

Pass Microsoft 98-364 Exam

Number: 98-364
Passing Score: 800
Time Limit: 120 min
File Version: 26.4



Pass Microsoft 98-364 Exam

Exam Name: Database Fundamentals

Examsheets**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 $A \rightarrow\!> B|C$, 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**

You manage a database named Customers, which includes a table named Orders. The Orders table is frequently queried, but only orders with a sales total of more than \$1000.00 are required in the query. You want to create an index to speed up these types of queries at the same time, ensuring the index is as small as possible. What type of index should you use?

- A. Non-clustered
- B. Filtered
- C. Clustered
- D. XML

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

A filtered index is used to filter out unwanted data in the index. These indexes are ideally suited to improve queries on data that is mostly NULL values. A filtered query only indexes the rows that have non-null data in the indexed column. The filtered index would be defined as a non-clustered index. Filtered indexes consume less space in the system, and also require less processing power to query. Moreover, filtered indexes are mainly created on columns configured as sparse columns. In this case, the filter within the CREATE INDEX statement could look something like WHERE SalesTotal > 1000.

Answer: C is incorrect. A clustered index is organized as a B-tree structure. The index consists of a root page, intermediate levels, and leaf levels. The leaf level nodes contain the data pages of the underlying table. The root and intermediate level nodes contain index pages that hold index rows. Each index row contains a key value and a pointer. A clustered index is useful when you want to optimize queries based on ranges since it orders the data based on the clustered index. However, the clustered index would include all of the rows and will never be smaller than a filtered index.



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: A is incorrect. A nonclustered index has the same B-tree structure as the clustered index. The index consists of a root page, intermediate levels, and a leaf level. The leaf level of a nonclustered index does not contain the actual data. It contains pointers to the data that is stored in the data pages. A nonclustered index does not physically rearrange the data. A non-clustered index can be used to improve performance of queries. However, a non-clustered index would include all of the rows and will never be smaller than a filtered index. While a filtered index must be a non- clustered index, all non-clustered indexes do not include WHERE clauses and are not filtered indexes. Answer: A is incorrect. A nonclustered index has the same B-tree structure as the clustered index. The index consists of a root page, intermediate levels, and a leaf level. The leaf level of a nonclustered index does not contain the actual data. It contains pointers to the data that is stored in the data pages. A nonclustered index does not physically rearrange the data. A non-clustered index can be used to improve performance of queries. However, a non-clustered index would include all of the rows and will never be smaller than a filtered index. While a filtered index must be a non- clustered index, all non-clustered indexes do not include WHERE clauses and are not filtered indexes. Answer: D is incorrect. The XML column provides great ability to retrieve data from within the XML column using various query methods. To support the querying of XML data, users can create indexes on these columns. There are four different types of XML indexes supported by SQL Server 2005/2008. There is a single primary XML index and three different flavors of secondary XML indexes. The primary XML index is a clustered index on an internal table known as the node table that users cannot use directly from their T-SQL statements. The primary XML index basically contains one row for each node in the XML instance. After the primary XML index has been created, users can create supplementary three kinds of secondary XML indexes. The secondary XML indexes help in specific types of XQuery processing. These are called the PATH, PROPERTY, and VALUE indexes. An XML index is only appropriate for XML data.

QUESTION 5

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 6

Which of the following are the main approaches in the database design? Each correct answer represents a complete solution. Choose three.

- A. Top-down approach
- B. Bottom-up approach
- C. Middle approach
- D. Inside-out approach

Correct Answer: ABD

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:

There are four different types of approaches in database design, which are as follows:

- 1.Top-down approach
- 2.Bottom-up approach
- 3.Inside-out approach
- 4.Mixed approach

QUESTION 7

You work as a Software Developer for ArchiTech Inc. You plan to create a well-designed database for the company. You implement normalization rules for designing the database. Choose the appropriate normalization rules from the Choose and Reorder window pane.



The screenshot shows a software interface for choosing normalization rules. It consists of two main panes:

- Normalization Rules**: An empty list box with scroll bars on the right.
- Rules**: A list box containing the following items:
 - A table in a database should avoid containing nullable columns.
 - A table in a database should avoid creating candidate key.
 - A table in a database should have a unique row identifier.
 - A table in a database should store information for a single entity.
 - A table in a database should avoid containing too many columns.
 - A table in a database should not contain repeating values.

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

Correct Answer:

Section: (none)
Explanation

Explanation/Reference:

Normalization Rules

A table in a database should have a unique row identifier.
A table in a database should store information for a single
A table in a database should avoid containing nullable colu
A table in a database should not contain repeating values

Rules

A table in a database should avoid containing too many cc
A table in a database should avoid creating candidate key

Explanation:

The following are some normalization rules that can be made for achieving a well-designed database: Each table in a database should have a unique row identifier, i.e. an ID column in which no two records can share the same ID value. Such column is known as the primary key column that serves as the unique row identifier for the table. For example, ProductID can be the primary key column for the Products table. Each table in a database should store information for a single type of entity. Too much information stored in one table can hinder the efficiency and reliability of data in




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



the table. For example, in a table, the columns that contain information of both sales order and customer details can lead to several problems. It is a best practice to store both these information in separate tables for better maintenance and avoidance of any query issues. A table in a database should avoid containing nullable columns. However, some columns of a table can be defined to allow null values. A table in a database should not contain repeated values or columns for a specific piece of information. For example, a product may be purchased from multiple vendors. If the product name and the names of all the vendors are stored in one table, it will contain columns for all the vendors with repeating product ID. This will create several problems. Therefore, it is a best practice to store product information in one table and the names of all vendors in another table. For creating relationships between the two tables, a common column such as ProductID should be placed in both the tables.

