David M. Kroenke and David J. Auer

Database Processing

Fundamentals, Design, and Implementation

Chapter Seven:

SQL for Database Construction and Application Processing

SQL Categories

- SQL statements can be divided into five categories:
  - **Data definition language (DDL)**
  - **Data manipulation language (DML)** statements
  - **SQL/Persistent Stored Modules (SQL/PSM)** statements
  - **Transaction control language (TCL)** statements
  - **Data control language (DCL)** statements

SQL DDL and DML

- **Data definition language (DDL)**
  - Statements used for creating tables, relationships, and other structures
- **Data manipulation language (DML)**
  - Statements used for queries and data modification
- **SQL/Persistent Stored Modules (SQL/PSM)**
  - Statements used to add procedural programming capabilities
    - Variables and Control-of-flow statements

Database Processing, 13e (Kroenke/Auer)
Chapter 1: Introduction

1) The purpose of a database is to help people keep track of things.
ANSWER: TRUE
Diff: 1 Page Ref: 3

2) In a database, each table stores data about a different type of thing.
ANSWER: TRUE
Diff: 1 Page Ref: 3

3) In a database, each row in a spreadsheet has data about a particular instance.
ANSWER: FALSE
Diff: 2 Page Ref: 3

4) In every database, not just the databases discussed in this book, table names are capitalized.
ANSWER: FALSE
Diff: 2 Page Ref: 3

5) A database shows data in tables and the relationships among the rows in those tables.
ANSWER: TRUE
Diff: 1 Page Ref: 4

6) Data is recorded facts and figures; information is knowledge derived from data.
ANSWER: TRUE
Diff: 1 Page Ref: 5

7) Databases record data in such a way that they can produce information.
ANSWER: TRUE
Diff: 2 Page Ref: 6

8) Enterprise Resource Planning (ERP) is an example of a data mining application.
ANSWER: FALSE
Diff: 2 Page Ref: 6-7

9) Databases are a key component of e-commerce order entry, billing, shipping and customer support.
ANSWER: TRUE
Diff: 2 Page Ref: 7

10) The largest databases in e-commerce are the order entry databases.
ANSWER: FALSE
Diff: 3 Page Ref: 7

11) The e-commerce companies use Web activity databases to determine which items on a Web page are popular and successful.
ANSWER: TRUE
Diff: 2 Page Ref: 7
12) Small databases typically have simple structures.
ANSWER: FALSE
Diff: 3 Page Ref: 8

13) A database system is typically defined as its four components: users, database applications, the DBMS and the databases.
ANSWER: TRUE
Diff: 1 Page Ref: 8

14) A database system as typically defined can be modified to include CODASYL.
ANSWER: FALSE
Diff: 1 Page Ref: 8

15) Applications are computer programs used directly by users.
ANSWER: TRUE
Diff: 1 Page Ref: 9

16) In a database system, applications write data to the database.
ANSWER: FALSE
Diff: 1 Page Ref: 9

17) Sequenced Query Language (SQL) is an internationally recognized standard language that is understood by all commercial database management system products.
ANSWER: FALSE
Diff: 2 Page Ref: 8

18) In database systems, indexes are held by the database.
ANSWER: TRUE
Diff: 3 Page Ref: 8-13 Fig 1-14

19) A database management system (DBMS) creates, processes and administers databases.
ANSWER: TRUE
Diff: 1 Page Ref: 11

20) The database management system (DBMS) is responsible for inserting, modifying, reading, and deleting data.
ANSWER: TRUE
Diff: 1 Page Ref: 11

21) The database application is responsible for concurrency control.
ANSWER: FALSE
Diff: 3 Page Ref: 12

22) The database management system (DBMS) is responsible for enforcing referential integrity constraints.
ANSWER: TRUE
Diff: 1 Page Ref: 12
23) Referential integrity constraints are rules about what data values are allowed in certain columns.
ANSWER: TRUE
Diff: 1 Page Ref: 12

24) A database is a self-describing collection of non-integrated tables.
ANSWER: FALSE
Diff: 2 Page Ref: 12

25) Integrated tables store both data and the relationships among the data.
ANSWER: TRUE
Diff: 2 Page Ref: 12

26) Microsoft Access is just a DBMS.
ANSWER: FALSE
Diff: 2 Page Ref: 14

27) Microsoft Access is a low-end product intended for individuals and small workgroups.
ANSWER: TRUE
Diff: 2 Page Ref: 14

28) The current DBMS engine in Microsoft Access is called ADE.
ANSWER: TRUE
Diff: 2 Page Ref: 14

29) In Microsoft Access, you can use the Oracle DBMS in place of the ADE DBMS.
ANSWER: FALSE
Diff: 2 Page Ref: 15

30) In an Enterprise-class database system, a database application interacts with the DBMS.
ANSWER: TRUE
Diff: 2 Page Ref: 15 Fig 1-16

31) In an Enterprise-class database system, a database application accesses the database data.
ANSWER: FALSE
Diff: 2 Page Ref: 15 Fig 1-16

32) In an Enterprise-class database system, business users interact directly with the DBMS, which directly accesses the database data.

ANSWER: FALSE
Diff: 2 Page Ref: 15 Fig 1-16

33) All database applications get and put database data by sending SQL statements to the DBMS.
ANSWER: TRUE
Diff: 2 Page Ref: 8

34) The DBMS ranked as having the “most power and features” in the text is IBM’s DB2.
ANSWER: FALSE
Diff: 3 Page Ref: 16 Fig 1-17

35) The DBMS ranked as being the “most difficult to use” in the text is Oracle Corporation’s Oracle Database.
ANSWER: TRUE
Diff: 3 Page Ref: 16 Fig 1-17

36) The DBMS ranked as being the “easiest to use” in the text is Microsoft’s SQL Server.
ANSWER: FALSE
Diff: 2 Page Ref: 16 Fig 1-17

37) The DBMS ranked as having the “least power and features” in the text is Microsoft Access.
ANSWER: TRUE
Diff: 2 Page Ref: 16 Fig 1-17

38) A database is called “self-describing” because it reduces data duplication.
ANSWER: FALSE
Diff: 2 Page Ref: 12

39) The description of a database’s structure that is stored within the database itself is called the “metadata.”
ANSWER: TRUE
Diff: 1 Page Ref: 12 Fig 1-14

40) In a database processing system, indexes are held by the database management system (DBMS).
ANSWER: TRUE
Diff: 3 Page Ref: 13 Fig 1-12

41) Database design is important, and fortunately it is simple to do.
ANSWER: FALSE
Diff: 1 Page Ref: 17

42) A database design may be part of a new systems development project.
ANSWER: TRUE
Diff: 1 Page Ref: 17-19 Fig 1-18

43) A database design is rarely a redesign of an existing database.
ANSWER: FALSE
Diff: 1 Page Ref: 17-19 Fig 1-18

44) Information systems that stored groups of records in separate files were called file processing systems.
ANSWER: TRUE
Diff: 2 Page Ref: 21-23 Fig 1-25

45) Data Language/I (DL/I) structured data relationships as a tree structure.
ANSWER: TRUE
Diff: 3 Page Ref: 21 Fig 1-25

46) The CODASYL DBTG mode structured data relationships as a tree structure.
ANSWER: FALSE
Diff: 3 Page Ref: 23 Fig 1-25
47) The relational model was first proposed in 1970 by E. F. Codd at IBM.
ANSWER: TRUE
Diff: 2 Page Ref: 23

48) The 1977 edition of this text contained a chapter on the relational model, and that chapter was reviewed by E. F. Codd.
ANSWER: TRUE
Diff: 3 Page Ref: 23

49) dBase was the first PC-based DBMS to implement true relational algebra on a PC.
ANSWER: FALSE
Diff: 3 Page Ref: 24

50) Paradox is the only major survivor of the “bloodbath of PC DBMS products.”
ANSWER: FALSE
Diff: 2 Page Ref: 24

51) Business organizations have resisted adopting object-oriented database systems because the cost of purchasing OODBMS packages is prohibitively high.
ANSWER: FALSE
Diff: 3 Page Ref: 24

52) Bill Gates has said that “XML is the lingua-franca of the Internet Age.”
ANSWER: TRUE
Diff: 2 Page Ref: 25

53) XML Web services allow database processing to be shared across the Internet.
ANSWER: TRUE
Diff: 2 Page Ref: 25

54) The NoSQL movement should really be called a NoRelational movement.
ANSWER: TRUE
Diff: 2 Page Ref: 24

55) Twitter and Facebook use NoSQL databases.
ANSWER: TRUE
Diff: 2 Page Ref: 25

56) The purpose of a database is to:
A) help people keep track of things.
B) store data in tables.
C) create tables of rows and columns.
D) maintain data on different things in different tables.
E) All of the above.
ANSWER: A
Diff: 2 Page Ref: 3
57) A database stores:
A) data.
B) relationships.
C) metadata.
D) A and B
E) A, B, and C
ANSWER: E
Diff: 2 Page Ref: 3-4 and 12

58) A database records:
A) facts.
B) figures.
C) information.
D) A and B
E) A, B, and C
ANSWER: D
Diff: 2 Page Ref: 3

59) A sales contact manager used by a salesperson is an example of a(n) _____.
A) single-user database application
B) multiuser database application
C) e-commerce database application
D) A or B
E) Any of A, B, or C
ANSWER: A
Diff: 2 Page Ref: 8 and Fig 1-5

60) A Customer Resource Management (CRM) system is an example of a(n) _____.
A) single-user database application
B) multiuser database application
C) e-commerce database application
D) A or B
E) Any of A, B, or C
ANSWER: B
Diff: 2 Page Ref: 8 and Fig 1-5

61) An online drugstore such as Drugstore.com is an example of a(n) _____.
A) single-user database application
B) multiuser database application
C) e-commerce database application
D) A or B
E) Any of A, B, or C
ANSWER: C
Diff: 1 Page Ref: 7 Fig 1-5
62) The industry standard supported by all major DBMSs that allows tables to be joined together is called _____.
A) Sequential Query Language (SQL)
B) Structured Question Language (SQL)
C) Structured Query Language (SQL)
D) Relational Question Language (RQL)
E) Relational Query Language (RQL)
ANSWER: C
Diff: 1 Page Ref: 8

63) A program whose job is to create, process and administer databases is called a _____.
A) Database Modeling System
B) Database Management System
C) Data Business Model System
D) Relational Model Manager
E) Data Business Management Service
ANSWER: B
Diff: 2 Page Ref: 8

64) Microsoft Access includes:
A) a DBMS.
B) an application generator.
C) a Web server.
D) A and B
E) A, B, and C
ANSWER: D
Diff: 2 Page Ref: 14 Fig 1-15

65) Microsoft Access may use which of the following DBMS engines?
A) ADE
B) SQL Server
C) Oracle
D) A and B
E) A, B, and C
ANSWER: D
Diff: 2 Page Ref: 14

66) Which of the following are basic components of an enterprise-class database system?
A) The user
B) The database application
C) The database management system (DBMS)
D) The database
E) All of the above
ANSWER: E
Diff: 1 Page Ref: 15-16 Fig 1-16
67) In an enterprise-class database system, _____.
A) the database application(s) interact(s) with the DBMS
B) the database application(s) access(es) the database data
C) the DBMS accesses the database data
D) A and B
E) A and C
ANSWER: E
Diff: 2 Page Ref: 15-16 Fig 1-16

68) In an enterprise-class database system, the database application _____.
A) creates queries
B) creates forms

C) creates reports
D) A and B
E) B and C
ANSWER: E
Diff: 2 Page Ref: 15-16 Fig 1-16

69) In an enterprise-class database system, reports are created by _____.
A) the user
B) the database application
C) the database management system (DBMS)
D) the database

E) All of the above
ANSWER: B
Diff: 2 Page Ref: 15-16 Fig 1-16

True false

Database Processing, 12e(Kroenke/Auer)Chapter 5: Data Modeling with the Entity-Relationship Model1) A data model is a plan for a database design.Answer:True

2) The method of constructing data models used in the text is the extended entity-relationship (E-R) model.Answer:True

3) An entity is something that users want to track.Answer:True

4) Entities of a given type are grouped into entity classes.Answer:True
5) An entity class is described by the structure of the entities in that class.Answer:True

6) An entity instance is the occurrence of a particular entity.Answer:True

7) An entity instance of an entity class is the representation of a particular entity and is described by the values of the attributes of the entity.Answer:True

8) An identifier of an entity instance is one or more attributes that name or identify entity instances.Answer:True

9) A compound identifier is an identifier consisting of two or more attributes.Answer:FALSE

10) In E-R modeling, an attribute describes the characteristics of an entity.Answer:True

11) In E-R modeling, entities within an entity class may have different attributes.Answer:FALSE

12) An identifier of an entity instance must consist of one and only one attribute.Answer:FALSE

13) A composite identifier is defined as a composite attribute that is an identifier.Answer:FALSE1

4) An identifier serves the same role for a table that a key does for an entity.Answer:FALSE

15) Entities can be associated with one another in relationships.Answer:True

16) The degree of a relationship is the number of entity classes in the relationship.Answer:True

17) A binary relationship is a relationship between two or more entities.Answer:FALSE

18) Relationships of degree two are referred to as binary relationships.Answer:True

19) E-R modeling recognizes both relationship classes and relationship instances.Answer:True

20) In today's E-R models, attributes of relationships are still used.Answer:FALSE

21) A single relationship class involves only one entity class.Answer:FALSE

22) A binary relationship is a relationship based on numerical entity instance identifiers.Answer:FALSE

23) The degree of a relationship is expressed as the relationship's maximum cardinality.Answer:FALSE

24) When transforming a data model into a relational design, relationships of all degrees are treated as combinations of binary relationships.Answer:True

25) The principle difference between an entity and a table is that you can express a relationship between entities without using foreign keys.Answer:True

26) When designing a database, first identify the entities, then determine the attributes, and finally create the relationships.Answer:FALSE

27) The notation 1:N shows the relationship's maximum cardinality.Answer:True

28) Relationships are classified by their cardinality.Answer:True

Exam Finished

Thank you for your time.

Test Bank for Database Processing 12th Edition Part 2

The CODASYL DBTG mode structured data relationships as a tree structure.

1. True
2. False

The relational model was first proposed in 1970 by E. F. Codd at IBM.

1. True
2. False

The 1977 edition of this text contained a chapter on the relational model, and that chapter was reviewed by E. F. Codd.

1. True

2. False

dBase was the first PC-based DBMS to implement true relational algebra on a PC.

1. True
2. False

Paradox is the only major survivor of the "bloodbath of PC DBMS products."

1. True
2. False

Business organizations have resisted adopting object-oriented database systems because the cost of purchasing OODBMS packages is prohibitively high.

1. True
2. False

Bill Gates has said that "XML is the lingua-franca of the Internet Age."

1. True
2. False

XML Web services allow database processing to be shared across the Internet.

1. True
2. False

The NoSQL movement should really be called a NoRelational movement.

1. True
2. False

Twitter and Facebook use NoSQL databases.

1. True
2. False

The purpose of a database is to:

1. A) help people keep track of things.
2. B) store data in tables.
3. C) create tables of rows and columns.
4. D) maintain data on different things in different tables.
5. E) All of the above.

A database stores:

1. A) data.
2. B) relationships.
3. C) metadata.
4. D) A and B
5. E) A, B and C

A database records:

1. A) facts.
2. B) figures.
3. C) information.
4. D) A and B
5. E) A, B and C

A sales contact manager used by a salesperson is an example of a(n) _____.

1. A) single-user database application
2. B) multiuser database application
3. C) e-commerce database application
4. D) A or B
5. E) Any of A, B or C

A Customer Resource Management (CRM) system is an example of a(n) _____.

1. A) single-user database application
2. B) multiuser database application
3. C) e-commerce database application
4. D) A or B
5. E) Any of A, B or C

An online drugstore such as Drugstore.com is an example of a(n) _____.

1. A) single-user database application
2. B) multiuser database application
3. C) e-commerce database application
4. D) A or B
5. E) Any of A, B or C

The industry standard supported by all major DBMSs that allows tables to be joined together is called _____.

1. A) Sequential Query Language (SQL)
2. B) Structured Question Language (SQL)

3. C) Structured Query Language (SQL)
4. D) Relational Question Language (RQL)
5. E) Relational Query Language (RQL)

A program whose job is to create, process and administer databases is called the _____.

1. A) Database Modeling System
2. B) Database Management System
3. C) Data Business Model System
4. D) Relational Model Manager
5. E) Data Business Management Service

Microsoft Access includes:

1. A) a DBMS.
2. B) an application generator.
3. C) a Web server.
4. D) A and B
5. E) A, B and C

Microsoft Access may use which of the following DBMS engines?

1. A) ADE
2. B) SQL Server
3. C) Oracle
4. D) A and B
5. E) A, B and C

Which of the following are basic components of an enterprise-class database system?

1. A) The user
2. B) The database application
3. C) The database management system (DBMS)
4. D) The database
5. E) All of the above.

In an enterprise-class database system _____.

1. A) the database application(s) interact(s) with the DBMS
2. B) the database application(s) access(es) the database data
3. C) the DBMS accesses the database data
4. D) A and B
5. E) A and C

In an enterprise-class database system, the database application _____.

1. A) creates queries
2. B) creates forms
3. C) creates reports
4. D) A and B
5. E) B and C

In an enterprise-class database system, reports are created by _____.

1. A) the user
2. B) the database application
3. C) the database management system (DBMS)
4. D) the database
5. E) All of the above.

In database systems, the DBMS _____.

1. A) inserts data
2. B) modifies data
3. C) reads data
4. D) deletes data
5. E) All of the above.

In database systems, the DBMS enforces rules about which data can be written in certain columns. The rules are know as _____.

1. A) data insertion control
2. B) data modification control
3. C) data reading control
4. D) concurrency control
5. E) referential integrity constraints

In database systems, the DBMS enforces rules which user can perform which action when. The rules are know as _____.

1. A) data insertion control
2. B) data modification control
3. C) data reading control
4. D) concurrency control
5. E) referential integrity constraints

A database is considered "self-describing" because _____.

1. A) all the users` data is in one place
2. B) it reduces data duplication
3. C) it contains a description of its own structure

4. D) it contains a listing of all the programs that use it
5. E) All of the above.

A database consists of integrated tables, which store _____.

1. A) data
2. B) relationships among the data
3. C) forms
4. D) A and B
5. E) All of the above.

In database systems, the database _____.

1. A) holds user data
2. B) holds metadata
3. C) holds indexes
4. D) holds stored procedures
5. E) All of the above.

A database may contain _____.

1. A) tables
2. B) metadata
3. C) triggers
4. D) stored procedures
5. E) All of the above.

A database may be designed _____.

1. A) from existing data
2. B) as a new systems development project
3. C) as a redesign of an existing database
4. D) A and B
5. E) A, B, and C

A database designed using spreadsheets from the Sales department is a database being designed _____.

1. A) from existing data
2. B) as a new systems development project
3. C) as a redesign of an existing database
4. D) A and B
5. E) A, B, and C

A database designed to implement requirements for a reporting application needed by the Sales department is a database being designed _____.

1. A) from existing non-database data
2. B) as a new systems development project
3. C) as a redesign of an existing database
4. D) A and B
5. E) A, B, and C

A database designed to combine two databases used by the Sales department is a database being designed _____.

1. A) from existing data
2. B) as a new systems development project
3. C) as a redesign of an existing database
4. D) A and B
5. E) A, B, and C

Database professionals use _____ as specific data sources for studies and analyses.

1. A) data marts
2. B) normalization
3. C) data models
4. D) entity-relationship data modeling
5. E) data migration

Database professionals use a set of principles called _____ to guide and assess database design.

1. A) data marts
2. B) normalization
3. C) data models
4. D) entity-relationship data modeling
5. E) data migration

A very popular development technique used by database professionals for database design is known as _____.

1. A) data marts
2. B) normalization
3. C) data models
4. D) entity-relationship data modeling
5. E) data migration

A very popular development technique used by database professionals to adopt a database design to a new or changing requirement is known as _____.

1. A) data marts
2. B) normalization

3. C) data models
4. D) entity-relationship data modeling
5. E) data migration

The predecessor(s) of database processing was (were) _____.

1. A) file managers
2. B) hierarchical models
3. C) network models
4. D) relational data model
5. E) All of the above were predecessors of database processing.

The relational model _____.

1. A) was first proposed in 1970
2. B) was developed by E. F. Codd
3. C) was developed at IBM
4. D) resulted in the DBMS product DB2
5. E) All of the above.

Modern microcomputer personal DBMS products _____.

1. A) are supplied by several well-established manufacturers
2. B) were essentially killed off by Microsoft Access
3. C) have poor response time
4. D) are not true DBMS products
5. E) are really just programming languages with generalized file-processing capabilities

Business organizations have resisted adopting object-oriented database management systems because _____.

1. A) object-oriented programming uses simplified data structures that fit easily into relational databases
2. B) the cost of purchasing OODBMS packages is prohibitively high
3. C) the cost of converting data from relational databases to OODBMSs is too high
4. D) most large organizations have older applications that are not based on object oriented programming
5. E) C and D

For database development, the most important Web-related technology to emerge in recent years is:

1. A) FTP.
2. B) HTTP.

3. C) XML.
4. D) OODBMS.
5. E) All of the above.

For database development, a challenge to the relational model and the use of SQL has recently developed. This challenge is known as:

1. A) the Web services movement.
2. B) the NoSQL movement.
3. C) the SOAP movement.
4. D) the NoSOAP movement.
5. E) the UDDI movement.

The purpose of a database is to help people keep track of things.

1. True
2. False

In a database, each table stores data about a different type of thing.

1. True
2. False

In a database, each row in a spreadsheet has data about a particular instance.

1. True
2. False

In every database, not just the databases discussed in this book, table names are capitalized.

1. True
2. False

A database shows data in tables and the relationships among the rows in those tables.

1. True
2. False

Data is recorded facts and figures; information is knowledge derived from data.

1. True
2. False

Databases record data in such a way that they can produce information.

1. True
2. False

Enterprise Resource Planning (ERP) is an example of a data mining application.

1. True
2. False

Databases are a key component of e-commerce order entry, billing, shipping and customer support.

1. True
2. False

The largest databases in e-commerce are the order entry databases.

1. True
2. False

The e-commerce companies use Web activity database to determine which items on a Web page are popular and successful.

1. True
2. False

Small databases typically have simple structures.

1. True
2. False

A database system is typically defined as the four components users, database applications, the DBMS and the databases.

1. True
2. False

A database system as typically defined can be modified to include CODASYL.

1. True
2. False

Applications are computer programs used directly by users.

1. True
2. False

In a database system, applications write data to the database.

1. True
2. False

Sequenced Query Language (SQL) is an internationally recognized standard language that is understood by all commercial database management system products.

1. True
2. False

In database systems, indexes are held by the database.

1. True
2. False

A database management system (DBMS) creates, processes and administers databases.

1. True
2. False

The database management system (DBMS) is responsible for inserting, modifying, reading, and deleting data.

1. True
2. False

The database application is responsible for concurrency control.

1. True
2. False

The database management system (DBMS) is responsible for enforcing referential integrity constraints.

1. True
2. False

Referential integrity constraints are rules about what data values are allowed in certain columns.

1. True
2. False

A database is a self-describing collection of non-integrated tables.

1. True

2. False

Integrated tables store both data and the relationships among the data.

1. True
2. False

Microsoft Access is just a DBMS.

1. True
2. False

Microsoft Access is a low-end product intended for individuals and small workgroups.

1. True
2. False

The current DBMS engine in Microsoft Access is called ADE.

1. True
2. False

In Microsoft Access, you can use the Oracle DBMS in place of the ADE DBMS.

1. True
2. False

In an Enterprise-class database system, a database application interacts with the DBMS.

1. True
2. False

In an Enterprise-class database system, a database application accesses the database data.

1. True
2. False

In an Enterprise-class database system, business users interact directly with the DBMS, which directly accesses the database data.

1. True
2. False

All database applications get and put database data by sending SQL statements to the DBMS.

1. True

2. False

The DBMS ranked as having the "most power and features" in the text is IBM's DB2.

1. True
2. False

The DBMS ranked as being the "most difficult to use" in the text is Oracle Corporation's Oracle Database.

1. True
2. False

The DBMS ranked as being the "easiest to use" in the text is Microsoft's SQL Server.

1. True
2. False

The DBMS ranked as having the "least power and features" in the text is Microsoft Access.

1. True
2. False

A database is called "self-describing" because it reduces data duplication.

1. True
2. False

The description of a database's structure that is stored within the database itself is called the "metadata."

1. True
2. False

In a database processing system, indexes are held by the database management system (DBMS).

1. True
2. False

Database design is important, but fortunately it is simple to do.

1. True
2. False

A database design may be a new systems development project.

1. True
2. False

A database design is rarely a redesign of an existing database.

1. True
2. False

Information systems that stored groups of records in separate files were called file processing systems.

1. True
2. False

Data Language/I (DL/I) structured data relationships as a tree structure.

1. True
2. False
3.

1) The SQL CREATE TABLE statement is used to name a new table and describe the table's columns. Answer: TRUE

4. Diff: 1 Page Ref: 281 2) The SQL keyword CONSTRAINT is used to define one of five types of constraints. Answer: TRUE
5. Diff: 1 Page Ref: 283 3) The SQL keyword PRIMARY KEY is used to designate the column(s) that are the primary key for the table. Answer: TRUE
6. Diff: 1 Page Ref: 283, 289 4) The SQL keyword CONSTRAINT is used to limit column values to specific values. Answer: TRUE

7. Diff: 2 Page Ref: 283, 289 5) The SQL keyword CONSTRAINT is used in conjunction with the SQL keywords PRIMARY KEY and FOREIGN KEY. Answer: TRUE
8. Diff: 1 Page Ref: 289-290 6) One advantage of using the CONSTRAINT command to define a primary key is that the database designer controls the name of the constraint. Answer: TRUE
9. Diff: 3 Page Ref: 289-290 7) The SQL keyword UNIQUE is used to define alternative keys. Answer: TRUE
10. Diff: 2 Page Ref: 290 8) If the table PRODUCT has a column PRICE, and PRICE has the data type Numeric (8,2), the value 98765 stored in that field will be displayed by the DBMS as 98765.00. Answer: FALSE
11. Diff: 3 Page Ref: 289 Fig 7-4 9) If the table ITEM has a column WEIGHT, and WEIGHT has the data type Numeric (7,2), the value 4321 with be displayed by the DBMS as 43.21. Answer: TRUE
12. Diff: 2 Page Ref: 289 Fig 7-4
13. 10) The SQL keyword CHECK is used to limit column values to specific values. Answer: TRUE
14. Diff: 2 Page Ref: 289, 292 11) The SQL keyword MODIFY is used to change the structure, properties or constraints of a table. Answer: FALSE
15. Diff: 1 Page Ref: 298 12) Data values to be added to a table are specified by using the SQL VALUES clause. Answer: TRUE

16. Diff: 2 Page Ref: 300-301 13) The SQL keyword DELETE is used to delete a table's structure. Answer: FALSE
17. Diff: 1 Page Ref: 298-299 14) When the correct SQL command is used to delete a table's structure, the command can only be used with a table that has already had its data removed. Answer: FALSE
18. Diff: 2 Page Ref: 298-299 15) One or more rows can be added to a table by using the SQL INSERT statement. Answer: TRUE
19. Diff: 2 Page Ref: 300-301 16) Unless it is being used to copy data from one table to another, the SQL INSERT statement can be used to insert only a single row into a table. Answer: TRUE
20. Diff: 3 Page Ref: 300-301 17) Rows in a table can be changed by using the SQL UPDATE statement. Answer: TRUE
21. Diff: 2 Page Ref: 307-308 18) The SQL SET keyword is used to specify a new value when changing a column value. Answer: TRUE
22. Diff: 3 Page Ref: 307-308 19) The SQL keyword MODIFY is used to change a column value. Answer: FALSE
23. Diff: 3 Page Ref: 307-308 20) Rows can be removed from a table by using the SQL DELETE statement. Answer: TRUE Diff: 2 Page Ref: 309

1) The SQL CREATE TABLE statement is used to name a new table and describe the table's columns. Answer: TRUE

Diff: 1 Page Ref: 281

2) The SQL keyword CONSTRAINT is used to define one of five types of constraints.
Answer: TRUE
Diff: 1 Page Ref: 283

3) The SQL keyword PRIMARY KEY is used to designate the column(s) that are the primary key for the table. Answer: TRUE
Diff: 1 Page Ref: 283, 289

4) The SQL keyword CONSTRAINT is used to limit column values to specific values.
Answer: TRUE
Diff: 2 Page Ref: 283, 289

5) The SQL keyword CONSTRAINT is used in conjunction with the SQL keywords PRIMARY KEY and FOREIGN KEY. Answer: TRUE
Diff: 1 Page Ref: 289-290

6) One advantage of using the CONSTRAINT command to define a primary key is that the database designer controls the name of the constraint. Answer: TRUE
Diff: 3 Page Ref: 289-290

7) The SQL keyword UNIQUE is used to define alternative keys. Answer: TRUE
Diff: 2 Page Ref: 290

8) If the table PRODUCT has a column PRICE, and PRICE has the data type Numeric (8,2), the value 98765 stored in that field will be displayed by the DBMS as 98765.00.
Answer: FALSE
Diff: 3 Page Ref: 289 Fig 7-4

9) If the table ITEM has a column WEIGHT, and WEIGHT has the data type Numeric (7,2), the value 4321 with be displayed by the DBMS as 43.21. Answer: TRUE
Diff: 2 Page Ref: 289 Fig 7-4

10) The SQL keyword CHECK is used to limit column values to specific values. Answer: TRUE
Diff: 2 Page Ref: 289, 292
11) The SQL keyword MODIFY is used to change the structure, properties or constraints of a table. Answer: FALSE
Diff: 1 Page Ref: 298

12) Data values to be added to a table are specified by using the SQL VALUES clause.
Answer: TRUE
Diff: 2 Page Ref: 300-301

13) The SQL keyword DELETE is used to delete a table's structure. Answer: FALSE

Diff: 1 Page Ref: 298-299

14) When the correct SQL command is used to delete a table's structure, the command can only be used with a table that has already had its data removed. Answer: FALSE
Diff: 2 Page Ref: 298-299

15) One or more rows can be added to a table by using the SQL INSERT statement.
Answer: TRUE
Diff: 2 Page Ref: 300-301

16) Unless it is being used to copy data from one table to another, the SQL INSERT statement can be used to insert only a single row into a table. Answer: TRUE
Diff: 3 Page Ref: 300-301

17) Rows in a table can be changed by using the SQL UPDATE statement. Answer: TRUE
Diff: 2 Page Ref: 307-308

18) The SQL SET keyword is used to specify a new value when changing a column value. Answer: TRUE
Diff: 3 Page Ref: 307-308

19) The SQL keyword MODIFY is used to change a column value. Answer: FALSE
Diff: 3 Page Ref: 307-308

20) Rows can be removed from a table by using the SQL DELETE statement. Answer: TRUE
Diff: 2 Page Ref: 309

21) An SQL virtual table is called a view.
Answer: TRUE
Diff: 1 Page Ref: 309

22) The SQL command CREATE USER VIEW is used to create a virtual table. Answer: FALSE
Diff: 2 Page Ref: 309
23) SQL views are constructed from SELECT statements.
Answer: TRUE
Diff: 2 Page Ref: 309-311

24) According to the SQL-92, statements used to construct views cannot contain the WHERE clause. Answer: FALSE
Diff: 2 Page Ref: 309

25) The SQL command SELECT is used to retrieve view instances. Answer: TRUE
Diff: 1 Page Ref: 309

26) The values in an SQL view are not always changeable through the view itself.
Answer: TRUE
Diff: 2 Page Ref: 318-319 Fig 7-19

27) SQL views can be used to hide columns.
Answer: TRUE
Diff: 1 Page Ref: 311-313 Fig 7-17

Chapter 8: Database Redesign

1) Database redesign is rarely needed because databases are usually built correctly the first time.
Answer: FALSE
Diff: 1 Page Ref: 367

2) In a real sense, information systems and organizations do not just influence each other, but rather they create each other. Answer: TRUE
Diff: 1 Page Ref: 367

3) A continuous circular process of changes in user behaviors and change in the information systems they use is a natural outcome of information system use. Answer: TRUE
Diff: 1 Page Ref: 367

4) The continuous circular process of changes is known as the Systems Development Life Cycle (SDLC). Answer: TRUE
Diff: 1 Page Ref: 367

5) Database redesign is equally difficult whether or not the database has data in it. Answer: FALSE
Diff: 1 Page Ref: 367

6) In the database redesign process, it is often useful to test whether certain conditions or assumptions are valid before proceeding with the redesign. Answer: TRUE
Diff: 1 Page Ref: 3

7) In the database redesign process, two SQL tools are useful for testing whether or not certain conditions or assumptions are valid: uncorrelated subqueries and EXISTS/NOT EXISTS.
Answer: FALSE
Diff: 2 Page Ref: 367

8) A correlated subquery looks very different from a noncorrelated subquery. Answer: FALSE

Diff: 2 Page Ref: 368

9) Correlated subqueries can be used to verify functional dependencies. Answer: TRUE
Diff: 2 Page Ref: 367-371

10) In a correlated subquery, the same table is used in the upper and lower SELECT statements.
Answer: TRUE
Diff: 2 Page Ref: 368-370
11) In the SQL statement:

SELECT S1.CustName, S1.SalesRepNo FROM SALES S1;

the "S1" is called an alias.
Answer: TRUE
Diff: 2 Page Ref: 369

12) In a correlated subquery, the DBMS can run the lower SELECT statement by itself and then send the results to the upper SELECT statement. Answer: FALSE
Diff: 2 Page Ref: 368-370

13) In a correlated subquery, the DBMS must run the lower SELECT statement as a process that is nested within the upper SELECT statement. Answer: TRUE
Diff: 1 Page Ref: 368-370

14) There is a common trap in writing a correlated subquery, which will cause no rows to ever be displayed in the results. Answer: TRUE
Diff: 2 Page Ref: 370

15) Although correlated subqueries are useful in database redesign, they cannot be used to verify functional dependencies. Answer: FALSE
Diff: 1 Page Ref: 370-371

16) When using queries with EXISTS and NOT EXISTS, the processing of the associated SELECT statements must be nested. Answer: TRUE
Diff: 2 Page Ref: 371

17) The use of a double nested set of NOT EXISTS SELECT statements can be used to find rows that meet some specified condition for every row in a table. Answer: TRUE
Diff: 3 Page Ref: 372

18) The use of a double nested set of NOT EXISTS SELECT statements is a famous pattern in SQL use. Answer: TRUE
Diff: 2 Page Ref: 372

19) EXISTS and NOT EXISTS are actually just another form of correlated subqueries. Answer: TRUE

Diff: 1 Page Ref: 371

20) Because EXISTS and NOT EXISTS are forms of correlated subqueries, the processing of the associated SELECT statements must be nested. Answer: TRUE
Diff: 2 Page Ref: 371
21) The EXISTS keyword will be true if any row in the subquery meets the condition. Answer: TRUE
Diff: 2 Page Ref: 373

22) The NOT EXISTS keyword will be true if any row in the subquery fails to meet the condition. Answer: FALSE
Diff: 2 Page Ref: 372

23) A double nested set of NOT EXISTS SELECT statements can be used to find rows that meet some specified condition for every row in a table. Answer: TRUE
Diff: 3 Page Ref: 372

24) When using a double nested set of NOT EXISTS SELECT statements, a row that does not match any row matches every row. Answer: TRUE
Diff: 3 Page Ref: 372...

