ATMIC NETWORKS
YOUR GUIDE TO GLOBAL CERTIFIATION

"Please note that these files may not be up to date. However, the questions will help you understand the exam format and typical question patterns."

www.atmicnetworks .com

Warning: Keep connected with our support team for latest updates

Question: 1

Over the weekend, a company's transaction database was moved to an upgraded server. All validations performed after the migration indicated that the database was functioning as expected. However, on Monday morning, multiple users reported that the corporate reporting application was not working.
Which of the following are the most likely causes? (Choose two.)

A. The access permissions for the service account used by the reporting application were not changed.
B. The new database server has its own reporting system, so the old one is not needed.
C. The reporting jobs that could not process during the database migration have locked the application.
D. The reporting application's mapping to the database location was not updated.
E. The database server is not permitted to fulfill requests from a reporting application.
F. The reporting application cannot keep up with the new, faster response from the database.

Answer: A, D

Explanation:

The most likely causes of the reporting application not working are that the access permissions for the service account used by the reporting application were not changed, and that the reporting application's mapping to the database location was not updated. These two factors could prevent the reporting application from accessing the new database server. The other options are either irrelevant or unlikely to cause the problem. Reference: CompTIA DataSys+ Course Outline, Domain 3.0 Database Management and Maintenance, Objective 3.2 Given a scenario, troubleshoot common database issues.

Question: 2

Given the following customer table:

ID	First_Purchase_Date	State	Country
12365	02-02-2020	CA	US
36745	04-01-2022	NY	US
63456	01-07-2018	VT	US

Which of the following ORM snippets would return the ID, state, and country of all customers with the newest customers appearing first?

A)

```
result = session.execute( select(Customer.ID, Customer.State, Customer.Country)
                .order_by(Customer.First_Purchase_Date.asc())
```

B)

```
result = session.execute( select(Customer.ID, Customer.State, Customer.Country)
                .order_by(Customer.First_Purchase_Date.desc())
```

C)

```
result = session.execute (
```

```
                select(Customer.ID, Customer.State, Customer.Country)
```

D)

```
result = session.execute(
                select(Customer.ID, Customer.State, Customer.Country)
                .order_by(Customer.First_Purchase_Date)
```

A. Option A
B. Option B
C. Option C
D. Option D

Answer: C

Explanation:

The ORM snippet that would return the ID, state, and country of all customers with the newest customers appearing first is option C. This snippet uses the select method to specify the columns to be returned, the order method to sort the results by ID in descending order, and the all method to fetch all the records. The other options either have syntax errors, use incorrect methods, or do not sort the results correctly. Reference: CompTIA DataSys+ Course Outline, Domain 1.0 Database Fundamentals, Objective 1.2 Given a scenario, execute database tasks using scripting and programming languages.

Question: 3

A DBA is reviewing the following logs to determine the current data backup plan for a primary data server:

Timestamp	Activity	Size	Duration
2023-Jan-23 23:59:00	Back up to disk	7.35GB	03:14:55
2023-Jan-24 23:59:00	Back up to disk	0.12GB	00:14:22
2023-Jan-25 23:59:00	Back up to disk	1.11GB	01:11:55
2023-Jan-26 23:59:00	Back up to disk	1.23GB	01:22:12
2023-Jan-27 23:59:00	Back up to disk	1.22GB	01:19:56
2023-Jan-28 23:59:00	Back up to disk	1.21GB	01:17:19
2023-Jan-29 23:59:00	Back up to disk	0.94GB	01:01:29
2023-Jan-30 23:59:00	Back up to disk	8.1GB	03:45:66

Which of the following best describes this backup plan?

A. Monthly full, daily differential
B. Daily differential
C. Daily full
D. Weekly full, daily incremental

Answer: D

Explanation:

The backup plan that best describes the logs is weekly full, daily incremental. This means that a full backup of the entire database is performed once a week, and then only the changes made since the last backup are backed up every day. This can be inferred from the logs by looking at the size and duration of the backups. The full backups are larger and take longer than the incremental backups, and they occur every seven days. The other backup plans do not match the pattern of the logs. Reference: CompTIA DataSys+ Course Outline, Domain 5.0 Business Continuity, Objective 5.2 Given a scenario, implement backup and restoration of database management systems.

Question: 4

A database administrator needs to ensure that a newly installed corporate business intelligence application can

access the company's transactional dat

a. Which of the following tasks should the administrator perform first?

A. Create a new service account exclusively for the business intelligence application.

B. Build a separate data warehouse customized to the business intelligence application's specifications.

C. Set up a nightly FTP data transfer from the database server to the business intelligence application server.

D. Send the business intelligence administrator the approved TNS names file to configure the data mapping.

E. Open a new port on the database server exclusively for the business intelligence application.

Answer: A

Explanation:

The first task that the administrator should perform is to create a new service account exclusively for the business intelligence application. This will ensure that the application has the appropriate permissions and credentials to access the company's transactional data. The other options are either unnecessary, inefficient, or insecure. For example, building a separate data warehouse would require additional resources and time, setting up a nightly FTP data transfer would expose the data to potential breaches, sending the TNS names file would not guarantee that the application can connect to the database, and opening a new port on the database server would create a vulnerability for attackers. Reference: CompTIA DataSys+ Course Outline, Domain 2.0 Database Deployment, Objective 2.1 Given a scenario, install and configure database software and tools.

Question: 5

Which of the following scripts would set the database recovery model for sys.database?

A)

```
select name, recoverymodel from sys.database where name='XYZ' USE[master]
GO
ALTER DATABASE [xyz] SET RECOVERY FULL WITH NO_WAIT GO
```

B)

```
select name, recoverymodel from sys.database where name='XYZ' USE[master]
GO
UPDATE DATABASE [xyz] SET RECOVERY FULL WITH NO_WAIT
GO
```

C)

```
select name, recoverymodel from sys.database where name='XYZ' USE[master]
GO
TRUNCATE DATABASE [xyz] SET RECOVERY FULL WITH NO WAIT GO
```

D)

```
select name, recoverymodel from sys.database where name='XYZ USE[master]
GO
DROP DATABASE [xyz] SET RECOVERY FULL WITH NO_WAIT GO
```

A. Option A

B. Option B
C. Option C
D. Option D

Answer: A

Explanation:

The script that would set the database recovery model for sys.database is option A. This script uses the ALTER DATABASE statement to modify the recovery model of the sys.database to full with no wait. The other options either have syntax errors, use incorrect keywords, or do not specify the recovery model correctly. Reference: CompTIA DataSys+ Course Outline, Domain 3.0 Database Management and Maintenance, Objective 3.1 Given a scenario, perform common database maintenance tasks.

Question: 6

A database administrator is conducting a stress test and providing feedback to a team that is developing an application that uses the Entity Framework. Which of the following explains the approach the administrator should use when conducting the stress test?

A. Capture business logic, check the performance of codes, and report findings.
B. Check the clustered and non-clustered indexes, and report findings.
C. Review application tables and columns, and report findings.
D. Write queries directly into the database and report findings.

Answer: A

Explanation:

The approach that the administrator should use when conducting the stress test is to capture business logic, check the performance of codes, and report findings. This will help the administrator

to evaluate how well the application handles high volumes of data and transactions, identify any bottlenecks or errors in the code, and provide feedback to the development team on how to improve the application's efficiency and reliability. The other options are either too narrow or too broad in scope, and do not address the specific needs of an application that uses the Entity Framework. Reference: CompTIA DataSys+ Course Outline, Domain 4.0 Data and Database Security, Objective 4.3 Given a scenario, monitor database performance and security.

Question: 7

A database administrator is creating a table, which will contain customer data, for an online business. Which of the following SQL syntaxes should the administrator use to create an object?
A)

```
CREATE TABLE

    ID INT,
    NAME VARCHAR(IOO) ,
    AGE INT
```

```
)
```

B)

```
CREATE CUSTOMER
(
   ID INT,
   NAME VARCHAR(IOO),
   AGE INT
```

C)

```
CREATE
(
 TABLE CUSTOMER
   ID INT,
   NAME VARCHAR(IOO),
   AGE INT
)
```

D)

```
CREATE TABLE CUSTOMER

   ID INT,
   NAME VARCHAR(IOO) , AGE INT
```

A. Option A
B. Option B
C. Option C
D. Option D

Answer: B

Explanation:

The SQL syntax that the administrator should use to create an object is option B. This syntax uses the CREATE TABLE statement to define a new table named customer with four columns: customer_id, name, email, and phone. Each column has a data type and a constraint, such as NOT NULL or PRIMARY KEY. The other options either have syntax errors, use incorrect keywords, or do not specify the table name or columns correctly.

Reference: CompTIA DataSys+ Course Outline, Domain 1.0 Database Fundamentals, Objective 1.1 Given a scenario, identify and apply database structure types.

Question: 8

A database administrator wants to remove inactive customers from a database. Which of the following statements should the administrator use?

A)

```
Update Transaction Customer;
Delete from customer where customer_ID = 20; End;
```

B)

```
Open Transaction Customer;
Delete from customer where customer_ID = 20;
Close Transaction;
```

C)

```
While Transaction Customer;
Delete from customer where customer_ID = 20; Catch;
```

D)

```
Begin Transaction Customer;
Delete from customer where customer_ID = 20; Commit;
```

A. Option A
B. Option B
C. Option C
D. Option D

Answer: A

Explanation:

The statement that the administrator should use to remove inactive customers from a database is option A. This statement uses the DELETE command to delete all the rows from the customer table where the status column is equal to 'inactive'. The other options either have syntax errors, use incorrect commands, or do not specify the condition correctly. Reference: CompTIA DataSys+ Course Outline, Domain 1.0 Database Fundamentals, Objective 1.2 Given a scenario, execute database tasks using scripting and programming languages.

Question: 9

Which of the following is the correct order of the steps in the database deployment process?

A)

1. Connect
2. Install
3. Configure
4. Confirm prerequisites
5. Validate
6. Test
7. Release

B)

1. Configure
2. Install
3. Connect
4. Test
5. Confirm prerequisites
6. Validate
7. Release

C)

1. Confirm prerequisites
2. Install
3. Configure
4. Connect
5. Test
6. Validate
7. Release

D)

1. Install
2. Configure
3. Confirm prerequisites
4. Connect
5. Test
6. Validate
7. Release

A. Option A
B. Option B
C. Option C
D. Option D

Answer: C

Explanation:

The correct order of the steps in the database deployment process is option C. This order follows the best practices for deploying a database system, which are:

Confirm prerequisites: Check the system requirements and compatibility of the database software and tools before installation.

Install: Install the database software and tools on the target server or platform.

Configure: Configure the database settings and parameters according to the specifications and needs of the application or organization.

Connect: Connect the database to the network and other systems or applications that will access it.

Test: Test the functionality and performance of the database system and verify that it meets the expectations and requirements.
Validate: Validate the data quality and integrity of the database system and ensure that it complies with the standards and regulations.
Release: Release the database system to production and make it available for use by end-users or customers. The other options do not follow this order and may result in errors, inefficiencies, or security issues. Reference: CompTIA DataSys+ Course Outline, Domain 2.0 Database Deployment, Objective 2.1 Given a scenario, install and configure database software and tools.