QUESTION 8

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" relationship 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 9

Which of the following is a design of the user interface and the application programs that use and process the database?

- A. Application Design
- B. Object Modeling
- C. Database Design
- D. Database Planning

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

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: C is incorrect. 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: D 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. 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.

QUESTION 10

Which of the following is used to give a default value to a column if the value for that column is unknown?

- A. Concurrency
- B. DEFAULT constraint
- C. Nested transaction
- D. Transaction

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 DEFAULT constraint is a constraint in which a default value is given to a column if the value for that column is unknown. If a user does not provide any value for that column, the default value is automatically inserted. If a default value is not provided, then NULL is inserted. If a column does not allow NULL value and a default value is also not assigned for that column, an error is sent by the database engine.

Answer: D is incorrect. A transaction is a logical unit of work, where related DML statements are grouped together. When a transaction is active, either all the statements in the transaction will complete successfully or none of them will run. This means that if any of the statements in the transaction fails due to any reason, the whole transaction will be rolled back. Transactions are ended by using the COMMIT or ROLLBACK statement. A perfect example would be a bank transaction, where a check deposit should cause a debit in the check issuer's account and a credit in the depositor's account. If any one operation fails, the bank should reverse the whole transaction, thus maintaining correct balances in both the accounts.

Answer: C is incorrect. A nested transaction occurs when a new transaction is started by an instruction that is already inside an existing transaction. The new transaction is said to be nested within the existing transaction. Nested transactions are implemented differently in databases. However, these nested transactions have in common that the changes are not made visible to any unrelated transactions until the outermost transaction has committed. This means that a commit in an inner transaction does not necessarily persist updates to the database. Answer: A is incorrect. Concurrency is the ability of several users to access data from the database simultaneously. Database concurrency increases when the number of simultaneous operations that the

database engine can support is large. In Microsoft SQL Server 2008 Compact 3.5, concurrency control is achieved by using locks to protect data. The locks control how multiple users can access and modify shared data simultaneously without conflicting with each other.

QUESTION 11

Consider the case of a distance education university. The university has many regional centers across the country. Each regional center has a head known as regional manager. Each regional center has allotted more than one study center in its region. Each regional center has a region code, which is unique and specifies a region. Each study center also has a study center code, which is also unique. What is the relationship between regional center and study center?

- A. One-to-one
- B. One-to-many
- C. There is no relationship.
- D. Many-to-many

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

It is given in the scenario that there is one regional center for each region, and each regional center has more than one study center. Therefore, one-to-many relationship

QUESTION 12

You work as a Database Administrator for DataOneWorld Inc. Management instructs you to remove an object from the relational database management system. Which of the following statements will you use to accomplish the task?

- A. DROP
- B. SELECT
- C. CREATE
- D. ALTER

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The DROP statement in SQL is used to remove an object from an RDBMS. The types of objects that can be dropped depend on which RDBMS is being used, but most support the dropping of tables, users, and databases. The syntax of the DROP TABLE statement is as follows:

`DROP TABLE "table_name"`



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: C is incorrect. The CREATE statement is used to make a new database, table, index, or stored query. The CREATE statement in SQL creates an object inside DBMS. The types of objects that can be created depend on which DBMS is being used, but most support the creation of tables, indexes, users, and databases.

Answer: D is incorrect. The ALTER statement in SQL changes the properties of an object inside DBMS. The types of objects that can be altered depend on which DBMS is being used. The syntax of the ALTER statement is as follows:

ALTER <specification> [Specification name]

Answer: B is incorrect. The SELECT statement is used to select data from a database. The result is stored in a result table called the result-set. The SELECT statement retrieves zero or more rows from one or more base tables, temporary tables, or views in a database.

QUESTION 13

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 14

You work as a Database Designer for Tritech Inc. You are designing a table named Info. The table contains the following attributes:

Login and Password

Name

ContactNumber

Address

Each new employee joining the company is assigned an E-mail Login and Password. This E-mail Login and Password is used by each employee to access his account on the company's Web site. The value of both the attributes should be unique and cannot be NULL. You design an artificial key EmpID in the Info table to serve as a primary key. Which of the following factors lead you to define an artificial key for use as a primary key in the Info table?

Each correct answer represents a part of the solution. Choose all that apply.

- A. The value of the primary key cannot be NULL.
- B. The value of the primary key must be persistent.
- C. The natural key contains highly sensitive data.
- D. The value of the primary key must be unique.