Database Processing, 13e (Kroenke/Auer)
Chapter 9: Managing Multiuser Databases

1) Database administration tasks have to be performed for single-user, personal databases. Answer: TRUE
Diff: 3 Page Ref: 396

2) Database administration is more important but less difficult in multiuser database systems than in single-user database systems. Answer: FALSE
Diff: 2 Page Ref: 396

3) In general, the overall responsibility of the DBA is to facilitate the development and use of the database system. Answer: TRUE
Diff: 1 Page Ref: 396

4) The DBA has to find a balance between the conflicting goals of maximizing availability of the database to users and protecting the database. Answer: TRUE
Diff: 2 Page Ref: 396

5) The DBA is responsible for managing changes to the database structure, but is rarely involved in the original design of the structure. Answer: FALSE
Diff: 2 Page Ref: 396

6) Changes in the database structure usually involve only one application. Answer: FALSE

Diff: 3 Page Ref: 396

7) The database is most vulnerable to failure after a change to its structure. Answer: TRUE
Diff: 3 Page Ref: 397

8) One important reason for documenting changes to the database structure is for diagnosing errors. Answer: TRUE
Diff: 1 Page Ref: 397

9) Concurrency control measures are taken to ensure that one user's work has absolutely no influence on another user's work. Answer: FALSE
Diff: 2 Page Ref: 397

10) A transaction is sometimes called atomic because it is performed as a unit. Answer: TRUE
Diff: 1 Page Ref: 398

11) A transaction is a series of actions to be taken on the database such that either all of them are performed successfully or none of them is performed at all. Answer: TRUE
Diff: 1 Page Ref: 398

12) A transaction is a group of alternative database actions, from which the database can choose to perform only one of them. Answer: FALSE
Diff: 1 Page Ref: 398

13) The size of a lock is referred to as the lock granularity. Answer: TRUE
Diff: 3 Page Ref: 401

14) An exclusive lock locks the item from access of any type. Answer: TRUE
Diff: 1 Page Ref: 401

15) An exclusive lock locks the item from change but not from read access. Answer: FALSE
Diff: 1 Page Ref: 402

16) Resource locking is one remedy to the lost update problem. Answer: TRUE
Diff: 2 Page Ref: 401

17) Locks placed by a command issued to the DBMS from the application program are called explicit locks. Answer: FALSE
Diff: 2 Page Ref: 401

18) Locks placed automatically by the DBMS are called implicit locks. Answer: TRUE
Diff: 2 Page Ref: 401

19) Explicit locks are locks that are placed automatically by the DBMS. Answer: FALSE
Diff: 1 Page Ref: 401

20) Locks with large granularity are easy for the DBMS to administer but frequently cause

conflicts. Answer: TRUE
Diff: 2 Page Ref: 401

21) Two-phased locking is a scheme for achieving serializability of transactions. Answer: TRUE
Diff: 3 Page Ref: 402

22) Two-phased locking has a growing phase and a shrinking phase. Answer: TRUE
Diff: 3 Page Ref: 402
23) In two-phase locking, all locks are obtained during the growing phase. Answer: TRUE
Diff: 2 Page Ref: 402

24) In two-phase locking, all locks are released during the contracting phase. Answer: FALSE
Diff: 2 Page Ref: 402

25) Resource locking must be carefully planned because most DBMS products cannot detect a deadlock condition. Answer: FALSE
Diff: 2 Page Ref: 402-403

26) Requiring all application programs to lock resources in the same order is one way of preventing a deadlock condition. Answer: TRUE
Diff: 3 Page Ref: 402-403

27) With optimistic locking, the assumption is made that no conflict will occur. Answer: TRUE
Diff: 1 Page Ref: 403

28) With pessimistic locking, the assumption is made that a conflict will occur....

chapter 9

Managing Multi-User Databases

True-False Questions

1. Database administration tasks have to be performed for single-user, personal databases.

   Answer: *True* ***Level:*** *hard*

   Page: *300*

2. Database administration is more important but less difficult in multi-user database systems than in single-user database systems.

Answer: *False* ***Level:*** *moderate*

Page: *300*

3. In general, the overall responsibility of the DBA is to facilitate the development and use of the database system.

Answer: *True* ***Level:*** *easy*

Page: *301*

4. The DBA has to find a balance between the conflicting goals of maximizing availability of the database to users and protecting the database.

Answer: *True* ***Level:*** *moderate*

Page: *301*

5. The DBA is responsible for managing changes to the database structure, but is rarely involved in the original design of the structure.

Answer: *False* ***Level:*** *moderate*

Page: *301*

6. Changes in the database structure usually involve only one application.

Answer: *False* ***Level:*** *hard*

Page: *301*

7. One important reason for documenting changes to the database structure is for diagnosing errors.

Answer: *True* ***Level:*** *easy*

Page: *301*

8. Concurrency control measures are taken to ensure that one user's work has absolutely no influence on another user's work.

Answer: *False* ***Level:*** *moderate*

Page: *302*

9. A transaction is a group of alternative database actions from which the database can choose to perform only one of them.

Answer: *False* ***Level:*** *easy*

Page: *303*

10. "Resource locking" is one remedy to the lost update problem.

Answer: *True* ***Level:*** *moderate*

Page: *306*

11. "Explicit locks" are locks that are placed automatically by the DBMS.

Answer: *False* ***Level:*** *easy*

Page: *307*

12. Locks with large granularity are easy for the DBMS to administer but frequently cause conflicts.

Answer: *True* ***Level:*** *moderate*

Page: *307*

13. In general, the boundaries of a transaction should correspond to the boundaries of the database view it is processing.

Answer: *True* ***Level:*** *hard*

Page: *308*

14. Resource locking must be carefully planned because most DBMS products cannot detect a deadlock condition.

Answer: *False* ***Level:*** *moderate*

Page: *309*

15. Resources are locked for a shorter amount of time with pessimistic locking because

the transaction is pre-processed.

Answer: *False* ***Level:*** *easy*

Page: *309*

16. In general, optimistic locking is the preferred technique for Internet databases.

Answer: *True* ***Level:*** *hard*

Page: *309*

17. A "dirty read" happens when one transaction reads a changed record that has not been committed to the database.

Answer: *True* ***Level:*** *hard*

Page: *312*

18. “Repeatable Reads” isolation is the most restrictive level of isolation.

Answer: *False* ***Level:*** *moderate*

Page: *312*

19. The goal of database security is to ensure that only authorized users can perform authorized activities at authorized times.

Answer: *True* ***Level:*** *easy*

Page: *304*

20. In regard to database security, neither the DBMS nor the database applications can enforce processing responsibilities.

Answer: *True* ***Level:*** *moderate*

Page: *305*

21. Processing responsibilities should be documented and encoded into manual procedures.

Answer: *True* ***Level:*** *easy*

Page: *305*

22. Processing rights may be implemented at the DBMS level.

Answer: *True* ***Level:*** *easy*

Page: *305*

23. All commercial DBMS products use some version of "username and password" as part of their security features.

Answer: *True* ***Level:*** *easy*

Page: *316*

24. The security provided by the DBMS often has to be augmented by additional security features within the application program.

Answer: *True* ***Level:*** *moderate*

Page: *319*

25. A "database save" is used to mark the end of a transaction.

Answer: *False* ***Level:*** *moderate*

Page: *320*

26. Reprocessing is normally the most convenient method for recovery after a system failure.

Answer: *False* ***Level:*** *easy*

Page: *320*

27. Rollforward and reprocessing are two different names for the same technique.

Answer: *False* ***Level:*** *hard*

Page: *320*

28. Both rollforward and rollback require the use of a log of transaction results.

Answer: *True* ***Level:*** *easy*

Page: *320*

29. The DBA should periodically analyze run-time statistics of database performance to help manage the DBMS.

Answer: *True* ***Level:*** *easy*

Page: *323*

30. A passive repository is preferred over an active repository because it requires less human intervention.

Answer: *False* ***Level:*** *moderate*

Page: *324*

Multiple Choice Questions

31. Which of the following is not a database administration responsibility of a DBA?

a.) managing the database structure

b.) managing data activity

c.) managing the DBMS

d.) maintaining the data repository

e.) All of the above are database administration responsibilities of a DBA.

Level: *easy*

Page: *300-301 [See Figure 9-1]*

32. Which of the following is true about making changes to the database structure?

a.) The DBA need not get input from users on the issue because it is a technical decision.

b.) Formal policies and procedures for requesting a change are not used because they are too limiting.

c.) Documentation of when the change was made, how it was made, and why it was made must be created.

d.) Changes do not produce unexpected results because the DBA will have investigated the change thoroughly before implementing it.

e.) If the database is properly designed, changes should not be necessary throughout the system's lifetime.

Level: *hard*

Page: *301*

33. The task of diagnosing errors due to changes in the database structure is eased by:

a.) formal policies for requesting changes.

b.) database structure change documentation.

c.) rollback analysis.

d.) configuration control.

e.) None of the above.

Level: *moderate*

Page: *301-302*

34. Measures that are taken to prevent one user's work from inappropriately influencing another user's work are called:

a.) concurrency control.

b.) checkpoint.

c.) database recovery.

d.) database logging.

e.) interleaving.

Level: *easy*

Page: *302*

35. A series of actions to be taken on the database such that either all actions are completed successfully, or none of them can be completed, is known as a(n):

a.) checkpoint.

b.) log.

c.) lock.

d.) transaction.

e.) concurrent.

Level: *easy*

Page: *303*

36. When two transactions are being processed against the database at the same time,

a.) they are called concurrent transactions.

b.) they are usually interleaved.

c.) they always result in a lost update problem.

d.) one must be rolled back.

e.) both a and b

Level: *easy*

Page: *303*

37. The situation that occurs when one user's changes to the database are lost by a second user's changes to the database is known as the:

a.) concurrent update problem.

b.) deadly embrace problem.

c.) inconsistent read problem.

d.) inconsistent write problem.

e.) deadlock problem.

Level: *hard*

Page: *303-306*

38. One remedy for the inconsistencies caused by concurrent processing is _____.

a.) lost updates

b.) checkpointing

c.) rollback

d.) resource locking

e.) concurrency

Level: *easy*

Page: *306*

39. A lock placed automatically by the DBMS is called a(n) _____ lock.

a.) exclusive

b.) explicit

c.) granular

d.) implicit

e.) shared

Level: *moderate*

Page: *307*

40. Which of the following is <u>not</u> true about locks?

a.) Locks with large granularity are easier for the DBMS to administer.

b.) Locks with small granularity cause more conflicts.

c.) Locks with large granularity produce fewer details for the DBMS to track.

d.) Locks may have a table-level granularity.

e.) Locks may have a database-level granularity.

Level: *hard*

Page: *307*

41. Which type of lock prevents all types of access to the locked resource?

a.) exclusive lock

b.) shared lock

c.) two-phased lock

d.) explicit lock

e.) implicit lock

Level: *easy*

Page: *307*

42. Which type of lock still allows other transactions to have read-only access to the locked resource?

a.) exclusive lock

b.) shared lock

c.) two-phased lock

d.) explicit lock

e.) implicit lock

Level: *easy*

Page: *307*

43. Which of the following is not true about two-phased locking?

a.) can make transactions serializable

b.) uses only shared locks

c.) has a growing phase

d.) has a shrinking phase

e.) cannot obtain a new lock once a lock has been released

Level: *hard*

Page: *307-308*

44. The situation that occurs when two users are each waiting for a resource that the other person has locked is known as a(n):

a.) lost update problem.

b.) deadlock.

c.) inconsistent read problem.

d.) inconsistent write problem.

e.) checkpoint.

Level: *moderate*

Page: *308*

45. Requiring all application programs to lock resources in the same order is a technique for preventing what problem?

a.) concurrent update

b.) lost update

c.) deadlock

d.) exclusive locks

e.) growing phase locking

Level: *hard*

Page: *308*

46. Locks that are placed assuming that a conflict will occur are called:

a.) dynamic locks.

b.) explicit locks.

c.) implicit locks.

d.) optimistic locks.

e.) pessimistic locks.

Level: *moderate*

Page: *309*

47. Locks that are placed assuming that a conflict will not occur are called:

a.) dynamic.

b.) explicit.

c.) implicit.

d.) optimistic.

e.) pessimistic.

Level: *moderate*

Page: *309*

48. Ensuring that all rows impacted by the actions of a transaction are protected from changes until the entire transaction is completed is called:

a.) statement level consistency.

b.) optimistic locking.

c.) transaction level consistency.

d.) durable transactions.

e.) ARID transactions.

Level: *hard*

Page: *312*

49. Which of the following is allowed by "Repeatable Read Isolation?"

a.) nonrepeatable reads

b.) dirty reads

c.) phantom reads

d.) a and b

e.) a, b and c

Level: *hard*

Page: *312-313 [See Figure 9-10]*

50. Which of the following is true of forward only cursors?

a.) Current values for each row are retrieved when the application accesses a row.

b.) All changes of any type from any source are visible.

c.) Changes made by the transaction are visible only if they occur on rows ahead of the cursor.

d.) Applications may scroll backward in the record set.

e.) It requires the greatest overhead of any cursor type.

Level: *moderate*

Page: *313-314 [See Figure 9-11]*

51. Which of the following cannot be enforced in the DBMS or application programs?

a.) processing rights

b.) security

c.) processing responsibilities

d.) cursors

e.) transaction isolation

Level: *moderate*

Page: *315*

52. Once processing rights have been defined, they may be implemented at any of these levels except:

a.) network.

b.) operating system.

c.) data.

d.) DBMS.

e.) application.

Level: *hard*

Page: *315*

53. Which of the following is not true of DBMS security features?

a.) Users may be assigned to one or more roles.

b.) A role may be assigned to only one user.

c.) Both users and roles can have many permissions.

d.) Objects have many permissions.

e.) Each permission pertains to one user or role and one object.

Level: *moderate*

Page: *316*

54. Recovering a database via reprocessing involves:

a.) restoring the database from the save and reprocessing all the transactions since the save.

b.) restoring the database from the save and reapplying all the changes made by transactions since the save.

c.) undoing the changes made by erroneous or partially processed transactions, and restarting the valid transactions that were in process at the time of the failure.

d.) recreating the database by reentering all of the data from the beginning, and then reprocessing all of the transactions.

e.) synchronizing the database and the transaction log by checkpointing.

Level: *easy*

Page: *320*

55. Recovering a database via rollforward involves:

a.) restoring the database from the save and reprocessing all the transactions since the save.

b.) restoring the database from the save and reapplying all the changes made by transactions since the save.

c.) undoing the changes made by erroneous or partially processed transactions, and restarting the valid transactions that were in process at the time of the failure.

d.) re-creating the database by re-entering all of the data from the beginning, and then reprocessing all of the transactions.

e.) synchronizing the database and the transaction log by checkpointing.

Level: *easy*

Page: *320*

56. Recovering a database via rollback involves:

a.) restoring the database from the save and reprocessing all the transactions since the save.

b.) restoring the database from the save and reapplying all the changes made by transactions since the save.

c.) undoing the changes made by erroneous or partially processed transactions, and restarting the valid transactions that were in process at the time of the failure.

d.) re-creating the database by re-entering all of the data from the beginning and, then reprocessing all of the transactions.

e.) synchronizing the database and the transaction log by checkpointing.

Level: *easy*

Page: *320*

57. Which of the following would <u>not</u> be contained in a transaction log?

a.) before-images

b.) type of operation

c.) pointers

d.) time of the action

e.) permissions

Level: *moderate*

Page: *320*

58. Which of the following would a DBA do in managing the DBMS?

a.) analyze system performance statistics

b.) investigate user complaints

c.) evaluate new DBMS product features

d.) tune DBMS product options to accommodate other software in use

e.) All of the above.

Level: *easy*

Page: *323*

59. Which of the following is not true of data repositories?

a.) They are usually created after the database has been implemented and optimized for performance.

b.) They may be virtual.

c.) They may contain metadata about database applications.

d.) They may contain metadata about users.

e.) They may contain metadata about web pages.

Level: *moderate*

Page: *324*

60. Which type of data repository is composed of metadata that is created automatically as the system components are created?

a.) passive

b.) dynamic

c.) active

d.) automatic

e.) summary

Level: *moderate*

Page: *324*

Fill in the Blank Questions

61. The overall responsibility of the *DBA* is to facilitate the development and use of the database.

Level: *easy*

Page: *301*

62. The database is most vulnerable to failure after a change to its *structure*.

Level: *hard*

Page: *301*

63. A(n) *transaction* is a series of actions to be taken on the database such that either all of them are performed successfully or none of them is performed at all.

Level: *easy*

Page: *303*

64. A transaction is sometimes called *atomic*, since it is performed as a unit.

Level: *easy*

Page: *303*

65. Locks placed automatically by the DBMS are called *implicit* locks.

Level: *moderate*

Page: *307*

66. Locks placed by a command issued to the DBMS from the application program are called *explicit* locks.

Level: *moderate*

Page: *307*

67. The size of a lock is referred to as the lock *granularity*.

Level: *hard*

Page: *307*

68. A(n) *exclusive* lock locks the item from access of any type.

Level: *easy*

Page: *307*

69. A(n) *shared* lock locks the item from change but not from read access.

Level: *easy*

Page: *307*

70. Two-phased locking is a scheme for achieving *serializability* of transactions.

Level: *hard*

Page: *307*

71. In two-phase locking, all locks are obtained during the *growing* phase.

Level: *moderate*

Page: *308*

72. In two-phase locking, all locks are released during the *shrinking* phase.

Level: *moderate*

Page: *308*

73. Requiring all application programs to lock resources in the same order is one way of preventing a *deadlock* condition.

Level: *hard*

Page: *308*

74. With *optimistic* locking, the assumption is made that no conflict will occur.

Level: *easy*

Page: *309*

75. With *pessimistic* locking, the assumption is made that a conflict will occur.

Level: *easy*

Page: *309*

76. The transaction *boundaries* are the essential information that the DBMS needs from the application programs to enforce different locking strategies.

Level: *hard*

Page: *309-310*

77. A(n) *durable* transaction is one for which all committed changes are permanent.

Level: *hard*

Page: *311*

78. Transaction level *consistency* means that all rows impacted by any actions in a transaction are protected from change during the entire transaction.

Level: *moderate*

Page: *312*

79. A(n) *dirty read* occurs when one transaction reads a changed record that has not been committed to the database.

Level: *hard*

Page: *312*

80. *Nonrepeatable reads* occur when a transaction rereads data it has previously read and finds modifications or deletions caused by a committed transaction.

Level: *hard*

Page: *312*

81. According to ANSI SQL, the *serializable* isolation level will not allow phantom reads, dirty reads, and nonrepeatable reads.

Level: *hard*

Page: *312*

82. A(n) *static* cursor processes a snapshot of the relation that was taken when the cursor was opened.

Level: *moderate*

Page: *313*

83. A(n) *keyset* cursor saves primary key values when the cursor is opened and retrieves the

values for each row as the application program accesses it.

Level: *moderate*

Page: *313*

84. __*Rollforward*__ is a method of database recovery that restores the database save and all valid transactions since the save was reapplied.

Level: *easy*

Page: *320*

85. To support rollforward and rollback recovery, transactions must be written to a(n) __*log*__ before they are applied to the database.

Level: *easy*

Page: *320*

86. Copies of each database record or page before it was changed by a transaction that are saved for use in database recovery are called __*before-images*__.

Level: *easy*

Page: *320*

87. Copies of each database record or page after it was changed by a transaction that are saved for use in database recovery are called *after-images*.

Level: *easy*

Page: *320*

88. A(n) *checkpoint* is a point of synchronization between the database and the transaction log.

Level: *moderate*

Page: *322*

89. A(n) active data repository is one in which the metadata is automatically created as the system components are created.

Level: *easy*

Page: *324*

90. A(n) passive data repository is one which requires a person to generate the metadata and place it in the repository.

Level: *easy*

Page: *324*

10) The Microsoft SQL Server Management Studio cannot verify SQL statements until they areexecuted.Answer:FALSEDiff: 1Page Ref: 10A-1411) Stored procedures and triggers in SQL Server may become confused while executing storedprocedures and triggers when SQL Server special words such asTRANSACTIONappear as tableor other names, even if they are enclosed in square brackets [ ].Answer:TRUEDiff: 2Page Ref: 10A-1212) The name TRANSACTION is so special to SQL Server that no stored procedures will workon a table with that name.Answer:TRUEDiff: 3Page Ref: 10A-1213) In the Microsoft SQL Server Management Studio, the starting

value for a field that has beenset as anIDENTITYin SQL Server is set in the seed (StartingValue) property.Answer:TRUEDiff: 3Page Ref: 10A-1214) In the Microsoft SQL Server Management Studio, the amount to add to the surrogate keyvalue when adding a new row is specified by the increment property.Answer:TRUEDiff: 3Page Ref: 10A-1215) When viewing a table in the GUI tools table design window in SQL Server, the primary keyis indicated by a black arrow.Answer:FALSEDiff: 2Page Ref: Fig 10A-1716) When viewing a table in the GUI tools table design window in SQL Server, the notation thatshould supply the values for a surrogate key is that theIdentityproperty of that column is set to"Yes."Answer:TRU

Database Processing, 13e(Kroenke/Auer)Online Chapter 10C:Managing Databases with MySQL 5.61) MySQL is an open source DBMS that only runs on Unix and Linux operating systems.Answer:FALSEDiff: 1Page Ref: 10C-32) Microsoft Windows users should install MySQL Community Edition using the MySQLInstaller for

Windows.Answer:TRUEDiff: 1Page Ref: 10C-33) The MySQL Workbench is included in the MySQL Server Community Edition installation.Answer:FALSEDiff: 1Page Ref: 10C-34) MySQL is configured using the MySQL Server Instance Configuration Wizard.Answer:TRUEDiff: 1Page Ref: 10C-55) For use with this book, it is recommended that MySQL be configured as a developer machine.Answer:TRUEDiff: 1Page Ref: 10C-66) For use with this book, it is recommended that MySQL be configured as a Decision Support(DSS)/OLAP machine.Answer:TRUEDiff: 1Page Ref: 10C-67) It is recommended that MySQL be used with the InnoDB storage engine.Answer:TRUEDiff: 1Page Ref: 10C-6-10C-78) To start working with MySQL Community Edition, use the command Start | All Programs |MySQL | MySQL Workbench CE 5.2.Answer:TRUEDiff: 1Page Ref: 10C-89) Users log into MySQL Workbench using the Connect to Database and Connect to MySQLServer dialog boxes.Answer:TRUEDiff: 1Page Ref: 10C-9-10C-1210) MySQL Community Edition uses operating system authentication.Answer:FALSEDiff: 1Page Ref: 10C-9-10C-1

11) Schema is the MySQL synonym for database.Answer:TRUEDiff: 1Page Ref: 10C-1212) Database objects are displayed in the Object Explorer window in the MySQL Workbench.Answer:TRUEDiff: 2Page Ref: 10C-10-10C-1313) A new MySQL database is created using the New Database dialog box.Answer:FALSEDiff: 2Page Ref: 10Cb-10-10B-1214) The MySQL Command Line Client is the MySQL command line utility.Answer:TRUEDiff: 2Page Ref: 10C-14-10C-1515) MySQL Workbench is the MySQL GUI application development utility.Answer:TRUEDiff: 2Page Ref: 10C-7, 10C-1616) MySQL Workbench is the MySQL GUI database administration utility.Answer:TRUEDiff: 2Page Ref: 10C-717) The MySQL Table Editor is accessible from the MySQL Workbench.Answer:TRUEDiff: 2Page Ref: 10C-21-10C-2418) SQL statements can be run individually or as part of a related group of SQL statementsknown as a script.Answer:TRUEDiff: 1Page Ref: 10C-4119) To run an SQL script, click the Execute SQL Script in Connected Server button on theMySQL Workbench.Answer:TRUEDiff: 1Page Ref: 10C-21 Fig10C-2020) To run an SQL script, click the Execute button on the Query Toolbar.Answer:FALSEDiff: 2Page Ref: 10C-21 Fig10C-2021) In MySQL, surrogate keys

that use a sequence maintained by MySQL should be given a datatype of IDENTITY.Answer:FALSEDiff: 1Page Ref: 10C-20

22) MySQL indexes can be created by using the GUI MySQL Table Editor or an SQL statement.Answer:TRUEDiff: 1Page Ref: 10C-25-10C-2

Database Processing, 13e(Kroenke/Auer)Online Chapter 10B:Managing Databases with Oracle Database 11g Release 21) Oracle Database provides only one method for creating an Oracle Database database.Answer:FALSEDiff: 2Page Ref: 10B-82) The Oracle Database Configuration Assistant is the easiest Oracle Database tool to use whenyou are creating an Oracle database.Answer:TRUEDiff: 1Page Ref: 10B-83) When creating a new database in Oracle Database, two new account passwords areautomatically created in the database.Answer:FALSEDiff: 2Page Ref: 10B-64) The Oracle Enterprise Manager 11g Database Control is a Web-based Oracle Databaseadministration utility.Answer:TRUEDiff: 1Page Ref: 10B-85) Instead of creating a new database instance in Oracle Database, it is

preferable to create a newtablespace.Answer:TRUEDiff: 1Page Ref: 10B-86) Each Oracle Database tablespace has one or more datafiles associated with it.Answer:TRUEDiff: 1Page Ref: 10B-10 Fig10B-97) An Oracle Database tablespace is a logical subdivision of an Oracle database instance that isused to group related table, view, and other similar objects.Answer:TRUEDiff: 1Page Ref: 10B-108) In Oracle Database security, a System Privilege is the right to general tasks such asperforming some action on the database data.Answer:TRUEDiff: 3Page Ref: 10B-14-10B-159) In Oracle Database security, a Role can have many system privileges.Answer:TRUE

10) Among its many functions, SQL*Plus is also a text editor.Answer:TRUEDiff: 2Page Ref: 10B-19-10B-2111) The Oracle Database GUI tool used for application development is the Oracle SQLDeveloper.Answer:TRUEDiff: 1Page Ref: 10B-2112) The default file extension for text files created for SQL Developer is .sql.Answer:TRUEDiff: 2Page Ref: 10B-2313) Oracle Database supports the ON UPDATE CASCADE constraint.Answer:FALSEDiff: 3Page Ref: 10B-2514) When using Oracle Database, asequenceis an

object that generates a sequential series ofunique numbers.Answer:TRUEDiff: 1Page Ref: 10B-33-10B-3415) Sequences are most often used to provide values for surrogate keys.Answer:TRUEDiff: 2Page Ref: 10B-3316) The NextVal method provides the next value in a sequence.Answer:TRUEDiff: 1Page Ref: 10B-3317) The CurrVal method provides the current value in a sequence.Answer:TRUEDiff: 1Page Ref: 10B-3318) With Oracle Database, using asequencedoes not guarantee valid surrogate key values.Answer:TRUEDiff: 3Page Ref: 10B-33-10B-3419) With Oracle Database, data must be entered at the Oracle Database command prompt.Answer:FALSEDiff: 2Page Ref: 10B-33-10B-3820) The Oracle Database FORMAT_DATE function can be used to make sure that OracleDatabase interprets date data correctly.Answer:FALSE

Chapter 12

Database Processing, 13e(Kroenke/Auer)Chapter 12:Big Data, DataWarehouse, and Business Intelligence Systems1) Big Data is the name given to the enormous

datasets generated by Web 2.0 applications.Answer:TRUEDiff: 1Page Ref: 534-5352) Business Intelligence (BI) systems are information systems that help users analyze and usedata.Answer:TRUEDiff: 1Page Ref: 5363) Business Intelligence (BI) systems support operational activities.Answer:FALSEDiff: 1Page Ref: 5364) Business Intelligence (BI) systems obtain data in three different ways.Answer:TRUEDiff: 1Page Ref: 536-537 Fig 12-25) Business Intelligence (BI) reporting systems are used to filter data, sort data, group data andmake simple calculations based on the data.Answer:TRUEDiff: 2Page Ref: 537 Fig 12-36) Business Intelligence (BI) reporting systems can analyze data using standard SQL.Answer:TRUEDiff: 1Page Ref: 5377) Business Intelligence (BI) reporting systems summarize the current status of businessactivities and compare that status with past events, but not with predicted future activities.Answer:FALSEDiff: 3Page Ref: 5378) Data mining uses sophisticated statistical and mathematical techniques to perform what-ifanalyses, to make predictions, and to facilitate decision making.Answer:TRUEDiff: 1Page Ref: 5379) Report delivery is more important for data mining than it is for reporting systems.Answer:FALSEDiff: 2Page Ref:

53710) A data warehouse is a database system that has data, programs and personnel specialized inBusiness Intelligence (BI) processing.Answer:TRUE

11) Data warehouses are populated with data prepared by data extraction, transformation andload (ETL) programs.Answer:TRUEDiff: 2Page Ref: 538 Fig 12-412) Data warehouse data are frequently denormalized.Answer:TRUEDiff: 1Page Ref: 53813) Data warehouses also store the data warehouse metadata.Answer:TRUEDiff: 3Page Ref: 538-53914) Data warehouses often include data purchased from outside vendors.Answer:TRUEDiff: 1Page Ref: 540 Fig 12-415) Metadata about the data's source, format, assumptions and constraints are kept in a datawarehouse metadata database.Answer:TRUEDiff: 2Page Ref: 538-53916) Problematic data are called "dirty data."Answer:TRUEDiff: 1Page Ref: 539-54017) A data mart is a collection of data that addresses a particular component of a functional areaof a business.Answer:TRUEDiff: 1Page Ref: 540-541 Fig 12-718) Operational databases store historical data.Answer:FALSEDiff: 1Page Ref: 542 Fig 12-819) Dimensional databases are used for analytical data

processing.Answer:TRUEDiff: 1Page Ref: 541 Fig 12-820) Dimensional databases use the star schema.Answer:TRUEDiff: 1Page Ref: 54221) Operational databases contain a fact table.Answer:FALSE

SQL is a(n) .

a.) data sublanguage
b.) product of IBM research
c.) national standard
d.) combination of a data definition language and a data manipulation language e.) All of the above.

e.) All of the above.
???