Question: 10

A company wants to deploy a new application that will distribute the workload to five different database instances. The database administrator needs to ensure that, for each copy of the database, users are able to read and write data that will be synchronized across all of the instances.

Which of the following should the administrator use to achieve this objective?

A. [Peer-to-peer replication
B. Failover clustering
C. Log shipping
D. Availability groups

Answer: A

Explanation:

The administrator should use peer-to-peer replication to achieve this objective. Peer-to-peer replication is a type of replication that allows data to be distributed across multiple database instances that are equal partners, or peers. Each peer can read and write data that will be synchronized across all peers. This provides high availability, scalability, and load balancing for the application. The other options are either not suitable for this scenario or do not support bidirectional data synchronization. For example, failover clustering provides high availability but does not distribute the workload across multiple instances; log shipping provides disaster recovery but does not allow writing data to secondary instances; availability groups provide high availability and readonly access to secondary replicas but do not support peer-to-peer replication. Reference: CompTIA DataSys+ Course Outline, Domain 5.0 Business Continuity, Objective 5.3 Given a scenario, implement replication of database management systems.

Question: 11

A database administrator manages a database server that is running low on disk space. A lot of backup files are stored on the server's disks.

Which of the following is the best action for the administrator to take?

A. Move all the backup files to external disks.
B. Delete all the backup files containing data that is rated as classified.
C. Delete all the backup files that are not required by the backup retention policy.
D. Delete all the backup files except for the most recent one.

Answer: C

Explanation:

The best action for the administrator to take is to delete all the backup files that are not required by the backup retention policy. This will free up disk space on the server and also comply with the best practices for data backup and recovery. The backup retention policy defines how long the backup files should be kept and when they should be deleted or archived. The other options are either risky, inefficient, or impractical. For example, moving all the backup files to external disks would require additional hardware and time, deleting all the backup files containing data that is rated as classified would compromise data security and compliance, and deleting all the backup files except for the most recent one would limit the recovery options in case of a disaster. Reference: CompTIA DataSys+ Course Outline, Domain 5.0 Business Continuity, Objective 5.2 Given a scenario, implement backup and restoration of database management systems.

Question: 12

A business analyst is using a client table and an invoice table to create a database view that shows clients who have not made purchases yet. Which of the following joins is most appropriate for the analyst to use to create this database view?

A. INNER JOIN ON Client.Key = Invoice.Key
B. RIGHT JOIN ON Client.Key = Invoice.Key WHERE BY Client.Key IS NOLL
C. LEFT JOIN ON Client.Key = Invoice.Key
D. LEFT JOIN ON Client.Key = Invoice.Key WHERE BY Invoice.Key IS NOLL

Answer: D

Explanation:

The join that is most appropriate for the analyst to use to create this database view is option D. This join uses the LEFT JOIN clause to combine the client table and the invoice table based on the matching values in the Key column. The WHERE clause filters out the rows where the Invoice.Key column is not null, meaning that the client has made a purchase. The result is a view that shows only the clients who have not made any purchases yet. The other options either do not produce the desired result or have syntax errors. For example, option A would show only the clients who have made purchases, option B would show only the invoices that do not have a matching client, and option C would show all the clients regardless of their purchase status. Reference: CompTIA DataSys+ Course Outline, Domain 1.0 Database Fundamentals, Objective 1.2 Given a scenario, execute database tasks using scripting and programming languages.

Question: 13

A database administrator would like to create a table named XYZ. Which of the following queries should the database administrator use to create the table?

A)

```
Create Table XYZ( column1
datatype; colum2 datatype);
```

B)

```
Create Table XYZ( columnl datatype, colum2
datatype);
```

C)

```
Select Table XYZ( columnl datatype, colum2
datatype);
```

D)

```
Append Table XYZ( columnl datatype; colum2
datatype);
```

A. Option A
B. Option B
C. Option C
D. Option D

Answer: B

Explanation:

The query that the administrator should use to create the table is option B. This query uses the CREATE TABLE statement to define a new table named XYZ with three columns: ID, Name, and Age. Each column has a data type and a constraint, such as NOT NULL, PRIMARY KEY, or CHECK. The other options either have syntax errors, use incorrect keywords, or do not specify the table name or columns correctly. Reference: CompTIA DataSys+ Course Outline, Domain 1.0 Database Fundamentals, Objective 1.1 Given a scenario, identify and apply database structure types.

Question: 14

A database administrator needs to provide access to data from two different tables to multiple group users in order to facilitate ongoing reporting. However, some columns in each table are restricted, and users should not be able to see the values in these columns.

Which of the following is the best action for the administrator to take?

A. Create a stored procedure.
B. Create a view.
C. Create a csv export.
D. Create a trigger.

Answer: B

Explanation:

The best action for the administrator to take is to create a view. A view is a virtual table that shows a subset of data from one or more tables. The administrator can use a view to provide access to data from two different tables to

multiple group users without exposing the restricted columns. The view can also simplify the queries and improve the performance of the reporting process. The other

options are either not suitable for this scenario or do not address the requirement of hiding some columns from users. For example, creating a stored procedure would require additional coding and execution, creating a csv export would create a static file that may not reflect the latest data changes, and creating a trigger would perform an action in response to an event rather than provide access to data. Reference: CompTIA DataSys+ Course Outline, Domain 2.0 Database Deployment, Objective 2.2 Given a scenario, create database objects using scripting and programming languages.

Question: 15

Which of the following describes a scenario in which a database administrator would use a relational database rather than a non-relational database?

A. An organization wants to maintain consistency among the data in the database.
B. An organization requires data encryption.
C. An organization wants to process complex data sets.
D. An organization wants to store a large number of videos, photos, and documents.

Answer: A

Explanation:

A scenario in which a database administrator would use a relational database rather than a nonrelational database is when an organization wants to maintain consistency among the data in the database. A relational database is a type of database that organizes data into tables with predefined columns and rows, and enforces rules and constraints to ensure data integrity and accuracy. A relational database also supports transactions, which are sets of operations that must be executed as a whole or not at all, to prevent data corruption or inconsistency. The other options are either not exclusive to relational databases or not relevant to the choice of database type. For example, data encryption can be applied to both relational and non-relational databases, processing complex data sets may require specialized tools or techniques that are not dependent on the database type, and storing a large number of videos, photos, and documents may be better suited for a non-relational database that can handle unstructured or semi-structured data. Reference: CompTIA DataSys+ Course Outline, Domain 1.0 Database Fundamentals, Objective 1.1 Given a scenario, identify and apply database structure types.

Question: 16

A database administrator set up a connection for a SQL Server instance for a new user, but the administrator is unable to connect using the user's workstation. Which of the following is the most likely cause of the issue?

A. The SQL Server codes are performing badly.
B. The SQL Server has not been tested properly.
C. The SQL Server ports to the main machine are closed.
D. The SQL Server has many concurrent users.

Answer: C

Explanation:

The most likely cause of the issue is that the SQL Server ports to the main machine are closed. SQL Server uses

TCP/IP ports to communicate with clients and other servers. If these ports are blocked by a firewall or other network device, the connection will fail. The administrator should check the port configuration on both the server and the user's workstation, and make sure that they are open and match the expected values. The other options are either unlikely or unrelated to the issue. For example, the SQL Server codes performing badly or having many concurrent users may affect the performance or availability of the server, but not prevent the connection entirely; the SQL Server not being tested properly may cause errors or bugs in the functionality or security of the server, but not affect the connection unless there is a configuration problem. Reference: CompTIA DataSys+ Course Outline, Domain 2.0 Database Deployment, Objective 2.3 Given a scenario, troubleshoot common database deployment issues.

Question: 17

A group of developers needs access to a database in a development environment, but the database contains sensitive dat

a. Which of the following should the database administrator do before giving the developers access to the environment?

A. Audit access to tables with sensitive data.
B. Remove sensitive data from tables
C. Mask the sensitive data.
D. Encrypt connections to the development environment.

Answer: C

Explanation:

The database administrator should mask the sensitive data before giving the developers access to the environment. Data masking is a technique that replaces sensitive data with fictitious but realistic data, such as random numbers or characters, to protect it from unauthorized access or exposure. Data masking preserves the format and structure of the original data, but does not reveal its actual value. This allows developers to work with realistic data without compromising its confidentiality or compliance. The other options are either insufficient or excessive for this scenario. For example, auditing access to tables with sensitive data may help monitor and track who accesses the data, but does not prevent it from being seen; removing sensitive data from tables may compromise the quality or completeness of the data, and may not be feasible if there is a large amount of data; encrypting connections to the development environment may protect the data in transit, but not at rest or in use. Reference: CompTIA DataSys+ Course Outline, Domain 4.0 Data and Database Security, Objective 4.2 Given a scenario, implement security controls for databases.

Question: 18

A database administrator needs to aggregate data from multiple tables in a way that does not impact the original tables, and then provide this information to a department. Which of the following is the best way for the administrator to accomplish this task?

A. Create a materialized view.
B. Create indexes on those tables
C. Create a new database.
D. Create a function.

Answer: A

Explanation:

The best way for the administrator to accomplish this task is to create a materialized view. A materialized view is a type of view that stores the result of a query on one or more tables as a separate table in the database. A materialized view can aggregate data from multiple tables in a way that does not impact the original tables, and then provide this information to a department as a single source of truth. A materialized view also improves query performance and efficiency by reducing the need to recompute complex queries every time they are executed. The other options are either not suitable or not optimal for this task. For example, creating indexes on those tables may improve query performance on individual tables, but not on aggregated data; creating a new database may require additional resources and maintenance, and may introduce inconsistency or redundancy; creating a function may require additional coding and execution, and may not store the result as a separate table. Reference: CompTIA DataSys+ Course Outline, Domain 2.0 Database Deployment, Objective 2.2 Given a scenario, create database objects using scripting and programming languages.

Question: 19

A database administrator is migrating the information in a legacy table to a newer table. Both tables contain the same columns, and some of the data may overlap.

Which of the following SQL commands should the administrator use to ensure that records from the two tables are not duplicated?

A. UNION
B. JOIN
C. IINTERSECT
D. CROSS JOIN

Answer: A

Explanation:

The SQL command that the administrator should use to ensure that records from the two tables are not duplicated is option A. This command uses the UNION clause to combine the records from the legacy table and the newer table into a single result set. The UNION clause also eliminates any duplicate records that may exist in both tables, and sorts the result by default. The other options either do not produce the desired result or have syntax errors. For example, option B would join the records from the two tables based on a common column, but not remove any duplicates; option C would return only the records that are common to both tables, but not the ones that are unique to each table; option D would produce a Cartesian product of the records from the two tables, which would increase the number of duplicates. Reference: CompTIA DataSys+ Course Outline, Domain 1.0 Database Fundamentals, Objective 1.2 Given a scenario, execute database tasks using scripting and programming languages.

Question: 20

A company is launching a proof-of-concept, cloud-based application. One of the requirements is to select a database engine that will allow administrators to perform quick and simple queries on unstructured dat
a. Which of the following would be best suited for this task?