Correct Answer:

Section: (none)

Explanation

Explanation/Reference:

Answer: D, A, B, and

C

Explanation:

The Name attribute cannot be used as a primary key, as two employees may have the same name in a company. For an attribute to be used as a primary key, it must have a unique value. The primary key attribute cannot contain NULL values. ContactNumber cannot be used as a primary key, for it is possible that some employees in the company may not have a ContactNumber, which will result in NULL values of ContactNumber for such employees. To be used as a primary key, a natural key must contain a unique value. This is a prerequisite for a primary key. In addition to this, if the natural key contains sensitive data, it cannot be used as a primary key. The primary key must be persistent, i.e. the value of the attribute must remain the same throughout the lifetime of the instance. Therefore, Address cannot be used as a primary key, as it is possible that an employee might change his house in future and this may lead to a change in the value of the attribute. In the given scenario, Login and Password of each employee is unique and cannot be NULL, but sharing this information with others is not at all appropriate, as these are classified as sensitive data. If one employee comes to know about the login and password of another employee, he can access his mail account and read his confidential mails or even send mails from his account to other people. Therefore, all these factors lead you to define an artificial key.



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**

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 16

Which of the following steps in database planning is used to record an object so that it is represented visually in the system?

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

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

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. 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 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.



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: 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 17

Remo works as a Database Designer for Tech Inc. He wants to create a table named Product. He issues the following query to create the Product table:
CREATE Product (

ProductID Char (10) NOT NULL,
OrderID Char (10) NULL
ProductName Varchar NOT NULL,
Primary key (OrderID, ProductID)

What are the errors in the above query?

Each correct answer represents a complete solution. Choose two.

- A. An attribute declared as a primary key cannot contain NULL values.
- B. Each attribute should be defined as a primary key separately.
- C. A table cannot have two primary keys.
- D. ProductName is declared as Varchar without specifying the width of the column.

Correct Answer:

Section: (none)

Explanation

Explanation/Reference:

Answer: D and A

Explanation:

When a column is declared as Varchar, it is mandatory to declare the maximum width of the column. The OrderID and ProductID columns are declared as the composite primary key attribute. Therefore, neither ProductID nor OrderID can contain NULL values, whereas in the declaration of OrderID, it is not specified as NOT NULL. Therefore, both answer options D and A are responsible for introducing errors in the query.



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 18

What is the relationship between the foreign key and primary key? Each correct answer represents a complete solution. Choose all that apply.

- A. A foreign key constraint works in conjunction with a primary key constraint to enforce referential integrity among related entities.
- B. A foreign key ties attribute(s) of an entity to the primary key of another entity, for the purpose of creating a dependency.
- C. There is no relationship between a primary key and a foreign key.
- D. A foreign key and a primary key create a link between two entities.

Correct Answer:

Section: (none)

Explanation

Explanation/Reference:

Answer: A and B

Explanation:

A foreign key ties attribute(s) of an entity to the primary key of another entity, for the purpose of creating a dependency. It is used to establish and enforce a relationship between data in two tables. A relationship between two tables is created by adding a column or columns holding the primary key values of the primary table to the related table. This column becomes a foreign key in the second table. You can create a foreign key by defining a FOREIGN KEY constraint, when you create or alter a table.

QUESTION 19

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 20

You work as a Database Designer for Netwell Inc. The database has a table named Project, which is depicted below:

ProjectID	ProjectName	TaskID	TaskName
A	Macwell	Design	A
Macwell	Coding	B	Vintech
Coding	B	Vintech	Design
B	Vintech	Design	Coding
Vintech	Coding	Coding	

Is the table given above in normalized form? If not, then which normal form (NF) is 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. The table is in normal form, as each cell contains only one value.

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

Correct Answer: BD

Section: (none)

Explanation

Explanation/Reference:

Explanation:

and The table given above violates 2NF, as TaskID and TaskName are not functionally dependent on the primary key attribute, i.e. ProjectID. TaskID functionally determines TaskName. Therefore, in order to make it in 2NF, you should split the parent table Project and create a table named Task with TaskID as the primary key attribute and migrate ProjectID to the Task table as the foreign key. The table is not in 2NF. Therefore, it is also not in 3NF. Answer: is incorrect. The table is in 1NF, as each row contains atomic values, i.e. each cell contains a single value. Reference: Database Management Systems, Contents: "Database Normalization"

QUESTION 21

Which of the following statements about rebuilding an index is true?

- A. The NOLOGGING and ONLINE keywords can never be used together.
- B. Without the ONLINE keyword the index is locked for any DML operation.
- C. The NOLOGGING and ONLINE keywords are always used together.
- D. Without the ONLINE keyword the index is locked for the SELECT operation.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The syntax of the REBUILD command has many options such as TABLESPACE, ONLINE, and NOLOGGING. The discussion is narrowed down to ONLINE and NOLOGGING options. NOLOGGING disables redo generation only for the index rebuild; all subsequent DML against the index will generate redo as normal. Without the ONLINE keyword the index is locked for DML, but not for SELECT. However, the rebuild still continues.