Which SQL keyword is used to name a new table and describe the table's columns? a.)
SET
b.) CREATE
c.) SELECT
d.) ALTER
e.) MODIFY

b.) CREATE
???

Which SQL keyword is used to change the structure, properties or constraints of a table? a.)
SET
b.) CREATE
c.) SELECT
d.) ALTER
e.) MODIFY

d.) ALTER
???

One advantage of using the ALTER command to define a primary is that the database designer controls the .
a.) name of the primary key
b.) name of the foreign key
c.) name of the constraint
d.) a and b
e.) a, b and c

c.) name of the constraint
???
If the table PRODUCT has a column PRICE that has the data type Numeric (8,2), the values stored for this attribute will be stored .
a.) as eight digits, a decimal point, and two more digits
b.) as six digits, a decimal point, and two more digits
c.) as ten digits with no stored decimal point
d.) as eight digits with no stored decimal point

e.) as six digits with no stored decimal point

d.) as eight digits with no stored decimal point
???

If the table PRODUCT has a column PRICE that has the data type Numeric (8,2), the value 12345 will be displayed by the DBMS as . a.) 123.45
b.) 12345
c.) 12345.00
d. 123450.00
e.) 00012345

a.) 123.45
???

All constraints are stored in . a.)
the parent table
b.) the child table
c.) an intersection table
d.) the database metadata
e.) a and b

d.) the database metadata
???

Which of the following illustrates the author's preferred style of defining a primary key? a.) CREATE TABLE CUSTOMER (CustomerID Integer Primary Key LastName Char(35) Not Null First Name Char(25) Null):
b.) CREATE TABLE CUSTOMER (CustomerID Integer Not Null
LastName Char(35) Not Null
First Name Char(25) Null
CONSTRAINT CustomerPK PRIMARY KEY (CustomerID);

c.) CREATE TABLE CUSTOMER (CustomerID Integer Not Null LastName Char(35) Not Null First Name Char(25) Null):

ALTER TABLE CUSTOMER
ADD CONSTRAINT CustomerPK PRIMARY KEY (CustomerID);
d.) either b or c
e.) Te author does not state a preference for how to define a primary key.

b.) CREATE TABLE CUSTOMER (CustomerID
Integer Not Null
LastName Char(35) Not Null

First Name Char(25) Null
CONSTRAINT CustomerPK PRIMARY KEY (CustomerID);
???

Which SQL keyword is used to delete a table's structure? a.)
DELETE
b.) DROP
c.) DISPOSE
d.) ALTER
e.) MODIFY

b.) DROP
???

When the correct SQL command is used to delete a table's structure, what happens to the data in the table?
a.) If the deleted table was a parent table, the data is added to the appropriate rows of the child table.
b.) If the deleted table was a child table, the data is added to the appropriate rows of the parent table.
c.) The data in the table is also deleted.
d.) Nothing because there was no data in the table - only an empty table can be deleted. e.) a and b

c.) The data in the table is also deleted.
???

In an SQL query, which SQL keyword is used to specify the table(s) to be used? a.)
EXISTS

b.) FROM
c.) SELECT
d.) SET
e.) WHERE

b.) FROM
???

In an SQL query, which SQL keyword must be used to remove duplicate rows from the result table? a.) DELETE
b.) DISTINCT
c.) NOT EXISTS
d.) UNIQUE
e.) KEY

b.) DISTINCT
□□□

In an SQL query, which SQL keyword is used to state the condition that specifies which rows are to be selected? a.) EXISTS
b.) FROM
c.) SELECT
d.) SET
e.) WHERE

e.) WHERE
□□□

In an SQL query, which SQL keyword is used to join two conditions that both must be true for the rows to be selected? a.) AND
b.) EXISTS
c.) HAVING
d.) IN
e.) OR

a.) AND
□□□
In an SQL query, which SQL keyword is used to determine if a column value is equal to any one of a set of values? a.) AND
b.) EXISTS
c.) HAVING
d.) IN
e.) OR

d.) IN
□□□

In an SQL query, which of the following symbols is used by ANSI SQL to represent a single unspecified character? a.) _ (underscore)
b.) ? (question mark)
c.) * (asterisk)
d.) % (percent)
e.) # (pound)

A) _ (underscore)
□□□

Given a table with the structure: EMPLOYEE (EmpNo, Name, Salary, HireDate), which of the following would find all employees whose name begins with the letter "S?" a.)
SELECT *
FROM EMPLOYEE
WHERE Name IN ['S']; b.)
SELECT EmpNo FROM
EMPLOYEE
WHERE Name LIKE 'S'; c.)
SELECT *
FROM Name
WHERE EMPLOYEE LIKE 'S*'; d.)
SELECT *
FROM EMPLOYEE
WHERE Name LIKE 'S%'; e.)
None of the above.

d.) SELECT *
FROM EMPLOYEE
WHERE Name LIKE 'S%';

□□□

In an SQL query, which SQL keyword is used to sort the result table by the values in one or more columns? a.) GROUP BY
b.) ORDER BY
c.) SELECT
d.) SORT BY
e.) WHERE

b.) ORDER BY
□□□

In an SQL query, which built-in function is used to compute the number of rows in a table? a.) AVG
b.) COUNT
c.) MAX
d.) MIN
e.) MEAN

b.) COUNT
□□□

In an SQL query, the built-in function COUNT works with columns containing data of which of the following data types? a.) Integer

b.) Numeric
c.) Char
d.) a and b
e.) a, b and c

e.) a, b and c
□□□

In an SQL query, which built-in function is used to total numeric columns? a.) AVG
b.) COUNT
c.) MAX
d.) MEAN
e.) SUM

e.) SUM

□□□

In an SQL query, which built-in function is used to compute the average value of numeric columns? a.) AVG
b.) MEAN
c.) MAX
d.) MIN
e.) SUM

a.) AVG
□□□

In an SQL query, which built-in function is used to obtain the largest value of numeric columns? a.) AVG
b.) COUNT
c.) MAX
d.) MIN
e.) SUM

c.) MAX
□□□

In an SQL query, which built-in function is used to obtain the smallest value of numeric columns? a.) AVG
b.) COUNT
c.) MAX
d.) MIN

e.) SUM

d.) MIN
□□□

In an SQL query, which SQL keyword is used with built-in functions to group together rows that have the same value in a specified column? a.) GROUP BY
b.) ORDER BY
c.) SELECT
d.) SORT BY
e.) DISTINCT SET

a.) GROUP BY
□□□

In an SQL query, which SQL keyword is used with GROUP BY to select groups meeting specified criteria? a.) AND
b.) EXISTS
c.) HAVING
d.) IN
e.) WHERE

c.) HAVING
□□□

Given a table with the structure: EMPLOYEE (EmpNo, Name, Salary, HireDate), which of the following is not a valid ANSI SQL command? a.) SELECT *
FROM EMPLOYEE
WHERE Name LIKE 'Ja%'; b.)
SELECT COUNT(*)
FROM EMPLOYEE
WHERE Salary < 30000;
c.) SELECT COUNT(EmpNo)
FROM EMPLOYEE;
d.) SELECT HireDate, COUNT(*)
FROM EMPLOYEE
WHERE Salary < 30000;
e.) SELECT HireDate, COUNT(*)
FROM EMPLOYEE GROUP
BY HireDate;

d.) SELECT HireDate, COUNT(*)
FROM EMPLOYEE
WHERE Salary < 30000;

□□□

Based on the tables above, which of the following ANSI SQL commands would return the average customer balance grouped by SalesRepNo? a.) SELECT AVG (Balance)
FROM CUSTOMER
WHERE SalesRepNo;

b.) SELECT AVG (Balance)
FROM CUSTOMER
GROUP BY SalesRepNo;
c.) SELECT AVG (Balance)
FROM CUSTOMER, SALESREP
WHERE SALESREP.SalesRepNo = CUSTOMER.SalesRepNo; d.)
SELECT AVG (Balance)
FROM CUSTOMER
ORDER BY SalesRepNo;
e.) SELECT AVG (Balance)
FROM CUSTOMER, SALESREP
WHERE CUSTOMER.SalesRepNo = CUSTOMER.SalesRepNo
HAVING SalesRepNo;

b.) SELECT AVG (Balance)
FROM CUSTOMER
GROUP BY SalesRepNo;
□□□

In an SQL query, which SQL keyword is used to implement a subquery? a.)
GROUP BY
b.) HAVING
c.) ORDER BY
d.) SELECT
e.) SORT BY

d.) SELECT
□□□

When one SQL query is embedded in the WHERE clause of another SQL query, this is referred to as a _____. a.) subset
b.) join
c.) WHERE Query
d.) subquery
e.) set query

d.) subquery
□□□

In an SQL query, which SQL keyword is used to specify the names of tables to be joined? a.) FROM

b.) HAVING
c.) JOIN
d.) SELECT
e.) WHERE

a.) FROM
□□□

In an SQL query, which SQL keyword is used to specify the condition(s) for a join operation? a.) FROM
b.) HAVING
c.) JOIN
d.) SELECT
e.) WHERE

e.) WHERE
□□□

Based on the tables above, which of the following commands in ANSI SQL would return only the name of the sales representative and name of the customer for each customer that has a balance greater than 400?
a.) SELECT *
FROM SALESREP, CUSTOMER
WHERE Balance > 400;
b.) SELECT DISTINCT RepName, CustName
FROM SALESREP, CUSTOMER
WHERE Balance > 400; c.)
SELECT *
FROM SALESREP, CUSTOMER
WHERE SALESREP.SalesRepNo = CUSTOMER.SalesRepNo
AND Balance > 400;
d.) SELECT RepName, CustName
FROM SALESREP, CUSTOMER
WHERE SALESREP.SalesRepNo = CUSTOMER.SalesRepNo
AND Balance > 400;
e.) SELECT RepName, CustName
FROM SALESREP, CUSTOMER

WHERE Balance > 400 GROUP
BY SalesRepNo;

d.) SELECT RepName, CustName FROM
SALESREP, CUSTOMER WHERE
SALESREP.SalesRepNo =
CUSTOMER.SalesRepNo
AND Balance > 400;
???

Regarding the interchangeability of subqueries and joins .
a.) A join can always be used as an alternative to a subquery, and a subquery can always be used as an alternative to a join.
b.) A join can sometimes be used as an alternative to a subquery, and a subquery can sometimes be used as an alternative to a join.
c.) A join can always be used as an alternative to a subquery, and a subquery can sometimes be used as an alternative to a join.
d.) A join can sometimes be used as an alternative to a subquery, and a subquery can always be used as an alternative to a join.
e.) A join can never be used as an alternative to a subquery, and a subquery can never be used as an alternative to a join.

b.) A join can sometimes be used as an alternative to a subquery, and a subquery can sometimes be used as an alternative to a join.
???

In an SQL query of two tables, which SQL keyword indicates that we want data from all the rows of one table to be included in the result, even if the row does not correspond to any data in the other table?
a.) LEFT JOIN
b.) RIGHT JOIN
c.) INCLUDE
d.) a and b
e.) a, b and c

d.) a and b
???

Which SQL keyword is used to add one or more rows to a table? a.)
DELETE
b.) INSERT
c.) SELECT
d.) SET

e.) UPDATE

b.) INSERT
□□□
Which SQL keyword is used to change one or more rows in a table? a.)
MODIFY
b.) INSERT
c.) SELECT
d.) CHANGE
e.) UPDATE

e.) UPDATE
□□□

Which SQL keyword is used to change a column value? a.)
CHANGE
b.) INSERT
c.) SELECT
d.) SET
e.) MODIFY

d.) SET
□□□

Which keyword is used to remove one or more rows from a table? a.)
DELETE
b.) INSERT
c.) ERASE
d.) SET
e.) UPDATE

a.) DELETE
□□□

Based on the tables above, which of the following commands in ANSI SQL would increase the balance of the Gonzales by $100 to a total of $450? a.) SELECT Gonzales
FROM CUSTOMER
INSERT VALUES PLUS (100) INTO Balance; b.)
SELECT Gonzales
FROM CUSTOMER
INSERT VALUES (450) INTO Balance;
c.) INSERT INTO CUSTOMER VALUES PLUS (100)
SELECT Balance
WHERE CustName = 'Gonzales';

d.) INSERT INTO CUSTOMER VALUES (450)

SELECT Balance
WHERE CustName = 'Gonzales'; e.)
UPDATE CUSTOMER
SET Balance = 450 WHERE
CustName = ''Gonzales';

e.) UPDATE CUSTOMER
SET Balance = 450
WHERE CustName = ''Gonzales';

The SQL CREATE TABLE statement is used to name a new table and describe the table's columns. T

The SQL keyword CONSTRAINT is used to define one of five types of constraints. T

The SQL keyword PRIMARY KEY is used to designate the column(s) that are the primary key for the table. T

The SQL keyword CONSTRAINT is used to limit column values to specific values. T

The SQL keyword CONSTRAINT is used in conjunction with the SQL keywords PRIMARY KEY and FOREIGN KEY. T

One advantage of using the CONSTRAINT command to define a primary key is that the database designer controls the name of the constraint.
T

The SQL keyword UNIQUE is used to define alternative keys.
T

If the table PRODUCT has a column PRICE, and PRICE has the data type Numeric (8,2), the value 98765 stored in that field will be displayed by the DBMS as 98765.00. F

If the table ITEM has a column WEIGHT, and WEIGHT has the data type Numeric (7,2), the value 4321 with be displayed by the DBMS as 43.21. T

The SQL keyword CHECK is used to limit column values to specific values. T

The SQL keyword MODIFY is used to change the structure, properties or constraints of a table. F

Data values to be added to a table are specified by using the SQL VALUES clause. T

The SQL keyword DELETE is used to delete a table's structure. F

When the correct SQL command is used to delete a table's structure, the command can only be used with a table that has already had its data removed. F

One or more rows can be added to a table by using the SQL INSERT statement. T
Unless it is being used to copy data from one table to another, the SQL INSERT statement can be used to insert only a single row into a table. T

Rows in a table can be changed by using the SQL UPDATE statement.
T

The SQL SET keyword is used to specify a new value when changing a column value. T

The SQL keyword MODIFY is used to change a column value.
F

Rows can be removed from a table by using the SQL DELETE statement.
T

An SQL virtual table is called a view.
T

The SQL command CREATE USER VIEW is used to create a virtual table. F

SQL views are constructed from SELECT statements. T

According to the SQL-92, statements used to construct views cannot contain the WHERE clause.
F

The SQL command SELECT is used to retrieve view instances. T

The values in an SQL view are not always changeable through the view itself.

SQL views can be used to hide columns.
T T

SQL views can be used to provide a level of insulation between data processed by applications and the data actually stored in the database tables. T

If the values in an SQL view are changeable through the view itself, the SQL command UPDATE is used to change the values. T

The values in an SQL view are always changeable through the view itself. F

SQL views are updatable when the view is based on a single table with no computed columns, and all non-null columns are present in the view. T

Because SQL statements are table-oriented, whereas programs are variable-oriented, the results of SQL statements used in programs are treated as pseudofiles.
T

A set of SQL statements stored in an application written in a standard programming language is called embedded SQL. T

Because SQL statements are table-oriented, whereas programs are variable-oriented, the results of SQL statements used in programs are accessed using an SQL cursor. T

A stored program that is attached to a table or view is called a stored procedure. F

SQL triggers use the ANSI SQL keywords BEFORE, INSTEAD OF, and AFTER. T

SQL triggers can be used with SQL operations INSERT, UPDATE, and DELETE. T

SQL triggers can be used when the DBMS receives an INSERT request.
T

SQL triggers are used for providing default values, validity checking, updating views, and performing referential integrity actions.
T

The Oracle DBMS supports the SQL BEFORE trigger.
T

The SQL Server DBMS supports the SQL BEFORE trigger.
F

SQL triggers can be used when the DBMS receives an insert request. T

To set a column value to an initial value that is selected according to some business logic, you would use the SQL DEFAULT constraint with the CREATE TABLE command. F

SQL triggers are created using the SQL ADD TRIGGER statement. F

If the values in an SQL view are not changeable through the view itself, you may still be able to update the view by using unique application logic. In this case, the specific logic is placed in an INSTEAD OF trigger. T

If a trigger is being written to enforce referential integrity actions, you cannot use an INSTEAD OF trigger. F

When a trigger is fired, the DBMS makes the appropriate data available to the trigger code. T

A stored program that is stored within the database and compiled when used is called a trigger. F

Stored procedures have the advantage of greater security, decreased network traffic, SQL optimized by the DBMS compiler, and code sharing. T

Unlike application code, stored procedures are never

distributed to the client computers. T

Because SQL stored procedures allow and encourage code sharing among developers, stored procedures give database application developers the advantages of less work, standardized processing, and specialization among developers. T

Which SQL keyword is used to name a new table and describe the table's columns?

A) SET
B) CREATE
C) SELECT
D) ALTER
E) CONSTRAINT

B

If the table PRODUCT has a column PRICE that has the data type Numeric (8,2), the value 12345 will be displayed by the DBMS as _____.

A) 123.45
B) 12345
C) 12345.00
D) 123450.00
E) 00012345 A

Which SQL keyword is used to impose restrictions on a table, data or relationship?

A) SET
B) CREATE
C) SELECT
D) ALTER

E) CONSTRAINT

E

One advantage of using the CONSTRAINT phrase to define a primary key is that the database designer controls the _____.

A) name of the primary key
B) name of the foreign key
C) name of the constraint
D) A and B
E) A, B, and C

C

Which of the following illustrates the authors' preferred style of defining a primary key? A)

```
CREATE TABLE CUSTOMER (
CustomerID Integer Primary Key
LastName Char(35) Not Null
First Name Char(25) Null
);
```

B) CREATE TABLE CUSTOMER (CustomerID Integer Not Null

```
LastName Char(35) Not Null
First Name Char(25) Null
CONSTRAINT CustomerPK PRIMARY KEY (CustomerID)
);
```

C) CREATE TABLE CUSTOMER (CustomerID Integer Not Null

```
LastName Char(35) Not Null
First Name Char(25) Null
);

ALTER TABLE CUSTOMER
ADD CONSTRAINT CustomerPK PRIMARY KEY (CustomerID);
```

D) either B or C
E) The authors do not demonstrate a preference for how to define a primary key. B

Given the SQL statement

```
CREATE TABLE SALESREP (
SalesRepNo int NOT NULL,
RepName char(35) NOT NULL,
HireDate date NOT NULL,

CONSTRAINT SalesRepPK PRIMARY KEY (SalesRepNo),
CONSTRAINT SalesRepAK1 UNIQUE (RepName)
);
```

we know that _____. A) RepName is the primary key B) RepName is a foreign key
C) RepName is a candidate key
D) RepName is a surrogate key
E) None of the above is true C

The SQL keyword used to limit column values to specific values is _____.
A) CONSTRAINT
B) CHECK
C) NOT NULL
D) UNIQUE
E) UPDATE
B

Which SQL keyword is used to change the structure, properties or constraints of a table?
A) SET

B) CREATE
C) SELECT
D) ALTER
E) CONSTRAINT D

Which SQL keyword is used to delete a table's structure?
A) DELETE
B) DROP
C) DISPOSE
D) ALTER
E) MODIFY B

When the correct SQL command is used to delete a table's structure, what happens to the data in the table?
A) If the deleted table was a parent table, the data is added to the appropriate rows of the child table.
B) If the deleted table was a child table, the data is added to the appropriate rows of the parent table. C) The data in the table is also deleted.
D) Nothing because there was no data in the table since only an empty table can be deleted.
E) A and B C

Which SQL keyword is used to add one or more rows of data to a table? A) DELETE
B) INSERT
C) SELECT
D) SET
E) UPDATE
B

Which SQL keyword is used to change one or more rows in a table? A) MODIFY
B) INSERT
C) SELECT
D) CHANGE
E) UPDATE
E

Which SQL keyword is used to change the values of an entire column? A) CHANGE
B) INSERT
C) SELECT
D) SET
E) MODIFY D

Which keyword is used to remove one or more rows from a table? A) DELETE
B) INSERT
C) ERASE
D) SET
E) UPDATE A

Based on the tables below, which of the following SQL statements would increase the balance of the Gonzales account by $100 to a total of $450?

GENERAL SALES DATABASE:

SALESREP
SalesRepNo RepName HireDate
654 Jones 01/02/2005
734 Smith 02/03/2007

345 Chen 01/25/2004
434 Johnson 11/23/2004

CUSTOMER
CustNo CustName Balance SalesRepNo
9870 Winston 500 345
8590 Gonzales 350 434
7840 Harris 800 654
4870 Miles 100 345

A) SELECT Gonzales
FROM CUSTOMER
INSERT VALUES PLUS (100) INTO Balance;
B) SELECT Gonzales
FROM CUSTOMER
INSERT VALUES (450) INTO Balance;
C) INSERT INTO CUSTOMER VALUES PLUS (100)
SELECT Balance
WHERE CustName = 'Gonzales';
D) INSERT INTO CUSTOMER VALUES (450)
SELECT Balance
WHERE CustName = 'Gonzales';
E) UPDATE CUSTOMER
SET Balance = 450
WHERE CustName = 'Gonzales';
E

An SQL virtual table is called _____.
A) a CHECK constraint
B) a view
C) embedded SQL
D) a trigger
E) a stored procedure
B

The SQL command used to create a virtual table is _____.

A) CREATE VTABLE

B) CREATE VIEW

C) VTABLE

D) VIEW

E) NEWLOOK

B

SQL views are constructed from _____.

A) CREATE statements

B) INSERT statements

C) UPDATE statements

D) SELECT statements

E) VIEW statements D

According to the SQL-92, statements used to construct views cannot contain _____.

A) the SELECT clause

B) the FROM clause

C) the WHERE clause

D) the ORDER BY clause

E) SQL view statements can use all of the listed clauses. D

Which SQL statement is used to retrieve view instances?

A) CREATE

B) DELETE

C) INSERT

D) SELECT

E) UPDATE D

SQL views are used _____.

A) to hide columns

B) to show results of computed columns
C) to hide complicated SQL statements
D) to provide a level of indirection between data processed by applications and the data actually stored in the database tables
E) SQL views are used for all of the above. E

If the values in an SQL view are changeable through the view itself, which SQL statement is used to change the values? A) CREATE
B) DELETE
C) INSERT
D) SELECT
E) UPDATE
E

SQL views are always updatable when _____.
A) the view is based on a single table with no computed columns, and all non-null columns are present in the view
B) the view is based on any number of tables, with or without computed columns, and the INSTEAD OF trigger is defined for the view
C) the view is based on multiple tables, the update is being done on the most subordinate table, and the rows of that table can be uniquely identified
D) A and B
E) A, B, and C
D

A set of SQL statements stored in an application written in a standard programming language is called _____.
A) a CHECK constraint

B) a view
C) embedded SQL
D) a trigger
E) a stored procedure C

Because SQL statements are set-oriented, whereas programs are element-oriented, the results of SQL statements used in programs are treated as _____.
A) tables
B) columns
C) rows
D) files
E) pseudofiles E

Because SQL statements are table-oriented, whereas programs are element-oriented, the results of SQL statements used in programs are accessed using _____.
A) standard programming tools
B) custom written programming tools
C) an SQL cursor
D) an SQL trigger
E) an SQL stored procedure C

A stored program that is attached to a table or view is called _____.
A) a CHECK constraint
B) a view
C) embedded SQL
D) a trigger
E) a stored procedure D

Which of the following is not an ANSI SQL trigger?
A) BEFORE UPDATE

B) INSTEAD OF UPDATE

C) BEFORE INSERT

D) INSTEAD OF CONSTRAINT

E) AFTER DELETE D

Which of the following is an SQL trigger Oracle supports?

A) BEFORE

B) INSTEAD OF

C) AFTER

D) B and C only

E) A, B, and C E

Which of the following is an SQL trigger Microsoft SQL Server supports?

A) BEFORE

B) INSTEAD OF

C) AFTER

D) B and C only

E) A, B, and C D

SQL triggers can be used when the DBMS receives a(n) _____ request.

A) INSERT

B) UPDATE

C) DELETE

D) A and B

E) A, B, and C

E

SQL triggers are used for _____.

A) validity checking

B) providing default values

C) updating views

D) A and B
E) A, B, and C
E

When a trigger is fired, the DBMS makes the appropriate data available to _____. A) the SQL interpreter
B) the application code
C) the embedded SQL code
D) the trigger code
E) the stored procedure code D

SQL triggers are created using _____.
A) the SQL CREATE TRIGGER statement
B) the SQL ADD TRIGGER statement
C) the SQL TRIGGER statement
D) the SQL ADD CONSTRAINT TRIGGER statement
E) the SQL CONSTRAINT TRIGGER statement A

To set a column value to an initial value that is selected according to some business logic, you would use:
A) the SQL DEFAULT constraint with the CREATE TABLE command.
B) an SQL view.
C) embedded SQL.
D) an SQL trigger.
E) an SQL stored procedure. D

If the values in an SQL view are not changeable through the view itself, you may still be able to update the view by using unique application logic. In this case, the specific logic is placed in _____. A) a BEFORE trigger
B) an INSTEAD OF trigger

C) an AFTER trigger
D) Depending on the specific logic, either A or B can be used.
E) Depending on the specific logic, any of A, B, or C can be used. B

A stored program that is attached to the database is called _____.
A) a CHECK constraint
B) a view
C) embedded SQL
D) a trigger
E) a stored procedure E

Stored procedures have the advantage of _____.
A) greater security
B) decreased network traffic
C) SQL optimized by the DBMS compiler
D) code sharing
E) All of the above E

Because SQL stored procedures allow and encourage code sharing among developers, stored procedures give database application developers the advantages of _____.
A) less work
B) standardized processing
C) specialization among developers
D) A and B
E) A, B, and C
E

Explain the essential format of the CREATE TABLE statement.

The essential format for the CREATE TABLE statement is:
CREATE TABLE tablename (column-description,
column-description, column-description,
. . .
optional table constraints
);
"Tablename" is the name that will be given to the newly created table. "Column-description" is a three-part description of each column to appear in the table. This description includes the name of the column, the column's data type, and an optional column constraint (either Primary Key, Null, or Not Null), in that order. The CONSTRAINT phrase can be used to set optional primary key, foreign key and referential integrity constraints for the table. All SQL statements must end with a semi-colon (;).

Explain how relationships are created using SQL.

In SQL, relationships are created using a FOREIGN KEY constraint. This has the format:

CONSTRAINT ConstraintNameFK FOREIGN KEY({ForeignKeyColumnInCurrentTable} REFERENCES {ReferencedTableName}(PrimaryKeyColumnInReferencedTable })

The constraint thus names the foreign key column in the current table and its corresponding primary key in a referenced table.

Discuss SQL data types.

Common examples of standard SQL data types are Char,

VarChar, Integer, and Numeric. The Char data type is for fixed-length character data. VarChar is for variable-length character data. Integer is for numeric data that are whole numbers only. Numeric is for numeric data that may include decimals. Char, VarChar, and Numeric must be qualified by a length specification to indicate the amount of storage space to be allocated for each data item. For example, Char(10) indicates fixed-length character data that is always stored as 10 characters.

Discuss what is meant by a data type of "Numeric (10,3)."

Numeric indicates a non-integer, decimal number in SQL Server. Oracle Database uses Number, and MySQL uses Decimal or Fixed. The (10,3) is in (n,d) format, where n is the total number of digits allowed, and d is the number of digits to the right of the decimal place. Thus, "10, 3" allows a maximum of ten digits, and the last three are considered to be to the right of the decimal place. For example, "1234567" would be read as "1234.567."