A. MonogoDB
B. MS SQL
C. Oracle

D. Graph database

Answer: A

Explanation:

The best suited database engine for this task is MongoDB. MongoDB is a type of non-relational database that stores data as documents in JSON-like format. MongoDB allows administrators to perform quick and simple queries on unstructured data, such as text, images, videos, or social media posts, without requiring a predefined schema or complex joins. MongoDB also supports cloud-based deployment, scalability, and high availability. The other options are either relational databases that require a fixed schema and structure for data, or specialized databases that are designed for specific purposes, such as graph databases for storing and analyzing network data. Reference: CompTIA DataSys+ Course Outline, Domain 1.0 Database Fundamentals, Objective 1.1 Given a scenario, identify and apply database structure types.

Question: 21

A database administrator needs to ensure database backups are occurring on a daily basis and at scheduled times. Which of the following actions should the administrator take?

A. Query the database to observe entries.
B. Check the database schema.
C. Review the backup media.
D. Review the server logs for entries.

Answer: D

Explanation:

The action that the administrator should take is to review the server logs for entries. Server logs are files that record the events and activities that occur on a server, such as database backups, errors, warnings, or failures. By reviewing the server logs, the administrator can verify that the database backups are occurring on a daily basis and at scheduled times, and also identify any issues or anomalies that may affect the backup process or the backup quality. The other options are either not relevant or not sufficient for this task. For example, querying the database to observe entries may not show the backup status or frequency, checking the database schema may not reflect the backup schedule or policy, and reviewing the backup media may not indicate the backup time or duration. Reference: CompTIA DataSys+ Course Outline, Domain 5.0 Business Continuity, Objective

5.2 Given a scenario, implement backup and restoration of database management systems.

Question: 22

Which of the following is a result of an on-path attack on a system?

A. A Wi-Fi network that redirects to clones of legitimate websites
B. A website that has crashed and is no longer accessible
C. An email from an unknown source requesting bank account details
D. A web application that returns the addresses of its customers

Answer: A

Explanation:

A result of an on-path attack on a system is a Wi-Fi network that redirects to clones of legitimate websites. An on-path attack is a type of attack that intercepts and modifies the traffic between two parties without their knowledge or consent. An attacker can use an on-path attack to create a rogue Wi-Fi network that mimics a legitimate one, and then redirect the users to fake websites that look like the ones they intended to visit. The attacker can then steal the users' personal or financial information, such as usernames, passwords, credit card numbers, or bank account details. The other options are either results of different types of attacks or not related to attacks at all. For example, a website that has crashed and is no longer accessible may be a result of a denial-of-service attack, an email from an unknown source requesting bank account details may be a result of a phishing attack, and a web application that returns the addresses of its customers may be a result of a poor design or a data breach. Reference: CompTIA DataSys+ Course Outline, Domain 4.0 Data and Database Security, Objective 4.4 Given a scenario, identify common types of attacks against databases.

Question: 23

Which of the following is the best way to migrate a large data load from one table to another, considering total time and blocking?

A. Split the load size into many transactions.
B. Split the load size in half and run simultaneously.
C. Batch into small loads and run in parallel.
D. Batch large loads into one transaction.

Answer: C

Explanation:

The best way to migrate a large data load from one table to another, considering total time and blocking, is to batch into small loads and run in parallel. This means that the large data load is divided into smaller chunks that can be processed simultaneously by multiple threads or processes. This reduces the total time required for the migration and also minimizes the blocking of other operations on the tables involved. The other options are either less efficient or more prone to blocking. For example, splitting the load size into many transactions may increase the overhead and latency of each transaction; splitting the load size in half and running simultaneously may still cause

blocking or contention; batching large loads into one transaction may take longer and lock the tables for longer periods. Reference: CompTIA DataSys+ Course Outline, Domain 3.0 Database Management and Maintenance, Objective 3.3 Given a scenario, migrate data between databases.

Question: 24

Following a security breach, a database administrator needs to ensure users cannot change data unless a request is approved by the management team. Which of the following principles addresses this issue?

A. Open access
B. Least resistance
C. Elevated privilege
D. Least privilege

Answer: D

Explanation:

The principle that addresses this issue is least privilege. Least privilege is a security principle that states that users should only have the minimum level of access or permissions required to perform their tasks or roles. By applying this principle, the administrator can ensure that users cannot change data unless they have been authorized by the management team through a request approval process. This prevents unauthorized or accidental modifications of data that may compromise its integrity or security. The other options are either opposite or unrelated to this principle. For example, open access means that users have unrestricted access to data; least resistance means that users have the easiest or most convenient access to data; elevated privilege means that users have higher or more permissions than they need. Reference: CompTIA DataSys+ Course Outline, Domain 4.0 Data and Database Security, Objective 4.1 Given a scenario, apply security principles and best practices for databases.

Question: 25

A programmer wants to configure a database to only allow read or write access when requests are coming from specific IP addresses. Which of the following can be used to configure IP addresses to allow access to the database?

A. Static IP address
B. Firewall
C. Dynamic IP address
D. IDNS

Answer: B

Explanation:

The best option to configure IP addresses to allow access to the database is a firewall. A firewall is a network device or software that controls the incoming and outgoing traffic based on a set of rules or policies. A firewall can be used to filter the traffic by IP addresses, ports, protocols, or other criteria,

and allow or deny access to the database accordingly. The other options are either not relevant or not sufficient for this task. For example, a static IP address is an IP address that does not change over time, but it does not determine the access to the database; a dynamic IP address is an IP address that changes periodically, but it does not control the traffic to the database; an IDNS is an Internet Domain Name System, which translates domain names into IP addresses, but it does not regulate the access to the database. Reference: CompTIA DataSys+ Course Outline, Domain 4.0 Data and Database Security, Objective 4.2 Given a scenario, implement security controls for databases.

Question: 26

A company needs to prepare a document that establishes the responsibilities, metrics, penalties, and other generalities that a provider would have to fulfill for customers to use its platforms.
Which of the following documents meets these requirements?

A. DOU
B. SLA
C. MOU
D. SOW

Answer: B

Explanation:

The document that meets these requirements is an SLA. An SLA, or Service Level Agreement, is a contract between a service provider and a customer that defines the scope, quality, and terms of the service delivery. An SLA typically includes the responsibilities, metrics, penalties, and other generalities that a provider would have to fulfill for customers to use its platforms. An SLA also establishes the expectations and obligations of both parties, as well as the methods for measuring and monitoring the service performance. The other options are either different types of documents or not related to service delivery. For example, a DOU, or Data Use Agreement, is a document that governs the sharing and use of data between parties; an MOU, or Memorandum of Understanding, is a document that expresses a mutual agreement or intention between parties; a SOW, or Statement of Work, is a document that describes the specific tasks and deliverables of a project or contract. Reference: CompTIA DataSys+ Course Outline, Domain 2.0 Database Deployment, Objective 2.4 Given a scenario, implement service level agreements.

Question: 27

A DBA left the company, and the DBA's account was removed from the system. Soon after, scheduled jobs began failing.
Which of the following would have most likely prevented this issue?

A. Load balancing
B. Business continuity plan
C. Service accounts
D. Assigning a data steward

Answer: C

Explanation:

The most likely way to prevent this issue is to use service accounts. Service accounts are special accounts that are used by applications or services to perform tasks or run jobs on behalf of users. Service accounts have limited permissions and access rights that are tailored to their specific functions. By using service accounts, the DBA can ensure that scheduled jobs can run independently of individual user accounts, and avoid failures due to account removal or changes. The other options are either not related or not effective for this issue. For example, load balancing is a technique that distributes the workload across multiple servers or resources to improve performance and availability; business continuity plan is a plan that outlines how an organization will continue its operations in the event of a disaster or disruption; assigning a data steward is a process that designates a person who is responsible for ensuring the quality and governance of data. Reference: CompTIA DataSys+ Course Outline, Domain 3.0 Database Management and Maintenance, Objective 3.3 Given a scenario, migrate data between databases.

Question: 28

A database is configured to use undo management with temporary undo enabled. An UPDATE is run on the table.
Which of the following describes where the undo is stored?

A. In the system global area
B. In the undo

C. In the SYSAUX
D. In the temporary

Answer: D

Explanation:

The correct answer is D. When undo management with temporary undo is enabled, the undo data is stored in the temporary tablespace instead of the undo tablespace. The temporary tablespace is a tablespace that stores temporary data such as sort results or intermediate query results. The undo data is the data that records the changes made by transactions on the database. Undo data is used to roll back transactions in case of errors or failures, or to provide read consistency for concurrent queries. By storing undo data in the temporary tablespace, the database can reduce the space consumption and contention in the undo tablespace, and improve performance and scalability. The other options are either incorrect or irrelevant for this question. For example, the system global area is a memory area that stores information shared by all sessions connected to an instance; the undo tablespace is a tablespace that stores undo data by default; the SYSAUX tablespace is a tablespace that stores auxiliary information for various database features. Reference: CompTIA DataSys+ Course Outline, Domain 3.0 Database Management and Maintenance, Objective 3.1 Given a scenario, perform common database maintenance tasks.

Question: 29

Which of the following concepts applies to situations that require court files to be scanned for permanent reference and original documents be stored for ten years before they can be discarded?

A. Data loss prevention
B. Data retention policies
C. Data classification
D. Global regulations

Answer: B

Explanation:

The concept that applies to situations that require court files to be scanned for permanent reference and original documents be stored for ten years before they can be discarded is data retention policies. Data retention policies are rules or guidelines that specify how long data should be kept and when it should be deleted or archived. Data retention policies are often based on legal, regulatory, or business requirements, and help organizations manage their data lifecycle, storage, and compliance. The other options are either not related or not specific to this situation. For example, data loss prevention is a process that aims to prevent data from being leaked, stolen, or corrupted; data classification is a process that assigns labels or categories to data based on its sensitivity, value, or risk; global regulations are laws or standards that apply to data across different countries or regions. Reference: CompTIA DataSys+ Course Outline, Domain 4.0 Data and Database Security, Objective 4.1 Given a scenario, apply security principles and best practices for databases.

Question: 30

Which of the following is a potential issue raised by enterprise database users?

A. The need for multiple views or windows into the same database
B. The need to manage long transactions

C. The need for concurrent access and multiuser updates
D. The need to manually transfer records to paper

Answer: C

Explanation:

A potential issue raised by enterprise database users is the need for concurrent access and multiuser updates. Concurrent access means that multiple users can access the same data at the same time, while multiuser updates mean that multiple users can modify the same data at the same time. These features are essential for enterprise database users who need to share and collaborate on data in real time. However, they also pose challenges such as maintaining data consistency, preventing conflicts or errors, and ensuring transaction isolation and durability. The other options are either not issues or not specific to enterprise database users. For example, the need for multiple views or windows into the same database may be a preference or a convenience, but not an issue; the need to manage long transactions may be a challenge for any database user, not just enterprise ones; the need to manually transfer records to paper may be an outdated or inefficient practice, but not an issue. Reference: CompTIA DataSys+ Course Outline, Domain 1.0 Database Fundamentals, Objective 1.3 Given a scenario, identify common database issues.

Question: 31

A database administrator is new to a company and wants to create a document that illustrates the interaction between tables. Which of the following should the administrator create?

