

Microsoft.Actualexams.98-364.v2014-08-23.by.EVA.40q

Number: 98-364
Passing Score: 800
Time Limit: 120 min
File Version: 12.5



Pass Microsoft 98-364 Exam

Exam Name: Database Fundamentals

Actual-Exams

QUESTION 1

John works as a Database Administrator for Bluewell Inc. The company has a SQL Server database. A table in the database has a candidate key and an attribute that is not a constituent of the candidate key. The non-key attribute depends upon the whole of the candidate key rather than just a part of it. Which of the following normal forms is represented in the scenario?

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

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Second normal form (2NF) is used in database normalization. A table that is in first normal form must meet additional criteria if it is to qualify for second normal form. Specifically, a 1NF table is in 2NF if and only if, given any candidate key and any attribute that is not a constituent of a candidate key, the non-key attribute depends upon the whole of the candidate key rather than just a part of it. Answer: C is incorrect. The 1NF is a normalization form in which each column in a row contains a single value, i.e., each attribute of the entity is single valued. Single valued attributes are also known as atomic attributes, as they cannot be decomposed into smaller units. There are mainly three kinds of attributes that prevent a table from being in the first normal form. They are as follows:

Composite attributes

Plural attributes (attributes that have more than one value) Attributes with complex data types

The table below is in 1 NF, as all the columns in each row contain a single value.

SSN	Name	Salary	DOB	Department
A101	Smith	20000	23/09/78	Sales
A102	Flora	23000	12/12/67	Sales
A103	Dinzi	12000	12/25/67	Production
A104	Henry	23000	23/5/78	Marketing
A105	Ricky	15000	12/09/68	Sales
A106	Esha	12500	05/06/77	Marketing
A107	Ema	34000	07/09/76	IT



For interactive and self-paced preparation of exam 98-364, try our practice exams. Practice exams also include self assessment and reporting features!
www.selftestengine.com



Answer: D is incorrect. Third normal form (3NF) is used in database normalization. A table is in 3NF if and only if the relation S (table) is in second normal form (2NF) and every non-prime attribute of S is non-transitively dependent on every key of S.

Answer: A is incorrect. Fourth normal form (4NF) is a normal form used in database normalization. Introduced by Ronald Fagin in 1977, 4NF is the next level of normalization after Boyce-Codd normal form (BCNF). Whereas the second, third, and Boyce-Codd normal forms are concerned with functional dependencies, 4NF is concerned with a more general type of dependency known as a multivalued dependency. Symbolically, it can be represented as follows:

If $A \twoheadrightarrow B|C$, then $A \rightarrow B$ and $A \rightarrow C$

Here, A, B, and C are attributes.

QUESTION 2

You work as a Database Designer for DataOneWorld Inc. The company has a SQL Server database. You are assigned the task of creating a data model of an enterprise based on a specific data model. The model to be created should be independent of a particular DBMS. Which of the following database designs will help you accomplish the task?

- A. Logical database design
- B. Conceptual database design
- C. Physical database design
- D. Application design

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Logical database design is the process of creating a data model of an enterprise based on a specific data model, but independent of a particular DBMS. In the logical database design stage, users can define tables and relationships that reflect the conceptual design. Answer: C is incorrect. Physical database design is the process of producing a description of the database implementation on secondary storage. It describes base relations, file organizations, and indexes used to achieve efficient access to data. It also describes any associated integrity constraints and security measures, tailored to a specific DBMS system. This involves the actual programming of the database, taking into account the features and limitations of the DBMS and the client. Answer: B is incorrect. Conceptual database design is the process of constructing a model for the data at a high level of abstraction without using DBMS. It consists of the construction of an Entity- Relationship schema, providing an optimal description of the user requirements. Answer: D is

incorrect. Application design is a design of the user interface and the application programs that use and process the database. Database and application design are parallel activities. Application design includes two important activities: transaction design and user interface design.



For interactive and self-paced preparation of exam 98-364, try our practice exams. Practice exams also include self assessment and reporting features! www.selftestengine.com



QUESTION 3

Which of the following steps in database planning helps to determine the requirements of the database through interviewing?

- A. Gathering information
- B. Identifying the relationship between objects
- C. Identify the objects
- D. Identifying the type of information for each objects
- E. Modeling the object

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Gathering information is the first step of database planning and requires a good understanding of what tasks are to be performed in the database. In case the database has to replace a paper-based or a manual system, the old system can provide all the necessary information. Moreover, all those who were involved with the previous system should be interviewed to determine what they did and what is now needed for the new database. The limitations, problems, and bottlenecks of the old system should also be identified.

Answer: C is incorrect. In the object identification process in database planning, it is necessary to identify the key objects or entities for database management. There are generally a few primary objects, and once these are identified, the related items become visible. Answer: E is incorrect. Object modeling is the third step in database planning in which objects are identified in the system. A user should record an object in such a way that it is represented visually in the system. He can use the database model as a reference during implementation of the database. Database developers use tools such as pencils, paper, word processing etc., as well as create software programs for data modeling of the database. Answer: D is incorrect. Identifying information for objects is the fourth step in database planning where the information for each object is stored. In a table, there are four columns that contain different types of information for objects:

1. Categorical column: In this column, data is classified and grouped, and a limited selection of data is stored.
2. Relational or referential column: In this column, a link is established between information in one table and related information in another table.
3. Identifier column: In this column, each item stored in the table is identified.
4. Raw data column: In this column, tangible information is stored and is determined by a source that is external to the database.

Answer: B is incorrect. Identifying relationships between objects is the last step in database planning, which is the ability to relate information about various items in the database. Isolated types of information can be stored separately, but the data can be combined according to the requirements of the database engine. Identifying the relationship between objects in the design process necessitates looking at the tables, determining how the tables are logically related, and adding relational columns to the tables that establish a link from one table to another.



For interactive and self-paced preparation of exam 98-364, try our practice exams. Practice exams also include self assessment and reporting features!
www.selftestengine.com



QUESTION 4

Which of the following database terms is described in the statement below? "It prevents the current database operation from reading or writing a data item while the data item is being accessed by another operation."

- A. Lock
- B. Encryption
- C. Constraint
- D. Deadlock

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Lock is the process to prevent the current database operations from reading and writing a data while that data item is being accessed by another operation. A lock is used when multiple users need to access a database concurrently. It prevents data from getting corrupt or invalidated, when multiple users try to write to the database. When a lock is applied, a single user can only modify that record to which the lock is applied. It gives the user exclusive access to the record until the lock is released. Locking also prevents reading of unfinished (uncommitted) data.



For interactive and self-paced preparation of exam 98-364, try our practice exams. Practice exams also include self assessment and reporting features!
www.selftestengine.com



Answer: B is incorrect. Encryption provides an additional security layer, protecting the data from unauthorized viewing with the help of an algorithm called cipher. Even if access to the database is obtained, it will not be easy to decipher encrypted data into a readable form. Answer: D is incorrect. A deadlock is a situation wherein two or more competing actions wait for the other to finish, and neither ever does.