Distinguish between Char and VarChar data types.

Char data type is fixed-length, so that no matter the actual length of the data entered it will always take exactly the same storage space. For example, Char(10) indicates that 10 characters will always be stored for each value of that column. If the actual data entered is less than the specified fixed- length, the data will be padded with blanks. VarChar data type is variable length so that only the amount of space actually needed to store the data is used. Although VarChar may be more efficient in its use of space, it is not always preferred. VarChar requires the storage of some extra data to indicate the length of the data values, plus it requires some extra processing by the DBMS to arrange the variable length data.

The following database will be used in this question:

GENERAL SALES DATABASE:

SALESREP
SalesRepNo RepName HireDate
654 Jones 01/02/2005
734 Smith 02/03/2007
345 Chen 01/25/2004
434 Johnson 11/23/2004

CUSTOMER
CustNo CustName Balance SalesRepNo
9870 Winston 500 345
8590 Gonzales 350 434
7840 Harris 800 654
4870 Miles 100 345

Explain the use of the SQL statement CREATE TABLE. Do NOT discuss the ALTER statement in your answer, but DO include an example based on the SALESREP table in the General Sales database.

The SQL statement CREATE TABLE forms that basis for all SQL table construction. CREATE TABLE is used to name tables and specify their structure, including column names and their associated data types. In addition, CREATE TABLE can be used to define primary keys, foreign keys, and to specify constraints on tables, columns and column values. There are five SQL keywords that specify these constraints: PRIMARY KEY, FOREIGN KEY, NULL/NOT NULL, UNIQUE and CHECK. PRIMARY KEY and FOREIGN KEY are used to specify keys. NULL and NOT NULL specify whether nulls are allowed as values in a column. UNIQUE specifies whether the column values have to be unique. CHECK is not discussed in this

chapter. The SQL keyword CONSTRAINT may be used to specify some of the table constraints. For example, to create the SALESREP table for the General Sales database, the following SQL statement may be use:

```
CREATE TABLE SALESREP (
SalesRepNo Integer Not Null,
RepName Char(25) Not Null,
HireDate Date Not Null,

CONSTRAINT SalesRepPK PRIMARY KEY (SalesRepNo)
);
```

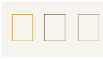

The following database will be used in this question:

GENERAL SALES DATABASE:

SALESREP
SalesRepNo RepName HireDate
654 Jones 01/02/2005
734 Smith 02/03/2007
345 Chen 01/25/2004
434 Johnson 11/23/2004

CUSTOMER
CustNo CustName Balance SalesRepNo
9870 Winston 500 345
8590 Gonzales 350 434
7840 Harris 800 654
4870 Miles 100 345

What is an SQL view, and what is it used for? Include an example based on the CUSTOMER table of the General Sales Database.

An SQL view is a virtual table constructed from database tables or other views. It is based on the SQL CREATE VIEW command and uses the SQL SELECT statement to construct the view. However, the ORDER BY clause cannot be used when creating a view. For example:

```
CREATE VIEW CustomerNameOnly AS
SELECT CustName
FROM CUSTOMER;
```

A view may be used to (1) hide columns or rows, (2) show the results of computed columns, (3) hide complicated SQL statements such as joins, (4) layer built-in functions, (5) provide a level of indirection between the data processed by applications and the actual table data,(6) assign different processing permissions to different views of the same table, and (7) assign different triggers to different views of the same table.

What is embedded SQL, and what considerations are necessary when using it in an application?

Embedded SQL are SQL statements used, or embedded, in program code, triggers or stored procedures. Applications are typically written in program code, using a programming language. There are two problems that arise. First, the results of SQL statements must be assigned to programming language variables. DBMS products typically provide the means of doing this. Second, SQL is table or set-oriented and SQL results use tables or sets of rows, whereas application programming languages are variable or row-oriented. This is resolved by treating SQL results as pseudofiles. A cursor is then used to move through the pseudofile one row at a time.

What are SQL triggers and how are they used?

An SQL trigger is a stored program that is attached to a table or view. The trigger is invoked by the DBMS whenever an insert, update or delete request is made on the table or view with the trigger. There are three commonly used triggers: BEFORE, INSTEAD OF, and AFTER (MS SQL server does not support BEFORE). This creates a set of nine possible trigger types: BEFORE + [INSERT or UPDATE or DELETE], INSTEAD OF + [INSERT or UPDATE or DELETE], and AFTER + [INSERT or UPDATE or DELETE]. Triggers are used (among other things) for (1) providing default values, (2) validity checking, (3) updating views, and (4) enforcing referential integrity actions.

What are SQL stored procedures and how are they used?

An SQL stored procedure is a stored program that is attached to a database instead of just a table or view. Stored procedures can receive input parameters and return results. They can be executed by any process that has permission with the database to use stored procedures. They can issue INSERT, UPDATE and DELETE commands. They are typically used by (1) database administrators to do common administrative tasks, and (2) database applications.

Chapter 7

When a surrogate key is used in the child table in an IDdependent relationship .

a.) the parent table becomes the holder of the foreign key for the relationship
b.) the surrogate key is allowed to contain null values
c.) the surrogate key is allowed to contain non-unique values
d.) the relationship changes to a non-identifying relationship
e.)All of the above.

d.) the relationship changes to a non-identifying relationship

Given the SQL statement

CREATE TABLE SALESREP(
SalesRepNo int NOT NULL,
RepName char(35) NOT NULL,
HireDate date NOT NULL,

CONSTRAINT SalesRepPK PRIMARY KEY (SalesRepNo),
CONSTRAINT SalesRepAK1 UNIQUE (RepName)):

we know that .
a.) RepName is the primary key
b.) RepName is a foreign key
c.) RepName is a candidate key
d.) RepName is a surrogate key
e.) none of the above is true
c.) RepName is a candidate key

The SQL keyword used to limit column values to specific values is .
a.) CONSTRAINT
b.) CHECK
c.) NOT NULL
d.) UNIQUE
e.) UPDATE
b.) CHECK

An SQL virtual table is called _____ .
a.) a CHECK constraint
b.) a view
c.) embedded SQL
d.) a trigger
e.) a stored procedure

b.) a view

The SQL command used to create a virtual table is . a.)
CREATE VTABLE
b.) CREATE VIEW
c.) VTABLE
d.) VIEW
e.) NEWLOOK
b.) CREATE VIEW

SQL views are constructed from .
a.) CREATE statements
b.) INSERT statements
c.) UPDATE statements
d.) SELECT statements
e.) VIEW statements
d.) SELECT statements

According to the SQL-92, statements used to construct
views cannot contain . a.)
the SELECT clause
b.) the FROM clause
c.) the WHERE clause
d.) the ORDER BY clause
e.) SQL view statements can use all of the listed clauses
d.) the ORDER BY clause

Which SQL statement is used to retrieve view instances?
a.) CREATE
b.) DELETE
c.) INSERT
d.) SELECT
e.) UPDATE

d.) SELECT

SQL views are used .
a.) to hide columns
b.) to show results of computed columns
c.) to hide complicated DQL statements
d.) to provide a level of indirection between data processed by applications
and the data actually stored in the database tables e.) SQL views are used for all of the above
e.) SQL views are used for all of the above

If the values in an SQL view are changeable through the view itself, which SQL statement is used to change the values?
a.) CREATE
b.) DELETE
c.) INSERT
d.) SELECT
e.) UPDATE
e.) UPDATE

When are the values in an SQL view changeable through the view itself? a) always
b.) not always, but the rules to determine modification status are simple
c.) not always, and the rules to determine modification status are complicated
d.) not always, and the rules to determine modification status are DBMS-dependent e.) c and d
e.) c and d

SQL views are always updateable when .
a.) the view is based on a single table with no computed columns,and all non-null columns are present in the view
b.) the view is based on any number of tables, with or without computed columns,and the INSTEAD OF trigger is defined for the view
c.) the view is based on multiple tables, the update is being done on the most subordinate table, and the rows of that table can be uniquely identified d.) a and b
e.) a, b and c
d.) a and b

How do SQL views compare to ANSI/SPARC external schemas /user views? a.)
they are identical
b.) SQL views are a subset of ANSI/SPARC user views
c.) ANSI/SPARC user views are a subset of SQL views
d.) they are completely different, with an unfortunate similarity in their names
e.) either b or c is correct depending on the particular database
b.) SQL views are a subset of ANSI/SPARC user views

A set of SQL statements stored in an application written in a standard programming language is called _____ . a.)
a CHECK constraint
b.) a view
c.) embedded SQL
d.) a trigger
e.) a stored procedure
c.) embedded SQL

Because SQL statements are set-oriented, whereas

programs are element-oriented, the results of SQL statements used in programs are treated as _____ . a.) tables
b.) columns
c.) rows
d.) files
e.) pseudofiles
e.) pseudofiles

Because SQL statements are set-oriented, whereas programs are element-oriented, the results of SQL statements used in programs are accessed using . a.) standard programming tools
b.) custom written programming tools
c.) an SQL cursor
d.) an SQL trigger
e.) an SQL stored procedure
c.) an SQL cursor

A stored program that is attached to a table or view is called _____ .
a.) a CHECK constraint
b.) a view
c.) embedded SQL
d.) a trigger
e.) a stored procedure
d.) a trigger

SQL triggers are used for .
a.) validity checking
b.) providing default values
c.) updating views
d.) a and b

e.) a, b and c

e.) a, b and c

Which of the following is not an ANSI SQL trigger?

a.) BEFORE UPDATE

b.) INSTEAD OF UPDATE

c.) BEFORE INSERT

d.) INSTEAD OF CONSTRAINT

e.) AFTER DELETE

d.) INSTEAD OF CONSTRAINT

Which of the following is an SQL trigger Oracle supports?

a.) BEFORE

b.) INSTEAD OF

c.) AFTER

d.) b and c only

e.) a, b and c

e.) a, b and c

Which of the following is an SQL trigger Microsoft SQL Server supports? a.) BEFORE

b.) INSTEAD OF

c.) AFTER

d.) b and c only

e.) a, b and c

d.) b and c only

SQL triggers can be used when the DBMS receives a _____request. a.) insert

b.) update

c.) delete

d.) a and 75

e.) a, b and 75

e.) a, b and 76

When a trigger is fired, the DBMS makes the appropriate data available to .

a.) the SQL interpreter

b.) the application code

c.) the embedded SQL code

d.) the trigger code

e.) the stored procedure code

d.) the trigger code

SQL triggers are created using .

a.) the SQL CREATE TRIGGER statement

b.) the SQL ADD TRIGGER statement

c.) the SQL TRIGGER statement

d.) the SQL ADD CONSTRAINT TRIGGER statement

e.) the SQL CONSTRAINT TRIGGER statement

a.) the SQL CREATE TRIGGER statement

To add a single initial value to a column, you would use:

a.) the SQL DEFAULT constraint with the CREATE TABLE command.

b.) an SQL view.

c.) embedded SQL.

d.) an SQL trigger.

e.) an SQL stored procedure.

a.) the SQL DEFAULT constraint with the CREATE TABLE command.

To set a column value to an initial value that is selected according to some business logic, you would use:

a.) the SQL DEFAULT constraint with the CREATE TABLE command.
b.) an SQL view.
c.) embedded SQL.
d.) an SQL trigger.
e.) an SQL stored procedure.

d.) an SQL trigger.

If the values in an SQL view are not changeable through the view itself, you may still be able to update the view by using unique application logic. In this case, the specific logic is placed in . a.) a BEFORE trigger
b.) an INSTEAD OF trigger
c.) an AFTER trigger
d.) depending on the specific logic, either a or b can be used
e.) depending on the specific logic, any of a, b or c can be used

b.) an INSTEAD OF trigger

If a trigger is being written to enforce referential integrity actions, you cannot use . a.)
a BEFORE trigger
b.) an INSTEAD OF trigger
c.) an AFTER trigger
d.) depending on the specific referential integrity action, either b or c may be disallowed
e.) depending on the specific referential integrity action, any of a, b or c may be disallowed c.) an AFTER trigger

A stored program that is attached to the database is called _____ .

a.) a CHECK constraint

b.) a view

c.) embedded SQL

d.) a trigger

e.) a stored procedure

e.) a stored procedure

Stored procedures have the advantage of .

a.) greater security

b.) decreased network traffic

c.) SQL optimized by the DBMS compiler

d.) code sharing

e.) All of the above.

e.) All of the above.

Unlike application code, stored procedures are never distributed to _____ . a.) the DBMS

b.) the client computers

c.) the network servers

d.) the database servers

e.) All of the above.

b.) the client computers

Because SQL stored procedures allow and encourage code sharing among developers, storedprocedures give database application developers the advantages of . a.) less work

b.) standardized processing

c.) specialization among developers

d.) a and b

e.) a, b and c

e.) a, b and c

Chapter 8

Database redesign is fairly easy when .

a.) information systems and organizations influence each other

b.) the design was done correctly the first time

c.) there is no data in the database

d.) good backups of the database are available

e.) All of the above.

c.) there is no data in the database

Which of the following is not a possible step in the database redesign process?

a.) checking whether certain conditions or assumptions about the data are valid

b.) reverse engineering the data model

c.) testing proposed changes

d.) maintaining backups of the existing database

e.) All of the above are possible steps in the database redesign process.

e.) All of the above are possible steps in the database redesign process.

In the database redesign process, before proceeding with the redesign it is often useful to .

a.) check whether certain conditions or assumptions about the data are valid

b.) find out why the design was not done properly the first time

c.) stop information systems and user from influencing each other

d.) set standards for user behavior

e.) All of the above.

a.) check whether certain conditions or assumptions about the data are valid

In the database redesign process, SQL tools that are useful for testing whether or not certain conditions or assumptions are valid are . a.) correlated subqueries
b.) EXISTS
c.) NOT EXISTS
d.) b and c
e.) a, b and c
e.) a, b and c

In a correlated subquery of a database that has tables T1 and T2, and if table T1 is used in the upper SELECT statements, then which table is used in the lower SELECT statement? a.) T1
b.) T2
c.) both T1 and T2
d.) either T1 or T2
e.) neither T1 nor T2
c.) both T1 and T2

In the SQL statements

```
SELECT C1.CustName, C1.SalesRepNo
FROM CUSTOMER C1;
```

the "C1" is called a(n) _____ . a.) term
b.) alias
c.) convention
d.) phrase
e.) label

b.) alias

When running a correlated subquery, the DBMS .
a.) runs the lower SELECT statement by itself and then
sends the results to the upper SELECT statement.
b.) runs the upper SELECT statement by itself and then
sends the results to the lower SELECT statement.
c.) alternates running the lower SELECT statement
with running the upper SELECT statement based on
each result of the lower SELECT statement
d.) either a or b may be used depending on the query
e.) None of the above describes how a correlated
subquery is run by the DBMS.
c.) alternates running the lower SELECT statement with running the upper SELECT statement based on each result of the lower SELECT statement

When running a correlated subquery, the DBMS always uses _____. a.) regular processing
b.) nested processing
c.) "quick and dirty" processing
d.) SQL-92 processing
e.) a form of processing the is specific to the DBMS product
b.) nested processing

Which of the following SQL statements is a correctly stated correlated subquery?

a.) SELECT C1.CustName, C1.SalesRepNo
FROM CUSTOMER C1
WHERE C1.SalesRepNo IN
(SELECT S1.SaleRepNo

```
FROM SALESREP S1 WHERE
S1.RepName = 'Smith');
b.) SELECT C1.CustName, C1.SalesRepNo
FROM CUSTOMER C1
WHERE C1.SalesRepNo IN
(SELECT S1.SaleRepNo
FROM SALESREP S1
WHERE S1.RepName = 'Smith') AND
C1.SalesRepNo=S1.SalesRepNo);
c.) SELECT C1.CustName, C1.SalesRepNo
FROM CUSTOMER C1
WHERE C1.SalesRepNo IN
(SELECT S1.SaleRepNo
FROM SALESREP S1
WHERE S1.RepName = 'Smith') AND
C1.SalesRepNo<>S1.SalesRepNo);
d.) SELECT C1.CustName, C1.SalesRepNo
FROM CUSTOMER C1
WHERE C1.SalesRepNo IN
(SELECT C2.SaleRepNo
FROM CUSTOMER C2
WHERE C1.SalesRepNo=C2.SalesRepNo);
AND C1.OrderNo<>C2.OrderNo);
e.) None of the above is a correctly stated correlated
    subquery.
d.) SELECT C1.CustName, C1.SalesRepNo
FROM CUSTOMER C1
WHERE C1.SalesRepNo IN
(SELECT C2.SaleRepNo
FROM CUSTOMER C2
WHERE C1.SalesRepNo=C2.SalesRepNo);
AND C1.OrderNo<>C2.OrderNo);
```

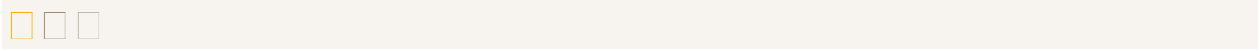

SQL queries that use EXISTS and NOT EXISTS are _____ . a.)
normal subqueries
b.) correlated subqueries
c.) uncorrelated subqueries
d.) constraint dependent subqueries
e.) constraint independent subqueries
b.) correlated subqueries

When running an SQL query that uses EXISTS or NOT
EXISTS, the DBMS always
uses _____.
a.) regular processing
b.) nested processing
c.) "quick and dirty" processing
d.) SQL-92 processing
e.) a form of processing the is specific to the DBMS product
b.) nested processing

When running an SQL query that uses EXISTS, the EXISTS
keyword will be true if .
a.) any row in the subquery meets the condition
b.) all rows in the subquery meet the condition
c.) no row in the subquery meets the condition
d.) any row in the subquery fails to meet the condition
e.) all rows in the subquery fail to meet the condition
a.) any row in the subquery meets the condition

When running an SQL query that uses NOT EXISTS, the NOT
EXISTS keyword will be true if .
a.) any row in the subquery meets the condition
b.) all rows in the subquery meet the condition
c.) no row in the subquery meets the condition
d.) any row in the subquery fails to meet the condition

e.) all rows in the subquery fail to meet the condition

e.) all rows in the subquery fail to meet the condition

The use of a double nested set of NOT EXISTS SELECT statements can be used to find rows that _____ some specified condition for _____ row(s) in a table.

a.) meet; any

b.) meet; every

c.) meet; no

d.) fails to meet; any

e.) fails to meet; every

b.) meet; every

A double nested set of NOT EXISTS SELECT statements is _____ .

a.) a famous pattern in SQL

b.) regularly used

c.) rarely used to nonexistent in the real world

d.) a and b

e.) a, b, and c

d.) a and b

The process of reading an actual database schema and producing a data model from that schema is called _____ . a.) data modeling

b.) data engineering

c.) reverse engineering

d.) schema modeling

e.) schema engineering

c.) reverse engineering

The data model produced by reverse engineering is a(n)

_____ .

a.) conceptual schema

b.) internal schema

c.) dependency graph

d.) table-relationship diagram

e.) entity-relationship diagram

d.) table-relationship diagram

Because of the need to know the functional dependencies in a database, it is a good idea to create a(n) _____ .

a.) conceptual schema

b.) internal schema

c.) dependency graph

d.) a table-relationship diagram

e.) an entity-relationship diagram

c.) dependency graph

Which of the following different copies of the database schema is/are typically used in the database redesign process? a.) small test database

b.) large test database

c.) operational database

d.) a and b

e.) a, b and c

e.) a, b and c

In order make sure the database redesign is working properly during the redesign testing process, a means must be created to _____ .

a.) reverse engineer all test databases

b.) graph dependencies in all test databases

c.) recover all test databases to their original state

d.) a and b

e.) a, b and c

c.) recover all test databases to their original state

To change a table name, we .

a.) use the SQL-92 RENAME TABLE command

b.) use the SQL-92 ALTER TABLENAME command

c.) use the SQL-92 MODIFY TABLENAME command

d.) create a new table, move the data and drop the old table

e.) None of the above is the correct way to change a table name.

d.) create a new table, move the data and drop the old table

When making any change to the database structure, we may need to check for effects of the change on . a.) data

b.) foreign keys

c.) constraints

d.) triggers

e.) All of the above may need to be checked.

e.) All of the above may need to be checked.

If it is necessary to copy a date when changing a table name, we can use the _____ to handle the task. a.) INSERT command by itself

b.) VALUES expression with the INSERT command

c.) SET expression with the INSERT command

d.) SET command by itself

e.) VALUES expression with the SET command

a.) INSERT command by itself

To add a NULL column to a table, we . a.)
use the REVISE TABLE command
b.) use the ALTER TABLE command
c.) use the MODIFY TABLE command
d.) create a new table with the NULL column, move the other data, and drop the old table
e.) None of the above is the correct way to add a NULL column.
b.) use the ALTER TABLE command

If a DEFAULT constraint is included when a new column is added to a table, the default value is
applied to .
a.) all existing rows at the time the column is added
b.) all new rows
c.) all new rows but only after the UPDATE command is issued
d.) a and b
e.) a and c
b.) all new rows

To add a NOT NULL column to a table, we . a.)
use the REVISE TABLE command
b.) use the ALTER TABLE command
c.) use the MODIFY TABLE command
d.) create a new NULL column, insert data values into every row,
and change the NULL constraint to NOT NULL
e.) None of the above is the correct way to add a NOT NULL column to a table.
d.) create a new NULL column, insert data values into every row, and change the NULL constraint to NOT NULL

When dropping a primary key column from a table, which of the following steps are included in the process [the order of the steps listed below is not relevant, only the steps themselves]?

a.) drop the primary key constraint from the table

b.) drop the primary key column from the table

c.) drop the foreign keys in other tables based on the primary key column d.) a and b

e.) a, b and c

e.) a, b and c

When changing column data types, which of the following data conversions will either usually or always succeed? a.) numeric to char or varchar

b.) date or money to char or varchar

c.) char or varchar numeric, date or money

d.) a and b

e.) a, b and c

d.) a and b

When changing the minimum cardinality on the parent side of the relationship from zero to one, the foreign key _____ .

a.) must be changed from NULL to NOT NULL

b.) must be changed from NOT NULL to NULL

c.) must be changed to a composite key

d.) must be changed to a surrogate key

e.) does not change

a.) must be changed from NULL to NOT NULL

Which of the following are difficulties when changing the maximum cardinality from 1:1 to 1:N? a.) preserving the existing tables

b.) preserving the existing relationships
c.) preserving the existing data
d.) a and b
e.) a, b and c
b.) preserving the existing relationships

When increasing cardinalities from 1:N to N:M, which of the following steps are included in the process [the order of the steps listed below is not relevant, only the steps themselves]? a.) create an intersection table
b.) populate the intersection table
c.) drop the old foreign key
d.) a and b
e.) a, b and c
e.) a, b and c

When decreasing cardinalities, there will always be _____ .
a.) relationship loss
b.) foreign key loss
c.) data loss
d.) a and b
e.) a, b and c
c.) data loss

When dropping tables and relationships, which of the following steps are included in the process [the order of the steps listed below is not relevant, only the steps themselves]?
a.) drop the foreign key constraints from the tables
b.) drop the tables
c.) drop the primary key constraints from the tables
d.) a and b

e.) a, b and c

d.) a and b

In order to have a normalized database, we may need to add _____ . a.) tables

b.) relationships

c.) triggers

d.) a and b

e.) a, b and c

d.) a and b

In a Cartesian join of tables T1 and T2 .

a.) every row of T1 is joined to matching rows in T2, and only these rows are in the resulting table

b.) every row of T1 is joined to matching rows in T2, and these rows together with any unmatched rows in T1 are in the resulting table

c.) every row of T1 is joined to matching rows in T2, and these rows together with any unmatched rows in T2 are in the resulting table

d.) every row of T1 is joined to every row in T2, and all these rows are in the resulting table

e.) None of the above correctly describes a Cartesian join.

d.) every row of T1 is joined to every row in T2, and all these rows are in the resulting table

Chapter 9

Which of the following is not a database administration responsibility of a DBA?

a.) managing the database structure

b.) managing data activity
c.) managing the DBMS
d.) maintaining the data repository
e.) All of the above are database administration responsibilities of a DBA.

e.) All of the above are database administration responsibilities of a DBA

Which of the following is true about making changes to the database structure?
a.) The DBA need not get input from users on the issue because it is a technical decision.
b.) Formal policies and procedures for requesting a change are not used because they are too limiting.
c.) Documentation of when the change was made, how it was made, and why it was made must be created.
d.) Changes do not produce unexpected results because the DBA will have investigated the change thoroughly before implementing it.
e.) If the database is properly designed, changes should not be necessary throughout the system's lifetime.

c.) Documentation of when the change was made, how it was made, and why it was made must be created.

The task of diagnosing errors due to changes in the database structure is eased by:
a.) formal policies for requesting changes.
b.) database structure change documentation.
c.) rollback analysis.
d.) configuration control.
e.) None of the above.

b.) database structure change documentation.

Measures that are taken to prevent one user's work from inappropriately influencing another user's work are called:

a.) concurrency control.

b.) checkpoint.

c.) database recovery.

d.) database logging.

e.) interleaving.

a.) concurrency control.

A series of actions to be taken on the database such that either all actions are completed successfully, or none of them can be completed, is known as a(n): a.) checkpoint.

b.) log.

c.) lock.

d.) transaction.

e.) concurrent.

d.) transaction.

When two transactions are being processed against the database at the same time,

a.) they are called concurrent transactions.

b.) they are usually interleaved.

c.) they always result in a lost update problem.

d.) one must be rolled back.

e.) both a and b

e.) both a and b

The situation that occurs when one user's changes to the database are lost by a second user's changes to the database is known as the:

a.) concurrent update problem.

b.) deadly embrace problem.

c.) inconsistent read problem.
d.) inconsistent write problem.
e.) deadlock problem.
a.) concurrent update problem.

One remedy for the inconsistencies caused by concurrent processing is _____. a.) lost updates
b.) checkpointing
c.) rollback
d.) resource locking
e.) concurrency
d.) resource locking

A lock placed automatically by the DBMS is called a(n) _____ lock. a.) exclusive
b.) explicit
c.) granular
d.) implicit
e.) shared
d.) implicit

Which of the following is not true about locks?
a.) Locks with large granularity are easier for the DBMS to administer.
b.) Locks with small granularity cause more conflicts.
c.) Locks with large granularity produce fewer details for the DBMS to track.
d.) Locks may have a table-level granularity.
e.) Locks may have a database-level granularity.
b.) Locks with small granularity cause more conflicts.

Which type of lock prevents all types of access to the locked resource?

a.) exclusive lock

b.) shared lock

c.) two-phased lock

d.) explicit lock

e.) implicit lock

a.) exclusive lock

Which type of lock still allows other transactions to have read-only access to the locked resource? a.) exclusive lock

b.) shared lock

c.) two-phased lock

d.) explicit lock

e.) implicit lock

b.) shared lock

Which of the following is not true about two-phased locking?

a.) can make transactions serializable

b.) uses only shared locks

c.) has a growing phase

d.) has a shrinking phase

e.) cannot obtain a new lock once a lock has been released

b.) uses only shared locks

The situation that occurs when two users are each waiting for a resource that the other person has locked is known as a(n):

a.) lost update problem.

b.) deadlock.

c.) inconsistent read problem.

d.) inconsistent write problem.

e.) checkpoint.

b.) deadlock.

Requiring all application programs to lock resources in the same order is a technique for preventing what problem? a.) concurrent update

b.) lost update

c.) deadlock

d.) exclusive locks

e.) growing phase locking

c.) deadlock

Locks that are placed assuming that a conflict will occur are called:

a.) dynamic locks.

b.) explicit locks.

c.) implicit locks.

d.) optimistic locks.

e.) pessimistic locks.

e.) pessimistic locks.

Locks that are placed assuming that a conflict will not occur are called:

a.) dynamic.

b.) explicit.

c.) implicit.

d.) optimistic.

e.) pessimistic

d.) optimistic.

Ensuring that all rows impacted by the actions of a

transaction are protected from changes until the entire transaction is completed is called: a.) statement level consistency.
b.) optimistic locking.
c.) transaction level consistency.
d.) durable transactions.
e.) ARID transactions.
c.) transaction level consistency.

Which of the following is allowed by "Repeatable Read Isolation?"
a.) nonrepeatable reads
b.) dirty reads
c.) phantom reads
d.) a and b
e.) a, b and c
c.) phantom reads

Which of the following is true of forward only cursors?
a.) Current values for each row are retrieved when the application accesses a row.
b.) All changes of any type from any source are visible.
c.) Changes made by the transaction are visible only if they occur on rows ahead of the cursor.
d.) Applications may scroll backward in the record set.
e.) It requires the greatest overhead of any cursor type.
c.) Changes made by the transaction are visible only if they occur on rows ahead of the cursor.

Which of the following cannot be enforced in the DBMS or application programs? a.) processing rights
b.) security
c.) processing responsibilities

d.) cursors

e.) transaction isolation

c.) processing responsibilities

Once processing rights have been defined, they may be implemented at any of these levels except: a.) network.

b.) operating system.

c.) data.

d.) DBMS.

e.) application.

c.) data.

Which of the following is not true of DBMS security features?

a.) Users may be assigned to one or more roles.

b.) A role may be assigned to only one user.

c.) Both users and roles can have many permissions.

d.) Objects have many permissions.

e.) Each permission pertains to one user or role and one object.

b.) A role may be assigned to only one user.

Recovering a database via reprocessing involves:

a.) restoring the database from the save and reprocessing all the transactions since the save.

b.) restoring the database from the save and reapplying all the changes made by transactions since the save.

c.)undoing the changes made by erroneous or partially processed transactions, and restarting the valid transactions that were in process at the time of the failure.

d.) recreating the database by reentering all of the data from the beginning, and then reprocessing all of the transactions.

e.) synchronizing the database and the transaction log by checkpointing.

a.) restoring the database from the save and reprocessing all the transactions since the save.

Recovering a database via rollforward involves:

a.) restoring the database from the save and reprocessing all the transactions since the save.

b.) restoring the database from the save and reapplying all the changes made by transactions since the save.

c.)undoing the changes made by erroneous or partially processed transactions, and restarting the valid transactions that were in process at the time of the failure.

d.) re-creating the database by re-entering all of the data from the beginning, and then reprocessing all of the transactions.

e.) synchronizing the database and the transaction log by checkpointing.

b.) restoring the database from the save and reapplying all the changes made by transactions since the save.

Recovering a database via rollback involves:

a.) restoring the database from the save and reprocessing all the transactions since the save.

b.) restoring the database from the save and reapplying all the changes made by transactions since the save.

c.)undoing the changes made by erroneous or partially processed transactions, and restarting the valid transactions that were in process at the time of the failure.

d.) re-creating the database by re-entering all of the data from the beginning and, then reprocessing all of the transactions.
e.) synchronizing the database and the transaction log by checkpointing.
c.) undoing the changes made by erroneous or partially processed transactions, and restarting the valid transactions that were in process at the time of the failure.