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**

You want to recover one or more tablespaces to a point in time older than the rest of the database. Which of the following are the recovery situations in which tablespace Point-in-Time Recovery (TSPITR) should be used? Each correct answer represents a part of the solution. Choose all that apply.

- A. To recover a tablespace that contains rollback segments.
- B. To recover a table that has become logically corrupted.
- C. To recover a DML statement that has affected only a subset of the database.
- D. To recover a tablespace on a very large database.

Correct Answer:

Section: (none)

Explanation

Explanation/Reference:

Answer: C, B, and A



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:

Following are the situations in which tablespace Point-in-Time Recovery (TSPITR) should be used:

To recover from an erroneous drop or truncate table operation. To recover a table that has become plausibly corrupted. To recover from an erroneous batch job or other DML statement that has affected only a subset of the database.

To recover one independent schema to a point dissimilar from the respite of a physical database (in cases where there are multiple independent schemas in individual tablespaces of one physical database). To recover a tablespace on a very large database (VLDB) instead of restoring the whole database from a backup and carrying out a complete database roll-forward.

Point-in-time recovery is a technique used to back up any database object or recover any database object to a particular target SCN. When the target SCN is specified using the date and time, it is known as time-based recovery. Following types of point-in-time recoveries are known:

Tablespace point-in-time recovery (TSPITR)

Database point-in-time recovery (DBPITR)

Answer: A is incorrect. Tablespace Point-in-Time Recovery cannot be used on the SYSTEM tablespace, an UNDO tablespace, or any tablespace that contains rollback segments.

QUESTION 23

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 24**

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 25

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 26

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	
	InquiryID
	ProductID
	LastInquiredOn
	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)



A list box titled "Primary Key Column(s)" containing no visible items. It includes scroll bars and standard Windows-style navigation buttons (up, down, first, last) at the top and bottom.

All Columns

ProductID
LastInquiredOn
SalesPersonID
InquiryID



A list box titled "All Columns" containing four items: ProductID, LastInquiredOn, SalesPersonID, and InquiryID. It includes scroll bars and standard Windows-style navigation buttons at the bottom.

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

Correct Answer:**Section: (none)****Explanation****Explanation/Reference:**

Primary Key Column(s)

InquiryID
ProductID

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 27

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 28

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 29

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 30

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.


C A:




C B:



C C:



C D:



- 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 31

Which of the following statements is true about referential integrity?

- A. It checks that no component of a primary key can have a null value and no duplicate entries can exist.
- B. It checks that the value of a primary key must be consistent throughout the life of an instance.
- C. It distinguishes between null values and zero entries.
- D. It checks that no record in a child table can exist if its corresponding record is not available in the parent table.

Correct Answer: D**Section: (none)****Explanation****Explanation/Reference:****Explanation:**

Referential integrity is enforced through a foreign key. There must be a reference of the foreign key in its originating table. For example, take two tables, Employee and Department. The DeptID attribute of the Department table is the primary key, and when it is exported to Employee table, it becomes the foreign key in the Employee table. The data value of DeptID in the Employee table cannot exist if it does not exist in the Department table.

Employee	
<u>EmpID</u>	Primary key
EmpName	
Address	
Contactnumber	
E-mail	
SSN	
<u>DeptID</u>	Foreign key

Department	
<u>DeptID</u>	Primary key
DepartmentName	



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



Referential integrity constraint is a feature in RDBMS (Relational Database Management System) that prevents a user from entering inconsistent data. There are various referential integrity rules that are applied on tables sharing a relationship. These are as follows:
 If the Employee table has a foreign key that is a primary key of the Department table, then the Employee table cannot have a value in the foreign key column if there is no corresponding value in the Department table. If a record is deleted from the Department table, the corresponding record must be deleted from the Employee table. This is known as cascading delete.

Answer: C is incorrect. According to an integrity constraint, a table can have only one primary key constraint. A primary key constraint imposes that the primary key attribute should be unique and cannot have NULL values. This ensures that the primary key value uniquely identifies each and every row in a table.

QUESTION 32

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 33

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 34

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 35

You are a database developer for a database named Customers hosted on a SQL Server 2008 server. Recently, several customers were deleted from the Customers database. To ensure this is not repeated in future, you have decided to create a DML trigger to prevent it. What code will create the trigger to meet your goals? Each correct answer represents a complete solution. Choose all that apply.