Answer: C is incorrect. A constraint enforces the integrity of a database. It defines the rules regarding the values allowed in the columns of a table. A constraint is the standard mechanism for enforcing integrity.

QUESTION 5

Which of the following can be used to populate a table? Each correct answer represents a complete solution. Choose all that apply.

- A. Data Pump
- B. SQL*Loader
- C. INSERT statement
- D. MERGE statement

Correct Answer:

Section: (none)

Explanation

Explanation/Reference:

Answer: C, B, A, and

D

Explanation:

The INSERT statement is a Data Manipulation Language (DML) statement that is used to add new rows of data to a specified database table.

Syntax:


```
INSERT INTO TableName [ (ColumnName [, ColumnName...]) ] VALUES ( value [, value...]);
```

where,

TableName specifies the name of the table.

ColumnName specifies the name of the column of the table, which is to be populated. value specifies the value for the corresponding column of the table. The number and datatypes of values specified in the VALUES clause must match the number and datatypes of corresponding columns specified in the INTO clause. If the columns are not specified, the values in the VALUES clause must be in the order in which the columns are defined in the table.

SQL*Loader (sqlldr) is a utility that is used for high performance data loads. The data can be in the form of a text file or embedded into a database. By using the information stored in the data file and the control file, sqlldr loads the data into a database. The architecture of SQL*Loader is given below:



For interactive and self-paced preparation of exam 98-364, try our practice exams. Practice exams also include self assessment and reporting features!
www.selftestengine.com




In the figure, 1 depicts the parameter file, 2 depicts the data file, 3 depicts the control file, 4 depicts the log file, 5 depicts the bad file, and 6 depicts the discard file. Note: While processing the information, SQL*Loader stores messages in the log file, bad rows in the bad file, and discarded rows in the discard file. Data Pump is a new feature introduced in Oracle 10g to move data and meta data between databases and to or from operating system files very efficiently. It provides parallel import and export utilities (impdp, expdp) on the command-line as well as the Web-based Oracle Enterprise Manager export/import interface. It is ideally beneficial for large databases and data warehousing environments. Oracle Data Pump facility runs on the server.

Following are some functions performed by Oracle Data Pump:

1. It is used to copy data from one schema to another between two databases or within a single database.
2. It can be used to extract a logical copy of the entire database, a list of tablespaces, a list of schemas or a list of tables. The MERGE statement is a Data Manipulation Language (DML) statement that is used to update or insert rows conditionally into a table. It selects rows from one table and updates or inserts them into another table. The decision whether to update or insert rows in the target table is based on a condition in the ON clause of the MERGE statement.

QUESTION 6

You work as a Database Designer for SoftTech Inc. You have developed the conceptual design of the database of an organization. Following is the ER-diagram, which you have designed. Now you have to convert the following diagram to a logical model.



Which of the following steps will you take to accomplish the task? Each correct answer represents a part of the solution. Choose all that apply.

- A. Now migrate the Start Date attribute of "works on" relationship to the Project table.
- B. Create two tables, Employee and Project.
- C. Create a primary key for both the tables.
- D. Migrate the primary key of the Employee table to the Project table to maintain referential integrity between the two tables.
- E. Create three tables, Employee, Project, and Manager.

Correct Answer:

Section: (none)

Explanation

Explanation/Reference:

Answer: B, C, D, and


A

Explanation:

While converting the conceptual model to a logical model, the following steps are to be followed:

Convert the entities defined in the ER-diagram to tables. The attributes of entities are represented as columns in the table. In the above mentioned ER-diagram, there will be two tables, Employee and Project. Define the primary keys for both the tables, EmpID for the Employee table and ProjectID for the

Project table. Migrate EmpID, i.e. the primary key of the Employee table, to the Project table to maintain referential integrity constraint between the two tables. Migrate the attribute Start Date of the "works on" relationship in the same direction as the primary key, i.e. in Project table. The process is shown below as a logical diagram:



Note: The relational model does not support "relationship" data structure like the hierarchical model; therefore, the information that was defined in the relationship attribute should be migrated to an existing or associative entity. The Start Date attribute of "works on" relationships is therefore migrated to the Project table.



For interactive and self-paced preparation of exam 98-364, try our practice exams. Practice exams also include self assessment and reporting features! www.selftestengine.com



QUESTION 7

You work as the Database Administrator for www.company.com.com. All servers on the www.company.com.com network run Windows Server 2003 and Windows Server 2000, and all client computers run Windows XP professional and Windows Vista. The www.company.com.com network area consists of a single Active Directory domain named www.company.com.com. The www.company.com.com network contains a Microsoft SQL Server 2005 database server named COMPANY-DB111, which you administer at a regular interval of time. COMPANY-DB111 contains and hosts three databases that support different types of business requirements. The network uses the DBA database that supports internal operations including maintaining data, accounting, and mailing. The network's regular users make use of stored procedures for the data values that need insertion into the tables. The stored procedures used in the network of the company are designed to access



For interactive and self-paced preparation of exam 98-364, try our practice exams. Practice exams also include self assessment and reporting features! www.selftestengine.com



SQL variant type parameters and then use the values to build and execute ad hoc query statements that are the part of the SQL statements. During a routine network check, you discover that there are several odd occurrences in the database. The odd occurrences include data deleted from tables and other unauthorized activity. You suspect a user is executing the unauthorized statements through the stored procedures. You are required to stop the unauthorized changes while having least impact on the stored procedures that are in use. Your explanation must use the least administrative effort. What should you do?

- A. The input parameters should be parsed to watch for and block any input including single quotes.
- B. The stored procedures used for the table should be replaced with ad hoc queries.
- C. The stored procedure should be customized to use type-specific and length-restricted parameters.
- D. All data handling activity on the table should be audited and recorded.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

You should consider modifying the stored procedures to use type-specific and length-restricted SQL parameters if there are problems generated from a SQL Injection attack when the malicious code is inserted in a query for execution on the server.

Answer: D is incorrect. This choice could be used for identifying the cause, but it will not help in stopping the attacks.

Answer: A is incorrect. This choice could be used, but the solution is incomplete. Answer: B is incorrect. You should not consider this explanation, as it is even less secure and more prone to a SQL injection attack.

QUESTION 8

You work as a Database Administrator for InfoTech Inc. The company has a database and they want you to create a design for a database that will support the enterprise's operations and objectives. Which of the following will you use to accomplish the task?

- A. Application design
- B. Database design
- C. Database planning
- D. Requirements collection and analysis

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Database design is the process of creating a design for a database that will support the enterprise's operations and objectives for the required database system. It represents data and relationships between data required by all major application areas and user groups. Answer: C is incorrect. Database planning is a management activity that allows the stages of the database application to be realized as efficiently and effectively as possible. It identifies the work to be done and the resources with which to do it. Database planning is integrated with the overall IS strategy of the organization.