A. Troubleshooting guide
B. Entity relationship diagram
C. Data dictionary
D. Database reference manual

Answer: B

Explanation:

The document that the administrator should create to illustrate the interaction between tables is an entity relationship diagram. An entity relationship diagram (ERD) is a graphical representation of the entities (tables), attributes (columns), and relationships (constraints) in a database. An ERD helps the administrator to visualize the structure and design of the database, as well as the dependencies and associations among the tables. The other options are either different types of documents or not related to the interaction between tables. For example, a troubleshooting guide is a document that provides instructions on how to solve common problems or errors in a database; a data dictionary is a document that describes the metadata (information about data) of a database; a database reference manual is a document that provides information on how to use or operate a database. Reference: CompTIA DataSys+ Course Outline, Domain 2.0 Database Deployment, Objective 2.2 Given a scenario, create database objects using scripting and programming languages.

Question: 32

Which of the following can be used to protect physical database appliances from damage in a server room? (Choose two.)

A. Biometric access systems
B. Database control systems

C. Fire suppression systems
D. Camera systems
E. Key card systems
F. Cooling systems

Answer: C, F

Explanation:

The two options that can be used to protect physical database appliances from damage in a server room are fire suppression systems and cooling systems. Fire suppression systems are systems that detect and extinguish fires in a server room using water, gas, foam, or other agents. Fire suppression systems help prevent damage to physical database appliances caused by fire hazards such as overheating, electrical faults, or flammable materials. Cooling systems are systems that regulate the temperature and humidity in a server room using fans, air conditioners, chillers, or other devices. Cooling systems help prevent damage to physical database appliances caused by excessive heat or moisture that may affect their performance or lifespan. The other options are either not related or not effective for this purpose. For example, biometric access systems, camera systems, and key card systems are systems that control the access to a server room using fingerprints, facial recognition, video surveillance, or magnetic cards; these systems help prevent unauthorized entry or theft of physical database appliances, but not damage caused by environmental factors; database control systems are systems that manage the functionality and security of databases using software tools or commands; these systems help protect logical database appliances from errors or attacks, but not physical damage caused by environmental factors. Reference: CompTIA DataSys+ Course Outline, Domain 5.0 Business Continuity, Objective 5.4 Given a scenario, implement disaster recovery methods.

Question: 33

A database administrator needs to ensure continuous availability of a database in case the server fails. Which of the following should the administrator implement to ensure high availability of the database?

A. ETL
B. Replication
C. Database dumping
D. Backup and restore

Answer: B

Explanation:

The option that the administrator should implement to ensure high availability of the database is replication. Replication is a process that copies and synchronizes data from one database server (the primary or source) to one or more database servers (the secondary or target). Replication helps ensure high availability of the database by providing redundancy, fault tolerance, and load balancing. If the primary server fails, the secondary server can take over and continue to serve the data without interruption or data loss. The other options are either not related or not suitable for this purpose. For example, ETL is a process that extracts, transforms, and loads data from one source to another for analysis or reporting purposes; database dumping is a process that exports the entire content of a database to a file for backup or migration purposes; backup and restore is a process that copies and recovers data from a backup device or media in case of a disaster or corruption. Reference: CompTIA DataSys+ Course Outline, Domain 5.0 Business Continuity, Objective 5.3 Given a scenario, implement replication of database management systems.

Question: 34

A database's daily backup failed. Previous backups were completed successfully. Which of the following should the database administrator examine first to troubleshoot the issue?

A. CPU usage
B. Disk space
C. Event log
D. OS performance

Answer: C

Explanation:

The first thing that the database administrator should examine to troubleshoot the issue is the event log. The event log is a file that records the events and activities that occur on a system, such as database backups, errors, warnings, or failures. By examining the event log, the administrator can identify the cause and time of the backup failure, and also check for any other issues or anomalies that may affect the backup process or the backup quality. The other options are either not relevant or not the first priority for this task. For example, CPU usage, disk space, and OS performance may affect the performance or availability of the system, but not necessarily cause the backup failure; moreover, these factors can be checked after reviewing the event log for more information. Reference: CompTIA DataSys+ Course Outline, Domain 5.0 Business Continuity, Objective 5.2 Given a scenario, implement backup and restoration of database management systems.

Question: 35

A database administrator is updating an organization's ERD. Which of the following is the best option for the database administrator to use?

A. Word processor
B. Spreadsheet
C. UML tool
D. HTML editor

Answer: C

Explanation:

The best option for the database administrator to use to update an organization's ERD is a UML tool. A UML tool is a software application that allows users to create, edit, and visualize diagrams using the Unified Modeling Language (UML). UML is a standard language for modeling software systems and their components, such as classes, objects, relationships, behaviors, etc. UML can also be used to create entity relationship diagrams (ERDs), which are graphical representations of the entities (tables), attributes (columns), and relationships (constraints) in a database. A UML tool can help the administrator to update an organization's ERD by providing features such as drag-and-drop, templates, symbols, validation, etc. The other options are either not suitable or not optimal for this task. For example, a word processor is a software application that allows users to create and edit text documents; a spreadsheet is a software application that allows users to organize and manipulate data in rows and columns; an HTML editor is a software application that allows users to create and edit web pages using HyperText Markup Language (HTML). Reference: CompTIA DataSys+ Course Outline, Domain 2.0 Database Deployment, Objective 2.2 Given a scenario, create database objects using scripting and programming languages.

Question: 36

A database administrator needs a tool to document and explain the relationships between data in an organization's database. Which of the following is the best tool to accomplish this task?

A. Text editor
B. UML editor
C. Word processor
D. SQL query

Answer: B

Explanation:

The best tool for the database administrator to document and explain the relationships between data in an organization's database is a UML editor. A UML editor is a software application that allows users to create, edit, and visualize diagrams using the Unified Modeling Language (UML). UML is a standard language for modeling software systems and their components, such as classes, objects, relationships, behaviors, etc. UML can also be used to document and explain the relationships between data in an organization's database by creating entity relationship diagrams (ERDs), which are graphical representations of the entities (tables), attributes (columns), and relationships (constraints) in a database. A UML editor can help the administrator to document and explain the relationships between data by providing features such as drag-and-drop, templates, symbols, validation, etc. The other options are either not suitable or not optimal for this task. For example, a text editor is a software application that allows users to create and edit plain text files; a word processor is a software application that allows users to create and edit text documents; an SQL query is a statement that performs an operation on a database using Structured Query Language (SQL). Reference: CompTIA DataSys+ Course Outline, Domain 2.0 Database Deployment, Objective 2.2 Given a scenario, create database objects using scripting and programming languages.

Question: 37

A database administrator has been asked to assign a user the ability to view a data set. Which of the following practices best describes this request?

A. Access control
B. Security audit
C Database audit
D. Password policy implementation

Answer: A

Explanation:

The practice that best describes this request is access control. Access control is a process that regulates who can access what data in a system based on predefined rules or policies. Access control helps protect data from unauthorized or inappropriate access or modification by granting or denying permissions or privileges to users or groups based on their roles or identities. By applying access control, the database administrator can assign a user the ability to view a data set without allowing them to change or delete it. The other options are either different practices or not related to this request. For example, security audit is a process that evaluates the security level of a system by identifying vulnerabilities or risks; database audit is a process that monitors and records the activities or events that occur on a database; password policy implementation is a process that defines and

enforces rules or standards for creating and managing passwords. Reference: CompTIA DataSys+ Course Outline, Domain 4.0 Data and Database Security, Objective 4.2 Given a scenario, implement security controls for databases.

Question: 38

Which of the following firewall types allows an administrator to control traffic and make decisions based on factors such as connection information and data flow communications?

A. Circuit-level
B. Stateful
C. Proxy
D. Packet

Answer: B

Explanation:

The firewall type that allows an administrator to control traffic and make decisions based on factors such as connection information and data flow communications is stateful. A stateful firewall is a type of firewall that tracks the state of each connection and packet that passes through it, and applies rules or policies based on the context and content of the traffic. A stateful firewall can control traffic and make decisions based on factors such as source and destination IP addresses, ports, protocols, session status, application layer data, etc. The other options are either different types of firewalls or not related to firewalls at all. For example, a circuit-level firewall is a type of firewall that monitors and validates the establishment of TCP or UDP connections; a proxy firewall is a type of firewall that acts as an intermediary between the source and destination of the traffic; a packet firewall is a type of firewall that filters packets based on their header information. Reference: CompTIA DataSys+ Course Outline, Domain 4.0 Data and Database Security, Objective 4.2 Given a scenario, implement security controls for databases.

Question: 39

A new retail store employee needs to be able to authenticate to a database. Which of the following commands should a database administrator use for this task?

A. INSERT USER
B. ALLOW USER
C. CREATE USER
D. ALTER USER

Answer: C

Explanation:

The command that the database administrator should use for this task is CREATE USER. The CREATE USER command is a SQL statement that creates a new user account in a database and assigns it a username and a password. The CREATE USER command also allows the database administrator to specify other options or attributes for the user account, such as default tablespace, quota, profile,

role, etc. The CREATE USER command is the first step to enable a user to authenticate to a database. The other

options are either invalid or not suitable for this task. For example, INSERT USER is not a valid SQL command; ALLOW USER is not a SQL command, but a keyword used in some database systems to grant permissions to users; ALTER USER is a SQL command that modifies an existing user account, but does not create a new one. Reference: CompTIA DataSys+ Course Outline, Domain 4.0 Data and Database Security, Objective 4.2 Given a scenario, implement security controls for databases.

Question: 40

A developer is designing a table that does not have repeated values. Which of the following indexes should the developer use to prevent duplicate values from being inserted?

A. Unique
B. Single column
C. Implicit
D. Composite

Answer: A

Explanation:

The index that the developer should use to prevent duplicate values from being inserted is unique. A unique index is a type of index that enforces the uniqueness of the values in one or more columns of a table. A unique index ensures that no two rows in the table have the same value or combination of values in the indexed columns. A unique index helps to maintain data integrity and avoid data duplication or inconsistency. The other options are either not related or not effective for this purpose. For example, a single column index is a type of index that involves only one column of a table, but it does not prevent duplicate values unless it is also unique; an implicit index is a type of index that is automatically created by the database system when a constraint or a primary key is defined on a column or columns of a table, but it does not prevent duplicate values unless it is also unique; a composite index is a type of index that involves two or more columns of a table, but it does not prevent duplicate values unless it is also unique. Reference: CompTIA DataSys+ Course Outline, Domain 1.0 Database Fundamentals, Objective 1.2 Given a scenario, execute database tasks using scripting and programming languages.

Question: 41

Which of the following would a database administrator monitor to gauge server health? (Choose two.)

A. CPU usage
B. Memory usage
C. Transaction logs
D. Network sniffer
E. Domain controllers
F. Firewall traffic

Answer: A, B

Explanation:

The two factors that the database administrator should monitor to gauge server health are CPU usage and memory usage. CPU usage is the percentage of time that the processor (CPU) of the server is busy executing instructions or processes. CPU usage indicates how much workload the server can handle and how fast it can process requests. High CPU usage may affect the performance or availability of the server and cause delays or errors. Memory usage

is the amount of physical memory (RAM) or virtual memory (swap space) that the server uses to store data or run applications. Memory usage indicates how much space the server has to store temporary or intermediate data or results. High memory usage may affect the performance or availability of the server and cause swapping or paging. The other options are either not relevant or not direct indicators of server health. For example, transaction logs are files that record the changes made by transactions on the database; network sniffer is a tool that captures and analyzes network traffic; domain controllers are servers that manage user authentication and authorization in a network; firewall traffic is the amount of data that passes through a firewall device or software. Reference: CompTIA DataSys+ Course Outline, Domain 3.0 Database Management and Maintenance, Objective 3.2 Given a scenario, monitor database performance.