Which of the following would not be contained in a transaction log?
a.) before-images
b.) type of operation
c.) pointers
d.) time of the action
e.) permissions
e.) permissions

Which of the following would a DBA do in managing the DBMS?
a.) analyze system performance statistics
b.) investigate user complaints
c.) evaluate new DBMS product features
d.) tune DBMS product options to accommodate other software in use
e.) All of the above.
e.) All of the above.

Which of the following is not true of data repositories?
a.) They are usually created after the database has been implemented and optimized for performance. b.) They may be virtual.

c.) They may contain metadata about database applications.

d.) They may contain metadata about users.

e.) They may contain metadata about web pages.

a.) They are usually created after the database has been implemented and optimized for performance.

Which type of data repository is composed of metadata that is created automatically as the system components are created?

a.) passive

b.) dynamic

c.) active

d.) automatic

e.) summary

c.) active

true

The SQL keyword CREATE is used to name a new table and describe the table's columns.

true

The SQL keyword CONSTRAINT is used to define one of five types of constraints.

true

The SQL keyword PRIMARY KEY is used to designate the column(s) that are the primary key for the table.

false

The SQL keyword CONSTRAINT is used to limit column values to specific values.

true

The SQL keyword UNIQUE is used to define alternative keys.

false

If the table PRODUCT has a column PRICE that has the data type Numeric (8,2), the value 98765 stored in that field will be displayed by the DBMS as 98765.00

false

The SQL keyword MODIFY is used to change the structure, properties or constraints of a table.

false

The SQL keyword DELETE is used to delete a table's structure.

false

When the correct SQL command is used to delete a table's structure, the command can only be used with a table that has already had its data removed.

true

Unless it is being used to copy data from one table to another, the SQL INSERT command can be used to insert only a single row into a table.

false

The SQL keyword MODIFY is used to change a column value.

true

The SQL syntax JOIN . . . ON can be used as alternate way of writing an SQL join statement.

true

Joins that show only matching rows from the joined tables in their results are called inner joins.

true

Joins that show the matching rows from the joined tables plus unmatched rows from one other table in their results are called outer joins.

false

Outer joins can be either up joins or down joins.

true

An SQL virtual table is called a view.

false

The SQL command CREATE USER VIEW is used to create a virtual table.

true

SQL views are constructed from SELECT statements.

false

According to the SQL-92, statements used to construct views cannot contain the WHERE clause.

true

The SQL command SELECT is used to retrieve view instances.

true

SQL views can be used to hide columns.

true

SQL views can be use used to provide a level of indirection between data processed by applications and the data actually stored in the database tables.

true

If the values in an SQL view are changeable through the view itself, the SQL command UPDATE is used to change the values.

false

The values in an SQL view are always changeable through the view itself.

true

SQL views are updateable when the view is based on a single table with no computed columns, and all non-null columns are present in the view.

true

A set of SQL statements stored in an application written in a standard programming language is called embedded SQL.

true

Because SQL statements are table-oriented, whereas programs are variable-oriented, the results of SQL statements used in programs are accessed using an SQL cursor.

false

A stored program that is attached to a table or view is called a stored procedure.

true

SQL triggers are used for providing default values, validity checking, updating views, and performing referential integrity actions.

true

The Oracle DBMS supports the SQL trigger BEFORE.

false

The SQL Server DBMS supports the SQL trigger BEFORE.

true

SQL triggers can be used when the DBMS receives an insert request.

false

To set a column value to an initial value that is selected according to some business logic, you would use the SQL DEFAULT constraint with the CREATE TABLE command.

false

SQL triggers are created using the SQL ADD TRIGGER statement.

true

If the values in an SQL view are not changeable through the view itself, you may still be able to update the view by using

unique application logic. In this case, the specific logic is placed in an INSTEAD OF trigger.

false

If a trigger is being written to enforce referential integrity actions, you cannot use an INSTEAD OF trigger.

false

A stored program that is stored within the database and complied when used is called a trigger.

true

Stored procedures have the advantage of greater security, decreased network traffic, SQL optimized by the DBMS compiler, and code sharing.

true

Unlike application code, stored procedures are never distributed to the client computers.

true

Because SQL stored procedures allow and encourage code sharing among developers, stored procedures give database application developers the advantages of less work, standardized processing, and specialization among developers.

CREATE

The SQL keyword _____ is used to name a new table and describe the table's columns.

43.21

If the table PRODUCT has a column PRICE, and PRICE has the data type Numeric (7,2), the value 4321 with be displayed by the DBMS as ____ .

CONSTRAINT

The SQL keyword _____ is used in conjunction with the SQL keywords PRMARY KEY and FOREIGN KEY.

name of the constraint

One advantage of using the CONSTRAINT command to define a primary key is that the database designer controls the _____.

UNIQUE

The SQL keyword _____ is used to define an alternate key.

CHECK

The SQL keyword _____ is used to limit column values to specific values.

ALTER

The SQL keyword _____ is used to change the structure, properties or constraints of a table.

DROP

The SQL keyword _____is used to delete a table's structure.

INSERT

One or more rows can be added to a table by using the _____ statement.

VALUES

Data values to be added to a table are specified by using the _____statement.

UPDATE

Rows in a table can be changed by using the_____statement.

SET

The SQL keyword_____is used to specify a new value when changing a column value.

DELETE

Rows can be removed from a table by using the SQL keyword _____.

JOIN...ON

In addition to the standard SQL join command structure, joins can be created using the SQL syntax _____.

view

An SQL virtual table is called a _____.

ORDER BY

According to SQL-92, statements used to construct views cannot contain the _____ clause.

CREATE VIEW

The SQL command _____ is used to create a virtual table.

SELECT

SQL views are constructed from _____ statements.

SELECT

columns

The SQL command _____ is used to retrieve view instances.

SQL views can be used to hide table _____ .

the view itself

The values in an SQL view are not always changeable through _____.

UPDATE

If the values in an SQL view are changeable through the view itself, the SQL command _____ is used to change the values.

updateable

When an SQL view is based on a single table with no computed columns, and all non-null columns are present in the view, then the view is _____.

embedded SQL

A set of SQL statements stored in an application written in a standard programming language is called _____.

pseduofiles

Because SQL statements are table-oriented, whereas programs are variable-oriented, the results of SQL statements used in programs are treated as _____.

SQL cursor

Because SQL statements are table-oriented, whereas programs are variable-oriented, the results of SQL statements used in programs are accessed using a(n) _____.

SQL trigger

A stored program that is attached to a table or view is called a(n) _____.

BEFORE, INSTEAD OF, AFTER

SQL triggers use the ANSI SQL keywords _____, _____, and _____.

INSERT, UPDATE, DELETE

SQL triggers can be used with SQL operations _____, _____, and _____.

BEFORE, INSTEAD OF, AFTER

The Oracle DBMS supports the SQL triggers _____, _____ , and _____.

BEFORE

The SQL Server DBMS does not support the SQL trigger _____.

triggers

SQL _____ can be used when the DBMS receives an INSERT request.

the trigger code

When a trigger is fired, the DBMS makes the appropriate data available to _____ .

an SQL trigger

To set a column value to an initial value that is selected according to some business logic, you would use _____ .

CREATE TRIGGER

SQL triggers are created using the SQL _____ statement.

INSTEAD OF

If the values in an SQL view are not changeable through the view itself, you may still be able to update the view by using unique application logic. In this case, the specific logic is placed in an _____ trigger.

stored procedure

A stored program that is attached to the database is called a _____ .

more

In terms of application security, stored procedures are _____ secure than application code.

never distributed

Unlike application code, stored procedures are _____to the client computers.

code sharing

SQL stored procedures give database application developers the advantages of less work, standardized processing, and specialization among developers because stored procedures allow and encourage _____ among developers.

38) SQL Server user built-in system functions start with the @@ symbol. Answer: TRUE Diff: 2 Page Ref: 10A-51

39) When using SQL Server, developers place explicit locks. Answer: FALSE Diff: 1 Page Ref: 10A-64-10A65 4 ScholarStock

40) SQL Server automatically makes determinations of whether or not to promote or demote a lock. Answer: TRUE Diff: 3 Page Ref: 10A-64-10A-65

41) In SQL Server, it is not possible to make dirty reads . Answer: FALSE Diff: 2 Page Ref: 10A-65-10A-67

42) The default cursor concurrency for a dynamic cursor in SQL Server is optimistic. Answer: TRUE Diff: 3 Page Ref: 10A-65-10A-67

43) The default transaction isolation level for SQL Server is Read Committed. Answer: TRUE Diff: 2 Page Ref: 10A-65-10A-67

44) In SQL Server, SCROLL_LOCK is a version of pessimistic locking. Answer: TRUE Diff: 1 Page Ref: 10A65-10A-67

45) Locking hints included in an SQL statement will override locking behavior based on explicit transaction isolation level statements. Answer: TRUE Diff: 2 Page Ref: 10A-65-10A-67

46) In SQL Server, locking behavior can be modified by providing locking hints in the WITH parameter of the FROM clause in SELECT statements. Answer: TRUE Diff: 2 Page Ref: 10A-65-10A-67

47) SQL Server security provides three modes of authentication. Answer: FALSE Diff: 1 Page Ref: 10A-67

48) The recommended SQL Server security is Windows-only security. Answer: TRUE Diff: 2 Page Ref: 10A-67

49) SQL Server security roles provide a simple way to control user privileges in a database. Answer: TRUE
Diff: 1 Page Ref: 10A-69-10A-70 5 ScholarStock

Background image of page 5

50) When recovering a database in SQL Server, it is possible to process the log to a particular point in time or to a transaction mark. Answer: TRUE Diff: 3 Page Ref: 10A-71-10A-73

51) In SQL Server, a differential backup makes a copy of the entire database. Answer: FALSE Diff: 2 Page Ref: 10A-72

52) In SQL Server, an interactive backup makes a copy of the changes that have been made since the last time a backup of the entire database was made. Answer: FALSE Diff: 1 Page Ref: 10A-72

53) SQL Server supports three recovery models: simple, complex and bulk logged. Answer: FALSE Diff: 3 Page Ref: 10A-73

54) SQL Server bulk-logged recovery differs from other log-based recovery methods because it does not include log entries for changes that cause large log entries. Answer: TRUE Diff: 2 Page Ref: 10A-73

55) In SQL Server, the full recovery model creates a log for every change to the database. Answer: FALSE Diff: 1 Page Ref: 10A-73

56) The SQL Server GUI interface database management program is the _____. A) Enterprise Manager B) Query Analyzer C) Microsoft SQL Server Management Studio D) A and B E) A, B, and C Answer: C Diff: 1 Page Ref: 10A-8

57) Which of the following is generally used to create a new database? A) Enterprise Manager B) Query Analyzer C) Microsoft SQL Server Management Studio D) A and B E) A, B, and C Answer: C Diff: 2 Page Ref: 10A-4-10A-7 6 ScholarStock

82) Which of the following is not a transaction isolation level supported by SQL Server? A) Pessimistic B) Read Committed C) Serializable D) Read Uncommitted E) Repeatable Read Answer: A Diff: 3 Page Ref:
10A-65 Fig10A-46

83) Which of the following is the default transaction isolation level for SQL Server? A) Pessimistic B) Read Committed C) Serializable D) Read Uncommitted E) Repeatable Read Answer: B Diff: 2 Page Ref:
10A-65

84) Which of the following is the most restrictive transaction isolation level supported by SQL Server? A) Pessimistic B) Read Committed C) Serializable D) Read Uncommitted E) Repeatable Read Answer: C Diff: 2 Page Ref: 10A-65 12 ScholarStock

85) Which of the following is a type of cursor concurrency that causes SQL Server to place an update lock on a row when the row is read? A) SCROLL_LOCK B) Optimistic C) Serializable D) Read only E) Forward-only Answer: A Diff: 3 Page Ref: 10A-65-10A-66

86) The default cursor concurrency setting for a static cursor in SQL Server is _____. A) SCROLL_LOCK B) Optimistic C) Serializable D) Read only E) Pessimistic Answer: D Diff: 2 Page Ref: 10A-65-10A-66

87) The default cursor concurrency setting for a keyset cursor in SQL Server is _____. A) SCROLL_LOCK B) Optimistic C) Serializable D) Read only E) Pessimistic Answer: B Diff: 2 Page Ref: 10A65-10A-66

88) Specifying locking hints in which parameter in an SQL clause can modify locking behavior in SQL Server? A) SELECT B) WITH C) LOCK D) TRANS E) HINT Answer: B Diff: 2 Page Ref: 10A-66-10A-67

89) Which type of backup in SQL Server will create a backup of the entire database? A) Full backup B) Transaction backup C) Database backup D) Complete backup E) DBMS backup Answer: D Diff: 1 Page Ref: 10A-71-10A-73 13 ScholarStock

90) Which of the following is a type of recovery model supported in SQL Server in which all database changes are logged? A) Logged B) Bulk logged C) Full D) Simple E) Complete Answer: C Diff: 2 Page Ref:
10A-73

91) Describe the database files that are automatically created when a new database is created in SQL Server. Answer: By default there are two database files created. The first is a data file, and the second is a log file for the database. However, other files can be created during the database creation process.
Diff: 1 Page Ref: 10A-4-10A-7

70) Which of the following is true of clustered indexes? A) They are normally slower than nonclustered indexes for retrieving data. B) Only one is allowed per table. C) They are normally slower than nonclustered indexes for updating data. D) They do not have data in the bottom level. E) A and C Answer: B Diff: 2 Page Ref: 10A-

---

71) Which type of index in SQL Server is faster for data retrieval? A) Clustered B) Nonclustered C) Unique D) Baseline E) There is no advantage to any particular index for data retrieval, only for data updating. Answer: A Diff: 2 Page Ref: 10A-20

72) Which of the following are ways of processing an SQL Server database? A) From an application written in a programming language such as C# or Visual Basic B) From text files processed by the SQL Server Query Analyzer C) From invoked SQL Server stored procedures D) From invoked SQL Server triggers E) All of the above Answer: E Diff: 1 Page Ref: 10A-33-10A-35

73) The language available in SQL Server that adds programming constructs to the SQL language is known as _____. A) PL/SQL B) SQLP C) C++ D) INTERACT-SQL E) TRANSACT-SQL Answer: E Diff: 1 Page Ref: 10A-35

74) Which of the following is known to be true if the command EXEC abcCust @Cost = 5 is used to invoke a stored procedure in the Microsoft SQL Server Management Studio? A) Cost = 5 is a CHECK constraint. B) The name of the stored procedure is abcCust. C) The name of the stored procedure is abc . D) Cust is given the value 5. E) abcCust is the name of a table or view in the database. Answer: B Diff: 3 Page Ref: 10A-40-10A-41 10 ScholarStock

75) A stored procedure can be stored in a text file and run with the Microsoft SQL Server Management Studio. If we are running the code that creates this procedure for the third time, it must start with

_____. A) STORED PROCEDURE B) CREATE PROCEDURE C) ALTER PROCEDURE D) MODIFY PROCEDURE E) RUN PROCEDURE Answer: C Diff: 2 Page Ref: 10A-40-10A-41

76) SQL Server user variables and parameters start with the symbol _____. A) /* B) # C) ## D) @ E) @@ Answer: D Diff: 1 Page Ref: 10A-35

77) Which of the following is not a type of trigger supported by SQL Server? A) BEFORE B) INSTEAD OF C) CONCURRENT D) AFTER E) A and C Answer: E Diff: 2 Page Ref: 10A-49

78) A view may not be assigned to a(n) _____. A) BEFORE trigger B) INSTEAD OF trigger C) CONCURRENT trigger D) AFTER trigger E) Any of the above Answer: D Diff: 1 Page Ref: 10A-49

79) SQL Server system functions start with the symbol _____. A) /* B) # C) ## D) @ E) @@ Answer: E Diff: 2 Page Ref: 10A-51 11 ScholarStock

80) For triggers on insert and update actions, the new column values are stored in _____. A) a table named new:TableName B) a table named old:TableName C) a pseudotable name inserted D) a pseudotable named deleted E) None of the above Answer: C Diff: 2 Page Ref: 10A-5

Which of the following is not a lock granularity supported by SQL Server?

Answer: All of these are supported.

SQL Server supports AFTER and INSTEAD OF triggers, but not BEFORE triggers. (T/F)	True	
	In SQL Server, insert and update triggers store old values in a pseudotable named deleted. (T/F)	~~False~~ TRUE!
	Each entity in the extended E-R model is represented as a table in the relational database design. (T/F)	True
	A surrogate key is appropriate when the primary key of a table contains a lengthy text field. (T/F)	True
	A null value is an attribute value that has been set to zero. (T/F)	False

When placing a foreign key for a 1:1 relationship, the key of either table can be used as the foreign key in the other table. (T/F)	True
Like all ID-dependent relationships, the parents of an association table are required. (T/F)	True
All identifying relationships are 1:N. (T/F)	True
When the parent entity has a surrogate key, the enforcement actions are the same for both parent and child. (T/F)	False
SQL Server AFTER triggers may be assigned to either tables or views. (T/F)	False
Inline user defined function can use input parameters. (T/F)	True
A view may not be assigned to a(n) ____. a) BEFORE trigger b) INSTEAD OF trigger c) CONCURRENT trigger d) AFTER trigger e) any of the above	d) AFTER trigger
Which of the following is not true about representing subtypes in a	b) all of the attributes of the supertype are added to the

relational database design? a) one table is created for the supertype and one for each subtype b) all of the attributes of the supertype are added to the subtype relations c) the key of the supertype is made the key of the subtypes d) a subtype and its supertype are representations of the same underlying table e) and instance of the supertype may be related to one instance each of several subtypes	subtype relations
A referential integrity constraint policy that insures that all rows containing a particular foreign key value in a table are eliminated from the table when the row containing the corresponding primary a) incremental updates b) incremental deletes c) controlled key adjustments d) cascading updates e) cascading deletes	e) cascading deletes
Which of the following is not a step in the database design process? a) create tables and columns from entities and attributes b) select primary keys c) represent relationships d) create constraints and triggers e) all of the above are steps in the	d) create constraints and triggers

database design process	
The identifier of the entity becomes the _____ of the corresponding table. a) primary key b) foreign key c) supertype d) subtype e) either a or b	a) primary key
One of the important properties of an attribute is whether or not it is _____. a) found in more than one entity b) required c) character or numeric d) subject to normalization e) subject to de-normalization	b) required
In many-to-many relationships in a relational database design _____. a) the key of the child is placed as a foreign key into the parent b) the key of the parent is placed as a foreign key into the child c) the keys of both tables are placed in a third table d) the keys of both tables are joined into a composite key e) c and d	e) c and d
When transforming an ID-dependent E-R data model relationship into a relational database design, and the child	c) the relationship changes to a non-ID-dependent relationship

entity is designed to use a surrogate key, then _____.

a) the parent entity must also use a surrogate key

b) the relationship remains an IDdependent relationship

c) the relationship changes to a non-ID-dependent relationship d) a and b

e) a and c

RAISEERROR sends a message, optionally logging it in the error log, but it does not affect batch execution unless ____, in which case the connection is closed. a) the state of the message is set to 19 b) the severity is 20 or more c) the disconnect option is turned on d) the error message has no log and a message_id of 5000 e) a and c	b) the severity is 20 or more
For triggers on insert and update actions, the new column values are stored in ____. a) a table named new:TableName b) a table named old: TableName c) a pseudotable name inserted d) a pseudotable name deleted e) none of the above	c) a pseudotable name inserted
You must ____ if you want a function to use an input parameter's default value. a) leave the parameter area blank	b) specify the DEFAULT keyword

b) specify the DEFAULT keyword
c) specify the NULL keyword
d) enter '%%%%'
e) you cannot have default values with inline user defined functions

SQL server does not support ____ triggers as Oracle does.	BEFORE
For update and delete triggers, old data values are stored in a pseudotable named _____.	deleted
A referential integrity constraint policy that insures that foreign key values in a table are correctly maintained when there is a change to the primary key value in the parent is called ____.	cascading updates
In 1:N relationships, the table on the "one" side is called the _____.	parent
In representing a M:N relationship in a relational database design, a(n) ___ table is created.	intersection
To represent a M:N relationship in a relational database design, in essence it is reduced to two ____ relationships.	1:N
The ____ system variable provides the error status of the most	@@ERROR
recently executed statement.	

For insert and update triggers, new data values are stored in a pseudotable named _____.	inserted
User-defined functions that return a table data type are referred to as _____.	table-valued functions

---

1) The SQL Server GUI program is the Enterprise Manager. Answer: FALSE Diff: 1 Page Ref: 10A-8

2) The Microsoft SQL Server Management Studio Manager can be used to create a new database. Answer: TRUE Diff: 1 Page Ref: 10A-4-10A-7

3) The SQL Server default when creating a new database is to create two data files and one log file for each database. Answer: FALSE Diff: 2 Page Ref: 10A-4-10A-7

4) In addition to the standard SQL Server 2012 tools, SQL statements such as CREATE TABLE can now also be executed by using Microsoft's PowerShell. Answer: TRUE Diff: 3 Page Ref: 10A-7-10A-8

5) When viewing the SQL Server list of tables in a database, dbo means "domain base object" and indicates a system table. Answer: FALSE Diff: 1 Page Ref: 10A-15

6) SQL Server 2012 provides two ways to create tables—graphically and through SQL code. Answer: TRUE Diff: 1 Page Ref: 10A-12 7) With SQL Server, the preferred way to create database structures is with SQL statements. Answer: TRUE Diff: 1 Page Ref: 10A-12

8) SQL Statements can be submitted to SQL Server from the Microsoft SQL Server Management Studio. Answer: TRUE Diff: 2 Page Ref: 10A-12

9) To use an SQL Server reserved word as a user identifier, it must be enclosed in square brackets [ ]. Answer: TRUE Diff: 3 Page Ref: 10A-12

10) The Microsoft SQL Server Management Studio cannot verify SQL statements until they are executed. Answer: FALSE Diff: 1 Page Ref: 10A-14

11) Stored procedures and triggers in SQL Server may become confused while executing stored procedures and triggers when SQL Server special words such as TRANSACTION appear as table or other names, even if they are enclosed in square brackets [ ]. Answer: TRUE Diff: 2 Page Ref: 10A-12

12) The name TRANSACTION is so special to SQL Server that no stored procedures will work on a table with that name. Answer: TRUE Diff: 3 Page Ref: 10A-12

13) In the Microsoft SQL Server Management Studio, the starting value for a field that has been set as an IDENTITY in SQL Server is set in the seed (StartingValue) property. Answer: TRUE Diff: 3 Page Ref: 10A-12 14) In the Microsoft SQL Server Management Studio, the amount to add to the surrogate key value when adding a new row is specified by the increment property. Answer: TRUE Diff: 3 Page Ref: 10A-12 15) When viewing a table in the GUI tools table design window in SQL Server, the primary key is indicated by a black arrow. Answer: FALSE Diff: 2 Page Ref: Fig 10A-17

16) When viewing a table in the GUI tools table design window in SQL Server, the notation that should supply the values for a surrogate key is that the Identity property of that column is set to "Yes." Answer: TRUE Diff: 2 Page Ref: Fig 10A-17

SQL Server 2008 provides two ways to create tables graphically and through SQL code.

Selected Answer: True

ALSO TRUE for SQL SERVER 2012

59) SQL statements can be passed to SQL Server using the _____. A) Enterprise Manager B) Query Analyzer C) Microsoft SQL Server Management Studio D) A and B E) A, B, and C Answer: C Diff: 1 Page Ref: 10A-8

60) SQL Server tables can be created and modified using _____. A) Enterprise Manager B) Query Analyzer C) Microsoft SQL Server Management Studio D) A and B E) A, B, and C Answer: C Diff: 1 Page Ref: 10A-12

61) When creating SQL Server tables in a database, SQL statements can now also be submitted by using _____. A) Internet Explorer B) Microsoft PowerShell C) Microsoft Word D) Microsoft Excel E) PL/SQL Answer: B Diff: 2 Page Ref: 10A-7-10A-8 7 ScholarStock

62) Anytime you want to use an SQL Server reserved word as a user identifier, enclose it in _____. A) { } B) [ ] C) ( ) D) " " E) < > Answer: B Diff: 1 Page Ref: 10A-12

63) The column(s) that is/are the primary key in an SQL table are indicated by _____ when viewed graphically in Enterprise Manager. A) the black arrow symbol B) the outline arrow symbol C) the key symbol D) the asterisk symbol E) being underlined Answer: C Diff: 1 Page Ref: Fig 10-17

64) What name is so special to SQL Server that no stored procedure will work on a table with that name, not even if it is enclosed in brackets? A) NAME B) UPDATE C) TRANSACTION D) SQL E) KEY Answer: C Diff: 2 Page Ref: 10A-12

65) When reviewing a table in the SQL Server GUI tools, which property should be set to "Yes" to indicate that the column is a surrogate key for which SQL Server should automatically supply values? A)
Identity B) Surrogate C) AutoIncrement D) AutoNumber E) Sequence Answer: A Diff: 3 Page Ref: Fig 1017 8 ScholarStock

66) Using only SQL Server tools, you can enter data into a table by _____. A) entering it into a table grid in the Microsoft SQL Server Management Studio B) using INSERT statements through the Microsoft SQL Server Management Studio C) using an SQL Server form D) A or B E) A, B, or C Answer: D Diff: 2 Page Ref: 10A-22

67) When creating a T-SQL statement, you should: A) predefine it graphically in Enterprise Manager. B) limit it to one table. C) limit it to two tables. D) end it with the semicolon normally used to terminate SQL statements. E) include the word "View" in the view name. Answer: D Diff: 2 Page Ref: 10A-35

68) Which of the following is true about indexes in SQL Server? A) SQL Server automatically creates indexes for columns appearing in WHERE clauses in queries. B) SQL Server pads all indexes. C) SQL Server supports filling up to a limit of 70 percent. D) SQL Server automatically creates indexes for foreign keys. E) All of the above. Answer: D Diff: 2 Page Ref: 10A-2

18) The developer can direct SQL Server to create indexes on non-key fields. Answer: TRUE Diff: 2 Page Ref: 10A-20 2 ScholarStock

Background image of page 2

19) SQL Server supports three types of indexes: clustered, unclustered, and random. Answer: FALSE Diff: 2 Page Ref: 10A-20

20) In an SQL Server clustered index, the data are stored in the bottom level of the index and in the same order as the index. Answer: TRUE Diff: 3 Page Ref: 10A-20

21) With SQL Server, a cluster index has the data with, and in the same order as, the bottom level of the index. Answer: TRUE Diff: 3 Page Ref: 10A-20

22) With SQL Server, only one cluster index is allowed per table. Answer: TRUE Diff: 2 Page Ref: 10A-20

23) With SQL Server, a nonclustered index does not contain data but has pointers to the data. Answer: TRUE Diff: 3 Page Ref: 10A-20

24) With SQL Server, the clustered index is faster for data retrieval. Answer: TRUE Diff: 2 Page Ref: 10A-

20

25) Data can be entered directly into a table in Enterprise Manager or via SQL INSERT statements submitted through the Microsoft SQL Server Management Studio. Answer: TRUE Diff: 1 Page Ref: 10A22

26) Database views can only be created in SQL Server through the use of SQL commands. Answer: FALSE Diff: 1 Page Ref: 10A-31

27) One means of processing an SQL Server database is to create application code in a programming language and invoke SQL Server DBMS commands from that program. Answer: TRUE Diff: 1 Page Ref:
10A-33-10A-35

28) One means of processing an SQL Server database is to save groups of SQL Server DBMS commands in a text file and then process this file in the Microsoft SQL Server Management Studio. Answer: TRUE Diff: 1 Page Ref: 10A-33-10A-35 3 ScholarStock

29) PL/SQL is a programming language for SQL Server that adds programming constructs to the SQL language. Answer: FALSE Diff: 2 Page Ref: 10A-35

30) SQL Server user variables and parameters start with the @ symbol. Answer: TRUE Diff: 1 Page Ref: 10A-35

31) When you use the Microsoft SQL Server Management Studio to run and rerun a set of SQL Server DBMS commands stored in a text file, the stored procedure must start with CREATE PROCEDURE. Answer: FALSE Diff: 3 Page Ref: 10A-40-10A-41

32) SQL Server supports AFTER and INSTEAD OF triggers, but not BEFORE triggers. Answer: TRUE Diff: 2 Page Ref: 10A-49

33) SQL Server AFTER triggers may be assigned to either tables or views. Answer: FALSE Diff: 2 Page Ref: 10A-49

34) SQL Server AFTER triggers may be used with insert and update actions, but not delete actions. Answer: FALSE Diff: 2 Page Ref: 10A-49

35) In SQL Server, triggers can roll back the transaction that caused them to be fired. Answer: TRUE Diff: 2 Page Ref: 10A-49-10A-50

36) In SQL Server, insert and update triggers store new values in a pseudotable named inserted . Answer: TRUE Diff: 2 Page Ref: 10A-50

1) Big Data is the name given to the enormous datasets generated by Web 2.0 applications. Answer: TRUE Diff: 1 Page Ref: 534-535

2) Business Intelligence (BI) systems are information systems that help users analyze and use data. Answer: TRUE Diff: 1 Page Ref: 536

3) Business Intelligence (BI) systems support operational activities. Answer: FALSE Diff: 1 Page Ref: 536

4) Business Intelligence (BI) systems obtain data in three different ways. Answer: TRUE Diff: 1 Page Ref: 536-537 Fig 12-2

5) Business Intelligence (BI) reporting systems are used to filter data, sort data, group data and make simple calculations based on the data. Answer: TRUE Diff: 2 Page Ref: 537 Fig 12-3

6) Business Intelligence (BI) reporting systems can analyze data using standard SQL. Answer: TRUE Diff: 1 Page Ref: 537

7) Business Intelligence (BI) reporting systems summarize the current status of business activities and compare that status with past events, but not with predicted future activities. Answer: FALSE Diff: 3 Page Ref: 537

8) Data mining uses sophisticated statistical and mathematical techniques to perform what-if analyses, to make predictions, and to facilitate decision making. Answer: TRUE Diff: 1 Page Ref: 537

9) Report delivery is more important for data mining than it is for reporting systems. Answer: FALSE Diff: 2 Page Ref: 537

10) A data warehouse is a database system that has data, programs and personnel specialized in Business Intelligence (BI) processing. Answer: TRUE Diff: 2 Page Ref: 538 1 ScholarStock

11) Data warehouses are populated with data prepared by data extraction, transformation and load (ETL) programs. Answer: TRUE Diff: 2 Page Ref: 538 Fig 12-4

12) Data warehouse data are frequently denormalized. Answer: TRUE Diff: 1 Page Ref: 538

13) Data warehouses also store the data warehouse metadata. Answer: TRUE Diff: 3 Page Ref: 538-539

14) Data warehouses often include data purchased from outside vendors. Answer: TRUE Diff: 1 Page Ref:
540 Fig 12-4

15) Metadata about the data's source, format, assumptions and constraints are kept in a data warehouse metadata database. Answer: TRUE Diff: 2 Page Ref: 538-539

16) Problematic data are called "dirty data." Answer: TRUE Diff: 1 Page Ref: 539-540

17) A data mart is a collection of data that addresses a particular component of a functional area of a business. Answer: TRUE Diff: 1 Page Ref: 540-541 Fig 12-7

18) Operational databases store historical data. Answer: FALSE Diff: 1 Page Ref: 542 Fig 12-8

19) Dimensional databases are used for analytical data processing. Answer: TRUE Diff: 1 Page Ref: 541 Fig 12-8

20) Dimensional databases use the star schema. Answer: TRUE Diff: 1 Page Ref: 542

21) Operational databases contain a fact table. Answer: FALSE Diff: 1 Page Ref: 5

68) Which of the following popular data mining techniques require special software? A) Decision tree analysis B) Logistic regression analysis C) Neural networks D) A and B E) A, B, and C Answer: E Diff: 3 Page Ref: J-22-J-23

69) Market basket analysis is: A) a data mart specific technique. B) a reporting technique. C) an RFM application. D) an OLAP application. E) a data mining technique. Answer: E Diff: 2 Page Ref: J-2

Here is Market Basket data for 1,000 transactions at a Drug Store. Use the table below to answer the following questions.