Answer: A is incorrect. Application design is a design of the user interface and the application programs that use and process the database. Database and application design are parallel activities. Application design includes two important activities: transaction design and user interface design. Answer: D is incorrect. Requirements collection and analysis is the process of collecting and analyzing information about the part of the organization that is to be supported by the database application, and using this information to identify the user's requirements of the new system.

QUESTION 9

Which of the following is a building working model of a database system?

- A. Prototyping
- B. Bottom-up approach
- C. Conceptual database design
- D. Top-down approach

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:



For interactive and self-paced preparation of exam 98-364, try our practice exams. Practice exams also include self assessment and reporting features!
www.selftestengine.com



Explanation:

Prototyping refers to the creation of a working model of the database system. The purpose of Prototyping is to suggest improvements or even add new features; to identify features of a system that work well, or are inadequate; to evaluate feasibility of a particular system design; to clarify a user's requirements.

Answer: D is incorrect. In a top-down approach, a schema is created through a series of successive refinements, starting from the first schema, which is gradually expanded by using appropriate modifications that increase the detail of the various concepts. Moving from one level to another, the schema is modified using some basic transformations called top-down transformation primitives.

Answer: B is incorrect. In a bottom-up approach, the preliminary specifications are broken down into smaller components so that each component describes a basic fragment of the specifications. These components are then represented by schemas, some of which may consist of single concepts. The schemas are later obtained to be integrated together until a final schema is obtained.

Answer: C is incorrect. Conceptual database design is the process of constructing a model for the data at a high level of abstraction

without using DBMS. It consists of the construction of an Entity-Relationship schema, providing an optimal description of the user requirements.

QUESTION 10

You are the database administrator for a MySQL database server. The network contains new and old (pre-4.1) clients. You configure another database server on the network. You allow the network clients to connect to the new server. Some users complain that when they try to connect to the server, they receive the following error:

ERROR 1251: The client does not support authentication protocol requested by the server; consider upgrading MySQL client

You do not want to upgrade any client or server right now, and want to enable all clients to connect to the new server successfully. Which of the following steps will you take to resolve the issue?

- A. Run the server with the --secure-auth option.
- B. Run the server with the --old-password option.
- C. Run the server with the --allow-old option.
- D. Run the server with the --enable-old option.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:



For interactive and self-paced preparation of exam 98-364, try our practice exams. Practice exams also include self assessment and reporting features!
www.selftestengine.com



Explanation:

According to the question, you are receiving Error 1251. Servers from MySQL 4.1 and later, use an enhanced authentication mechanism for providing greater security. This mechanism provides better passwords than in the older version.

When client programs try to connect to the server, they do not understand this mechanism and produce the above stated error. In order to resolve the issue, you will have to run the server with the --old-password option. Starting the server with this option will provide support to the older clients.

Answer: A is incorrect. This option will enforce the server to reject connection attempts for any client that has a password in the old format.

Answer: D and C are incorrect. No such options are available with the server.

QUESTION 11

Mark works as a Database Administrator for DataWorld Inc. The company has a SQL Server database. Management instructs him to ensure that no

inconsistent data is entered in the database. Which of the following will help him to accomplish the task?

- A. Referential integrity
- B. Abstraction
- C. Encryption
- D. Authentication
- E. Native auditing

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Referential integrity is a feature or property that prevents users or applications from entering inconsistent data which, when satisfied, requires every value of one column of a table to exist as a value of another column in a different or the same table.



For interactive and self-paced preparation of exam 98-364, try our practice exams. Practice exams also include self assessment and reporting features!
www.selftestengine.com



Answer: D is incorrect. Authentication is the process of verifying the identity of a user or software connecting to a database. During the process of authentication, the user must submit sufficient information validated either by the database itself, or by the operating system within which the database is installed.

Answer: B is incorrect. Abstraction is a method of coding that enables a user to focus on the coding rather than the syntax for a specific database API. A user can use generic methods to access a database. In order to perform database abstraction, a user requires additional code or libraries. The benefit of database abstraction is that it provides a uniform interface for database access.

Answer: E is incorrect. Native auditing is the process of extracting trails on a regular basis so that they can be transferred to a designated security system where the database administrators do not have access. This ensures a certain level of segregation of duties and provides evidence that the native audit trails were not modified by authenticated administrators. Answer: C is incorrect. Encryption provides an additional security layer, protecting the data from unauthorized viewing with the help of an algorithm called cipher. Even if access to the database is obtained, it will not be easy to decipher encrypted data into a readable form.

QUESTION 12

You work as a database developer for www.company.com.com. The company has a database named www.company.com that is hosted on a SQL Server 2008 server. The database includes a table named Employees, which contains the details of the employees. You want to ensure that anytime any data in the Employees table is modified with an UPDATE statement, they are recorded in a table named EmployeeChanges. The EmployeeChanges

table includes columns to record what the change was, who made the change, and when they made the change. What should you create to meet this need?

- A. A DDL trigger that queries the updated table
- B. A DML trigger that queries the inserted table
- C. A DDL trigger that queries the inserted table
- D. A DML trigger that queries the updated table

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

A Data Manipulation Language (DML) trigger can be used to capture DML statements (such as UPDATE, INSERT, and DELETE). DML triggers can use the inserted and deleted tables, which are virtual tables that hold data from the last INSERT, UPDATE, or DELETE statement.

An INSERT will record the inserted data in the inserted table. A DELETE will record the deleted data in the deleted table. An UPDATE will record the original row in the deleted table, and the modified row in the inserted table. The GetDate() function can be used to retrieve the current date.

Answer: D is incorrect. There is no such thing as an updated table. The inserted and deleted tables are used to record updated data.

Answer: A and C are incorrect. DDL triggers will capture Data Definition Language statements (such as CREATE, DROP, and ALTER) but will not capture statements that modify data.

QUESTION 13

Linda works as a Database Designer for Lion Inc. She has been given an assignment to design the database of a publishing company. The database has a table named Author, which has a composite key, AuthorID and TitleID. Royalty and LiteraryAgent are the other attributes. The functional dependencies are such that AuthorID + TitleID-> Royalty (i.e. Royalty is functionally dependent on AuthorID and TitleID) and AuthorID-> LiteraryAgent (i.e. LiteraryAgent is functionally dependent on AuthorID). Which of the following is true about this table?

- A. It violates 4NF.
- B. It violates 2NF.
- C. It violates 1NF.
- D. There is no violation.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

It violates the 2NF. According to 2NF, all the attributes that are not part of the primary key should be functionally dependent on the entire key attributes (composite key). Here, Literary Agent is functionally dependent on one of the composite key attributes, i.e. AuthorID. To make this in 2NF, the Author

table should be decomposed into two tables such that the first table contains AuthorID, TitleID, and Royalty as its attributes, and the second table contains AuthorID and LiteraryAgent as its attributes. This will make the two tables in 2NF.