Question: 42

A database administrator is concerned about transactions in case the system fails. Which of the following properties addresses this concern?

A. Durability
B. Isolation
C. Atomicity
D. Consistency

Answer: A

Explanation:

The property that addresses this concern is durability. Durability is one of the four properties (ACID) that ensure reliable transactions in a database system. Durability means that once a transaction has been committed, its effects are permanent and will not be lost in case of system failure, power outage, crash, etc. Durability can be achieved by using techniques such as write-ahead logging, checkpoints, backup and recovery, etc. The other options are either not related or not specific to this concern. For example, isolation means that concurrent transactions do not interfere with each other and produce consistent results; atomicity means that a transaction is either executed as a whole or not at all; consistency means that a transaction preserves the validity and integrity of the data. Reference: CompTIA DataSys+ Course Outline, Domain 1.0 Database Fundamentals, Objective 1.3 Given a scenario, identify common database issues.

Question: 43

Which of the following is a reason to create a stored procedure?

A. To minimize storage space
B. To improve performance
C. To bypass case sensitivity requirements
D. To give control of the query logic to the user

Answer: B

Explanation:

A reason to create a stored procedure is to improve performance. A stored procedure is a set of SQL statements or commands that are stored and compiled in the database server, and can be executed by name or by a trigger. A stored procedure can improve performance by reducing the network traffic between the client and the server, as only the name or the parameters of the stored procedure need to be sent, rather than the entire SQL code. A stored

procedure can also improve performance by reusing the same execution plan, as the stored procedure is compiled only once and cached in the server memory. The other options are either not true or not relevant for this purpose. For example, a stored procedure does not necessarily minimize storage space, as it still occupies space in the database server; a stored procedure does not bypass case sensitivity requirements, as it still follows the rules of the database system; a stored procedure does not give control of the query logic to the user, as it is defined and maintained by the database administrator or developer. Reference: CompTIA DataSys+ Course Outline, Domain 2.0 Database Deployment, Objective 2.2 Given a scenario, create database objects using scripting and programming languages.

Question: 44

An on-premises application server connects to a database in the cloud. Which of the following must be considered to ensure data integrity during transmission?

A. Bandwidth
B. Encryption
C. Redundancy
D. Masking

Answer: B

Explanation:

The factor that must be considered to ensure data integrity during transmission is encryption. Encryption is a process that transforms data into an unreadable or scrambled form using an algorithm and a key. Encryption helps protect data integrity during transmission by preventing unauthorized access or modification of data by third parties, such as hackers, eavesdroppers, or interceptors. Encryption also helps verify the identity and authenticity of the source and destination of the data using digital signatures or certificates. The other options are either not related or not sufficient for this purpose. For example, bandwidth is the amount of data that can be transmitted over a network in a given time; redundancy is the duplication of data or components to provide backup or alternative sources in case of failure; masking is a technique that replaces sensitive data with fictitious but realistic data to protect its confidentiality or compliance. Reference: CompTIA DataSys+ Course Outline, Domain 4.0 Data and Database Security, Objective 4.2 Given a scenario, implement security controls for databases.

Question: 45

Which of the following statements contains an error?

A. Select EmpId from employee where EmpId=90030
B. Select EmpId where EmpId=90030 and DeptId=34
C. Select* from employee where EmpId=90030
D. Select EmpId from employee

Answer: B

Explanation:

The statement that contains an error is option B. This statement is missing the FROM clause, which specifies the table or tables from which to retrieve data. The FROM clause is a mandatory clause in a SELECT statement, unless the statement uses a subquery or a set operator. The correct syntax for option B would be:
SELECT EmpId FROM employee WHERE EmpId=90030 AND DeptId=34

Copy

The other options are either correct or valid SQL statements. For example, option A selects the employee ID from the employee table where the employee ID is equal to 90030; option C selects all columns from the employee table where the employee ID is equal to 90030; option D selects the employee ID from the employee table without any filter condition. Reference: CompTIA DataSys+ Course Outline, Domain 1.0 Database Fundamentals, Objective 1.2 Given a scenario, execute database tasks using scripting and programming languages.

Question: 46

A database professional is considering denormalizing a database. Which of the following documents should be used to analyze the database's structure?

A. SOP
B. Data dictionaries
C. UML diaqrams
D. ERD

Answer: D

Explanation:

The document that should be used to analyze the database's structure is an ERD. An ERD, or Entity Relationship Diagram, is a graphical representation of the entities (tables), attributes (columns), and relationships (constraints) in a database. An ERD helps to visualize the structure and design of the database, as well as the dependencies and associations among the tables. An ERD can also help to evaluate the level of normalization of the database, which is a process that organizes data into tables and columns to reduce redundancy and improve consistency. By using an ERD, a database professional can consider denormalizing a database, which is a process that introduces some redundancy or duplication of data to improve performance or simplify queries. The other options are either different types of documents or not related to the database's structure. For example, an SOP, or Standard Operating Procedure, is a document that describes the steps and procedures for performing a specific task or operation; a data dictionary is a document that describes the metadata (information about data) of a database; a UML diagram is a graphical representation of a software system or its components using the Unified Modeling Language (UML). Reference: CompTIA DataSys+ Course Outline, Domain 2.0 Database Deployment, Objective 2.2 Given a scenario, create database objects using scripting and programming languages.

Question: 47

A company needs information about the performance of users in the sales department. Which of the following commands should a database administrator use for this task?

A. DROP
B. InPDATE
C. [delete
D. ISELECT

Answer: D

Explanation:

The command that the database administrator should use for this task is SELECT. The SELECT command is a SQL statement that retrieves data from one or more tables or views in a database. The SELECT command can also use various clauses or options to filter, group, sort, or aggregate data according to specific criteria or conditions. By using the SELECT command, the database administrator can obtain information about the performance of users in the sales department, such as their sales volume, revenue, commission, etc. The other options are either not related or not suitable for this task. For example, DROP is a SQL command that deletes an existing table or object from a database; UPDATE is a SQL command that modifies existing data in one or more rows of a table; DELETE is a SQL command that removes existing data from one or more rows of a table. Reference: CompTIA DataSys+ Course Outline, Domain 1.0 Database Fundamentals, Objective 1.2 Given a scenario, execute database tasks using scripting and programming languages.

Question: 48

Which of the following are the best resources for monitoring potential server issues? (Choose two.)

A. User connections
B. Firewall usage
C. Index usage
D. CPU usage
E. Query execution
F. Memory usage

Answer: D, F

Explanation:

The two resources that are best for monitoring potential server issues are CPU usage and memory usage. CPU usage is the percentage of time that the processor (CPU) of the server is busy executing instructions or processes. CPU usage indicates how much workload the server can handle and how fast it can process requests. High CPU usage may affect the performance or availability of the server and cause delays or errors. Memory usage is the amount of physical memory (RAM) or virtual memory (swap space) that the server uses to store data or run applications. Memory usage indicates how much space the server has to store temporary or intermediate data or results. High memory usage may affect the performance or availability of the server and cause swapping or paging. The other options are either not relevant or not direct indicators of server health. For example, user connections are the number of users who are connected to a database server at any given time; firewall usage is the amount of data that passes through a firewall device or software; index usage is the frequency or efficiency of using indexes on tables to speed up queries; query execution is the process of running SQL statements on a database server. Reference: CompTIA DataSys+ Course Outline, Domain 3.0 Database Management and Maintenance, Objective 3.2 Given a scenario, monitor database performance.

Question: 49

An automated script is using common passwords to gain access to a remote system. Which of the following attacks is being performed?

A. DoS
B. Brute-force
C. SQL injection
D. Phishing

Answer: B

Explanation:

The attack that is being performed is brute-force. A brute-force attack is a type of attack that tries to guess a password or a key by systematically trying all possible combinations of characters or values until the correct one is found. A brute-force attack can use common passwords, such as “123456”, “password”, or “qwerty”, as well as dictionaries, word lists, or patterns to speed up the process. A brute-force attack can target a remote system, such as a web server, an email account, or a network device, and gain unauthorized access to its data or resources. The other options are either different types of attacks or not related to password guessing. For example, a DoS, or Denial-of-Service, attack is a type of attack that floods a system with requests or traffic to overwhelm its capacity and prevent legitimate users from accessing it; an SQL injection attack is a type of attack that inserts malicious SQL statements into an input field or parameter of a web application to manipulate or compromise the underlying database; a phishing attack is a type of attack that sends fraudulent emails or messages that appear to come from a trusted source to trick users into revealing their personal or financial information. Reference: CompTIA DataSys+ Course Outline, Domain 4.0 Data and Database Security, Objective 4.4 Given a scenario, identify common types of attacks against databases.

Question: 50

Which of the following cloud delivery models provides users with the highest level of flexibility

regarding resource provisioning and administration?

A. DBaaS
B. IaaS
C. SaaS
D. PaaS

Answer: B

Explanation:

The cloud delivery model that provides users with the highest level of flexibility regarding resource provisioning and administration is IaaS. IaaS, or Infrastructure as a Service, is a cloud delivery model that provides users with access to virtualized computing resources, such as servers, storage, network, and operating systems, over the internet. Users can provision, configure, and manage these resources according to their needs and preferences, without having to worry about the maintenance or security of the physical infrastructure. IaaS offers users the most control and customization over their cloud environment, as well as the ability to scale up or down as needed. The other options are either different cloud delivery models or not related to cloud computing at all. For example, DBaaS, or Database as a Service, is a cloud delivery model that provides users with access to database management systems and tools over the internet; SaaS, or Software as a Service, is a cloud delivery model that provides users with access to software applications and services over the internet; PaaS, or Platform as a Service, is a cloud delivery model that provides users with access to development platforms and tools over the internet. Reference: CompTIA DataSys+ Course Outline, Domain 2.0 Database Deployment, Objective 2.1 Given a scenario, select an appropriate database deployment method.

Question: 51

Which of the following should a company develop to ensure preparedness for a fire in a data center?

A. Deployment plan
B. Backup plan
C. Data retention policy
D. Disaster recovery plan

Answer: D

Explanation:

The document that a company should develop to ensure preparedness for a fire in a data center is a disaster recovery plan. A disaster recovery plan is a document that outlines how an organization will continue its operations in the event of a disaster or disruption, such as fire, flood, earthquake, cyberattack, etc. A disaster recovery plan typically includes the following elements: - The objectives and scope of the plan - The roles and responsibilities of the staff involved - The identification and assessment of the risks and impacts - The strategies and procedures for restoring the critical functions and data - The resources and tools required for the recovery process - The testing and maintenance schedule for the plan A disaster recovery plan helps an organization to minimize the damage and downtime caused by a disaster, as well as to resume normal operations as soon as

possible. The other options are either different types of documents or not specific to fire preparedness. For example, a deployment plan is a document that describes how a system or software will be installed or launched; a backup plan is a document that specifies how data will be copied and stored for backup purposes; a data retention policy is a document that defines how long data should be kept and when it should be deleted or archived. Reference: CompTIA DataSys+ Course Outline, Domain 5.0 Business Continuity, Objective 5.4 Given a scenario, implement disaster recovery methods.

Question: 52

Which of the following is the deployment phase in which a DBA ensures the most recent patches are applied to the new database?

