



Free Questions for DS0-001 by dumpsheet

Shared by Phelps on 12-12-2023

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

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

Options:

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.

Question 2

Question Type: MultipleChoice

Which of the following commands is part of DDL?

Options:

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 3

Question Type: MultipleChoice

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

Options:

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 4

Question Type: MultipleChoice

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

Options:

- 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 5

Question Type: MultipleChoice

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

Options:

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 6

Question Type: MultipleChoice

Which of the following describes the purpose of a snapshot?

Options:

- 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 7

Question Type: MultipleChoice

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

Options:

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 8

Question Type: MultipleChoice

Which of the following is an advantage of creating indexes?

Options:

- 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.

To Get Premium Files for DS0-001 Visit

<https://www.p2pexams.com/products/ds0-001>

For More Free Questions Visit

<https://www.p2pexams.com/comptia/pdf/ds0-001>