For interactive and self-paced preparation of exam 98-364, try our practice exams. Practice exams also include self assessment and reporting features!
www.selftestengine.com



QUESTION 14

Which of the following is a management activity that allows the stages of the database application to be realized as efficiently and effectively as possible?

- A. Database planning
- B. Identifying information for objects
- C. Object identification
- D. Requirements collection and analysis

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Database planning is a management activity that allows the stages of the database application to be realized as efficiently and effectively as possible. It identifies the work to be done and the resources with which to do it. Database planning is integrated with the overall IS strategy of the organization.

Answer: B is incorrect. Identifying information for objects is the fourth step in database planning where the information for each object is stored. In a table, there are four columns that contain different types of information for objects:

1.Categorical column: In this column, data is classified and grouped, and a limited selection of data is stored.

2.Relational or referential column: In this column, a link is established between information in one table and related information in another table.

3.Identifier column: In this column, each item stored in the table is identified. 4.Raw data column: In this column, tangible information is stored and is determined by a source that is external to the database.

Answer: D is incorrect. Requirements collection and analysis is the process of collecting and analyzing information about the part of the organization that is to be supported by the database application, and using this information to identify the user's requirements of the new system.

Answer: C is incorrect. In the object identification process in database planning, it is necessary to identify the key objects or entities for database management. There are generally a few primary objects, and once these are identified, the related items become visible.



For interactive and self-paced preparation of exam 98-364, try our practice exams. Practice exams also include self assessment and reporting features!
www.selftestengine.com



QUESTION 15

You work as a Database Administrator for Bell Ceramics Inc. An employee of the company has fired a query, including a DML statement, such as INSERT, against a table named Sales. You notice that the DML statement has not executed. What will you do to resolve the issue? Each correct answer represents a complete solution. Choose two.

- A. Clean up all the uncommitted transactions on the Sales table.
- B. Add more space to the tablespace and increase the users quota on the tablespace.
- C. Provide redundant network paths from the client computer to the server along with additional listener connections on the Oracle server and redundant network cards on the Oracle server.
- D. Provide an appropriate privilege or create views on the Sales table, and grant privileges on the view.

Correct Answer:

Section: (none)

Explanation

Explanation/Reference:

Answer: D and B

Explanation:

If an Oracle database fails due to any DML statements issued in a database, then the type of failure is categorized under statement failures. To overcome the statement failures, appropriate privileges on the table and views should be provided. The issue can also be resolved by adding more space to the tablespace or by increasing the users' quota on the tablespace. Answer: A is incorrect. Cleaning up all the uncommitted transactions in the Sales table will not resolve the issue if an Oracle database faces inconsistencies due to statement failures.

Answer: C is incorrect. The redundant network paths from the client computer to the server along with additional listener connections on the Oracle server and the redundant network cards on the Oracle server will resolve the issue only if the Oracle database fails due to network inconsistencies.

QUESTION 16

You work as a Database Developer for GraySoft Inc. You are designing a database to record information about potential customers. You create a table named ProductInquiry for the database.

The table is given below:

ProductInquiry	
<input type="checkbox"/>	InquiryID
<input type="checkbox"/>	ProductID
<input type="checkbox"/>	LastInquiredOn
<input type="checkbox"/>	SalesPersonID

You want to uniquely identify the product in which a customer is interested. Choose the columns on which a PRIMARY KEY constraint needs to be created to accomplish this.



For interactive and self-paced preparation of exam 98-364, try our practice exams. Practice exams also include self assessment and reporting features!
www.selftestengine.com



Primary Key Column(s)

All Columns

ProductID

LastInquiredOn

SalesPersonID

InquiryID

- A.
- B.
- C.

D.

Correct Answer:

Section: (none)

Explanation

Explanation/Reference:

Primary Key Column(s)

Primary Key Column(s)
InquiryID
ProductID

All Columns

All Columns
SalesPersonID
LastInquiredOn

Explanation: Create a PRIMARY KEY constraint on the InquiryID and ProductID columns to uniquely identify the product in which a customer is interested.

The ProductInquiry table is working as a junction table to create a many-to-many relationship between the Products table and the Inquiries table. A customer can be interested in many products, while a product can also be inquired by many customers. However, the inquiry information of one product by a particular customer should be stored only once. For the next subsequent inquiry for the same product by the same customer, only the LastInquiredOn column should be updated. The correct entity relationship diagram for this scenario is given below:



For interactive and self-paced preparation of exam 98-364, try our practice exams. Practice exams also include self assessment and reporting features!
www.selftestengine.com



QUESTION 17

Which of the following is the process of extracting trails on a regular basis so that they can be transferred to a designated security system where the database administrators do not have access?

- A. Native auditing
- B. Abstraction
- C. Database security
- D. Encryption

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Native auditing is the process of extracting trails on a regular basis so that they can be transferred to a designated security system where the database administrators do not have access. This ensures a certain level of segregation of duties and provides evidence that the native audit trails were not modified by authenticated administrators. Answer: B is incorrect. Abstraction is a method of coding that enables a user to focus on the coding rather than the syntax for a specific

database API. A user can use generic methods to access a database. In order to perform database abstraction, a user requires additional code or libraries. The benefit of database abstraction is that it provides a uniform interface for database access.

Answer: C is incorrect. Database security refers to the system, processes, and procedures that protect a database from unintended activities, such as authenticated misuse, malicious attacks, or inadvertent mistakes made by authorized individuals or processes. Database security can begin with the process of creating and publishing appropriate security standards for the database environment. These standards may include specific controls for the various relevant database platforms. Answer: D is incorrect. Encryption provides an additional security layer, protecting the data from unauthorized viewing with the help of an algorithm called cipher. Even if access to the database is obtained, it will not be easy to decipher encrypted data into a readable form.

QUESTION 18

You manage a large database named Sales. The Sales database contains a table named OrderDetails, which is a heavily transacted table with frequent inserts. Indexes in the table often become fragmented due to excessive page splitting. You want to minimize the amount of fragmentation due to page splits. What should you do?



For interactive and self-paced preparation of exam 98-364, try our practice exams. Practice exams also include self assessment and reporting features!
www.selftestengine.com



- A. Update the statistics on the indexes.
- B. Change the fillfactor for the indexes to 100.
- C. Change the fillfactor for the indexes to 60.
- D. Change the fillfactor for the indexes to 0.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

FILLFACTOR specifies a percentage that indicates how much free space will be in the leaf level of each index page. When a clustered index is created on a table, SQL Server does not stuff data wherever it finds space, but it physically rearranges data in order. SQL Server cannot rearrange data without page split if it does not find free space at an index page. To help SQL Server accomplish this, there is a need to leave a little space at the leaf level of each page on a clustered index. This free space is called FILLFACTOR. The fillfactor determines how much empty space is left in the indexes. Once a page is full, it must split to accommodate additional data.