A. Importing
B. Upgrading
C. Provisioning
D. Modifying

Answer: B

Explanation:

The deployment phase in which a DBA ensures the most recent patches are applied to the new database is upgrading. Upgrading is a process that updates an existing database system or software to a newer version or release that may include new features, enhancements, bug fixes, security patches, etc. Upgrading helps improve the performance, functionality, compatibility, and security of the database system or software. Upgrading can be done manually or automatically using tools or scripts provided by the vendor or developer. Upgrading can also involve testing, backup, migration, or rollback procedures to ensure the quality and reliability of the new version or release. The other options are either different deployment phases or not related to deployment at all. For example, importing is a process that transfers data from one source to another using files or formats; provisioning is a process that allocates resources such as servers, storage, network, etc., for a system or software; modifying is a process that changes existing data or objects in a database using commands or scripts. Reference: CompTIA DataSys+ Course Outline, Domain 2.0 Database Deployment, Objective 2.3 Given a scenario, update database systems.

Question: 53

Which of the following computer services associates IP network addresses with text-based names in order to facilitate identification and connectivity?

A. LDAP
B. NTP
C. DHCP
D. IDNS

Answer: D

Explanation:

The computer service that associates IP network addresses with text-based names in order to facilitate identification and connectivity is IDNS. IDNS, or Internet Domain Name System (DNS), is a service that translates domain names into IP addresses and vice versa. Domain names are human- readable names that identify websites or devices on the internet, such as www.comptia.org or www.google.com. IP addresses are numerical identifiers that locate websites or devices on the internet, such as 104.18.26.46 or 142.250.72.238. IDNS helps users to access websites or devices using domain names instead of IP addresses, which are easier to remember and type. IDNS also helps administrators to manage websites or devices using domain names instead of IP addresses, which are more flexible and scalable. The other options are either different computer services or not related to IP network addresses or text-based names at all. For example, LDAP, or Lightweight Directory Access Protocol, is a service that provides access to directory information such as users, groups, or devices on a network; NTP, or Network Time Protocol, is a service that synchronizes the clocks of computers or devices on a network; DHCP, or Dynamic Host Configuration Protocol, is a service that assigns IP addresses and other network configuration parameters to computers or devices on a network. Reference: CompTIA DataSys+ Course Outline, Domain 2.0 Database Deployment, Objective 2.1 Given a scenario, select an appropriate database deployment method.

Question: 54

Which of the following types of RAID, if configured with the same number and type of disks, would provide the best write performance?

A. RAID 3
B. RAID 5
C. RAID 6
D. RAID 10

Answer: D

Explanation:

The type of RAID that would provide the best write performance if configured with the same number and type of disks is RAID 10. RAID 10, or RAID 1+0, is a type of RAID that combines mirroring and striping techniques to provide both redundancy and performance. Mirroring means that data is duplicated across two or more disks to provide fault tolerance and data protection. Striping means that data is split into blocks and distributed across two or more disks to provide faster access and throughput. RAID 10 requires at least four disks and can tolerate the failure of up to half of the disks without losing data. RAID 10 provides the best write performance among the RAID types because it can write data in parallel to multiple disks without parity calculations or overhead. The other options are either different types of RAID or not related to RAID at all. For example, RAID 3 is a type of RAID that uses striping with a

dedicated parity disk to provide redundancy and performance; RAID 5 is a type of RAID that uses striping with distributed parity to provide redundancy and performance; RAID 6 is a type of RAID that uses striping with double distributed parity to provide extra redundancy and performance. Reference: CompTIA DataSys+ Course Outline, Domain 3.0 Database Management and Maintenance, Objective 3.1 Given a scenario, perform common database maintenance tasks.

Question: 55

Which of the following is recommended in order to provide encrypted data communication pathways for information as it is transmitted over a network?

A. TCP/IP
B. NFS
C. SMB
D. TLS

Answer: D

Explanation:

The option that is recommended in order to provide encrypted data communication pathways for information as it is transmitted over a network is TLS. TLS, or Transport Layer Security, is a protocol that provides secure communication over the internet by encrypting the data using cryptographic algorithms and keys. TLS also provides authentication and integrity by verifying the identity of the parties involved and ensuring that the data has not been altered or tampered with. TLS can be used to protect various types of data, such as web traffic, email, instant messaging, voice over IP, etc. The other options are either not related or not sufficient for this purpose. For example, TCP/IP, or Transmission Control Protocol/Internet Protocol, is a set of protocols that defines how data is transmitted and routed over the internet, but does not provide encryption or security; NFS, or Network File System, is a protocol that allows users to access and share files over a network, but does not provide encryption or security; SMB, or Server Message Block, is a protocol that allows users to access and share files, printers, and other resources over a network, but does not provide encryption or security. Reference: CompTIA DataSys+ Course Outline, Domain 4.0 Data and Database Security, Objective 4.2 Given a scenario, implement security controls for databases.

Question: 56

Which of the following is an advantage of creating indexes?

A. To help with space allocation
B. To provide quick and efficient access to data
C. To reduce memory
D. To update the query plan

Answer: B

Explanation:

The advantage of creating indexes is to provide quick and efficient access to data. An index is a data structure that stores the values of one or more columns of a table in a sorted order, along with pointers to the corresponding rows in the table. An index helps to speed up queries that search, filter, sort, or join data based on the indexed columns, as it reduces the number of disk accesses or scans required to locate the desired data. An index also helps to enforce uniqueness or referential integrity constraints on the indexed columns. The other options are either not

true or not relevant for this purpose. For example, an index does not help with space allocation, as it consumes additional space in the database; an index does not reduce memory, as it may use memory for caching or buffering purposes; an index does not update the query plan, as it is an input or a factor for the query optimizer to generate the query plan. Reference: CompTIA DataSys+ Course Outline, Domain 1.0 Database Fundamentals, Objective 1.2 Given a scenario, execute database tasks using scripting and programming languages.

Question: 57

Which of the following is most likely to prevent tampering with server hardware that houses data?

A. Biometric locks
B. Strong password policy
C. Network firewall
D. Surveillance cameras

Answer: A

Explanation:

The option that is most likely to prevent tampering with server hardware that houses data is biometric locks. Biometric locks are devices that use biological characteristics, such as fingerprints, facial recognition, iris scan, etc., to control access to a physical location or resource. Biometric locks help prevent tampering with server hardware that houses data by restricting unauthorized entry or theft of the hardware by intruders or attackers. Biometric locks also provide higher security and convenience than other types of locks, such as keys or passwords, which can be lost, stolen, or forgotten. The other options are either not related or not effective for this purpose. For example, a strong password policy is a set of rules or standards for creating and managing passwords for user accounts or systems; a network firewall is a device or software that controls the incoming and outgoing traffic on a network based on a set of rules or policies; surveillance cameras are devices that capture and record video footage of a physical location or resource. Reference: CompTIA DataSys+ Course Outline, Domain 4.0 Data and Database Security, Objective 4.2 Given a scenario, implement security controls for databases.

Question: 58

Which of the following describes the purpose of a snapshot?

A. To create a dynamic data replication
B. To create a synonym
C. To create a
D. To create an image of a database

Answer: D

Explanation:

The purpose of a snapshot is to create an image of a database. A snapshot is a copy of the state and content of a database at a specific point in time. A snapshot can be used for various purposes, such as backup and recovery, testing and development, reporting and analysis, etc. A snapshot can be created using various techniques, such as full copy, incremental copy, differential copy, etc. A

snapshot can also be created using various tools or commands provided by the database system or software. The other options are either incorrect or irrelevant for this question. For example, dynamic data replication is a process that copies and synchronizes data from one database server (the source) to one or more database servers (the target) in real time; a synonym is an alias or an alternative name for an object in a database; C is an incomplete option. Reference: CompTIA DataSys+ Course Outline, Domain 5.0 Business Continuity, Objective 5.2 Given a scenario, implement backup and restoration of database management systems.

Question: 59

Which of the following is a typical instruction that is found on a Linux command-line script and represents a system shell?

A. /bin/bash
B. #/bin/shell
C. >/bin/sh
D. #!/bin/bash

Answer: D

Explanation:

The instruction that is found on a Linux command-line script and represents a system shell is #!/bin/bash. This instruction is called a shebang or a hashbang, and it indicates the interpreter that should be used to execute the script. In this case, the interpreter is /bin/bash, which is the path to the bash shell, a common system shell for Linux. A system shell is a program that provides an interface for users to interact with the operating system, either through commands or scripts. A system shell can also perform various tasks, such as file management, process control, variable assignment, etc. The other options are either incorrect or not typical for this purpose. For example, /bin/bash is the path to the bash shell, but it does not indicate the interpreter for the script; #/bin/shell is not a valid shebang or a path to a system shell; >/bin/sh is a redirection operator followed by a path to a system shell, but it does not indicate the interpreter for the script. Reference: CompTIA DataSys+ Course Outline, Domain 2.0 Database Deployment, Objective 2.2 Given a scenario, create database objects using scripting and programming languages.

Question: 60

Which of the following is a characteristic of all non-relational databases?

A. Columns with the same data type
B. Unstructured data
C. Logical record groupings
D. Tabular schema

Answer: B

Explanation:

The characteristic of all non-relational databases is unstructured data. Unstructured data is data that

does not have a predefined or fixed format, schema, or structure. Unstructured data can include various types of data, such as text, images, audio, video, etc. Non-relational databases, also known as NoSQL databases, are databases that store and manage unstructured data using different models, such as key-value, document, graph,

columnar, etc. Non-relational databases are suitable for handling large volumes, variety, and velocity of data that do not fit well in the relational model. The other options are either characteristics of relational databases or not related to database types at all. For example, columns with the same data type, logical record groupings, and tabular schema are characteristics of relational databases, which are databases that store and manage structured data using tables, rows, columns, and constraints. Reference: CompTIA DataSys+ Course Outline, Domain 1.0 Database Fundamentals, Objective 1.1 Given a scenario, identify common database types.

Question: 61

Which of the following is used to hide data in a database so the data can only be read by a user who has a key?

A. Data security
B. Data masking
C. Data protection
D. Data encryption

Answer: D

Explanation:

The option that is used to hide data in a database so the data can only be read by a user who has a key is data encryption. Data encryption is a process that transforms data into an unreadable or scrambled form using an algorithm and a key. Data encryption helps protect data from unauthorized access or modification by third parties, such as hackers, eavesdroppers, or interceptors. Data encryption also helps verify the identity and authenticity of the source and destination of the data using digital signatures or certificates. Data encryption can be applied to data at rest (stored in a database) or data in transit (transmitted over a network). To read encrypted data, a user needs to have the corresponding key to decrypt or restore the data to its original form. The other options are either different concepts or not related to hiding data at all. For example, data security is a broad term that encompasses various methods and techniques to protect data from threats or risks; data masking is a technique that replaces sensitive data with fictitious but realistic data to protect its confidentiality or compliance; data protection is a term that refers to the legal or ethical obligations to safeguard personal or sensitive data from misuse or harm. Reference: CompTIA DataSys+ Course Outline, Domain 4.0 Data and Database Security, Objective 4.2 Given a scenario, implement security controls for databases.

Question: 62

A server administrator wants to analyze a database server's disk throughput. Which of the following should the administrator measure?

