BY CustomerID followed by ORDER BY SUM(Amount) DESC.

Answer: (A)

Q46.

Solution

Concept:

E-waste (Electronic Waste) contains toxic substances like lead, mercury, and cadmium. Improper disposal can lead to environmental degradation and health hazards. Sustainable IT practices focus on the ”3 Rs”: Reduce, Reuse, and Recycle through official channels.

Solution:

1. Landfills and burning are hazardous and release toxins into the soil and air. 2. Uncertified dealers often use unsafe methods to extract metals. 3. Refurbishing extends the lifecycle of the product (Reuse), and authorized recycling ensures that toxic materials are handled safely (Recycle). 4. This approach minimizes the environmental footprint of technology.

Final Answer: Refurbishing and using authorized recycling vendors is the most sustainable method.

Answer: (C)

Q47.

Solution

Concept:

The index of a DataFrame is the set of labels used to identify rows. By default, it is a RangeIndex (0, 1, 2...). To use a specific column as the index, the `set_index()` method is used.

Solution:

1. `df.set_index('Date')` returns a new DataFrame with the new index but does not modify the original one. 2. To make the change permanent on the existing object without reassignment, the `inplace=True` parameter must be passed. 3. `reindex()` is used to change the order of the existing index, not to promote a column to index status.

Final Answer: The correct command is df.set$_index('Date', inplace = True)$.

Answer: (B)

Q48.

Solution

Concept:

SQL provides different types for string storage. `CHAR` is a fixed-length character type, while `VARCHAR` is a variable-length character type.

Solution:

1. If you store "ABC" in `CHAR(10)`, SQL will use 10 bytes of storage, padding the remaining 7 spaces with blanks. 2. If you store "ABC" in `VARCHAR(10)`, SQL will only use 3 bytes (plus a small amount of overhead for length), saving storage space. 3. `CHAR` is generally faster for very short, consistent data (like 'M'/'F' or state codes), while `VARCHAR` is more efficient for names and descriptions.

Final Answer: CHAR is fixed length and pads with spaces, while VARCHAR is variable length.

Answer: (C)

Q49.

Solution

Concept:

Matplotlib's `pyplot` module (often imported as `plt`) provides functions for data visualization. For categorical comparison, a bar chart is ideal. Labels are added using specific functions for each axis.

Solution:

1. `plt.bar()` is the function to create vertical bar charts. It takes data for x and y axes. 2. The `color` parameter allows customization of the bar aesthetics. 3. `plt.ylabel()` is used to set the text label for the vertical axis. 4. Option D is incorrect because `barplot` is a Seaborn function, and `yaxis` is not the standard label function.

Final Answer: The correct syntax is plt.bar(regions, sales, color='green'); plt.ylabel('Sales').

Answer: (A)

Q50.

Solution

Concept:

Intellectual Property Rights (IPR) protect the creations of the mind. Common violations include Copyright Infringement, Trademark Infringement, and Plagiarism.

Solution:

1. Phishing and Hacking are cybercrimes focused on stealing data or gaining unauthorized access. 2. Spamming refers to sending unsolicited bulk messages. 3. Plagiarism is the specific act of taking someone else's work or ideas and passing them off as one's own without appropriate acknowledgement. It is a fundamental ethical and legal violation in the realm of IPR.

Final Answer: Plagiarism is a violation of Intellectual Property Rights.

Answer: (B)

Answer Key

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