- A. CREATE TRIGGER trgDeleteCustomer
ON dbo.Customers
BEFORE DELETE
AS
RAISERROR ('Customers cannot be deleted. An error has been logged', 16, 10) WITH LOG ROLLBACK TRANSACTION
- B. CREATE TRIGGER trgDeleteCustomer
ON dbo.Customers
AFTER DELETE
AS
RAISERROR ('Customers cannot be deleted. An error has been logged', 16, 10) WITH LOG ROLLBACK TRANSACTION
- C. CREATE TRIGGER trgDeleteCustomer
ON dbo.Customers
AFTER DELETE
AS



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



IF (SELECT COUNT(*) FROM DELETED) > 1
BEGIN

```
RAISERROR ('Customers cannot be deleted. An error has been logged', 16, 10) WITH LOG ROLLBACK TRANSACTION  
END
```

D. CREATE TRIGGER trgDeleteCustomer
ON dbo.Customers
AFTER DELETE
AS
IF (SELECT COUNT(*) FROM DELETED) > 0
BEGIN
RAISERROR ('Customers cannot be deleted. An error has been logged', 16, 10) WITH LOG ROLLBACK TRANSACTION
END

Correct Answer:

Section: (none)

Explanation

Explanation/Reference:

Answer: D and B

Explanation:

In the code with answer option D, the deleted table is queried to see if it holds any data. The deleted table is a virtual table that exists immediately after the DELETE statement in a trigger. It will hold the deleted row or rows. The COUNT(*) is counting the number of rows in the table, and if any rows are present, the trigger will raise the error and roll back the transaction. This method is useful if you want to prevent more than a specific count (such as more than 5 customers) from being deleted.

QUESTION 36

JDBC API 2.0 defines five levels of transaction isolation for database concurrency control. Which of the following is the lowest level of transaction isolation?

- A. TRANSACTION_SERIALIZABLE
- B. TRANSACTION_NONE
- C. TRANSACTION_READ_COMMITTED
- D. TRANSACTION_READ_UNCOMMITTED
- E. TRANSACTION_REPEATABLE_READ

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

There are five levels of transaction isolation provided by JDBC 2.0 Connection interface. The higher the transaction level, the more care is taken to avoid database conflicts. On the other hand, the higher the level of isolation, the slower the application executes (due to increased



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



locking overhead and decreased concurrency between users). The developer must balance the need for performance with the need for data consistency when making a decision about what isolation level to use. However, the level that can actually be supported depends on the capabilities of the underlying DBMS.

The highest transaction isolation level specifies that if one transaction is operating on a database, no other transaction can access that part of the database.

The transaction isolation levels in terms of concurrency control in ascending order are as follows:

TRANSACTION_NONE

TRANSACTION_READ_COMMITTED

TRANSACTION_READ_UNCOMMITTED

TRANSACTION_REPEATABLE_READ

TRANSACTION_SERIALIZABLE

The lowest transaction isolation level, i.e. TRANSACTION_NONE, specifies that transactions are not at all supported.

A developer can change a transaction isolation level by using the setTransactionIsolation() method of the Connection interface.

QUESTION 37

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 38

Which of the following DML SQL statements support usage of correlated sub-queries? Each correct answer represents a complete solution. Choose all that apply.

- A. INSERT
- B. UPDATE
- C. DELETE

Correct Answer:

Section: (none)

Explanation

Explanation/Reference:

Answer: B and C

Explanation:

A correlated sub-query is a sub-query where a column of a table referred in the outer parent query is referenced in the sub-query. Two tables are involved in a correlated sub-query and both have data in them. The two tables are linked to each other using a common column. In the UPDATE and

DELETE SQL statements the table can have data and can be linked with a common column. However, the same is not true for INSERT statements as in the INSERT statement the table does not contain any data, and hence, cannot be correlated with the inner query table. The best way to work with correlated sub-queries is to use table and/or column aliases to clearly understand which columns are being referenced from which tables. Without aliases Oracle will return an ambiguity error in some cases.

Answer: A is incorrect. The INSERT statements do not support correlated sub-queries. They support other sub-queries like INSERT INTO (SELECT FROM).

Correlated sub-queries with INSERT will not work because for a correlated sub-query to work there has to be a common column based on which rows in the table are linked to the columns of the table used in the sub-query. When using the INSERT state, the table does not contain any data; hence, a correlation cannot occur. You may however, INSERT data first into a table using INSERT INTO (SELECT FROM) and then use UPDATE statements and correlated sub-queries to update the data from other tables.



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 39

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 40

A table named Emp is given below that contains data of an organization.

	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

How many records will be displayed if the following SQL query is executed?
 SELECT *
 FROM Emp
 WHERE Name BETWEEN 'D' AND 'F'

- A. 2
- B. A character column cannot be used in the BETWEEN operator.
- C. 3
- D. 4

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:

The table given below shows the result of the query. The character column is compared against a string using the BETWEEN operator, which is equivalent to $\geq 'D'$ AND $\leq 'F'$. The name Flora will not be included because Flora is $> F$.

	SSN	Name	Salary	DOB	Department
	A103	Dinzi	12000	12/25/67	Production
	A106	Esha	12500	5/6/77	Marketing
	A107	Ema	34000	7/9/76	IT

QUESTION 41

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 42

Mark works as a Database Designer for Reon Inc. He is assigned the task to create a database for the company. He issues the following query to create the database.

CREATE DATABASE '24342'

What will be the output of the query?

- A. A database will be created.
- B. An error will be generated stating that a database name must begin with an alphabet and all the alphabets should be in uppercase.
- C. An error will be generated stating that a database name must begin with an alphabet.
- D. An error will be generated stating that a database name should be a combination of numerals and alphabets.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

A database name can begin with a digit, provided it is enclosed within single quotes. The digits within quotes are treated as a string. It is not necessary that all the alphabets in a database name should be in uppercase.