However, excessive splitting causes the indexes to be fragmented and can significantly impact performance. The fillfactor can be modified with a statement like this:

`ALTER INDEX ALL ON dbo.OrderDetails REBUILD WITH (FILLFACTOR = 60);` FILLFACTOR =60 tells SQL Server to fill the page to 60 percent full with 30 percent free space. Answer: D and B are incorrect. Fillfactor values of 0 and 100 are both the same. This indicates that the indexes fill the pages with only

one empty row. As soon as a second insert is performed, the page will split. Answer: A is incorrect. Statistics are used by the database engine to determine how useful an index may be for a query, but updating the index would not reduce the fragmentation.

QUESTION 19

You are the database administrator of a MySQL server that runs on a Windows server. All clients are local clients. For security, you want to disable connections from the remote clients. Which of the following steps will you take to accomplish the task?

- A. Start the server with the --disable-networking option.
- B. Start the server with the --shared-memory option.
- C. Start the server with the --skip-networking option.
- D. Start the server with the --secure-auth option.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:



For interactive and self-paced preparation of exam 98-364, try our practice exams. Practice exams also include self assessment and reporting features!
www.selftestengine.com






Explanation:

According to the question, all the clients are local and you want to disable connections from the remote clients. In order to accomplish this task, you will have to start the server with the --skip-networking option. This will disable the TCP/IP connections (the only type of connection that can be made by the remote clients).

Answer: B and D are incorrect. These options will not disable remote connections. Answer: A is incorrect. No such option is available with a MySQL server.

QUESTION 20

Consider the case of an organization having many departments. Each department has a department head known as senior manager. Several employees are working in a department under the senior manager. Identify the most appropriate ER-diagram that represents the relationship between department and employees.



- A.
- B.
- C.
- D.

Correct Answer:

Section: (none)

Explanation

Explanation/Reference:



For interactive and self-paced preparation of exam 98-364, try our practice exams. Practice exams also include self assessment and reporting features!
www.selftestengine.com



Explanation:

The relationship between department and employees is one-to-many and that of employees and department is one-to-one, i.e. an employee can work in only one department at a time.

In the figure, a single line across the right end of the relationship represents one-to-one cardinality. The crowfoot at the left end of the relationship signifies an optional cardinality of any number of, or several.

Answer: C is incorrect. A single line across the left end of the relationship represents one-to-one cardinality. The crowfoot at the right end of the relationship signifies an optional cardinality of any number of, or several. According to this figure, an employee can work in several departments and a department has a single employee, which is not possible. Answer: B is incorrect. A single line between the two entities means one-to-one cardinality. Answer: A is incorrect. A single line across the left end of the relationship represents one-to-one cardinality. However, in this case, a department can have several employees, not just one.

QUESTION 21

Which of the following are DML commands? Each correct answer represents a complete solution.

Choose all that apply.

- A. SELECT INTO
- B. DELETE
- C. UPDATE
- D. ALTER
- E. SELECT
- F. INSERT

Correct Answer:

Section: (none)

Explanation

Explanation/Reference:

Answer: A, C, B, and F

Explanation:

A data-manipulation language (DML) is a family of languages to access or manipulate data that is organized by the appropriate data model.

DML is used for the following operations:

Retrieval of information stored in the database.

Insertion of new information into the database.

Deletion of information from the database.

Modification of information stored in the database.

Currently the most popular data manipulation language is that of SQL, which is used to retrieve and manipulate data in a Relational database.

In SQL, the following statements are used for database manipulation:

SELECT ... INTO

INSERT

UPDATE

DELETE

Each SQL DML statement is a declarative command.

Answer: E is incorrect. The purely read-only SELECT query statement is classed with the 'SQL-data' statements and hence it is considered to be outside of DML. The SELECT ... INTO form is considered to be DML because it manipulates or modifies data. Answer: D is incorrect. ALTER is a DDL command.



For interactive and self-paced preparation of exam 98-364, try our practice exams. Practice exams also include self assessment and reporting features!
www.selftestengine.com



QUESTION 22

Which of the following commands can be used to change the authorization type in DB2 9? Each correct answer represents a complete solution. Choose all that apply.

- A. UPDATE DATABASE CONFIGURATION
- B. SET AUTHORISATION
- C. RESET DATABASE CONFIGURATION
- D. UPDATE DBM CFG

Correct Answer:

Section: (none)

Explanation

Explanation/Reference:

Answer: D and A

Explanation:

The UPDATE DBM CFG or UPDATE DATABASE CONFIGURATION command can be used to change the authorization type in DB2. Both UPDATE DBM CFG, UPDATE DATABASE CONFIGURATION are equivalent commands and used to put the desired changes on the DB2 9 database configuration file.

Answer: B is incorrect. This is not a valid DB2 command. Answer: C is incorrect. It is used to reset database configuration.

QUESTION 23

Smith works as a Database Administrator for DWorlds Inc. The management instructs him to plan a database where identifying the key objects or entities for database management is important. Which of the following database planning steps will help him to accomplish the task?

- A. Object identification
- B. Gathering information
- C. Identifying the relationships between objects
- D. Object modeling

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

In the object identification process in database planning, it is necessary to identify the key objects or entities for database management. There are generally a few primary objects, and once these are identified, the related items become visible.

Answer: B is incorrect. Gathering information is the first step of database planning and requires a good understanding of what tasks are



For interactive and self-paced preparation of exam 98-364, try our practice exams. Practice exams also include self assessment and reporting features!
www.selftestengine.com



to be performed in the database. In case the database has to replace a paper-based or a manual system, the old system can provide all the necessary information. Moreover, all those who were involved with the previous system should be interviewed to determine what they did and what is now needed for the new database. The limitations, problems, and bottlenecks of the old system should also be identified.

Answer: D is incorrect. Object modeling is the third step in database planning in which objects are identified in the system. A user should record an object in such a way that it is represented visually in the system. He can use the database model as a reference during implementation of the database. Database developers use tools such as pencils, paper, word processing etc., as well as create software

programs for data modeling of the database.

Answer: C is incorrect. Identifying relationships between objects is the last step in database planning, which is the ability to relate information about various items in the database. Isolated types of information can be stored separately, but the data can be combined according to the requirements of the database engine. Identifying the relationship between objects in the design process necessitates looking at the tables, determining how the tables are logically related, and adding relational columns to the tables that establish a link from one table to another.

QUESTION 24

Which of the following provides reliable units of work that allow correct recovery from failures and keeps a database consistent even in cases of system failure?

- A. Database security
- B. Concurrency control
- C. Two-phase commit
- D. Database transaction

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

A database transaction comprises a unit of work performed within a database management system against a database, and treated in a coherent and reliable way independent of other transactions. Transactions in a database environment have two main purposes:

1.To provide reliable units of work that allow correct recovery from failures and to keep a database consistent even in cases of system failure, such as when execution stops and many operations in a database transaction remain uncompleted.