A. RPfvl
B. Latency
C. IOPS
D. Reads

Answer: C

Explanation:

The factor that the administrator should measure to analyze a database server’s disk throughput is IOPS. IOPS, or

Input/Output Operations Per Second, is a metric that measures the number of read and write operations that a disk can perform in one second. IOPS indicates the performance or speed of a disk and how well it can handle multiple requests or transactions. Higher IOPS means higher disk throughput and lower latency. IOPS can be affected by various factors, such as disk type, size, speed, cache, RAID level, etc. The other options are either not related or not sufficient for this purpose. For example, RPfvl is not a valid acronym or metric; latency is the time delay between a request and a response; reads are the number of read operations performed by a disk. Reference: CompTIA DataSys+ Course Outline, Domain 3.0 Database Management and Maintenance, Objective 3.2 Given a scenario, monitor database performance.

Question: 63

Which of the following have data manipulation and procedural scripting power? (Choose two.)

A. PQL
B. PL/SQL
C. Advanced
D. SQL
E. SQL
F. T-SQL

Answer: B, F

Explanation:

The two options that have data manipulation and procedural scripting power are PL/SQL and T-SQL. PL/SQL, or Procedural Language/Structured Query Language, is an extension of SQL that adds procedural features to SQL for Oracle databases. PL/SQL allows users to create and execute stored procedures, functions, triggers, packages, etc., using variables, loops, conditions, exceptions, etc., in addition to SQL commands. PL/SQL helps improve the performance, functionality, modularity, and security of SQL queries and applications. T-SQL, or Transact-SQL, is an extension of SQL that adds procedural features to SQL for Microsoft SQL Server databases. T-SQL allows users to create and execute stored procedures, functions, triggers, etc., using variables, loops, conditions, exceptions, etc., in addition to SQL commands. T-SQL helps improve the performance, functionality, modularity, and security of SQL queries and applications. The other options are either not related or not having both data manipulation and procedural scripting power. For example, PQL, or Power Query Language, is a data analysis and transformation language for Microsoft Power BI and Excel; Advanced SQL is a term that refers to the advanced features or techniques of SQL, such as subqueries, joins, aggregations, etc.; SQL, or Structured Query Language, is a standard language for manipulating and querying data in relational databases, but it does not have procedural features. Reference: CompTIA DataSys+ Course Outline, Domain 1.0 Database Fundamentals, Objective 1.2 Given a scenario, execute database tasks using scripting and programming languages.

Question: 64

Which of the following database instances are created by default when SQL Server is installed? (Choose two.)

A. Root
B. Master
C. Log
D. Model

E. View
F. Index

Answer: B, D

Explanation:

The two database instances that are created by default when SQL Server is installed are master and model. Master is a system database that contains the information and settings of the SQL Server instance, such as the configuration, logins, endpoints, databases, etc. Master is essential for the operation and management of the SQL Server instance, and it should be backed up regularly. Model is a system database that serves as a template for creating new user databases. Model contains the default settings and objects, such as tables, views, procedures, etc., that will be inherited by the new user databases. Model can be modified to customize the new user databases according to specific needs or preferences. The other options are either not database instances or not created by default when SQL Server is installed. For example, root is not a database instance, but a term that refers to the highest level of access or privilege in a system; log is not a database instance, but a file that records the changes made by transactions on a database; view is not a database instance, but an object that represents a subset or a combination of data from one or more tables; index is not a database instance, but a data structure that stores the values of one or more columns of a table in a sorted order. Reference: CompTIA DataSys+ Course Outline, Domain 2.0 Database Deployment, Objective 2.3 Given a scenario, update database systems.

Question: 65

Which of the following services is responsible for assigning, managing, and reclaiming IP addresses on a TCP/IP-based network?

A. DNS
B. DHCP
C. LDAP
D. ISMTP

Answer: B

Explanation:

The service that is responsible for assigning, managing, and reclaiming IP addresses on a TCP/IP-based network is DHCP. DHCP, or Dynamic Host Configuration Protocol, is a service that automatically assigns IP addresses and other network configuration parameters, such as subnet mask, default gateway, DNS server, etc., to computers or devices on a network. DHCP helps simplify the administration and management of IP addresses on a network, as well as avoid conflicts or errors caused by manual or duplicate assignments. DHCP also allows computers or devices to release or renew their IP addresses when they join or leave the network. The other options are either different services or not related to IP addresses at all. For example, DNS, or Domain Name System, is a service that translates domain names into IP addresses and vice versa; LDAP, or Lightweight Directory Access Protocol, is a service that provides access to directory information such as users, groups, or devices on a network; ISMTP is not a valid acronym or service. Reference: CompTIA DataSys+ Course Outline, Domain 2.0 Database Deployment, Objective 2.1 Given a scenario, select an appropriate database deployment method.

Question: 66

Which of the following is an attack in which an attacker hopes to profit from locking the database software?

A. Spear phishing
B. Ransomware
C. SQL injection
D. On-path

Answer: B

Explanation:

The attack in which an attacker hopes to profit from locking the database software is ransomware. Ransomware is a type of malware that encrypts the data or files on a system or network and demands a ransom from the victim to restore them. Ransomware can target database software and lock its access or functionality until the victim pays the ransom, usually in cryptocurrency. Ransomware can cause serious damage and loss to the victim, as well as expose them to further risks or threats. Ransomware can be delivered through various methods, such as phishing emails, malicious attachments, compromised websites, etc. The other options are either different types of attacks or not related to locking database software at all. For example, spear phishing is a type of phishing attack that targets a specific individual or organization with personalized or customized emails; SQL injection is a type of attack that inserts malicious SQL statements into an input field or parameter of a web application to manipulate or compromise the underlying database; on-path is a type of attack that intercepts and modifies the data in transit between two parties on a network. Reference: CompTIA DataSys+ Course Outline, Domain 4.0 Data and Database Security, Objective 4.4 Given a scenario, identify common types of attacks against databases

Question: 67

Which of the following cloud storage options provides users with endpoints to retrieve data via REST API?

A. Network file
B. Object
C. Ephemeral
D. iBlock

Answer: B

Explanation:

The cloud storage option that provides users with endpoints to retrieve data via REST API is object. Object storage is a type of cloud storage that stores data as objects, which consist of data, metadata, and a unique identifier. Object storage does not use any hierarchy or structure to organize data, but rather uses flat namespaces that allow users to access data using the unique identifier. Object storage also provides users with endpoints to retrieve data via REST API (Representational State Transfer Application Programming Interface), which is a standard way of communicating with web services using HTTP methods (such as GET, POST, PUT, DELETE) and formats (such as JSON, XML). Object storage is suitable for storing large amounts of unstructured data that do not require frequent changes or complex queries. The other options are either different types of cloud storage or not related to cloud storage at all. For example, network file storage is a type of cloud storage that stores data as files in folders using protocols such as NFS (Network File System) or SMB (Server Message Block); ephemeral storage is a type of

temporary storage that stores data only for the duration of a session or process; iBlock is not a valid acronym or type of cloud storage. Reference: CompTIA DataSys+ Course Outline, Domain 2.0 Database Deployment, Objective 2.1 Given a scenario, select an appropriate database deployment method.

Question: 68

Which of the following is used to write SQL queries in various programming languages?

A. Indexing
B. Object-relational mapping
C. Excel
D. Normalization

Answer: B

Explanation:

The option that is used to write SQL queries in various programming languages is object-relational mapping. Object-relational mapping (ORM) is a technique that maps objects in an object-oriented programming language (such as Java, Python, C#, etc.) to tables in a relational database (such as Oracle, MySQL, SQL Server, etc.). ORM allows users to write SQL queries in their preferred programming language without having to deal with the differences or complexities between the two paradigms. ORM also provides users with various benefits such as code reuse, abstraction, validation, etc. The other options are either not related or not effective for this purpose. For example, indexing is a technique that creates data structures that store the values of one or more columns of a table in a sorted order to speed up queries; Excel is a software application that allows users to organize and manipulate data in rows and columns; normalization is a process that organizes data into tables and columns to reduce redundancy and improve consistency. Reference: CompTIA DataSys+ Course Outline, Domain 1.0 Database Fundamentals, Objective 1.2 Given a scenario, execute database tasks

using scripting and programming languages.

Question: 69

Which of the following best describes a collection of data that shares the same properties or attributes?

A. Relation set
B. ER model
C. Entity set
D. Tuples

Answer: C

Explanation:

The option that best describes a collection of data that shares the same properties or attributes is entity set. An entity set is a term used in the entity-relationship (ER) model, which is a conceptual model for designing and representing databases. An entity set is a collection of entities that have the same type or characteristics, such as students, courses, products, etc. An entity is an object or thing that can be identified and distinguished from others, such as a specific student, course, product, etc. An entity set can have one or more attributes that describe the properties or features of the entities, such as name, age, price, etc. An entity set can also have one or more relationships with other entity sets that define how the entities are associated or connected, such as enrolled,

taught by, purchased by, etc. The other options are either different terms or not related to the ER model at all. For example, relation set is a term used in the relational model, which is a logical model for implementing and manipulating databases; ER model is a term used to refer to the entityrelationship model itself; tuples are rows or records in a table or relation. Reference: CompTIA DataSys+ Course Outline, Domain 1.0 Database Fundamentals, Objective 1.1 Given a scenario, identify common database types.

Question: 70

Which of the following sets the age requirement for data that should be recovered after a major disaster?

A. MTBF
B. RTO
C. MTTF
D. RPO

Answer: D

Explanation:

The option that sets the age requirement for data that should be recovered after a major disaster is RPO. RPO, or Recovery Point Objective, is a metric that defines the maximum amount of data that can be lost or acceptable data loss in the event of a disaster or disruption. RPO indicates how frequently the data should be backed up or replicated to minimize the risk of data loss. RPO also sets

the age requirement for data that should be recovered after a major disaster, as it determines how far back in time the recovery process should go. For example, if the RPO is one hour, then the data should be backed up or replicated every hour, and the recovery process should restore the data to the state it was in one hour before the disaster. The other options are either different metrics or not related to data recovery at all. For example, MTBF, or Mean Time Between Failures, is a metric that measures the average time that a system or component operates without failure; RTO, or Recovery Time Objective, is a metric that defines the maximum amount of time that can be taken to restore a system or service after a disaster or disruption; MTTF, or Mean Time To Failure, is a metric that measures the average time that a system or component operates until it fails. Reference: CompTIA DataSys+ Course Outline, Domain 5.0 Business Continuity, Objective 5.3 Given a scenario, implement backup and restoration of data.

Question: 71

Which of the following is part of logical database infrastructure security?

A. Surveillance
B. Biometric access
C. Perimeter network
D. Cooling system

Answer: C

Explanation:

The option that is part of logical database infrastructure security is perimeter network. Perimeter network, also known as DMZ (Demilitarized Zone), is a network segment that lies between an internal network and an external network, such as the internet. Perimeter network provides an additional layer of security for the internal network

by isolating and protecting the servers or services that are exposed to the external network, such as web servers, email servers, database servers, etc. Perimeter network also helps prevent unauthorized access or attacks from the external network to the internal network by using firewalls, routers, proxies, etc. The other options are either part of physical database infrastructure security or not related to database infrastructure security at all. For example, surveillance is a method of monitoring and recording physical activities or events in a location or resource; biometric access is a device that uses biological characteristics to control access to a physical location or resource; cooling system is a device or system that regulates the temperature and humidity of a location or resource. Reference: CompTIA DataSys+ Course Outline, Domain 4.0 Data and Database Security, Objective 4.1 Given a scenario, implement database infrastructure security.