QUESTION 43

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 44

You work as a MySQL Database Administrator. Your server crashed and it was recovered and made online within half an hour. What will you do to ensure that the table structure and contents are in a proper state?

- A. Perform a table optimization operation.
- B. Perform a table repair operation.
- C. Perform a table analysis operation.
- D. Perform a table check operation.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

While working on MySQL database maintenance, a table check performs an integrity check to ensure that the table structure and contents are in a proper state. A table check can be performed for the MyISAM and InnoDB tables.

QUESTION 45

James works as an Administrator for Softech Inc. The company has a SQL Server 2005 computer named SQL1. The company has six departments;

Accounting, Sales, Marketing, Manufacturing, Production, and Research. The user accounts of each department are located in a domain group named after the department. James wants to create a separate database on SQL1 for each department which will store the department related information. Users in each department must be able to read and update data in the database belonging to their respective departments. No user must be able to access the database belonging to another department. James wants to configure access to the database to meet these requirements. What should James do to accomplish the task?

- A. James should create a Windows Authentication login for each domain group and configure the logins as database users for the appropriate database. Add each database user to the db_ddladmin database role.
- B. James should create a Windows Authentication login for each domain group and configure the logins as database users for the appropriate database. Add each database user to the db_datareader and db_datawriter database roles.
- C. James should create a Windows Authentication login for each domain user and configure the logins as database users for the appropriate database. Add each database user to the db_ddladmin database role.
- D. James should create a Windows Authentication login for each domain user and configure the logins as database users for the appropriate database. Add each database user to the db_datareader and db_datawriter database roles.

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:

In order to accomplish the task James should create a Windows Authentication login for each domain group and configure the logins as database users for the appropriate database. Add each database user to the db_datareader and db_datawriter database roles.

As the criterion in the question states that the users of each department belong to a group named after that department. These groups can be used to configure access to the database. Each department user should be assigned db_datareader and db_datawriter database roles for their respective databases. This will prevent a user in one department from accessing another department's database.

Answer: D is incorrect. Create a Windows Authentication login for each domain user, and configure it as a database user for the appropriate database. Then add each database user to the db_datareader and db_datawriter database roles. But it would require less administrative effort to create a Windows Authentication login for each domain group when compared to that created for each domain user.

Answer: A is incorrect. James could use the domain group to configure access to the database but the database user should not be added to the db_ddladmin database role. The db_ddladmin is a fixed database role. A user who is a member of this database role can add, modify, or drop objects at the database level but cannot issue the GRANT, REVOKE, or DENY statements. The db_ddladmin role exists in each database. As the scenario in the question states that the user can only read and update the data present in database, thus, assigning user

to the db_ddladmin database role will not serve the purpose. Answer: C is incorrect. James should not create a Windows Authentication login for each user, but for each group as this task will be accomplished with least administrative effort. Moreover, the database should be added to the db_datareader and db_datawriter database roles as they have to only read and update the data present in the database.

QUESTION 46

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



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 47

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> /ust/share/zoneinfo mysql_tzinfo_to_sql | mysql -u root mysql
- B. shell> /ust/share/zoneinfo mysql_tz_to_sql | mysql -u root mysql
- C. shell> mysql_tz_to_sql /ust/share/zoneinfo | mysql -u root mysql
- D. shell> mysql_tzinfo_to_sql /ust/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 /ust/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 48**

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 49

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 50

You manage a database named Exams accessed via a website. The website is regularly accessed by users in different countries. You are redesigning a table named History that records when users have purchased and taken online exams. You need one of the columns in the History table to record what time of day that a user has taken an exam. Additionally, you need this column to have time zone awareness. What data type should you use?

- A. datetimeoffset
- B. datetime
- C. Use two different columns with one column of datetime, and the second column indicating the UTC offset
- D. datetime2

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:**

The datetimeoffset data type combines a date with a time of day that has time zone awareness. The time is based on a 24-hour clock but also includes a UTC offset. For the datetimeoffset data type, both the UTC and local datetime values will be validated during update, insert, convert, arithmetic, or assign operations. The recognition of any invalid UTC or local datetime value will result in an invalid value error. For example, the 9999-12-31 10:10:00 datetime value is valid in UTC, but it overflows in local time when the time zone offset is +13:50.

QUESTION 51

Fill in the blank with the appropriate word.

The model for database management is a database model based on first-order predicate logic.

- A. relational

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation: The Relational model describes a database as a collection of predicates over a finite set of predicate variables, describing constraints on the possible values and combinations of values. The content of the database at any given time is a finite (logical) model of the database, i.e., a set of relations, one per predicate variable, such that all predicates are satisfied. A request for information from the database or a database query is also a predicate.

QUESTION 52

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 53

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 $A \rightarrow\!> B|C$, then $A \rightarrow B$ and $A \rightarrow C$

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 54

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 55

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 56

Mark works as a Database Administrator for DataOneWorld Inc. The company has a SQL Server database. Management instructs him to verify the identity of a user or software that is connecting to the database. Which of the following will help him to accomplish the task?