2.To provide isolation between multiple programs accessing a database concurrently. Answer: B is incorrect. Concurrency control in database management systems ensures that database transactions are performed concurrently without the concurrency violating the data integrity of a database. Executed transactions should follow the ACID rules. This mechanism is required to manage multiple users accessing the same groups of resources with the least possible overhead, either optimized



For interactive and self-paced preparation of exam 98-364, try our practice exams. Practice exams also include self assessment and reporting features!
www.selftestengine.com



for a specific task, or general purpose.

Answer: A is incorrect. Database security refers to the system, processes, and procedures that protect a database from unintended

activities, such as authenticated misuse, malicious attacks, or inadvertent mistakes made by authorized individuals or processes. Database security can begin with the process of creating and publishing appropriate security standards for the database environment. These standards may include specific controls for the various relevant database platforms. Answer: C is incorrect. A two-phase commit is a feature of transaction processing systems, enabling databases to be returned to the pre-transaction state if some error condition occurs. A single transaction can update many different databases. A two-phase commit strategy is designed to ensure that either all the databases are updated or none of them, so that the databases remain synchronized.

QUESTION 25

Which of the following are the types of prototyping strategies? Each correct answer represents a complete solution. Choose two.

- A. Requirements prototyping
- B. Physical prototyping
- C. Revolutionary prototyping
- D. Evolutionary prototyping

Correct Answer:

Section: (none)

Explanation

Explanation/Reference:

Answer: D and A

Explanation:

There are two types of prototyping strategies which are as follows:

Requirements prototyping: It is used to determine the requirements of a proposed database system and then the prototype is discarded.

Evolutionary prototyping: It is used for the same purpose, but the prototype is not discarded and with further development becomes the working database system.

QUESTION 26

Which of the following commands cannot be rolled back? Each correct answer represents a complete solution. Choose two.

- A. DELETE
- B. TRUNCATE
- C. UPDATE
- D. COMMIT

Correct Answer:

Section: (none)

Explanation

Explanation/Reference:

Answer: D and B

Explanation:

COMMIT statement is a transaction control statement that ends the current database transaction and makes all changes made in that transaction permanent. It also releases all locks on the transaction. The TRUNCATE statement is a Data Definition Language (DDL) statement that is used to remove all records from a table. It also releases the storage space used by the table. The TRUNCATE statement cannot be rolled back.

Syntax:

TRUNCATE TABLE TableName;

Therefore, both the commands, COMMIT and TRUNCATE, can never be rolled back. Note: The TRUNCATE command includes a built-in COMMIT command. Answer: C and A are incorrect. The SQL UPDATE statement allows a user to update the current values in a database with new values. It contains three main clauses. The first clause specifies the table that requires an update, the second clause specifies the column name and the data values that will be inserted in place of the current values, and the third clause specifies the criteria for the records that will be updated. An example query is given below:

Update Employee{first clause}

Set Fname="Roger"{second clause}

Lname="Pressman"



For interactive and self-paced preparation of exam 98-364, try our practice exams. Practice exams also include self assessment and reporting features!
www.selftestengine.com



WHERE Empid=1{third clause}

The DELETE statement is a Data Manipulation Language (DML) statement that is used to delete records from a table, based on the condition specified in the WHERE clause. If the WHERE clause is omitted, all records in the table are deleted.

The DELETE statement can be rolled back.

Syntax:

DELETE [FROM] TableName

[WHERE condition];

The above two DML commands can therefore be rolled back.

QUESTION 27

What is the ideal time to back up dynamic log files?

- A. When the server is running.
- B. You can back up dynamic log files any time.

- C. When the server is stopped.
- D. When the server is shut down.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Dynamic files such as the logs that the server changes as it runs are best backed up when the server is stopped.

Answer: D, C, and A are incorrect. These are not the appropriate statements for backing up dynamic log files.



For interactive and self-paced preparation of exam 98-364, try our practice exams. Practice exams also include self assessment and reporting features!
www.selftestengine.com



QUESTION 28

Two tables, namely, OrderItem and Item are depicted below in which all attributes mentioned therein respectively, refer to column heading labels.

OrderItem	
<u>OrderNumber(PK)</u>	
<u>ItemNumber(FK)</u>	
QuantityOrdered	
Price	
TotalAmount	

Item	
<u>ItemNumber(PK)</u>	
ItemName	
ItemPrice	

If the ItemName attribute is moved from the Item table to the OrderItem table, then will it violate any normalization form? If yes, then which of the following normalization forms will be violated?



For interactive and self-paced preparation of exam 98-364, try our practice exams. Practice exams also include self assessment and reporting features!
www.selftestengine.com



- A. It will not violate any normal form (NF).
- B. 1NF
- C. 3NF
- D. 2NF

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

It will violate 2NF, as ItemName is functionally dependent on ItemNumber, which is the primary key of the Item table. ItemNumber is the foreign key of the OrderItem table, and OrderNumber is the primary key of the OrderItem table. According to 2NF, all the non-key attributes in a table must be functionally dependent on the primary key column. Migrating ItemName to the OrderItem table will violate 2NF, since ItemName (a non-key attribute) is not functionally dependent on OrderNumber, which is the primary key of the OrderItem table.

QUESTION 29

You are the database administrator for a MySQL database. The database server is installed on a Unix system. The time zone files for the system are located at /usr/share/zoneinfo. You need to ensure that the system and MySQL time zones are based on the same information. Which of the following statements will you use to accomplish the task?

- A. shell> /usr/share/zoneinfo mysql_tzinfo_to_sql | mysql -u root mysql
- B. shell> /usr/share/zoneinfo mysql_tz_to_sql | mysql -u root mysql
- C. shell> mysql_tz_to_sql /usr/share/zoneinfo | mysql -u root mysql
- D. shell> mysql_tzinfo_to_sql /usr/share/zoneinfo | mysql -u root mysql

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

According to the question, you are required to ensure that the system and MySQL time zones are based on the same information. To ensure this, on operating systems that have their own time zone files, it is best to use them for loading the MySQL time zone tables. Furthermore, the question also provides information that the system time zone files are located at /usr/share/zoneinfo location. In order to accomplish this task, you will have to use the mysql_tzinfo_to_sql program to convert the file contents into SQL statements that can be loaded into MySQL by the mysql program. Use the following command for this:

shell> mysql_tzinfo_to_sql /usr/share/zoneinfo | mysql -u root mysql Answer: A is incorrect. Wrong syntax is used.
Answer: C and B are incorrect. No such command is available with MySQL.



For interactive and self-paced preparation of exam 98-364, try our practice exams. Practice exams also include self assessment and reporting features!
www.selftestengine.com



QUESTION 30

You work as a Database Designer for Hitech Inc. You are designing a database named Project. The company takes projects from outside and divides them into tasks (design, coding, etc.). Each task is assigned to only one employee. You want to use a primary key constraint to uniquely define each row in a table named Task. The tables are given below:

Employee : Table	
	Field Name
	EmployeeID
	SocialSecurityNumber
	Name
	JoiningDate

Task : Table	
	Field Name
	TaskID
	ProjectID
	EmployeeID
	SocialSecurityNumber

Project : Table	
	Field Name
	ProjectID
	ProjectName
	ProjectStartingDate
	ProjectEndDate
	EmployeeID

Which of the following attributes can be a primary key in the Task table?