1,000 Transactions	Drug A 500	Drug B 300	Drug C 200
Drug A	0	100	125
Drug B	100	25	75
Drug C	125	75	0
No Additional Drug	200	100	0

70) Refer to the table above. What is the definition of support for Drug A and Drug B? A) The probability that Drug A will be purchased. B) The probability that Drug B will be purchased. C) The probability that both Drug A and Drug B will be purchased. D) The probability that Drug A will be purchased given that Drug B has been purchased. E) The probability that Drug B will be purchased given that Drug A has been purchased. Answer: C Diff: 2 Page Ref: J25-J26 Fig J-18

71) Refer to the table above. What is the support for Drug A and Drug B? A) .025 B) .075 C) .100 D) .125 E) .200 Answer: C Diff: 3 Page Ref: J25-J26 Fig J-18

72) Refer to the table above. What is the definition of confidence for Drug B with reference to Drug A? A) The probability that Drug A will be purchased. B) The probability that Drug B will be purchased. C) The probability that both Drug A and Drug B will be purchased. D) The probability that Drug A will be purchased given that Drug B has been purchased. E) The probability that Drug B will be purchased given that Drug A has been purchased. Answer: E Diff: 2 Page Ref: J25-J26 Fig J-18 11 ScholarStock

Background image of page 11

QUESTION 73 not visible in the pool..

A has been purchased? A) .200 B) .250 C) .333 D) .400 E) .500 Answer: A Diff: 3 Page Ref: J25-J26 Fig J-18

74) Refer to the table above. What is the lift for Drug B being purchased given that Drug A has been purchased? A) .200 B) .250 C) .333 D) .500 E) .667 Answer: E Diff: 3 Page Ref: J25-J26 Fig J-18

75) What are Business Intelligence (BI) systems? Answer: Business Intelligence (BI) systems are information systems used by managers and other business professionals to analyze past and current activities and to predict future events. BI systems do not support the recording and processing of operational data — this is left to transaction processing systems. Instead, BI systems are management support systems use that produce information for assessment, analysis, planning and control. There are two main categories of BI systems: reporting systems and data mining applications. Diff: 2 Page Ref: J-4J_____

Database Processing, 12e (Kroenke/Auer)

Chapter 13: Database Processing for Business Intelligence Systems

1) Business Intelligence (BI) systems are information systems that help users analyze and use data.

Answer: TRUE

Diff: 1 Page Ref: 549

2) Business Intelligence (BI) systems support operational activities.

Answer: FALSE

Diff: 1 Page Ref: 549

3) Business Intelligence (BI) systems obtain data in three different ways.

Answer: TRUE

Diff: 1 Page Ref: 555 Fig 13-1

4) Business Intelligence (BI) reporting systems are used to filter data, sort data, group data and make simple calculations based on the data.

Answer: TRUE

Diff: 2 Page Ref: 550 Fig 13-2

5) Business Intelligence (BI) reporting systems can analyze data using standard SQL.

Answer: TRUE

Diff: 1 Page Ref: 550

6) Business Intelligence (BI) reporting systems summarize the current status of business activities and compare that status with past events, but not with predicted future activities.

Answer: FALSE

Diff: 3 Page Ref: 550

7) Data mining uses sophisticated statistical and mathematical techniques to perform what-if analyses, to make predictions, and to facilitate decision making.

Answer: TRUE

Diff: 1 Page Ref: 550

8) Report delivery is more important for data mining than it is for reporting systems.

Answer: FALSE

Diff: 2 Page Ref: 551

9) Problematic data are called "dirty data."

Answer: TRUE

Diff: 1 Page Ref: 552

10) A data warehouse is a database system that has data, programs and personnel specialized in Business Intelligence (BI) processing.

Answer: TRUE

Diff: 2 Page Ref: 551

11) Data warehouses are populated with data prepared by data extraction, transformation and load (ETL) programs.

Answer: TRUE

Diff: 2 Page Ref: 551-552 Fig 13-3

12) Data warehouse data are frequently denormalized.

Answer: TRUE

Diff: 1 Page Ref: 551

13) Data warehouses also store the data warehouse metadata.

Answer: TRUE

Diff: 3 Page Ref: 552

14) Data warehouses often include data purchased from outside vendors.

Answer: TRUE

Diff: 1 Page Ref: 554 Fig 13-3

15) Metadata about the data's source, format, assumptions and constraints are kept in a data warehouse metadata database.

Answer: TRUE

Diff: 2 Page Ref: 552

16) A data mart is a collection of data that addresses a particular component of a functional area of a business.

Answer: TRUE

Diff: 1 Page Ref: 554 Fig 13-6

17) Business Intelligence (BI) Reporting systems are intended to create meaningful information from disparate data sources and to deliver that information to the proper users on a timely basis.

Answer: TRUE

Diff: 2 Page Ref: 563

18) RFM analysis is a way of analyzing and ranking customers based on online survey data.

Answer: FALSE

Diff: 2 Page Ref: 563

19) In a common form of RFM analysis, customers are sorted in five groups and given an associated score depending on their group.

Answer: TRUE

Diff: 1 Page Ref: 563

20) In a common form of RFM analysis, customers with an R score of 5 are in the 20% of customers who have the most recent orders.

Answer: FALSE

Diff: 2 Page Ref: 563

21) In RFM analysis, R stands for "how recently."

Answer: TRUE

Diff: 2 Page Ref: 563

22) In RFM analysis, F stands for "how frequently."

Answer: TRUE

Diff: 2 Page Ref: 563

23) In RFM analysis, M stands for "how much money."

Answer: TRUE

Diff: 2 Page Ref: 563

24) In a common form of RFM analysis, a score of 1 is "high" or "good" while a score of 5 is "low" or "bad."

Answer: TRUE

Diff: 3 Page Ref: 563

25) In a common form of RFM analysis, an RFM score of {5 1 1} means that the customer orders frequently and orders items of high monetary value but has not ordered anything for some time.

Answer: TRUE

Diff: 3 Page Ref: 563-564

26) A reporting system does not maintain a database of metadata.

Answer: FALSE

Diff: 2 Page Ref: 567-568 Fig13-24

27) A digital dashboard is an electronic display customized for an individual user.

Answer: TRUE

Diff: 2 Page Ref: 568

28) Report modes include dynamic and static.

Answer: FALSE

Diff: 3 Page Ref: 568 Fig 13-25

29) Reports that do not change once prepared are called static reports.

Answer: TRUE

Diff: 1 Page Ref: 568 Fig 13-25

30) Reports that rely on the most current data when they are prepared are called dynamic reports.

Answer: TRUE

Diff: 1 Page Ref: 568 Fig 13-25

31) A report that is sent to users on a predetermined schedule is called a push report.

Answer: FALSE

Diff: 2 Page Ref: 569

32) A report that is sent to users only upon the user's request is called on-demand report.

Answer: FALSE

Diff: 2 Page Ref: 569

33) Report systems include three functions: authoring, management and delivery.

Answer: TRUE

Diff: 1 Page Ref: 569

34) Report authoring includes responsibility for delivery of the report.

Answer: FALSE

Diff: 2 Page Ref: 569-572

35) Report management defines who gets the report, how often, and how it is to be delivered.

Answer: TRUE

Diff: 2 Page Ref: 569-572

36) Report management systems include the ability to create associated user accounts and groups.

Answer: TRUE

Diff: 2 Page Ref: 569-572

37) OLAP provides the ability to sum, count, average and perform other simple arithmetic operations on groups of data.

Answer: TRUE

Diff: 1 Page Ref: 572

38) An OLAP cube is limited to three axes.

Answer: FALSE

Diff: 3 Page Ref: 572

39) The term drill down refers to the capability of seeing the data in smaller and smaller units.

Answer: TRUE

Diff: 2 Page Ref: 576

40) Data mining is the application of mathematical and statistical techniques to find patterns and relationships that can be used to classify and predict future outcomes.

Answer: TRUE

Diff: 1 Page Ref: 577

41) Most data mining techniques are simple and easy to use.

Answer: FALSE

Diff: 1 Page Ref: 578

42) In unsupervised data mining, explanations are created after patterns are found.

Answer: TRUE

Diff: 3 Page Ref: 578

43) In unsupervised data mining, a model or hypothesis is created ahead of time.

Answer: FALSE

Diff: 2 Page Ref: 578

44) In supervised data mining, statistical techniques are used to estimate the parameters of the model.

Answer: TRUE

Diff: 2 Page Ref: 580

45) Cluster analysis is a commonly used technique in supervised data mining.

Answer: FALSE

Diff: 3 Page Ref: 578-580

46) As specified in the text, three popular data mining techniques are cluster analysis, logistic regression analysis, and neural networks.

Answer: FALSE

Diff: 3 Page Ref: 580

47) Market basket analysis is a data mining technique for determining which sets of products customers tend to purchase at the same time.

Answer: TRUE

Diff: 2 Page Ref: 580

48) Market basket analysis as described in the text uses the technique of conditional probabilities.

Answer: TRUE

Diff: 3 Page Ref: 580-581

49) In market basket analysis, support is defined as the probability that two items will be purchased together.

Answer: TRUE

Diff: 2 Page Ref: 581

50) In market basket analysis, confidence is the probability that one particular product will be purchased.

Answer: FALSE

Diff: 2 Page Ref: 581

51) In market basket analysis, lift is confidence divided by support.

Answer: FALSE

Diff: 3 Page Ref: 581

52) Business Intelligence (BI) systems do which of the following?

A) Analyze current and past activities

B) Predict future events

C) Record and process transactions

D) A and B

E) A, B, and C

Answer: D

Diff: 1 Page Ref: 549

53) Business Intelligence (BI) systems fall into which of the following categories?

A) Processing

B) Reporting

C) Data mining

D) A and B

E) B and C

Answer: E

Diff: 1 Page Ref: 550

54) Business Intelligence (BI) systems obtain their data by which of the following means?

A) Read and process data from an operational database

B) Process extracts from operational databases

C) Process data purchased from data vendors

D) A and B

E) A, B and C

Answer: E

Diff: 2 Page Ref: 550 Fig 13-1

55) Business Intelligence (BI) reporting systems can do which of the following operations?

A) Filter data

B) Group data

C) Modify data

D) A and B

E) A, B and C

Answer: D

Diff: 2 Page Ref: 550

56) One Business Intelligence (BI) reporting system that uses extensions to SQL is: A)

cluster analysis.

B) OLAP.

C) regression analysis.

D) RFM analysis.

E) A, B and C

Answer: B

Diff: 2 Page Ref: 530

57) Data mining applications are used to accomplish which of the following tasks?

A) Perform what-if analysis

B) Make predictions

C) Facilitate decision making

D) A and B

E) A, B and C

Answer: E

Diff: 1 Page Ref: 530

58) Which of the following are true about data mining applications? A)

They use sophisticated mathematical techniques.

B) They use sophisticated statistical techniques.

C) Their report delivery is more important than report delivery for reporting systems.

D) A and B

E) A, B and C

Answer: D

Diff: 3 Page Ref: 550-551

59) Which of the following is a reason that operational data are difficult to read?

A) Dirty data

B) Missing values

C) Nonintegrated data

D) A and B

E) A, B and C

Answer: E

Diff: 2 Page Ref: 552-554 Fig 13-4

60) We have obtained access to the company's operational data. In one record, we find that a customer's age has been recorded as "337." This is an example of _____.

A) dirty data

B) inconsistent data

C) nonintegrated data

D) a "wrong format" problem

E) a "too much data" problem Answer: A

Diff: 2 Page Ref: 552-554

61) We have obtained access to the company's operational data. We examine 50 records for customers with phone numbers that should use the current area code of 345. Of these 50 records, we find 10 that still use an older area code of 567. This is an example of _____.

A) dirty data

B) inconsistent data

C) nonintegrated data

D) a "wrong format" problem

E) a "too much data" problem

Answer: B

Diff: 3 Page Ref: 552-554

62) We have obtained access to the company's operational data. We have been asked to produce a report with an item by item analysis of sales, but the only sales figure available is the total sale value for each order. This is an example of _____.

A) dirty data

B) inconsistent data

C) nonintegrated data

D) a "wrong format" problem

E) a "too much data" problem Answer: D

Diff: 3 Page Ref: 552-554

63) A data warehouse database differs from an operational database because: A)

data warehouse data are not stored in tables.

B) data warehouse databases do not have metadata.

C) data warehouse data are often denormalized.

D) B and C

E) A, B and C

Answer: C

Diff: 2 Page Ref: 551

64) Which of the following are components of a data warehouse?

A) Data extract, transform, and load (ETL) preparation programs

B) Data warehouse data

C) Data warehouse metadata

D) B and C

E) A, B and C

Answer: E

Diff: 2 Page Ref: 551-552 Fig 13-3

65) A data mart differs from a data warehouse in that: A)

it has a smaller database.

B) it deals with a particular component or functional area of the business.

C) data mart users do not have the data management expertise of data warehouse employees.

D) A and B

E) A, B and C

Answer: E

Diff: 1 Page Ref: 554-555 Fig 13-6

66) A Business Intelligence (BI) reporting system:

A) creates meaningful information from disparate data sources.

B) delivers information to users on a timely basis.

C) uses statistical procedures to predict future events.

D) A and B

E) A, B and C

Answer: D

Diff: 2 Page Ref: 563

67) RFM analysis analyzes and ranks customers based on: A)

their purchasing patterns.

B) their income status.

C) their residential location.

D) A and B

E) A, B and C

Answer: A

Diff: 2 Page Ref: 563

68) The "R" in RFM analysis stands for _____.

A) rank

B) recent

C) relationship

D) readiness E) random

Answer: B

Diff: 1 Page Ref: 563

69) The "M" in RFM analysis stands for _____.

A) money

B) mostly

C) modest

D) modern

E) modem

Answer: A

Diff: 2 Page Ref: 563

70) RFM scores commonly range from _____, with _____ being the "high" or "most desirable" (from the vendors point of view) score.

A) 0 to 5; 0

B) 0 to 5; 5

C) 1 to 5; 1

D) 1 to 5; 5

E) 1 to 10; 10

Answer: C

Diff: 2 Page Ref: 563

71) We have done an RFM analysis on our customer data. John Smith has a score of {5 1 1}. This means that John:

A) has ordered recently, and orders a lot when he orders.

B) hasn't ordered recently, but orders a lot when he orders.

C) has ordered recently, but doesn't order a lot when he orders.

D) hasn't ordered recently, and doesn't order a lot when he orders.

E) None of the above is correct.

Answer: B

Diff: 2 Page Ref: 563-564

72) The metadata maintained by a reporting system database describes _____.

A) reports

B) users

C) groups

D) roles

E) All of the above.

Answer: E

Diff: 2 Page Ref: 567-568

73) The reports generated by a reporting system can be classified as _____.

A) static

B) dynamic

C) fluid

D) A and B

E) A, B and C

Answer: D

Diff: 1 Page Ref: 568 Fig 13-25

74) The reports generated by a reporting system can be delivered in which of the following media?

A) Paper

B) Web portal

C) Digital dashboard

D) E-mail

E) All of the above.

Answer: E

Diff: 1 Page Ref: 568-569 Fig13-25

75) A report generated by a reporting system is delivered to the appropriate users via an electronic display customized for each user. This system uses which of the following media?

A) Paper

B) Web portal

C) Digital dashboard

D) E-mail

E) None of the above.

Answer: C

Diff: 2 Page Ref: 568-569

76) A report generated by a reporting system is delivered to the appropriate users via a printed report. This system uses which of the following report modes?

A) Push

B) Pull

C) Query

D) OLAP

E) None of the above.

Answer: A

Diff: 2 Page Ref: 569 Fig 13-25 77)

OLAP stands for _____.

A) OnLine Analytical Processing

B) OffLine Analytical Processing

C) OnLine Analysis Process

D) OffLine Analysis Process

E) Old, Lazy And Particular

Answer: A

Diff: 2 Page Ref: 550

78) In OLAP, the data item of interest is called a _____.

A) level

B) dimension

C) measure

D) member

E) slice

Answer: C

Diff: 1 Page Ref: 572

79) In OLAP, the characteristic of a measure is called a _____.

A) level

B) dimension

C) measure

D) member

E) slice

Answer: B

Diff: 2 Page Ref: 572

80) An OLAP cube is called that because some products show OLAP displays on _____ axes.

A) one

B) two

C) three

D) four

E) a number that varies depending upon the analysis being done

Answer: C

Diff: 2 Page Ref: 572

81) The term drill down means the user wants to _____.

A) summarize data

B) get older data

C) sort

D) get more details

E) aggregate data

Answer: D

Diff: 1 Page Ref: 576

82) Data mining techniques are used to find patterns and relationships that can be used to _____.

A) report

B) classify

C) predict

D) B and C

E) A, B and C

Answer: D

Diff: 2 Page Ref: 577

83) Which of the following is a common data mining technique?

A) Cluster analysis

B) Regression analysis

C) RFM analysis

D) A and B

E) A, B and C

Answer: D

Diff: 2 Page Ref: 578-580

84) Which of the following is a supervised data mining technique?

A) Cluster analysis

B) Regression analysis

C) RFM analysis

D) A and B

E) A, B and C

Answer: B

Diff: 2 Page Ref: 580

85) Which of the following popular data mining techniques require special software?

A) Decision tree analysis

B) Logistic regression analysis

C) Neural networks

D) A and B

E) A, B and C

Answer: E

Diff: 3 Page Ref: 580

86) Market basket analysis is _____.

A) a reporting technique

B) an OLAP application

C) a data mining technique

D) A and B

E) A, B and C

Answer: C

Diff: 2 Page Ref: 580

Here is Market Basket data for 1,000 transactions at a Drug Store. Use the table below to answer the following questions.

1,000

Transactions	Drug A
500	Drug B

	300	Drug C	
	200		
Drug A 0	100	125	
Drug B 100	25	75	
Drug C 125	75	0	
No Additional Drug	200	100	0

87) Refer to the table above. What is the definition of support for Drug A and Drug B? A) The probability that Drug A will be purchased.

B) The probability that Drug B will be purchased.

C) The probability that both Drug A and Drug B will be purchased.

D) The probability that Drug A will be purchased given that Drug B has been purchased.

E) The probability that Drug B will be purchased given that Drug A has been purchased.

Answer: C

Diff: 2 Page Ref: 580-582 Fig13-37

88) Refer to the table above. What is the support for Drug A and Drug B?

A) .025

B) .075

C) .100

D) .125

E) .200

Answer: C

Diff: 3 Page Ref: 580-582 Fig13-37

89) Refer to the table above. What is the definition of confidence for Drug B with reference to Drug A?

A) The probability that Drug A will be purchased.

B) The probability that Drug B will be purchased.

C) The probability that both Drug A and Drug B will be purchased.

D) The probability that Drug A will be purchased given that Drug B has been purchased.

E) The probability that Drug B will be purchased given that Drug A has been purchased.

Answer: E

Diff: 2 Page Ref: 580-582 Fig13-37

90) Refer to the table above. What is the confidence of Drug B being purchased given that Drug A has been purchased?

A) .200

B) .250

C) .333

D) .400

E) .500

Answer: A

Diff: 3 Page Ref: 580-582 Fig13-37

91) Refer to the table above. What is the lift for Drug B being purchased given that Drug A has been purchased? A) .200

B) .250

C) .333

D) .500

E) .667

Answer: E

Diff: 3 Page Ref: 580-582 Fig13-37

92) What are Business Intelligence (BI) systems?

Answer: Business Intelligence (BI) systems are information systems used by managers and other business professionals to analyze past and current activities and to predict future events. BI systems do not support the recording and processing of operational data • this is left to transaction processing systems. Instead, BI systems are management support systems that produce information for assessment, analysis, planning and control. There are two main categories of BI systems: reporting systems and data mining applications.

Diff: 1 Page Ref: 549-550

93) Why is operational data sometimes unusable for Business Intelligence (BI) use?

Answer: Data in operational databases can suffer from a number of problems. These include: (1) Dirty data • data that has problems with it, for example an age of "323;" (2) Missing values data values that are unknown, for example a person's age; (3) Inconsistent data old data values that needed to be updated may not have been for example a ZIP code in an area that was split into two new ZIP codes; (4) Data not integrated when data from two or more data sets is used, the data may be from two or more different DBMSs; (5) Data in the wrong format data values may have been recorded at an inappropriate value for the needed analysis for example, distance may have been recorded in miles when we need meters; (6) Too much data it is possible to simply have very large data sets in terms of records or fields.

Diff: 2 Page Ref: 552-554 Fig 13-4

94) What is a data warehouse?

Answer: A data warehouse is a nonoperational database intended for use by Business Intelligence (BI) applications. A data warehouse is a database system that has data, programs and personnel who specialize in the preparation of data for Business Intelligence processing. Database in data warehouses are frequently denormalized to speed up BI processes. The components of a data warehouse include programs for data extraction/cleaning/preparation, data warehouse DBMSs, data warehouse metadata and data warehouse data. The data warehouse metadata database records the data warehouse data's source, format, assumptions, constraints and other relevant facts. Diff: 2 Page Ref: 551-552 Fig 13-3

95) Compare the characteristics of an operational database and a dimensional database.

Answer: An operational database is used for structured transaction data processing, while a dimensional database is used for unstructured analytical data processing. An operational database uses current data, while a dimensional database uses both current and historical data. With an operational database, data are inserted, modified and deleted by users, while in a dimensional database data are loaded and updated systemically by the data warehouse administrators. Diff: 2 Page Ref: 555 Fig 13-7

96) What is a slowly changing dimension?

Answer: A slowly changing dimension is an attribute or set of attributes that can change over time, but typically does not change that often. Address and phone number are examples of slowly changing dimensions.

Diff: 2 Page Ref: 555

97) What is a star schema?

Answer: The star schema describes the arrangement of the tables in a dimensional database. A fact table is the center of the star, while dimension tables are connected to the fact table and can be illustrated as the points of the star.

Diff: 2 Page Ref: 555-560

98) What is a conformed dimension?

Answer: A conformed dimension is a dimension table in a data warehouse that is used in more than one star schema. Since the same table is reused, we can be sure that the data is consistent within each start schema.

Diff: 2 Page Ref: 560-561 Fig13-17

99) What is a reporting system?

Answer: A reporting system is one of the main categories of a Business Intelligence (BI) system. Reporting systems are used to filter data, sort data, group data and make simple calculations based on the data. These systems summarize the current status of business activities and compare the current status to past or predicted results. They are also used to classify entities such as customers, employees, products, etc. Report delivery is a critical aspect of reporting systems. Diff: 1 Page Ref: 563 Fig 13-2

100) What is an Excel PivotTable?

Answer: A PivotTable is the reporting mechanism Excel uses to display OLAP reports. Both measures (a dimensional database "fact") and dimensions are displayed and can be rearranged. Drill-down is also available.

Diff: 2 Page Ref: 572-577

101) What is a data mining application?

Answer: A data mining application uses sophisticated statistical and mathematical techniques to do what-if analyses, forecast predicted values and facilitate decision making. In unsupervised data mining, analysts run tools against the raw data without a prior hypothesis or theory • after the analysis hypothesis and theories are created to explain the results. In supervised data mining, the analysts create the theory and hypotheses first, and then run tools on the raw data to calculate parameter values for the theoretical models. Report delivery is not as critical in data mining as it is in reporting systems. Diff: 1 Page Ref: 550-51, 577-580

1) Database redesign is rarely needed because databases are usually built correctly the first time.	Answer: FALSE
2) In a real sense, information systems and organizations do not just influence each other, but rather they create each other.	Answer: TRUE
3) A continuous circular process of changes in user behaviors and change in the information systems they use is a natural outcome of information system use.	Answer: TRUE
4) The continuous circular process of changes is known as the Systems Development Life Cycle (SDLC).	Answer: TRUE
5) Database redesign is equally difficult whether or not the database has data in it.	Answer: FALSE
6) In the database redesign process, it is often useful to test whether certain conditions or assumptions are valid before proceeding with the redesign.	Answer: TRUE
7) In the database redesign process, two SQL tools	Answer: FALSE

are useful for testing whether or not certain conditions or assumptions are valid: uncorrelated subqueries and EXISTS/NOT EXISTS.	
8) A correlated subquery looks very different from a noncorrelated subquery.	Answer: FALSE
9) Correlated subqueries can be used to verify functional dependencies.	Answer: TRUE
10) In a correlated subquery, the same table is used in the upper and lower SELECT statements.	Answer: TRUE
11) In the SQL statement: SELECT S1.CustName, S1.SalesRepNo FROM SALES S1; the "S1" is called an alias.	Answer: TRUE
12) In a correlated subquery, the DBMS can run the lower SELECT statement by itself and then send the results to the upper SELECT statement.	Answer: FALSE
13) In a correlated subquery, the DBMS must run the lower SELECT statement as a process that is nested within the upper SELECT statement.	Answer: TRUE
14) There is a common trap in writing a correlated subquery, which will cause no rows to ever be displayed in the results.	Answer: TRUE
15) Although correlated subqueries are useful in database redesign, they cannot be used to verify functional dependencies.	Answer: FALSE

16) When using queries with EXISTS and NOT EXISTS, the processing of the associated SELECT statements must be nested.	Answer: TRUE
17) The use of a double nested set of NOT EXISTS SELECT statements can be used to find rows that meet some specified condition for every row in a table.	Answer: TRUE
18) The use of a double nested set of NOT EXISTS SELECT statements is a famous pattern in SQL use.	Answer: TRUE
19) EXISTS and NOT EXISTS are actually just another form of correlated subqueries.	Answer: TRUE
20) Because EXISTS and NOT EXISTS are forms of correlated subqueries, the processing of the associated SELECT statements must be nested.	Answer: TRUE
21) The EXISTS keyword will be true if any row in the subquery meets the condition.	Answer: TRUE
22) The NOT EXISTS keyword will be true if any row in the subquery fails to meet the condition.	Answer: FALSE
23) A double nested set of NOT EXISTS SELECT statements can be used to find rows that meet some specified condition for every row in a table.	Answer: TRUE
24) When using a double nested set of NOT EXISTS SELECT statements, a row that does not match any row matches every row.	Answer: TRUE

25) The use of a double nested set of NOT EXISTS SELECT statements is so rare that even if you are a professional database developer you will probably never see it used.	Answer: FALSE
26) There is no good SQL command that can be used to change table names.	Answer: TRUE
27) The process of reading an actual database schema and producing a data model from that schema is called reverse engineering.	Answer: TRUE
28) The data model produced by reverse engineering is a true conceptual schema.	Answer: FALSE
29) The data model produced by reverse engineering may include some entities that should not appear in the data model.	Answer: TRUE
30) The design produced by reverse engineering may be described as a table-relationship diagram.	Answer: TRUE
31) The authors refer to the data model produced by reverse engineering as the RE data model.	Answer: TRUE
32) Because of the need to know the functional dependencies in a database, it is a good idea to create a dependency graph.	Answer: TRUE
33) Dependency graphs are graphical displays like bar charts.	Answer: FALSE

34) Typically, there are at least four different copies of the database schema used in the redesign process.	Answer: FALSE
35) A means must be created to recover all test databases to their original state during the redesign testing process.	Answer: TRUE
36) Even if an organization has a very large database, it will be possible to make a complete backup copy of the operational database prior to making structure changes.	Answer: FALSE
37) SQL contains an SQL command RENAME TABLENAME that can be used to change table names.	Answer: FALSE
38) SQL Server 2012 contains a system-stored procedure named sp_rename that can be used to change table names.	Answer: TRUE
39) Changing table names is complicated by the fact that constraints and triggers are often associated with the table and will also need to be changed.	Answer: TRUE
40) In order to minimize the need to change table names, some organizations have a policy that no user or application should ever employ the true name of a table, but use views as table aliases instead.	Answer: TRUE
41) To add a NULL column to a table, we simply use the MODIFY TABLE statement.	Answer: FALSE

42) If a DEFAULT constraint is included when a new column is added to a table, the default value is only applied to new rows and not to the existing rows at the time the new column is added.	Answer: TRUE
43) To add a NOT NULL column to a table, we first add a NULL column, then we insert values into every row, and finally we change the NULL constraint to NOT NULL.	Answer: TRUE
44) To drop a nonkey column from a table, no preliminary steps are needed and we can simply use the ALTER TABLE DROP COLUMN statement.	Answer: TRUE
45) To drop a foreign key column from a table, no preliminary steps are needed and we can simply use the ALTER TABLE DROP COLUMN statement.	Answer: FALSE
46) To drop a primary key column from a table the primary key constraint must first be dropped, but this does not require that related foreign keys based on the column be dropped.	Answer: FALSE
47) To drop a constraint, no preliminary steps are needed and we can simply use the ALTER TABLE DROP CONSTRAINT statement.	Answer: TRUE
48) Converting date, money or other more specific data types to char or varchar will usually succeed.	Answer: TRUE
49) To change the minimum cardinality on the parent side from zero to one, the foreign key, which would have been NULL, must be changed to NOT NULL.	Answer: TRUE

50) Depending on the DBMS, when changing the minimum cardinality on the parent side from zero to one, the foreign key constraint that defines the relationship may have to be dropped before the change is made and re-added afterwards.	Answer: TRUE

51) There are several difficulties with increasing cardinalities from 1:1 to 1:N, one of which is preserving the existing relationships.	Answer: FALSE
52) When increasing cardinalities from 1:N to N:M, we basically create a new intersection table, fill it with data and drop the old foreign key.	Answer: TRUE
53) When decreasing cardinalities, there will always be data loss.	Answer: TRUE
54) Adding new tables and relationships to a database is difficult.	Answer: FALSE
55) Deleting tables and relationships is basically a matter of dropping foreign key constraints and then dropping the tables.	Answer: TRUE
56) Database redesign is fairly easy when _____. A) information systems and organizations influence each other B) the design was done correctly the first time C) there is no data in the database D) good backups of the database are available E) All of the above	Answer: C
57) Which of the following is not a possible step in the database redesign process?	Answer: E