- A. Referential integrity
- B. Abstraction
- C. Encryption
- D. Authentication

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:**

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: A is incorrect. 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.

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 57

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 58

Peter works as a Database Administrator for DataOneWorld Inc. The company has a SQL Server database. Peter has to construct an Entity-Relationship schema design for the data he will use in the database. Which of the following database designs will he use to accomplish the task?

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

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

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.



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. 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. 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: C is incorrect. 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.

QUESTION 59

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 60

Which of the following approaches begins with the identification of only a few important concepts and, based on these, the design proceeds, spreading outward 'radially'?

- A. Inside-out approach
- B. Bottom-up approach
- C. Mixed approach
- D. Top-down approach

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Inside-out approach can be regarded as a particular type of bottom-up approach. It begins with the identification of only a few important concepts and, based on these, the design proceeds, spreading outward 'radially'. First the concepts nearest to the initial concepts are represented, and a user then moves towards those further away by means of 'navigation' through the specification.

Answer: C is incorrect. In a mixed approach, the designer decomposes the requirements into a number of components, as in the bottom-up approach, but not to the extent where all the concepts are separated. At the same time, a skeleton schema is defined containing the main concepts of the application. This skeleton schema gives a unified view of the whole design and helps the integration of schemas developed separately. Then the designer examines separately these main concepts and can proceed with gradual refinements following the top-down approach or extending a portion with concepts that are not yet represented following the bottom-up approach.

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: 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.



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 61**

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

In the table, you deleted the record of Mark Smith. Which of the following will change as a result?

- A. Degree of the table
- B. Cardinality of the table
- C. Domain of the table
- D. Attribute of the table

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

In tables, the number of rows (or tuples) is called the cardinality. When a record (or row) of the table is deleted (or added), the cardinality of the table

changes. Answer: A and D When a column of the table is changed, the degree (or attribute is changed). 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.

Answer: C is incorrect. A Domain is a set of allowable values for one or more attributes. It is also known as the data domain which refers to all the unique values which a data element may contain. The Domain boundary can be defined by rules such as the data type with an enumerated list of values. Check constraint is also a simple example of defining the boundary of a domain.

QUESTION 62

Which of the following files contain configuration information that must be restored after a MySQL database server has crashed? Each correct answer represents a complete solution. Choose two.

- A. My.ini
- B. My.cnf
- C. Master.info
- D. Binary log files

Correct Answer:

Section: (none)

Explanation

Explanation/Reference:

Answer: B and A



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:

My.cnf and my.ini files are the option files used by a MySQL server. These files contain configuration information that must be restored after a crash. Answer: D is incorrect. The binary logs store updates that have been made after the backup was made. Answer: C is incorrect. Replication slave servers create a master.info file that contains information needed for connecting to the master server.

QUESTION 63

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.

QUESTION 64

Which of the following commands are used to terminate a transaction? Each correct answer represents a complete solution. Choose all that apply.

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

Correct Answer:

Section: (none)

Explanation

Explanation/Reference:

Answer: D, C, and A

Explanation:

The COMMIT, ROLLBACK, and TRUNCATE commands are used to terminate a transaction. While the first two do so explicitly, TRUNCATE does it implicitly. Answer: B is incorrect. 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];`

QUESTION 65

You are modifying a database that is used as a backend for a web site. The web site developers want to include positional data that can be used by the web site users to determine store locations information that is stored in the Stores table of the database. You are asked to add a column to the Stores table. What data type should you use?

- A. varchar
- B. nvarchar
- C. geography
- D. GPS

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:

The geography data type is used to hold spatial data. Geography is a .NET common language runtime (CLR) data type that can store ellipsoidal (round-earth) data such as GPS latitude and longitude coordinates. The geometry data type is a planar spatial data type that is implemented as a common language runtime (CLR) data type in SQL Server. This type represents data in a Euclidean (flat) coordinate system.

QUESTION 66

Philandrio Inc. is one of the largest drug manufacturing companies in the world. It has many offices worldwide. Each office has a manpower of approximately 14,000. The managers themselves are employees in the company. You are building a database for the company's Human Resources Department (HRD). You want to track the following attributes for each employee, and at the same time, eliminate duplicate entries and minimize data storage wherever possible.

FirstName
MiddleName
LastName
EmployeeIdentityNumber
Address
Date_of_Hire
Department

Salary

ManagerIdentityNumber

Which of the following is the appropriate table detail sufficient to track the above attributes, correctly and efficiently?

- A. Only one table, Employee
- B. Three tables, Employee, Manager, and a table to link both these entities
- C. Two tables, Employee and Manager
- D. Either two or three tables

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:

Only one table is sufficient to keep track of all the above attributes of each employee, because all these attributes refer to the Employee entity. However, there can be many employees working under one manager, which requires a one- to-many relationship between Manager and Employee.