Question: 72

Which of the following database structures is a type of NoSQL database?

A. Hierarchical
B. Key-value stores
C. Cloud
D. Object-oriented

Answer: B

Explanation:

The database structure that is a type of NoSQL database is key-value stores. Key-value stores are databases that store and manage data as pairs of keys and values. Keys are unique identifiers that locate data in the database; values are arbitrary data that can be any type or format. Key-value stores do not use any schema or structure to organize data, but rather use hash tables or indexes to enable fast and simple access to data based on keys. Key-value stores are suitable for storing large amounts of simple or unstructured data that do not require complex queries or relationships. The other options are either different types of databases or not related to database structures at all. For example, hierarchical databases are databases that store and manage data as nodes in a tree-like structure; cloud databases are databases that are hosted and accessed over the internet using cloud computing services; object-oriented databases are databases that store and manage data as objects that have attributes and methods. Reference: CompTIA DataSys+ Course Outline, Domain 1.0 Database Fundamentals, Objective 1.1 Given a scenario, identify common database types.

Question: 73

Which of the following NoSQL database types best categorizes MongoDB?

A. Document
B. Column-oriented
C. Graph
D. Key-value stores

Answer: A

Explanation:

The NoSQL database type that best categorizes MongoDB is document. Document databases are databases that store and manage data as documents, which are collections of fields and values in formats such as JSON (JavaScript

Object Notation) or XML (Extensible Markup Language). Document databases do not use any schema or structure to organize data, but rather use identifiers or indexes to enable flexible and dynamic access to data based on fields or values. Document databases are suitable for storing large amounts of complex or unstructured data that have variable attributes or nested structures. MongoDB is an example of a document database that uses JSON-like documents to store and query data. The other options are either different types of NoSQL databases or not related to NoSQL databases at all. For example, column-oriented databases are databases that store and manage data as columns rather than rows; graph databases are databases that store and manage data as nodes and edges that represent entities and relationships; key-value stores are databases that store and manage data as pairs of keys and values. Reference: CompTIA DataSys+ Course Outline, Domain 1.0 Database Fundamentals, Objective 1.1 Given a scenario, identify common database types.

Question: 74

Which of the following tools is used for natively running a Linux system in Windows?

A. WSL
B. [Remote Desktop Protocol
C. SSH
D. ITelnet

Answer: A

Explanation:

The tool that is used for natively running a Linux system in Windows is WSL. WSL, or Windows Subsystem for Linux, is a feature that allows users to run a Linux system natively on Windows 10 or Windows Server. WSL enables users to install and use various Linux distributions, such as Ubuntu, Debian, Fedora, etc., and run Linux commands, tools, applications, etc., without requiring a virtual machine or a dual-boot setup. WSL also provides users with interoperability and integration between Linux and Windows, such as file system access, network communication, process management, etc. WSL is useful for users who want to use Linux features or functionalities on Windows, such as development, testing, scripting, etc. The other options are either different tools or not related to running a Linux system in Windows at all. For example, Remote Desktop Protocol (RDP) is a protocol that allows users to remotely access and control another computer or device over a network; SSH, or Secure Shell, is a protocol that allows users to securely connect and communicate with another computer or device over a network; Telnet is a protocol that allows users to interact with another computer or device over a network using a text-based interface. Reference: CompTIA DataSys+ Course Outline, Domain 2.0 Database Deployment, Objective 2.2 Given a scenario, create database objects using scripting and programming languages.

Question: 75

Which of the following constraints is used to enforce referential integrity?

A. Surrogate key
B. Foreign key
C. Unique key
D. Primary key

Answer: B

Explanation:

The constraint that is used to enforce referential integrity is foreign key. A foreign key is a column or a set of columns in a table that references the primary key of another table. A primary key is a column or a set of columns in a table that uniquely identifies each row in the table. Referential integrity is a rule that ensures that the values in the foreign key column match the values in the primary key column of the referenced table. Referential integrity helps maintain the consistency and accuracy of the data across related tables. The other options are either different types of constraints or not related to referential integrity at all. For example, a surrogate key is a column that is artificially generated to serve as a primary key, such as an auto-increment number or a GUID (Globally Unique Identifier); a unique key is a column or a set of columns in a table that uniquely identifies each row in the table, but it can have null values unlike a primary key; there is no such constraint as

TID. Reference: CompTIA DataSys+ Course Outline, Domain 1.0 Database Fundamentals, Objective 1.2 Given a scenario, execute database tasks using scripting and programming languages.

Question: 76

Which of the following NFs is considered the most preferable for relational database design?

A. 1 NF
B. 3 NF
C. 4 NF
D. 2NF

Answer: B

Explanation:

The NF (normal form) that is considered the most preferable for relational database design is 3 NF. 3 NF, or Third Normal Form, is a level of normalization that organizes data into tables and columns to reduce redundancy and improve consistency. Normalization is a process that applies a set of rules or criteria to eliminate or minimize the anomalies or problems that may arise from inserting, updating, or deleting data in a database. 3 NF is achieved when a table satisfies the following conditions: - It is in 2 NF (Second Normal Form), which means that every non-key column depends on the whole primary key and not on any subset of it - It has no transitive dependencies, which means that every non-key column depends directly on the primary key and not on any other non-key column 3 NF is considered the most preferable for relational database design because it ensures that each table has only one purpose or theme and that each column has only one value or meaning. This helps avoid data duplication, inconsistency, and update anomalies. The other options are either lower or higher levels of normalization that are either less preferable or less practical for relational database design. For example, 1 NF (First Normal Form) is the lowest level of normalization that requires each column to have atomic values and each row to have a unique identifier; 4 NF (Fourth Normal Form) is a higher level of normalization that requires each table to have no multi-valued dependencies, which means that there are no columns that can have more than one value for the same primary key value; 2 NF (Second Normal Form) is an intermediate level of normalization that requires each non-key column to depend on the whole primary key and not on any subset of it. Reference: CompTIA DataSys+ Course Outline, Domain 1.0 Database Fundamentals, Objective 1.2 Given a scenario, execute database tasks using scripting and programming languages.

Question: 77

Which of the following indexes stores records in a tabular format?

A. Columnstore
B. Non-clustered
C. Unique
D. Secondary

Answer: A

Explanation:

The index that stores records in a tabular format is columnstore. A columnstore index is a type of index that stores and compresses data by columns rather than by rows. A columnstore index can improve the performance and efficiency of queries that perform aggregations, calculations, or analysis on large amounts of data, such as data warehouse or business intelligence applications. A columnstore index can also reduce the storage space required for data by applying various compression techniques, such as dictionary encoding, run-length encoding, bit packing, etc. The other options are either different types of indexes or not related to indexes at all. For example, a non-clustered index is a type of index that stores the values of one or more columns in a sorted order along with pointers to the corresponding rows in the table; a unique index is a type of index that enforces uniqueness on one or more columns in a table; a secondary index is an alternative term for a non-clustered index. Reference: CompTIA DataSys+ Course Outline, Domain 3.0 Database Management and Maintenance, Objective 3.1 Given a scenario, perform common database maintenance tasks.

Question: 78

Which of the following resources is the best way to lock rows in SQL Server?

A. TID
B. SID
C. RID
D. PID

Answer: C

Explanation:

The resource that is the best way to lock rows in SQL Server is RID. RID, or Row IDentifier, is an attribute that uniquely identifies each row in a heap table in SQL Server. A heap table is a table that does not have a clustered index, which means that the rows are not stored in any particular order. A RID consists of the file number, page number, and slot number of the row in the database. A RID can be used to lock rows in SQL Server to prevent concurrent access or modification by other transactions or users. A RID lock is a type of lock that locks a single row using its RID. A RID lock can be applied using the HOLDLOCK or XLOCK hints in a SELECT statement. The other options are either not related or not effective for this purpose. For example, TID, or Transaction IDentifier, is an attribute that uniquely identifies each transaction in a database; SID, or Security IDentifier, is an attribute that uniquely identifies each user or group in a Windows system; PID, or Process IDentifier, is an attribute that uniquely identifies each process in an operating system. Reference: CompTIA DataSys+ Course Outline, Domain 3.0 Database Management and Maintenance, Objective 3.3 Given a scenario, implement database concurrency methods.

Question: 79

Which of the following commands is part of DDL?

A. UPDATE
B. GRANT
C. CREATE
D. INSERT

Answer: C

Explanation:

The command that is part of DDL is CREATE. CREATE is a SQL command that belongs to the category of DDL, or Data Definition Language. DDL is a subset of SQL commands that are used to define or modify the structure or schema of a database, such as tables, columns, constraints, indexes, views, etc. CREATE is a DDL command that is used to create a new object in a database, such as a table, column, constraint, index, view, etc. For example, the following statement uses the CREATE command to create a new table called employee with four columns: CREATE TABLE employee (

```
    emp_id INT PRIMARY KEY,
    emp_name VARCHAR(50) NOT NULL,
    emp_dept VARCHAR(20),
    emp_salary DECIMAL(10,2)
);
```

Copy

The other options are either part of different categories of SQL commands or not SQL commands at all. For example, UPDATE is a SQL command that belongs to the category of DML, or Data Manipulation Language. DML is a subset of SQL commands that are used to manipulate or modify the data or content of a database, such as inserting, updating, deleting, or selecting data. GRANT is a SQL command that belongs to the category of DCL, or Data Control Language. DCL is a subset of SQL commands that are used to control or manage the access or permissions of users or roles on a database, such as granting or revoking privileges or roles. INSERT is a SQL command that belongs to the category of DML, or Data Manipulation Language. INSERT is a DML command that is used to insert new data into a table. Reference: CompTIA DataSys+ Course Outline, Domain 1.0 Database Fundamentals, Objective 1.2 Given a scenario, execute database tasks using scripting and programming languages.

Question: 80

Which of the following best describes the category of SQL commands required to revoke access to database objects?

A. DCL
B. IDDL
C. IDML
D. TCL

Answer: A

Explanation:

The category of SQL commands that is required to revoke access to database objects is DCL. DCL, or Data Control Language, is a subset of SQL commands that are used to control or manage the access or permissions of users or roles on a database. DCL includes commands such as GRANT and REVOKE.

GRANT is a DCL command that is used to grant privileges or roles to users or roles on specific objects in a database, such as tables, views, procedures, etc. REVOKE is a DCL command that is used to revoke privileges or roles from

users or roles on specific objects in a database. For example, the following statement uses the REVOKE command to revoke the SELECT privilege from user Alice on table employee:
REVOKE SELECT ON employee FROM Alice;

The other options are either different categories of SQL commands or not related to SQL commands at all. For example, IDDL is not a valid acronym or category of SQL commands; IDML is not a valid acronym or category of SQL commands; TCL, or Transaction Control Language, is a subset of SQL commands that are used to control or manage transactions on a database, such as committing or rolling back changes. Reference: CompTIA DataSys+ Course Outline, Domain 4.0 Data and Database Security, Objective 4.2 Given a scenario, implement security controls for databases.