A) Checking whether certain conditions or assumptions about the data are valid B) Reverse engineering the data model C) Testing proposed changes D) Maintaining backups of the existing database E) All of the above are possible steps in the database redesign process.	
58) In the database redesign process, before proceeding with the redesign it is often useful to _____. A) check whether certain conditions or assumptions about the data are valid B) find out why the design was not done properly the first time C) stop information systems and users from influencing each other D) set standards for user behavior E) All of the above	Answer: A
59) In the database redesign process, SQL tools that are useful for testing whether or not certain conditions or assumptions are valid are _____. A) correlated subqueries B) EXISTS C) NOT EXISTS D) B and C E) A, B, and C	Answer: E
60) In a correlated subquery of a database that has tables TableOne and TableTwo, if table TableOne is used in the upper SELECT statements, then which table is used in the lower SELECT statement? A) TableOne B) TableTwo C) both TableOne and TableTwo D) either TableOne or TableTwo	Answer: A

E) neither TableOne nor TableTwo	
61) In the SQL statements SELECT C1.CustName, C1.SalesRepNo / FROM CUSTOMER C1; the "C1" is called a(n) _____. A) term B) alias C) convention D) phrase E) label	Answer: B
62) When running a correlated subquery, the DBMS _____. A) runs the lower SELECT statement by itself and then sends the results to the upper SELECT statement. B) runs the upper SELECT statement by itself and then sends the results to the lower SELECT statement. C) alternates running the lower SELECT statement with running the upper SELECT statement based on each result of the lower SELECT statement D) Either A or B may be used depending on the query. E) None of the above describes how a correlated subquery is run by the DBMS.	Answer: C
63) When running a correlated subquery, the DBMS always uses _____. A) regular processing B) nested processing C) "quick and dirty" processing D) SQL-92 processing E) a form of processing that is specific to the DBMS product	Answer: B

65) SQL queries that use EXISTS and NOT EXISTS are _____. A) normal subqueries B) correlated subqueries C) uncorrelated subqueries D) constraint dependent subqueries E) constraint independent subqueries	Answer: B
66) When running an SQL query that uses EXISTS, the EXISTS keyword will be true if _____. A) any row in the subquery meets the condition B) all rows in the subquery meet the condition C) no row in the subquery meets the condition D) any row in the subquery fails to meet the condition E) all rows in the subquery fail to meet the condition	Answer: A
67) When running an SQL query that uses NOT EXISTS, the NOT EXISTS keyword will be true if _____. A) any row in the subquery meets the condition B) all rows in the subquery meet the condition C) no row in the subquery meets the condition D) any row in the subquery fails to meet the condition E) all rows in the subquery fail to meet the condition	Answer: E
68) A double nested set of NOT EXISTS SELECT statements is _____. A) a famous pattern in SQL B) regularly used C) rarely used to nonexistent in the real world D) A and B E) A and C	Answer: D

69) In a double nested set of NOT EXISTS SELECT statements, _____. A) if a row does not match any row, then it matches every row B) if a row matches any row, then it matches every row C) if a row does not match any row, then it does not match every row D) if a row matches any row, then it does not match every row E) if a row does not match any row, then referential integrity has been broken	Answer: A
70) The process of reading an actual database schema and producing a data model from that schema is called _____. A) data modeling B) data engineering C) reverse engineering D) schema modeling E) schema engineering	Answer: C
71) The data model produced by reverse engineering is not truly a logical model because it will contain tables for _____. A) strong entities B) weak non-ID-dependent entities C) ID-dependent entities D) intersection tables E) supertype/subtype tables	Answer: D
72) The data model produced by reverse engineering is a(n) _____. A) conceptual schema B) internal schema C) dependency graph D) table-relationship diagram	Answer: D

E) entity-relationship diagram

73) Because of the need to know the functional dependencies in a database, it is a good idea to create a(n) _____. A) conceptual schema B) internal schema C) dependency graph D) table-relationship diagram E) entity-relationship diagram	Answer: C
74) A dependency graph should include: A) tables B) views C) triggers D) A and B E) A, B, and C	Answer: E
75) Which of the following different copies of the database schema is/are typically used in the database redesign process? A) Small test database B) Large test database C) Operational database D) A and B E) A, B, and C	Answer: E
76) In order to make sure the database redesign is working properly during the redesign testing process, a means must be created to _____. A) reverse engineer all test databases B) graph dependencies in all test databases C) recover all test databases to their original state D) A and B E) A, B, and C	Answer: C

77) To change a table name, we _____. A) use the SQL RENAME TABLE command B) use the SQL ALTER TABLENAME command C) use the SQL MODIFY TABLENAME command D) create a new table, move the data, and drop the old table E) None of the above is the correct way to change a table name.	Answer: D
78) When making any change to the database structure, we may need to check for effects of the change on _____. A) data B) foreign keys C) constraints D) triggers E) All of the above may need to be checked.	Answer: E
79) To add a NULL column to a table, we _____. A) use the REVISE TABLE command B) use the ALTER TABLE command C) use the MODIFY TABLE command D) create a new table with the NULL column, move the other data, and drop the old table E) None of the above is the correct way to add a NULL column.	Answer: B
80) If a DEFAULT constraint is included when a new column is added to a table, the default value is applied to _____. A) all existing rows at the time the column is added B) all new rows C) all new rows but only after the UPDATE command is issued D) A and B E) A and C	Answer: B

81) To add a NOT NULL column to a table, we _____.

Answer: D

A) use the REVISE TABLE command
B) use the ALTER TABLE command
C) use the MODIFY TABLE command
D) create a new NULL column, insert data values into every row, and change the NULL constraint to NOT NULL
E) None of the above is the correct way to add a NOT NULL column to a table.

82) When dropping a nonkey column from a table, which of the following steps is (are) included in the process? (The order of the steps listed below is not relevant, only the steps themselves.) A) Drop any column constraints from the table B) Drop the column from the table C) Drop any foreign keys constraints based on the column D) A and B E) A, B, and C	Answer: B
83) When dropping a primary key column from a table, which of the following steps is (are) included in the process? (The order of the steps listed below is not relevant, only the steps themselves.) A) Drop the primary key constraint from the table B) Drop the primary key column from the table C) Drop the foreign keys in other tables based on the primary key column D) A and B E) A, B, and C	Answer: E
84) When dropping a foreign key column from a table, which of the following steps is (are) included in the process? (The order of the steps listed below is not relevant, only the steps themselves.) A) Drop the foreign key constraint from the table	Answer: D

B) Drop the foreign key column from the table C) Drop the primary key in the other table referenced by the referential integrity constraint
D) A and B
E) A, B, and C

85) Changing cardinalities _____. A) never occurs in database redesign B) rarely occurs in database redesign C) commonly occurs in database redesign D) always occurs in database redesign E) cannot be done once a database is implemented	Answer: C
86) When changing column data types, which of the following data conversions will either usually or always succeed? A) Numeric → char or varchar B) Date or money → char or varchar C) Char or varchar → numeric, date or money D) A and B E) A, B, and C	Answer: D
87) When changing the minimum cardinality on the parent side of the relationship from zero to one, the foreign key _____. A) must be changed from NULL to NOT NULL B) must be changed from NOT NULL to NULL C) must be changed to a composite key D) must be changed to a surrogate key E) does not change	Answer: A
88) Which of the following are difficulties when changing the maximum cardinality from 1:1 to 1:N? A) Preserving the existing tables B) Preserving the existing relationships C) Preserving the existing data D) A and B	Answer: B
E) A, B, and C	

89) When increasing cardinalities from 1:N to N:M, which of the following steps is/are included in the process? (The order of the steps listed below is not relevant, only the steps themselves.) A) Create an intersection table B) Populate the intersection table C) Drop the old foreign key D) A and B E) A, B, and C	Answer: E
90) When decreasing cardinalities, there will always be _____. A) relationship loss B) foreign key loss C) data loss D) A and B E) A, B, and C	Answer: C
91) When dropping tables and relationships, which of the following steps is/are included in the process? (The order of the steps listed below is not relevant, only the steps themselves.) A) Drop the foreign key constraints from the tables B) Drop the tables C) Drop the primary key constraints from the tables D) A and B E) A, B, and C	Answer :D

In a relational database design, all relationships are expressed by _____. Creating a foreign key

The identifier of the entity becomes the _____ of the corresponding table Primary Key

Which of the following is not true about surrogate keys? They are non-unique within a table

Which of the following is not a step in the database design process?

Each attribute of an entity becomes a(n) _____ of a table. Create constraints and triggers Column

The first step in transforming an extended E-R model into a relational database design is to _____ Create a table for each entity

A surrogate key should be considered when _____.

The ideal primary key is _____. The key contains a lengthy text field All of the above

When representing a 1:1 relationship in a relational database design _____.

The key of either table may be placed as a foreign key into the other

To which of the following actions are referential integrity constraints not applied? Create

The Entity-Relationship (E-R) Model may contain which of the following items:
Tables, Attributes, Identifiers, and Relationships

T/F: When the key of one table is placed into a second table to represent a relationship, the key is called a "relational key" in the second table
False

T/F: A foreign key is a key that does not belong in any table. False

T/F: When creating a table in the relational database design from an entity in the extended E-R model, the attributes of the entity become the rows of the table.

T/F: Data types are consistent across all DBMS products.
False
False

T/F: An intersection table can have additional attributes besides the keys of its parent tables. False

T/F: It is easy to enforce the referential integrity actions for M-M relationships. False

T/F: When creating a relational database design from E-R diagrams, first create a relation for each relationship.
False

T/F: To represent a 1:N relationship in a relational database design, an intersection table is created
False

T/F: In a 1:N relationship, the term "parent" refers to the table on the "many" side of the relationship. False

T/F: A null value is an attribute value that has been set to zero
False

In Microsoft Visio 2010, when the Entity object is placed in a diagram, it is labeled as a(n): Table

T/F: One must always use a discriminator with Supertype/Subtype entities. False
(if not, code must be created)

When defining column properties, limitations on data values may contain: All
of these

T/F: An entity instance of an entity class is the representation of a particular entity and is described by the values of the attributes of the entity True

T/F: An entity class is described by the structure of the entities in that class.

True

T/F: In E-R modeling, an attribute describes the characteristics of an entity

True

T/F: A data model is a plan for a database design.

T/F: An entity is something that users want to track

True True

T/F: An identifier of an entity instance is one or more attributes that name or identify entity instances.

T/F: Entities of a given type are grouped into entity classes.

True True

T/F: An entity instance is the occurrence of a particular entity. True

T/F: The method of constructing data models used in the text is the extended entity-relationship (E-R) model True

T/F: Entities can be associated with one another in relationships True

From which movie were quotes or paraphrases used in

reference to our Past, Present & Future database activities?

Matrix

T/F: One of the most important reasons to learn to use Microsoft Visio is that it provides a convenient method of creating the data models and the database designs discussed in our text

True

T/F: When creating an E-R Diagram, one must occassionally ask questions of the data owner to gain a better understand of how the E-R Diagram must look

True

Microsoft Visio 2010 can be used to create

Both: Data models & database designs

T/F: In a database design, many-to-many (N:M) relationships are broken down into two one-to-many (1:N) identifying relationships between three ID-dependent entities. True

A hash mark across the relationship line near an entity indicates _____.

A minimum cardinality of "required"

A circle across the relationship line near an entity indicates _____.

A minimum cardinality of "optional"

For a relationship to be considered a binary relationship it

must satisfy which of the following conditions? It must involve exactly two entity classes

A composite attribute is an attribute that _____

Consists of a group of attributes

Entities of a given type are grouped into an _____

Entity class

Minimum cardinality refers to _____

Whether or not an instance of one entity class is required to be related to an instance of another entity class

The characteristics of a thing are described by its _____

Attributes

Maximum cardinality refers to _____.

The most instances of one entity class that can be involved in a relationship instance with another entity class

Which of the following is not a key element of an E-R model?

Objects

The occurrence of a particular entity is called an _____

Entity instance

T/F: The Microsoft Visio tool provides a way of showing IE

The principles for database redesign include:

Crow's Foot
symbols. True A,B,
and C

Which of the following is not a possible step in the
database redesign process?
All of the above are possible steps in the database redesign
process.

In the database redesign process, SQL tools that are useful
for testing whether or not certain conditions
or assumptions are valid are _____. A, B,
and C

Correlated Subqueries may be:
Thought of as recursive nesting

When running an SQL query that uses NOT EXISTS, the NOT
EXISTS keyword will be true if _____.
Any row in the subquery meets the condition

In a correlated subquery of a database that has tables
TableOne and TableTwo, and if table TableOne is used in
the upper SELECT statements, then which table is used
in the lower SELECT statement? TableOne

When running an SQL query that uses EXISTS, the EXISTS
keyword will be true if _____
Any row in the subquery meets the condition

Database redesign is fairly easy when _____
There is no data in the database

In a double nested set of NOT EXISTS SELECT statements,

_____

If a row does not match any row, then it matches every row

SQL queries that use EXISTS and NOT EXISTS are _____.

Correlated subqueries

In the SQL statements

SELECT C1.CustName, C1.SalesRepNo
FROM CUSTOMER C1;

the "C1" is called a(n) _____. Alias

In the database redesign process, before proceeding with the redesign it is often useful to _____

Check whether certain conditions or assumptions about the data are valid

The process of reading an actual database schema and producing a data model from that schema is called _____.

Reverse engineering

The data in SQL views is constructed using _____

SELECT statements

The SQL command used to create a virtual table is

Which SQL keyword is used to delete a table's structure?

CREATE VIEW DROP

Which SQL keyword is used to add one or more rows of data to a table?

INSERT

Given the SQL statement

```
CREATE TABLE SALESREP (
SalesRepNo int NOT NULL,
RepName char(35) NOT NULL,
HireDate date NOT NULL,

CONSTRAINT SalesRepPK PRIMARY KEY (SalesRepNo),
CONSTRAINT SalesRepAK1 UNIQUE (RepName)
);
```

we know that SalesRepNo is the _____. Primary Key

Given the SQL statement

```
CREATE TABLE SALESREP (
SalesRepNo int NOT NULL,
RepName char(35) NOT NULL,
HireDate date NOT NULL,

CONSTRAINT SalesRepPK PRIMARY KEY (SalesRepNo),
CONSTRAINT SalesRepAK1 UNIQUE (RepName)
);
```

we know that _____ RepName is a candidate key

Which SQL keyword is used to change the values of an

entire column?

Set

One advantage of using the CONSTRAINT phrase to define a primary is that the database designer controls the _____

Name of the constraint

When the correct SQL command is used to delete a table's structure, what happens to the data in the table? The data in the table is also deleted.

Which SQL keyword is used to impose restrictions on a table, data, or relationship?

CONSTRAINT

An SQL virtual table is called _____

A view

SQL views may be used to _____

Do all of the above

Which SQL statement is used to retrieve view instances?

SELECT

T/F: The SQL TRUNCATE TABLE statement only removes the part of the data from a table specified in the SQL WHERE clause False

T/F: Omitting the SQL WHERE clause from the SQL DELETE statement will delete all rows from a table. This has the same effect as the SQL TRUNCATE statement True

T/F: The SQL DROP TABLE Statement always removes the entire table and all of its content
False (Not if it is a parent table with a child table relationship established.)

T/F: Any Join The Is Not Explicitly An Inner Join Is An Outer Join.
False

T/F: Correlated subqueries can be used to verify functional dependencies.
True

T/F: The continuous circular process of changes is known as the Systems Development Life Cycle (SDLC). True

T/F: The use of a double nested set of NOT EXISTS SELECT statements can be used to find rows that meet some specified condition for every row in a table
True

T/F: To drop a constraint, no preliminary steps are needed and we can simply use the ALTER TABLE DROP CONSTRAINT statement. True

T/F: Database redesign is rarely needed because databases are usually built correctly the first time False

T/F: In a real sense, information systems and organizations do not just influence each other, but rather they create each other True

T/F: A continuous circular process of changes in user behaviors and change in the information systems they use is a natural outcome of information system use True

T/F: The process of reading an actual database schema and producing a data model from that schema is called reverse engineering. True

T/F: The author refers to the data model produced by reverse engineering as the RE data model True

T/F: The use of a double nested set of NOT EXISTS SELECT statements is a famous pattern in SQL use True

T/F: In a correlated subquery, the same table is used in the upper and lower SELECT statements. True

T/F: In the database redesign process, it is often useful to test whether certain conditions or assumptions are valid before proceeding with the redesign
True

T/F: The SQL keyword UNIQUE is used to define alternative keys
True

T/F: Rows can be removed from a table by using the SQL DELETE statement. True

T/F: The SQL keyword MODIFY is used to change the structure, properties or constraints of a table. False

T/F: The SQL keyword CHECK is used to limit column values to specific values. True

T/F: The SQL keyword CONSTRAINT is used in conjunction with the SQL keywords PRIMARY KEY and FOREIGN KEY True

T/F: When the correct SQL command is used to delete a table's structure, the command can only be used with a table that has already had its data removed. False

T/F: The SQL keyword CONSTRAINT is used to define one of five types of constraints (plus 1 optional) True

T/F: One or more rows can be added to a table by using the SQL INSERT statement.

True

T/F: The SQL CREATE statement is used to name a new table and describe the table's columns

True

T/F: The SQL keyword DROP is used to delete a table's structure. True

T/F: The SQL command SELECT is used to retrieve view instances True

T/F: The SQL keyword PRIMARY KEY is used to designate the column(s) that are the primary key for the table

True

chapter 9

Managing Multi-User Databases

True-False Questions

1. Database administration tasks have to be performed for single-user, personal databases.

Answer: *True* ***Level:*** *hard*

Page: *300*

2. Database administration is more important but less difficult in multi-user database systems than in single-user database systems.

Answer: *False* ***Level:*** *moderate*

Page: *300*

3. In general, the overall responsibility of the DBA is to facilitate the development and use of the database system.

Answer: *True* ***Level:*** *easy*

Page: *301*

4. The DBA has to find a balance between the conflicting goals of maximizing availability of the database to users and protecting the database.

Answer: *True* ***Level:*** *moderate*

Page: *301*

5. The DBA is responsible for managing changes to the database structure, but is rarely involved

in the original design of the structure.

Answer: *False* ***Level:*** *moderate*

Page: *301*

6. Changes in the database structure usually involve only one application.

Answer: *False* ***Level:*** *hard*

Page: *301*

7. One important reason for documenting changes to the database structure is for diagnosing errors.

Answer: *True* ***Level:*** *easy*

Page: *301*

8. Concurrency control measures are taken to ensure that one user's work has absolutely no influence on another user's work.

Answer: *False* ***Level:*** *moderate*

Page: *302*

9. A transaction is a group of alternative database actions from which the database can choose to perform only one of them.

Answer: *False* ***Level:*** *easy*

Page: *303*

10. Resource locking is one remedy to the lost update problem.

Answer: *True* ***Level:*** *moderate*

Page: *306*

11. "Explicit locks" are locks that are placed automatically by the DBMS.

Answer: *False* ***Level:*** *easy*

Page: *307*

12. Locks with large granularity are easy for the DBMS to administer but frequently cause conflicts.

Answer: *True* ***Level:*** *moderate*

Page: *307*

13. In general, the boundaries of a transaction should correspond to the boundaries of the database view it is processing.

Answer: *True* ***Level:*** *hard*

Page: *308*

14. Resource locking must be carefully planned because most DBMS products cannot detect a deadlock condition.

Answer: *False* ***Level:*** *moderate*

Page: *309*

15. Resources are locked for a shorter amount of time with pessimistic locking because the transaction is pre-processed.

Answer: *False* ***Level:*** *easy*

Page: *309*

16. In general, optimistic locking is the preferred technique for Internet databases.

Answer: *True* ***Level:*** *hard*

Page: *309*

17. A "dirty read" happens when one transaction reads a changed record that has not been committed to the database.

Answer: *True* ***Level:*** *hard*

Page: *312*

18. “Repeatable Reads” isolation is the most restrictive level of isolation.

Answer: *False* ***Level:*** *moderate*

Page: *312*

19. The goal of database security is to ensure that only authorized users can perform authorized activities at authorized times.

Answer: *True* ***Level:*** *easy*

Page: *304*

20. In regard to database security, neither the DBMS nor the database applications can enforce processing responsibilities.

Answer: *True* ***Level:*** *moderate*

Page: *305*

21. Processing responsibilities should be documented and encoded into manual procedures.

Answer: *True* ***Level:*** *easy*

Page: *305*

22. Processing rights may be implemented at the DBMS level.

Answer: *True* ***Level:*** *easy*

Page: *305*

23. All commercial DBMS products use some version of “username and password” as part of their security features.

Answer: *True* ***Level:*** *easy*

Page: *316*

24. The security provided by the DBMS often has to be augmented by additional security features within the application program.

Answer: *True* ***Level:*** *moderate*

Page: *319*

25. A “database save” is used to mark the end of a transaction.

Answer: *False* ***Level:*** *moderate*

Page: *320*

26. Reprocessing is normally the most convenient method for recovery after a system failure.

Answer: *False* ***Level:*** *easy*

Page: *320*

27. Rollforward and reprocessing are two different names for the same technique.

Answer: *False* ***Level:*** *hard*

Page: *320*

28. Both rollforward and rollback require the use of a log of transaction results.

Answer: *True* ***Level:*** *easy*

Page: *320*

29. The DBA should periodically analyze run-time statistics of database performance to help manage the DBMS.

Answer: *True* ***Level:*** *easy*

Page: *323*

30. A passive repository is preferred over an active repository because it requires less human intervention.

Answer: *False* ***Level:*** *moderate*

Page: *324*

Multiple Choice Questions

31. Which of the following is not a database administration responsibility of a DBA?

a.) managing the database structure

b.) managing data activity

c.) managing the DBMS

d.) maintaining the data repository

e.) **All of the above are database administration responsibilities of a DBA.**

Level: *easy*

Page: *300-301 [See Figure 9-1]*

32. Which of the following is true about making changes to the database structure?

a.) The DBA need not get input from users on the issue because it is a technical decision.

b.) Formal policies and procedures for requesting a change are not used because they are too limiting.

c.) **Documentation of when the change was made, how it was made, and why it was made must be created.**

d.) Changes do not produce unexpected results because the DBA will have investigated the change thoroughly before implementing it.

e.) If the database is properly designed, changes should not be necessary throughout the system's lifetime.

Level: *hard*

Page: *301*

33. The task of diagnosing errors due to changes in the database structure is eased by:

a.) formal policies for requesting changes.

b.) **database structure change documentation.**

c.) rollback analysis.

d.) configuration control.

e.) None of the above.

Level: *moderate*

Page: *301-302*

34. Measures that are taken to prevent one user's work from inappropriately influencing another user's work are called:

a.) **concurrency control.**

b.) checkpoint.

c.) database recovery.

d.) database logging.

e.) interleaving.

Level: *easy*

Page: *302*

35. A series of actions to be taken on the database such that either all actions are completed successfully, or none of them can be completed, is known as a(n):

a.) checkpoint.

b.) log.

c.) lock.

d.) **transaction.**

e.) concurrent.

Level: *easy*

Page: *303*

36. When two transactions are being processed against the database at the same time,

a.) they are called concurrent transactions.

b.) they are usually interleaved.

c.) they always result in a lost update problem.

d.) one must be rolled back.

e.) **both a and b**

Level: *easy*

Page: *303*

37. The situation that occurs when one user's changes to the database are lost by a second user's changes to the database is known as the:

a.) **concurrent update problem.**

b.) deadly embrace problem.

c.) inconsistent read problem.

d.) inconsistent write problem.

e.) deadlock problem.

Level: *hard*

Page: *303-306*

38. One remedy for the inconsistencies caused by concurrent processing is _____.

a.) lost updates

b.) checkpointing

c.) rollback

d.) **resource locking**

e.) concurrency

Level: *easy*

Page: *306*

39. A lock placed automatically by the DBMS is called a(n) _____ lock.

a.) exclusive

b.) explicit

c.) granular

d.) **implicit**

e.) shared

Level: *moderate*

Page: *307*

40. Which of the following is <u>not</u> true about locks?

a.) Locks with large granularity are easier for the DBMS to administer.

b.) **Locks with small granularity cause more conflicts.**

c.) Locks with large granularity produce fewer details for the DBMS to track.

d.) Locks may have a table-level granularity.

e.) Locks may have a database-level granularity.

Level: *hard*

Page: *307*

41. Which type of lock prevents all types of access to the locked resource?

a.) **exclusive lock**

b.) shared lock

c.) two-phased lock

d.) explicit lock

e.) implicit lock

Level: *easy*

Page: *307*

42. Which type of lock still allows other transactions to have read-only access to the locked resource?

a.) exclusive lock

b.) **shared lock**

c.) two-phased lock

d.) explicit lock

e.) implicit lock

Level: *easy*

Page: *307*

43. Which of the following is not true about two-phased locking?

a.) can make transactions serializable

b.) **uses only shared locks**

c.) has a growing phase

d.) has a shrinking phase

e.) cannot obtain a new lock once a lock has been released

Level: *hard*

Page: *307-308*

44. The situation that occurs when two users are each waiting for a resource that the other person has locked is known as a(n):

a.) lost update problem.

b.) **deadlock.**

c.) inconsistent read problem.

d.) inconsistent write problem.

e.) checkpoint.

Level: *moderate*

Page: *308*

45. Requiring all application programs to lock resources in the same order is a technique for preventing what problem?

a.) concurrent update

b.) lost update

c.) **deadlock**

d.) exclusive locks

e.) growing phase locking

Level: *hard*

Page: *308*

46. Locks that are placed assuming that a conflict will occur are called:

a.) dynamic locks.

b.) explicit locks.

c.) implicit locks.

d.) optimistic locks.

e.) **pessimistic locks.**

Level: *moderate*

Page: *309*

47. Locks that are placed assuming that a conflict will not occur are called:

a.) dynamic.

b.) explicit.

c.) implicit.

d.) **optimistic.**

e.) pessimistic.

Level: *moderate*

Page: *309*

48. Ensuring that all rows impacted by the actions of a transaction are protected from changes until the entire transaction is completed is called:

a.) statement level consistency.

b.) optimistic locking.

c.) **transaction level consistency.**

d.) durable transactions.

e.) ARID transactions.

Level: *hard*

Page: *312*

49. Which of the following is allowed by "Repeatable Read Isolation?"

a.) nonrepeatable reads

b.) dirty reads

c.) **phantom reads**

d.) a and b

e.) a, b and c

Level: *hard*

Page: *312-313 [See Figure 9-10]*

50. Which of the following is true of forward only cursors?

a.) Current values for each row are retrieved when the application accesses a row.

b.) All changes of any type from any source are visible.

c.) **Changes made by the transaction are visible only if they occur on rows ahead of the cursor.**

d.) Applications may scroll backward in the record set.

e.) It requires the greatest overhead of any cursor type.

Level: *moderate*

Page: *313-314 [See Figure 9-11]*

51. Which of the following cannot be enforced in the DBMS or application programs?

a.) processing rights

b.) security

c.) **processing responsibilities**

d.) cursors

e.) transaction isolation

Level: *moderate*

Page: *315*

52. Once processing rights have been defined, they may be implemented at any of these levels except:

a.) network.

b.) operating system.

c.) **data.**

d.) DBMS.

e.) application.

Level: *hard*

Page: *315*

53. Which of the following is not true of DBMS security features?

a.) Users may be assigned to one or more roles.

b.) **A role may be assigned to only one user.**

c.) Both users and roles can have many permissions.

d.) Objects have many permissions.

e.) Each permission pertains to one user or role and one object.

Level: *moderate*

Page: *316*

54. Recovering a database via reprocessing involves:

a.) **restoring the database from the save and reprocessing all the transactions since the save.**

b.) restoring the database from the save and reapplying all the changes made by transactions since the save.

c.) undoing the changes made by erroneous or partially processed transactions, and restarting the valid transactions that were in process at the time of the failure.

d.) recreating the database by reentering all of the data from the beginning, and then reprocessing all of the transactions.

e.) synchronizing the database and the transaction log by checkpointing.

Level: *easy*

Page: *320*

55. Recovering a database via rollforward involves:

a.) restoring the database from the save and reprocessing all the transactions since the save.

b.) **restoring the database from the save and reapplying all the changes made by transactions since the save.**

c.) undoing the changes made by erroneous or partially processed transactions, and restarting the valid transactions that were in process at the time of the failure.

d.) re-creating the database by re-entering all of the data from the beginning, and then reprocessing all of the transactions.

e.) synchronizing the database and the transaction log by checkpointing.

Level: *easy*

Page: *320*

56. Recovering a database via rollback involves:

a.) restoring the database from the save and reprocessing all the transactions since the save.

b.) restoring the database from the save and reapplying all the changes made by transactions since the save.

c.) **undoing the changes made by erroneous or partially processed transactions, and restarting the valid transactions that were in process at the time of the failure.**

d.) re-creating the database by re-entering all of the data from the beginning and, then reprocessing all of the transactions.

e.) synchronizing the database and the transaction log by checkpointing.

Level: *easy*

Page: *320*

57. Which of the following would <u>not</u> be contained in a transaction log?

a.) before-images

b.) type of operation

c.) pointers

d.) time of the action

e.) **permissions**

Level: *moderate*

Page: *320*

58. Which of the following would a DBA do in managing the DBMS?

a.) analyze system performance statistics

b.) investigate user complaints

c.) evaluate new DBMS product features

d.) tune DBMS product options to accommodate other software in use

e.) **All of the above.**

Level: *easy*

Page: *323*

59. Which of the following is <u>not</u> true of data repositories?

a.) **They are usually created after the database has been implemented and optimized for performance.**

b.) They may be virtual.

c.) They may contain metadata about database applications.

d.) They may contain metadata about users.

e.) They may contain metadata about web pages.