- A. TaskID
- B. ProjectID
- C. EmployeeID
- D. SocialSecurityNumber (SSN)

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The TaskID attribute uniquely defines each row in the Task table. It is given in the scenario that one task is assigned to only one employee. Therefore, in order to avoid a task being given to two employees, TaskID should be entered just once (unique).

QUESTION 31

Which of the following is the process of creating a design for the database that will support the enterprise's operations and objectives for the required database system?

- A. Identifying relationships between objects
- B. Application design
- C. Database planning
- D. Database design

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:



For interactive and self-paced preparation of exam 98-364, try our practice exams. Practice exams also include self assessment and reporting features!
www.selftestengine.com



Explanation:

Database design is the process of creating a design for a database that will support the enterprise's operations and objectives for the required database system. It represents data and relationships between data required by all major application areas and user groups.

Answer: C is incorrect. Database planning is a management activity that allows the stages of the database application to be realized as efficiently and effectively as possible. It identifies the work to be done and the resources with which to do it. Database planning is integrated with the overall IS strategy of the organization.

Answer: B is incorrect. Application design is a design of the user interface and the application programs that use and process the database. Database and application design are parallel activities. Application design includes two important activities: transaction design and user interface design.

Answer: A is incorrect. Identifying relationships between objects is the last step in database planning, which is the ability to relate

information about various items in the database. Isolated types of information can be stored separately, but the data can be combined according to the requirements of the database engine. Identifying the relationship between objects in the design process necessitates looking at the tables, determining how the tables are logically related, and adding relational columns to the tables that establish a link from one table to another.

QUESTION 32

You are designing a database for a sports hostel. In the hostel, a player can be enrolled for multiple games. You are working on a relation shown in the image below:

Game_ID	P_ID	P_Age	P_DOB
001	3011	25	2-2-1991
001	3012	23	3-4-1993
003	5078	24	3-3-1992
004	3011	25	2-2-1991
005	7023	24	12-8-1992

Which of the following candidate keys would best serve as the primary key?

- A. P_ID and P_DOB
- B. Game_ID and P_ID
- C. Game_ID and P_Age
- D. P_ID and P_Age

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:



For interactive and self-paced preparation of exam 98-364, try our practice exams. Practice exams also include self assessment and reporting features!
www.selftestengine.com



Explanation:

According to the question, you are required to identify the keys that would best serve as the primary key for the given relation. The primary key must be unique for each row. The only combination of the candidate keys that makes it unique for each row of the table is Game_ID and

P_ID.

QUESTION 33

You work as a Database Administrator for Dowtech Inc. A table named StudentInformation exists in a database. The table has columns named student_names, student_marks, and student_courses. Some students who have opted for various courses have got different marks in them. This implies that the table has two many-to-many relationships. You have to ensure that these two relationships are not represented in a single table and that the student_name column exists in both tables. Which of the following normal forms will you use to represent these relationships in the tables?

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

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Fourth normal form (4NF) is a normal form used in database normalization. Introduced by Ronald Fagin in 1977, 4NF is the next level of normalization after Boyce-Codd normal form (BCNF). Whereas the second, third, and Boyce-Codd normal forms are concerned with functional dependencies, 4NF is concerned with a more general type of dependency known as a multivalued dependency. Symbolically, it can be represented as follows:

If $A \twoheadrightarrow B|C$, then $A \twoheadrightarrow B$ and $A \twoheadrightarrow C$

Here, A, B, and C are attributes.

Answer: D is incorrect. First normal form exists when all the columns in a table are atomic i.e., only a single value is allowed in each column. Also, there should not be a group value like more than one phone number in a phone number column. If a table violates the first normal rule, following steps should be taken to normalize the database:

To make column atomic, divide them into multiple columns. If there is group value, create separate tables for the column and relate new table with a foreign key. Answer: C is incorrect. Second normal form (2NF) is used in database normalization. A table that is in first normal form must meet additional criteria if it is to qualify for second normal form. Specifically, a 1NF table is in 2NF if and only if, given any candidate key and any attribute that is not a constituent of a candidate key, the non-key attribute depends upon the whole of the candidate key rather than just a part of it.

Answer: B is incorrect. Third normal form (3NF) is used in database normalization. A table is in 3NF if and only if the relation S (table) is in second normal form (2NF) and every non-prime attribute of S is non-transitively dependent on every key of S.



For interactive and self-paced preparation of exam 98-364, try our practice exams. Practice exams also include self assessment and reporting features!
www.selftestengine.com

**QUESTION 34**

Which of the following is an ANSI approved language for communicating and managing data in a relational database management system (RDBMS)?

- A. SQL
- B. QL
- C. DDL
- D. XML

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Structured Query Language (SQL) is used to communicate and manage data in a relational database management system (RDBMS).

Structured Query Language (SQL) is the most common language used to access relational databases.

SQL was developed by IBM in the

1970's, and SQL version 3 was officially defined by the American National Standards Institute (ANSI).

Most DBMS's conform to the ANSI SQL-92

standard. It is a nonprocedural language and supports many database systems that are in use today, such as MySQL, PostgreSQL, Berkeley DB, Oracle, DB2, Sybase, Informix, Microsoft SQL, etc. Answer: C is incorrect. Data Definition Language (DDL) is a part of SQL that allows a user to create objects (such as tables, stored procedures, indexes, and triggers) within a database. It also allows a user to delete already created objects.

Answer: B is incorrect. Query Language (QL) is a set of commands used to retrieve, modify, and update data from the database; it is a language for formulating queries for a given dataset.

Answer: D is incorrect. Extensible Markup Language (XML) is a general-purpose specification for creating custom markup languages. It is classified as an extensible language because it allows the user to define the mark-up elements. It is used to aid information systems in sharing structured data, especially via the Internet, to encode documents, and to serialize data. XML is recommended by the World Wide Web Consortium (W3C). It is a fee-free open standard. The recommendation specifies lexical grammar and parsing requirements.



For interactive and self-paced preparation of exam 98-364, try our practice exams. Practice exams also include self assessment and reporting features!
www.selftestengine.com



QUESTION 35

Examine the following table:

P_ID	First_Name	Last_Name	Score
001	Sharon	Hayes	25
002	James	Welch	23
003	Rick	Stuart	24
004	Tim	May	25
005	Mark	Smith	24

What is the degree of this table?

- A. 20
- B. 25
- C. 5
- D. 4

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The question being asked is the degree of the displayed table. The degree of a table refers to the number of columns that a table has. The number of columns contained in a table is referred to by the term degree of the table; if the degree of a table is increased by 1, it means that a new column is added to the table. The degree of the table denotes the number of columns or attributes. Since the table displayed has four attributes, the degree of this table is four (4). Answer: C, A, and B are incorrect. These are not the appropriate values for the degree of the table displayed.