To retrieve the name of the manager of a particular employee, a self join should be performed. A self join is a join of a table to itself. This table appears twice in the FROM clause and is followed by table aliases that qualify column names in the join condition.

Example:

```
SELECT A.EmployeeIdentityNumber, A.LastName, B.LastName FROM Employees A, Employees B  
WHERE A.ManagerIdentityNumber=B.EmployeeIdentityNumber;
```

QUESTION 6

You work as a Database Designer for Realtech Inc. You want to create two tables, Product and Order, such that a row in the Order table does not exist without a corresponding row in the Product table. The structure of the Product and Order tables with their respective columns is given below:

Product

ProductID

ProductName

Date_of_manufacturing

SupplierID

ProductPrice
ProductQty

Order
OrderID
OrderDate
OrderQuantity
OrderPrice

Based on this information, while maintaining the integrity of this rule, what will you do to accomplish the task?

- A. Create the Product table, and then create the Order table. Create a third table that will be used to join the Product and Order tables by using FOREIGN KEY constraints.
- B. Create the Product table, and then create the Order table that has a primary key referenced by the foreign key in the Product table.
- C. Create the Product table, and then create the Order table. Create a primary key in both the tables and form a one-to-one relationship.
- D. Create the Product table, and then create the Order table that has a foreign key constraint that references the primary key in the Product table.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The Product table and the Order table will have a one-to-many relationship. To form this relationship, you will have to add a column in the Order table, and create a foreign key constraint on that column such that it references the primary key in the Product table. Suppose ProductID is designated as the primary key of the Product table, then it will have to be migrated to the Order table to maintain the referential integrity constraint.



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 68

Which of the following are the steps of database planning life cycle? Each correct answer represents a part of the solution. Choose three.

- A. Application Design
- B. Data conversion and loading

- C. Conceptual database design
- D. Operational Maintenance

Correct Answer:

Section: (none)

Explanation

Explanation/Reference:

Answer: A, D, and B



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:

There are eleven steps in the database planning life cycle, which are as follows:

- 1.Database planning
- 2.System definition
- 3.Requirements collection and analysis
- 4.Database design
- 5.DBMS selection
- 6.Application design
- 7.Prototyping
- 8.Implementation
- 9.Data conversion and loading
- 10.Testing
- 11.Operational maintenance

QUESTION 69

You manage a SQL Server 2008 server that hosts a database named Sales. Unauthorized odifications to some tables within the Sales database have resulted in some views no longer working. You need to implement a method that will ensure that all modifications to any objects with the Sales database are logged. What should you do?

- A. Create a DDL trigger with database scope.
- B. Create a DML trigger with database scope.
- C. Create a DML trigger with server scope.

D. Create a DDL trigger with server scope.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

A Database Definition Language (DDL) trigger with database scope can be used to audit all DDL statements on any database objects. A DDL statement (such as CREATE, ALTER, and drop) is used to create, alter, or drop (remove or delete) objects within the database. Once the event is captured, the DDL trigger can be programmed to log events, or if desired, even roll back unauthorized events.

Answer: D is incorrect. A Database Definition Language (DDL) trigger with server scope can be used to audit all DDL statements at the server level, such as the creation, modification, or deletion of a database. However, DDL triggers at the server level will not capture DDL statements executed on objects within a database.

Answer: B and C are incorrect. Data Manipulation Language (DML) statements (such as UPDATE, INSERT, or DELETE) are used to update, insert, or delete data within a database. A DML trigger will not capture modifications to the objects. Further, DML triggers would be created with database scope since it works on data within a database.



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 70

Roger is the HR Manager of Neon Inc. He wants to generate a report of all the employees working in Department ID 30 from the Emp table on the following attributes:

Fname

Lname

Salary

All the records should be sorted in ascending order of the Lname attribute. Which of the following queries will he use to get the required output?

- A. SELECT Fname Lname Salary FROM Emp WHERE DeptId=30 ORDER BY Lname
- B. SELECT Fname, Lname, Salary FROM Emp ORDER BY Lname WHERE DeptId=30
- C. SELECT Fname Lname Salary FROM Emp ORDER BY Lname WHERE DeptId=30
- D. SELECT Fname, Lname, Salary FROM Emp WHERE DeptId=30 ORDER BY Lname

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The first part of the query, i.e. SELECT Fname, Lname, Salary selects the Fname, Lname, and Salary columns (attributes). The second part, i.e. FROM Emp selects the Emp table, and the third part, i.e. WHERE DeptId=30 specifies the condition that DeptId is 30. The ORDER BY clause is used to sort the name in ascending order of Lname. In the ORDER BY clause, ascending is the default order. Therefore, if no order is given, the column is sorted in ascending order. Answer: C is incorrect. The first part, i.e. SELECT Fname Lname Salary has an error, as each column name must be separated by a comma, and the ORDER BY clause should be after the WHERE clause and not before it.

Answer: B is incorrect. The ORDER BY clause should be used after the WHERE clause and not before it.

Answer: A is incorrect. The first part of the query, i.e. SELECT Fname Lname Salary is erroneous, as each column name must be separated by a comma.



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