Level: *moderate*

Page: *324*

60. Which type of data repository is composed of metadata that is created automatically as the system components are created?

a.) passive

b.) dynamic

c.) **active**

d.) automatic

e.) summary

Level: *moderate*

Page: *324*

Fill in the Blank Questions

61. The overall responsibility of the *DBA* is to facilitate the development and use of the database.

Level: *easy*

Page: *301*

62. The database is most vulnerable to failure after a change to its *structure*.

Level: *hard*

Page: *301*

63. A(n) *transaction* is a series of actions to be taken on the database such that either all of them are performed successfully or none of them is performed at all. ***Level:*** *easy*

Page: *303*

64. A transaction is sometimes called *atomic*, since it is performed as a unit.

Level: *easy*

Page: *303*

65. Locks placed automatically by the DBMS are called *implicit* locks.

Level: *moderate*

Page: *307*

---

2) Two-phased locking has a growing phase and a shrinking phase.Answer:TRUEDiff:

3Page Ref: 40223) In two-phase locking, all locks are obtained during the growing phase.Answer:TRUEDiff: 2Page Ref: 4022

4) In two-phase locking, all locks are released during the contracting phase.Answer:FALSEDiff: 2Page Ref: 4022

5) Resource locking must be carefully planned because most DBMS products cannot detect adeadlock condition.Answer:FALSEDiff: 2Page Ref: 402-4032

6) Requiring all application programs to lock resources in the same order is one way ofpreventing a deadlock condition.Answer:TRUEDiff: 3Page Ref: 402-4032

7) With optimistic locking, the assumption is made that no conflict will occur.Answer:TRUEDiff: 1Page Ref: 4032

8) With pessimistic locking, the assumption is made that a conflict will occur.Answer:TRUEDiff: 1Page Ref: 4032

9) Resources are locked for a shorter amount of time with pessimistic locking because thetransaction is pre-processed.Answer:FALSEDiff: 1Page Ref: 403-4043

0) In general, optimistic locking is the preferred technique for Internet databases.Answer:TRUEDiff: 3Page Ref: 403-4043ScholarStock

31) The transaction boundaries are the essential information that the DBMS needs from theapplication programs to enforce different locking strategies.Answer:TRUEDiff: 3Page Ref: 404-405

32) A durable transaction is one for which all committed changes are permanent.Answer:TRUEDiff: 3Page Ref: 406

33) Transaction level consistency means that all rows impacted by any actions in a transaction are protected from change during the entire transaction.Answer:TRUEDiff: 2Page Ref: 406-407

34) Adirty readhappens when one transaction reads a changed record that has not beencommitted to the database.Answer:TRUEDiff: 3Page Ref: 407-408 Fig 9-11

35)Nonrepeatable readsoccur when a transaction rereads data it has previously read and findsmodifications or deletions caused by a committed transaction.Answer:TRUEDiff: 3Page Ref: 407408 Fig 9-11

36)Repeatable Readisolation is the most restrictive level of isolation.Answer:FALSEDiff: 2Page Ref: 408 Fig 9-12

37) According to ANSI SQL, the serializable isolation level will not allow phantom reads, dirty reads, and nonrepeatable reads.Answer:TRUEDiff: 3Page Ref: 408 Fig 9-12

38) A static cursor processes a snapshot of the relation that was taken when the cursor was opened.Answer:TRUEDiff: 2Page Ref: 408-410 Fig 9-13

39) A dynamic cursor saves primary key values when the cursor is opened and retrieves thevalues for each row as the application program accesses it.Answer:FALSEDiff: 2Page Ref: 408-410 Fig 9-13

40) The goal of database security is to ensure that only authorized users can perform authorizedactivities at authorized times.Answer:TRUEDiff: 1Page

66. Locks placed by a command issued to the DBMS from the application program

are called *explicit* locks.

Level: *moderate*

Page: *307*

67. The size of a lock is referred to as the lock *granularity* .

Level: *hard*

Page: *307*

68. A(n) *exclusive* lock locks the item from access of any type.

Level: *easy*

Page: *307*

69. A(n) *shared* lock locks the item from change but not from read access.

Level: *easy*

Page: *307*

70. Two-phased locking is a scheme for achieving *serializability* of transactions.

Level: *hard*

Page: *307*

71. In two-phase locking, all locks are obtained during the *growing* phase.

Level: *moderate*

Page: *308*

72. In two-phase locking, all locks are released during the *shrinking* phase.

Level: *moderate*

Page: *308*

73. Requiring all application programs to lock resources in the same order is one way of preventing a *deadlock* condition.

Level: *hard*

Page: *308*

74. With *optimistic* locking, the assumption is made that no conflict will occur.

Level: *easy*

Page: *309*

75. With *pessimistic* locking, the assumption is made that a conflict will occur.

Level: *easy*

Page: *309*

76. The transaction *boundaries* are the essential information that the DBMS needs from the application programs to enforce different locking strategies.

Level: *hard*

Page: *309-310*

77. A(n) *durable* transaction is one for which all committed changes are permanent.

Level: *hard*

Page: *311*

78. Transaction level *consistency* means that all rows impacted by any actions in a transaction are protected from change during the entire transaction.

Level: *moderate*

Page: *312*

79. A(n) *dirty read* occurs when one transaction reads a changed record that has not been committed to the database.

Level: *hard*

Page: *312*

80. *Nonrepeatable reads* occur when a transaction rereads data it has previously read and finds modifications or deletions caused by a committed transaction.

Level: *hard*

Page: *312*

81. According to ANSI SQL, the *serializable* isolation level will not allow phantom reads, dirty reads, and nonrepeatable reads.

Level: *hard*

Page: *312*

82. A(n) *static* cursor processes a snapshot of the relation that was taken when the cursor was opened.

Level: *moderate*

Page: *313*

83. A(n) *keyset* cursor saves primary key values when the cursor is opened and retrieves the values for each row as the application program accesses it.

Level: *moderate*

Page: *313*

84. *Rollforward* is a method of database recovery that restores the database save and all valid transactions since the save was reapplied.

Level: *easy*

Page: *320*

85. To support rollforward and rollback recovery, transactions must be written to a(n) *log* before they are applied to the database.

Level: *easy*

Page: *320*

86. Copies of each database record or page before it was changed by a transaction that are saved for use in database recovery are called *before-images*.

Level: *easy*

Page: *320*

87. Copies of each database record or page after it was changed by a transaction that are saved for use in database recovery are called *after-images*.

Level: *easy*

Page: *320*

88. A(n) *checkpoint* is a point of synchronization between the database and the transaction log.

Level: *moderate*

Page: *322*

89. A(n) active data repository is one in which the metadata is automatically created as the system components are created.

Level: *easy*

Page: *324*

90. A(n) passive data repository is one which requires a person to generate the metadata and place it in the repository.

Level: *easy* ***Page:***

324

---

42) Processing responsibilities should be documented and encoded into manual procedures. Answer: TRUE Diff: 1 Page Ref: 411

43) Processing rights may be implemented at the DBMS level. Answer: TRUE Diff: 1 Page Ref: 411

44) All commercial DBMS products use some version of "username and password" as part of their security features. Answer: TRUE Diff: 1 Page Ref: 412

45) The security provided by the DBMS often has to be augmented by additional security features within the application program. Answer: TRUE Diff: 2 Page Ref: 414

46) A "database save" is used to mark the end of a transaction. Answer: FALSE Diff: 2 Page Ref: 416

47) A checkpoint is a point of synchronization between the database and the transaction log. Answer: TRUE Diff: 2 Page Ref: 417

48) Copies of each database record or page before being changed by a transaction that are saved for use in database recovery are called before images. Answer: TRUE Diff: 1 Page Ref: 416

49) Copies of each database record or page after it was changed by a transaction that are saved for use in database recovery are called ghost images. Answer: FALSE Diff: 1 Page Ref: 416

50) Reprocessing is normally the most convenient method for recovery after a system failure. Answer: FALSE Diff: 1 Page Ref: 416 5 ScholarStock

Background image of page 5

51) Rollforward is a method of database recovery that restores the database save and all valid transactions since the save was reapplied. Answer: TRUE Diff: 1 Page Ref: 416

52) Rollforward and reprocessing are two different names for the same technique. Answer: FALSE Diff: 3 Page Ref: 416

53) To support rollforward and rollback recovery, transactions must be written to a script before they are applied to the database. Answer: FALSE Diff: 1 Page Ref: 416

54) Both rollforward and rollback require the use of a log of transaction results. Answer: TRUE Diff: 1 Page Ref: 416

55) The DBA should periodically analyze run-time statistics of database performance to help manage the DBMS. Answer: TRUE Diff: 1 Page Ref: 418

56) An active data repository is one in which the metadata is automatically created as the system components are created. Answer: TRUE Diff: 1 Page Ref: 419-420

57) A standby data repository is one which requires a person to generate the metadata and place it in the repository. Answer: FALSE Diff: 1 Page Ref: 419-420

58) A passive data repository is preferred over an active repository because it requires less human intervention. Answer: FALSE Diff: 2 Page Ref: 419-420

59) Which of the following is not a database administration responsibility of a DBA? A) Managing the database structure B) Managing data activity C) Managing the DBMS D) Maintaining the data repository E) All of the above are database administration responsibilities of a DBA. Answer: E Diff: 1 Page Ref: 396397 Fig 9-

---

- **Front**

- **Back**

The recovery technique in which the database is returned to a known state and then all valid transactions are reapplied to the database is known as:	rollforward.
Ensuring that each SQL statement independently processes consistent rows is known as:	statement level consistency.
An ACID transaction is one that is:	All of these (isolated, consistent, atomic, durable)
When distributed databases break the database into sections and store the sections on different servers, this is known as:	partitioning.
What concurrent processing problem occurs when a transaction reads a changed record that has not been committed to the database?	Dirty reads.
The purpose of concurrency control is to:	ensure that one user's work does not interfere with another's.
Rollback requires the use of logs, but rollforward does not.	False

What is the least restrictive isolation level that will prevent nonrepeatable reads?	Repeatable Read
With optimistic locking, locks are first issued, then the transaction is processed, and then the locks are released.	False
Preventing multiple applications from obtaining copies of the same record when the record is about to be changed is called:	resource locking.
When distributed databases create copies of the database on different servers, this is known as	replication.
A point of synchronization between the database and the transaction log is called a(n):	checkpoint
When using the Microsoft Access 2010 Security Options dialog box, the default setting in the dialog box enables some specific capabilities of Access 2010	False
What is the least restrictive isolation level that will prevent dirty reads?	Read Committed
Which of the following is not true of database recovery through reprocessing?	Reprocessing will always return the database to its exact previous state.

Which of the following is not a means of	All of these are means of processing a database. (Programs coded in C#, Internet applications

processing databases?	using ASP, Internet applications using JSP, Applications invoking stored procedures)
When a transaction in such a way that either all of the transaction actions are completed or none of them will be, the transaction is said to be:	atomic.
What is the least restrictive isolation level that will prevent phantom read problems?	Serializable.
When Microsoft Access 2010 compiles VBA code in a database, the database file will be a(n):	*.accde file.
The term DBA refers to the:	database adminstrator.
Which type of lock assumes that no conflicts will occur?	Optimistic locks
Transactions should be written to the log before they are applied to the database itself.	True.
What concurrent processing problem occurs when a transaction rereads data and finds new rows that were inserted by a different transaction since the prior read?	Phantom reads.
Which of the following is true about a shared lock?	It allows reads to the locked item.
Whether a lock applies to data at the record level, page level, table level, or database level is referred to as:	lock granularity.

59) SQL statements can be passed to SQL Server using the _____. A) Enterprise Manager B) Query Analyzer C) Microsoft SQL Server Management Studio D) A and B E) A, B, and C Answer: C Diff: 1 Page Ref: 10A-8

60) SQL Server tables can be created and modified using _____. A) Enterprise Manager B) Query Analyzer C) Microsoft SQL Server Management Studio D) A and B E) A, B, and C Answer: C Diff: 1 Page Ref: 10A-12

61) When creating SQL Server tables in a database, SQL statements can now also be submitted by using _____. A) Internet Explorer B) Microsoft PowerShell C) Microsoft Word D) Microsoft Excel E) PL/SQL Answer: B Diff: 2 Page Ref: 10A-7-10A-8 7 ScholarStock

62) Anytime you want to use an SQL Server reserved word as a user identifier, enclose it in _____. A) { } B) [ ] C) ( ) D) " " E) < > Answer: B Diff: 1 Page Ref: 10A-12

63) The column(s) that is/are the primary key in an SQL table are indicated by _____ when viewed graphically in Enterprise Manager. A) the black arrow symbol B) the outline arrow symbol C) the key symbol D) the asterisk symbol E) being underlined Answer: C Diff: 1 Page Ref: Fig 10-17

64) What name is so special to SQL Server that no stored procedure will work on a table with that name, not even if it is enclosed in brackets? A) NAME B) UPDATE C) TRANSACTION D) SQL E) KEY Answer: C Diff: 2 Page Ref: 10A-12

65) When reviewing a table in the SQL Server GUI tools, which property should be set to "Yes" to indicate that the column is a surrogate key for which SQL Server should automatically supply values? A) Identity B) Surrogate C) AutoIncrement D) AutoNumber E) Sequence Answer: A Diff: 3 Page Ref: Fig 1017 8 ScholarStock

66) Using only SQL Server tools, you can enter data into a table by _____. A) entering it into a table grid in the Microsoft SQL Server Management Studio B) using INSERT statements through the Microsoft SQL Server Management Studio C) using an SQL Server form D) A or B E) A, B, or C Answer: D Diff: 2 Page Ref: 10A-22

67) When creating a T-SQL statement, you should: A) predefine it graphically in Enterprise Manager. B) limit it to one table. C) limit it to two tables. D) end it with the semicolon normally used to terminate SQL statements. E) include the word "View" in the view name. Answer: D Diff: 2 Page Ref: 10A-35

68) Which of the following is true about indexes in SQL Server? A) SQL Server automatically creates indexes for columns appearing in WHERE clauses in queries. B) SQL Server pads all indexes. C) SQL Server

supports filling up to a limit of 70 percent. D) SQL Server automatically creates indexes for foreign keys. E) All of the above. Answer: D Diff: 2 Page Ref: 10A-20

69) Which type(s) of index(s) in SQL Server has/have a bottom level that does not contain data but has pointers to the data?

Nonclustered

SQL stands for Standard Query Language.

F

SQL includes a data definition language, a data manipulation language, and SQL/Persistent stored modules.

T

SQL is only a data manipulation language (DML).

F

The American National Standards Institute (ANSI) maintains the standards for SQL.

T

SQL was developed by IBM in the late 1970s.

T

SQL is not a complete programming language. Rather it is a data sublanguage.

T

In addition to being a data sublanguage, SQL is also a programming language, like Java or C#.

F

SQL commands can be embedded in application programs.

T

SQL, although very popular, has never become a national standard.

F

The SQL keyword SELECT is used to specify the columns to be listed in the query results.
Answer: TRUE

T

The SQL keyword WHERE is used to specify the table(s) that contain(s) the data to be retrieved.

F

The SQL keyword FROM is used to specify the table to be used.

T

SQL can only query a single table.

F

SQL statements end with a colon.

F

The columns to be obtained by an SQL command are listed after the FROM keyword.

F

To remove duplicate rows from the result of a query, specify the SQL DISTINCT keyword.

T

To obtain all columns, use an asterisk (*) wildcard character instead of listing all the column names.

T

The SQL WHERE clause contains the condition that specifies which rows are to be selected.

T

The result of an SQL SELECT operation can contain duplicate rows.

T

To have SQL automatically eliminate duplicate rows from a result, use the keyword DISTINCT with the FROM keyword.

F

An asterisk (*) following the SELECT verb means that all columns are to be displayed.

T

The WHERE clause contains the condition that specifies which columns are to be selected.

F

The rows of the result table can be sorted by the values in one or more columns.

T

Sorting is specified by the use of the SORT BY phrase.

F

To sort the rows of the result table, the ORDER BY clause is specified.

F

□□□

Columns can be sorted in descending sequence by using the SQL DESC keyword.

T
□□□

A WHERE clause can contain only one condition.

F
□□□

When two conditions must both be true for the rows to be selected, the conditions are separated by the SQL AND keyword.

T
□□□

To refer to a set of values needed for a condition, use the SQL IN operator.

T
□□□

To exclude one or more values using a condition, the SQL OUT keyword must be used.

F
□□□
To refer to a set of values in a condition, the values are placed inside parentheses ( ) and separated by commas.

T
□□□

The condition in WHERE clauses can refer to a set of values by using the IN operator.

T
□□□

The SQL keyword LIKE is used in SQL expressions to select on partial values.

T
□□□

The SQL wildcard character "%" represents a series of one or more unspecified characters.

T

□□□

The SQL wildcard character "#" indicates a single, unspecified character in a specific location in an SQL query.

F
□□□

The Microsoft Access wildcard character "*" (asterisk) indicates a sequence of one or more unspecified characters in a Microsoft Access SQL query.

T
□□□

The Microsoft Access wildcard character "_" (underscore) indicates a single, unspecified character in a specific location in a Microsoft Access SQL query.

F
□□□

The SQL built-in function ADDUP totals values in numeric columns.

F

□□□

The SQL built-in function AVG computes the average of values in numeric columns.

T
□□□

The SQL built-in function MOST obtains the largest value in a numeric column.

T
□□□

The SQL built-in function MIN obtains the smallest value in a numeric column.

T
□□□

The SQL built-in function COUNT computes the number of rows in a query.

T
□□□

The built-in function SUM can be used with any column.

F

The clause SELECT COUNT (*) results in a table with a single row and a single column.

T

Arithmetic in SQL statements is limited to the operations provided by the built-in functions.

F

The SQL keyword GROUP BY instructs the DBMS to group together those rows that have the same value in a column.

T

A WHERE clause can contain another SELECT statement enclosed in parentheses.

T

A SELECT statement used in a WHERE clause is called a subquery.

T

A nested SELECT statement (one that appears within the WHERE clause of another SQL statement) is called a subquery and must be enclosed in parentheses.

T

Only two tables can be queried by using a subquery.

F

An alternative to combining tables by a subquery is to use a join.

T

When people use the term "join" they normally mean an "equijoin." T

Two or more tables are joined by giving the table names in the WHERE clause and specifying the equality of the respective column names as a condition in the GROUP BY clause.

F

The names of tables to be joined in an SQL query are listed in the FROM clause.

T

In an SQL query, a join operation is achieved by specifying the equality of the respective column names as a condition in the WHERE clause.

T

Every subquery can be alternatively expressed by a join.

F

While many subqueries can be alternatively written as joins, correlated subqueries do work that cannot be duplicated as a join. T

The SQL syntax JOIN . . . ON can be used as alternate way of writing an SQL join statement.

T

Joins that show only matching rows from the joined tables in their results are called inner joins.

T

Joins that show the matching rows from the joined tables plus unmatched rows from one other table in their results are called outer joins.

T

Outer joins can be either up joins or down joins.

F

???

SQL is a ________.

A) data sublanguage

B) product of IBM research

C) national standard

D) combination that includes a data definition language and a data manipulation language

E) All of the above

E

???

When making an SQL query, we are using SQL as a(n) ________.

A) DDL

B) DML

C) embedded language

D) SET

E) WHERE

B

???

In an SQL query, which SQL keyword actually creates the query?

A) EXISTS

B) FROM

C) SELECT

D) SET

E) WHERE

C

???

In an SQL query, which SQL keyword is used to specify the table(s) to be used?

A) EXISTS

B) FROM

C) SELECT

D) SET

E) WHERE B

???

In an SQL query, which SQL keyword must be used to remove duplicate rows from the result table?
A) DELETE
B) DISTINCT
C) NOT EXISTS
D) UNIQUE
E) KEY

B
□□□

In an SQL query, which of the following symbols is used by ANSI SQL to represent all the columns in a single table? A) _ (underscore)
B) ? (question mark)
C) * (asterisk)
D) % (percent)
E) # (pound)

C
□□□

In an SQL query, which SQL keyword is used to state the condition that specifies which rows are to be selected? A) EXISTS
B) FROM
C) SELECT
D) SET
E) WHERE

E
□□□

In an SQL query, which SQL keyword is used to join two conditions that both must be true for the rows to be selected?
A) AND
B) EXISTS
C) HAVING
D) IN
E) OR

A
□□□

In an SQL query, which SQL keyword is used to determine if a column value is equal to any one of a set of values? A) AND
B) EXISTS

C) HAVING

D) IN
E) OR

D
□□□

In an SQL query, which of the following symbols is used by ANSI SQL to represent a single unspecified character? A) _ (underscore)
B) ? (question mark)
C) * (asterisk)
D) % (percent)
E) # (pound)

A
□□□

In an SQL query, which of the following symbols is used by Microsoft Access to represent a single unspecified character? A) _ (underscore)
B) ? (question mark)
C) * (asterisk)
D) % (percent)
E) # (pound)

B
□□□

In an SQL query, which SQL keyword is used to sort the result table by the values in one or more columns? A) GROUP BY
B) ORDER BY
C) SELECT
D) SORT BY
E) WHERE

B
□□□

Given a table with the structure: EMPLOYEE (EmpNo, Name, Salary, HireDate), which of the following would find all employees whose name begins with the letter "S" using standard SQL?
A) SELECT *

FROM EMPLOYEE
WHERE Name IN ['S'];

B) SELECT EmpNo FROM EMPLOYEE
WHERE Name LIKE 'S';
C) SELECT *
FROM Name
WHERE EMPLOYEE LIKE 'S*';
D) SELECT *
FROM EMPLOYEE WHERE
Name LIKE 'S%'; E) None
of the above.

D
□□□

Given a table with the structure: EMPLOYEE (EmpNo, Name, Salary, HireDate), which of the following would find all employees whose name begins with the letter "S" using Microsoft Access?
A) SELECT *
FROM EMPLOYEE
WHERE Name IN ['S'];
B) SELECT EmpNo FROM EMPLOYEE
WHERE Name LIKE 'S';
C) SELECT *
FROM Name
WHERE EMPLOYEE LIKE 'S*';
D) SELECT *
FROM EMPLOYEE WHERE
Name LIKE 'S%'; E) None
of the above.

C
□□□

In an SQL query, which built-in function is used to total numeric columns?
A) AVG
B) COUNT
C) MAX
D) MEAN
E) SUM

E
□□□

In an SQL query, which built-in function is used to compute the average value of numeric columns?
A) AVG

B) MEAN
C) MAX
D) MIN
E) SUM A

□□□

In an SQL query, which built-in function is used to obtain the largest value of numeric columns?
A) AVG
B) COUNT
C) MAX
D) MIN
E) SUM C

□□□

In an SQL query, which built-in function is used to obtain the smallest value of numeric columns?
A) AVG
B) COUNT
C) MAX
D) MIN
E) SUM

D
□□□

In an SQL query, the built-in functions SUM and AVG work with columns containing data of which of the following data types?
A) Integer
B) Numeric
C) Char
D) A and B
E) A, B, and C

D

□□□

In an SQL query, which built-in function is used to compute the number of rows in a table?
A) AVG
B) COUNT
C) MAX
D) MIN
E) MEAN

B

In an SQL query, the built-in function COUNT works with columns containing data of which of the following data types?
A) Integer
B) Numeric
C) Char
D) A and B
E) A, B, and C

E

In an SQL query, which SQL keyword is used with built-in functions to group together rows that have the same value in a specified column?
A) GROUP BY
B) ORDER BY
C) SELECT
D) SORT BY
E) DISTINCT SET A

In an SQL query, which SQL keyword is used with GROUP BY to select groups meeting specified criteria?
A) AND
B) EXISTS
C) HAVING
D) IN
E) WHERE

C

Given a table with the structure: EMPLOYEE (EmpNo, Name, Salary, HireDate), which of the following is not a valid ANSI SQL command?
A) SELECT *
FROM EMPLOYEE
WHERE Name LIKE 'Ja%';
B) SELECT COUNT(*)
FROM EMPLOYEE
WHERE Salary < 30000;
C) SELECT COUNT(EmpNo)

FROM EMPLOYEE;
D) SELECT HireDate, COUNT(*)
FROM EMPLOYEE
WHERE Salary < 30000;
E) SELECT HireDate, COUNT(*)
FROM EMPLOYEE
GROUP BY HireDate;

D
???

Based on the tables below, which of the following ANSI SQL commands would return the average customer balance grouped by SalesRepNo?

GENERAL SALES DATABASE:

SALESREP
SalesRepNo RepName HireDate
654 Jones 01/02/2005
734 Smith 02/03/2007
345 Chen 01/25/2004
434 Johnson 11/23/2004

CUSTOMER
CustNo CustName Balance SalesRepNo
9870 Winston 500 345
8590 Gonzales 350 434
7840 Harris 800 654
4870 Miles 100 345

A) SELECT AVG (Balance)

FROM CUSTOMER
WHERE SalesRepNo;
B) SELECT AVG (Balance)
FROM CUSTOMER
GROUP BY SalesRepNo;
C) SELECT AVG (Balance)
FROM CUSTOMER, SALESREP
WHERE SALESREP.SalesRepNo = CUSTOMER.SalesRepNo;
D) SELECT AVG (Balance)
FROM CUSTOMER

ORDER BY SalesRepNo;
E) SELECT AVG (Balance)
FROM CUSTOMER, SALESREP
WHERE CUSTOMER.SalesRepNo = CUSTOMER.SalesRepNo
HAVING SalesRepNo;

B
???

Based on the tables below, which of the following commands in ANSI SQL would return only the name of the sales representative and the name of the customer for each customer that has a balance greater than 400?

GENERAL SALES DATABASE:

SALESREP
SalesRepNo RepName HireDate
654 Jones 01/02/2005
734 Smith 02/03/2007
345 Chen 01/25/2004
434 Johnson 11/23/2004

CUSTOMER
CustNo CustName Balance SalesRepNo
9870 Winston 500 345
8590 Gonzales 350 434
7840 Harris 800 654
4870 Miles 100 345

A) SELECT *
FROM SALESREP, CUSTOMER
WHERE Balance > 400;

B) SELECT DISTINCT RepName, CustName
FROM SALESREP, CUSTOMER
WHERE Balance > 400;
C) SELECT *
FROM SALESREP, CUSTOMER
WHERE SALESREP.SalesRepNo = CUSTOMER.SalesRepNo
AND Balance > 400;
D) SELECT RepName, CustName
FROM SALESREP, CUSTOMER

WHERE SALESREP.SalesRepNo = CUSTOMER.SalesRepNo
AND Balance > 400;
E) SELECT RepName, CustName
FROM SALESREP, CUSTOMER
WHERE Balance > 400
GROUP BY SalesRepNo;

D
□□□

In an SQL query, which SQL keyword is used to implement a subquery?
A) GROUP BY
B) HAVING
C) ORDER BY
D) SELECT
E) SORT BY

D
□□□

When one SQL query is embedded in another SQL query, this is referred to as a _____.
A) subset
B) join
C) WHERE Query
D) subquery
E) set query

D
□□□

When one SQL query is embedded in another SQL query, the second SQL query is embedded in the _____ of the first query.
A) SELECT

B) FROM
C) WHERE
D) GROUP BY
E) HAVING

C
□□□

When one SQL query is embedded in another SQL query, the first SQL query can still contain an SQL _____ clause.

A) FROM THE
B) WHERE THE
C) ORDER BY
D) GROUP BY
E) C and D E

□□□

In an SQL query, which SQL keyword is used to specify the names of tables to be joined?
A) FROM
B) HAVING
C) JOIN
D) SELECT E) WHERE A

□□□

In an SQL query, which SQL keyword is used to specify the column names to be used in a join?
A) FROM
B) HAVING
C) JOIN
D) SELECT
E) WHERE

E
□□□

Regarding the interchangeability of subqueries and joins,
A) a join can always be used as an alternative to a subquery, and a subquery can always be used as an alternative to a join.

B) a join can sometimes be used as an alternative to a subquery, and a subquery can sometimes be used as an alternative to a join.
C) a join can always be used as an alternative to a subquery, and a subquery can sometimes be used as an alternative to a join.
D) a join can sometimes be used as an alternative to a subquery, and a subquery can always be used as an alternative to a join.
E) a join can never be used as an alternative to a subquery, and a subquery can never be used as an alternative to a join.

B
□□□

In an SQL query of two tables, which SQL keyword indicates that we want data from all the rows of one table to be included in the result, even if the row does not correspond to any data in the other table? A) LEFT JOIN
B) RIGHT JOIN
C) INCLUDE
D) A and B
E) A, B, and C

D
□□□

What is SQL?

Structured Query Language (SQL) is used to create and use databases, tables and relationships. SQL is divided into two categories: SQL statements for database definition and SQL statements for database processing (querying and updating). The database definition commands are referred to as a data definition language (DDL), and the database query and update commands are referred to as a data manipulation language (DML). SQL was developed by IBM, and is endorsed as a national standard by the American National Standards Institute (ANSI). Although a newer standard, SQL3, exists, the most widely implemented version of SQL is the ANSI SQL-92 standard. SQL is not a full-featured programming language, but rather it is considered to be a data sublanguage.
□□□

Explain why it is important to learn SQL.

Most modern DBMS products support SQL as a standardized data language. These products usually provide graphical tools to perform the tasks associated with SQL, but there are some tasks that cannot be performed using these graphical tools. SQL is text-oriented, and SQL code must be written in order to embed SQL commands within program applications.

□□□

Briefly describe subqueries and joins. Explain when each is not an acceptable alternative for the other.

Subqueries and joins are both methods for retrieving data from multiple tables. Subqueries involve nesting one SELECT statement within another. The nested SELECT is used as part of a condition in the WHERE clause of the first SELECT statement. The nested SELECT statement can return a set of records from one table, which are then used in a logical operator within the parent SELECT query. A join combines records from each table into concatenated records containing the fields of both tables. The records are concatenated based on matching values in similar columns in the two tables. Subqueries cannot be used in situations where the results to be displayed include attributes from more than one table. Joins cannot be used as an alternative to a correlated subquery.
□□□

The following database will be used in this question:

GENERAL SALES DATABASE:

SALESREP
SalesRepNo RepName HireDate
654 Jones 01/02/2005
734 Smith 02/03/2007
345 Chen 01/25/2004
434 Johnson 11/23/2004

CUSTOMER
CustNo CustName Balance SalesRepNo
9870 Winston 500 345
8590 Gonzales 350 434
7840 Harris 800 654
4870 Miles 100 345

Explain the use of the of SQL keyword SELECT. Include an example based on the CUSTOMER table in the General Sales database.

The SQL keyword SELECT forms the basis for all SQL querying. Although SELECT technically specifies columns to be used in a query, it is always the first SQL keyword or command in an SQL phrase that includes at least the SQL keyword FROM and often the SQL keyword WHERE. The SQL keyword WHERE is used to specify which tables are used in the query, while the SQL keyword WHERE is used to specify conditions or constraints imposed on the query. For example,

```
SELECT *
FROM CUSTOMER
WHERE CustNo = 34567;
```

is a query that returns all information in the table CUSTOMER about the customer whose customer number is 34567.
□□□

The following database will be used in this question:

GENERAL SALES DATABASE:

SALESREP
SalesRepNo RepName HireDate
654 Jones 01/02/2005

734 Smith 02/03/2007
345 Chen 01/25/2004
434 Johnson 11/23/2004

CUSTOMER
CustNo CustName Balance SalesRepNo
9870 Winston 500 345
8590 Gonzales 350 434
7840 Harris 800 654
4870 Miles 100 345

Explain the use of the SQL keyword LIKE. Include an example based on the CUSTOMER table from the General Sales database.

The LIKE keyword is used in the WHERE clause of an SQL query to select rows based on partial values. Through the use of wildcard characters that can represent one or more unspecified characters, the LIKE operator can search for a given string of characters within a column value. The ANSI wildcard "%" represents one or more unspecified characters, and the wildcard " " (underscore) represents any single unspecified character. For example,

```
SELECT *
FROM CUSTOMER
WHERE CustName LIKE 'H%';
```

is a query that returns all information in the table CUSTOMER about customers whose CustName starts with the letter H.