QUESTION 36

A company named Rel Inc. has many authorized dealers across the country who sell their products. The Sales Manager of the company wants to see the details of the authorized dealers, including the name, region, and total sales in ascending order of sales. Which of the following queries should be issued to get the desired output?

- A. `SELECT MAX (Totalsales)`
`FROM Dealer`
- B. `SELECT`
`FROM Dealer`
- C. `SELECT *`
`FROM Dealer`
`WHERE MAX (Totalsales)`

D. SELECT Name, Region, Totalsales
FROM Dealer
ORDER BY Totalsales

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:



For interactive and self-paced preparation of exam 98-364, try our practice exams. Practice exams also include self assessment and reporting features!
www.selftestengine.com



Explanation:

The first part of the query (i.e. SELECT Name, Region, Totalsales) selects the Name, Region, and Totalsales attributes. The second part (i.e. FROM Dealer) specifies the table name, and the third part (i.e. ORDER BY Totalsales) gives Totalsales in ascending order. In the ORDER BY clause, ascending is the default order.

Answer: B is incorrect. The SELECT clause does not specify the attributes. Answer: A is incorrect. The query fails to provide the details from the table. It will provide only the maximum sales from the table.

Answer: C is incorrect. MAX is an aggregate function. An aggregate function cannot be in the WHERE clause.

QUESTION 37

You work as the Database Administrator for a MySQL database server. In an hour or so, you are able to bring a corrupted server online. You execute the CHECK TABLE command and find that some of the InnoDB tables need to be repaired. You restart the server and discover that auto-recover has failed. Which of the following steps will you take to recover the corrupted InnoDB tables?

- A. Run the ANALYZE TABLE command.
- B. Run the REPAIR TABLE command.
- C. Recover the tables from the last backup.
- D. Run the OPTIMIZE TABLE command.
- E. Restart the server with the innodb_force_recovery option.

Correct Answer: E

Section: (none)

Explanation

Explanation/Reference:

Explanation:

According to the question, the corrupted InnoDB tables are not auto-recovered. In such a case, you will have to restart the server with the innodb_force_recovery option. Further, you can dump the table, restart the server normally, and recover the tables from the dump file.

QUESTION 38

You have developed a stored procedure named usp_GetEmp that accepts an employee number as a parameter and retrieves the details about the employee from the CurrentEmp table of a database named Employees. You have tested it, and it works exactly as you expected. Later, another employee tries to use the stored procedure and receives the following error: "The SELECT permission was denied on the object 'CurrentEmp', database 'Employees.schema 'dbo' ".

What should you do to resolve the problem?



For interactive and self-paced preparation of exam 98-364, try our practice exams. Practice exams also include self assessment and reporting features! www.selftestengine.com



- A. Modify usp_GetEmp to include the With Execute As Owner clause.
- B. Modify usp_GetEmp to include the With Execute As Caller clause.
- C. Grant the employee the SELECT permission on the CurrentEmp table.
- D. Grant the employee the SELECT permission on the Employees database.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The Execute As OWNER clause can be used to identify what permissions a stored procedure will have when it is executed. The Execute As Owner clause results in the stored procedure running with the same permissions as the owner of the stored procedure, or the account that created the stored procedure. The EXECUTE AS permission is placed in the line right after the CREATE PROCEDURE line as follows:

CREATE PROCEDURE

WITH EXECUTE AS

Since you created the stored procedure and it worked when you tested it, it will work if it is executed with your permissions.

Answer: C is incorrect. While it may be possible to grant the SELECT permission to the table, it would not help if another user executes the stored procedure. The next user will have the same problem. Additionally, a common method of protecting databases is not to grant access to the tables directly, but instead grant access via stored procedures or views. Answer: B is incorrect. The SELECT permission is granted to tables or views, not entire databases.

QUESTION 39

Which of the following statements is true regarding a composite key in a table?

- A. It is a primary key.
- B. It is a primary or foreign key that consists of two or more attributes of the table.
- C. It is a combination of the primary key and the foreign key.
- D. It is a foreign key.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:



For interactive and self-paced preparation of exam 98-364, try our practice exams. Practice exams also include self assessment and reporting features!
www.selftestengine.com



Explanation:

A composite key is a primary or foreign key that consists of two or more attributes of the table. It is a key that consists of two or more attributes that uniquely identify an entity occurrence. Answer: C, D, and A are incorrect. These statements do represent a composite key.

QUESTION 40

You are a database developer for a database named Sales hosted on a SQL Server 2008 server. The Sales database includes the Exams and ExamQuestions tables. If any questions are modified for an exam, you need to ensure that the ExamModifiedDate in the Exams table is updated with the date and time of the update. The ExamQuestions table includes an INSTEAD OF trigger that records when and who last modified each question. Therefore, you decide to create a trigger. You only want the trigger to fire if a constraint violation has not occurred. Which of the following code will create the trigger to meet your goals?

- A.

```
CREATE TRIGGER trgExamQuestionsModified
ON dbo.ExamQuestions
INSTEAD OF UPDATE NOT FOR REPLICATION
AS
UPDATE Exams
SET ExamModifiedDate = GetDate()
FROM inserted
WHERE inserted.ExamID = Exams.ExamID
```

- B. CREATE TRIGGER trgExamQuestionsModified
ON dbo.ExamQuestions
AFTER UPDATE NOT FOR REPLICATION
AS
UPDATE Exams
SET ExamModifiedDate = GetDate()
FROM inserted
WHERE inserted.ExamID = Exams.ExamID
- C. CREATE TRIGGER trgExamQuestionsModified
ON dbo.ExamQuestions
AFTER UPDATE FOR REPLICATION
AS
UPDATE Exams
SET ExamModifiedDate = GetDate()
FROM inserted
WHERE inserted.ExamID = Exams.ExamID
- D. CREATE TRIGGER trgExamQuestionsModified
ON dbo.Exams
AFTER UPDATE NOT FOR REPLICATION
AS
UPDATE Exams
SET ExamModifiedDate = GetDate()
FROM inserted
WHERE inserted.ExamID = Exams.ExamID

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:



For interactive and self-paced preparation of exam 98-364, try our practice exams. Practice exams also include self assessment and reporting features!
www.selftestengine.com



Explanation:

The syntax The DML trigger is created on the ExamQuestions table as it should be. AFTER UPDATE is used to ensure that the trigger will only fire after constraints are checked and succeed. In other words, if the UPDATE fails because it does not meet the constraint checks,

the trigger will not fire. In contrast, the INSTEAD OF trigger will fire even if the UPDATE would violate constraints since constraints are not checked.

While replication is not specifically mentioned in the question, you can specify NOT FOR REPLICATION if the table(s) are involved in replication. The NOT FOR REPLICATION clause indicates that the trigger should not be executed when a replication agent modifies the table that is involved in the trigger.