

KillTest

Higher Quality, Better Service!



Q&A

<http://www.killtest.com>

We offer free update service for one year.

Exam : **70-431**

Title : Microsoft SQL Server 2005
Implementation &
Maintenance

Version : Demo

1. You are preparing for a new installation of SQL Server 2005. You need to select the protocols that client computers might use to connect to the server. Which two protocols can you use to achieve this goal?

(Each correct answer presents a complete solution. Choose two.)

- A. Named Pipes
- B. TCP/IP
- C. Shared Memory
- D. Virtual Interface Adapter (VIA)
- E. Multiprotocol

Answer: AB

2. Application developers in your company create an assembly that contains a CLR function. This CLR function reads data from a spreadsheet, performs some calculations, and returns the data to a SQL Server 2005 computer. You need to register the assembly with SQL Server 2005 by using the CREATE ASSEMBLY statement and the least privileged security permission set. Which permission set should you use?

- A. Default
- B. SAFE
- C. EXTERNAL_ACCESS
- D. UNSAFE

Answer: C

3. Your company has two SQL Server 2005 computers named SQL1 and SQL2. A user in your company named Eric writes many ad hoc queries against the company databases. Eric has access to the Customer database on SQL1. He does not have access to the Sales database on SQL2. You need to ensure that Eric can write queries that join information from both servers. What should you do first?

- A. Create a linked server on SQL1 to SQL2. Configure the linked server to use impersonation.
- B. Create a linked server on SQL1 to SQL2. Configure the linked server to use mapped logins.
- C. Instruct Eric to write the queries on SQL2 by using the OPENQUERY statement and specifying SQL1 as the server name.
- D. Instruct Eric to specify SQL Server object names by using four-part notation.

Answer: B

4. You configure a new SQL Server 2005 computer to use TCP/IP with all default settings. Your corporate policy requires that each server use a firewall. You find that you can connect to the SQL Server instance from the local computer. However, client computers cannot connect to the SQL Server instance. You need to identify the most likely cause of the connection issues. What should you do first?

- A. Ensure that port 1433 is open in your firewall.
- B. Ensure that port 443 is open in your firewall.
- C. Ensure that client computers connect by using Shared Memory protocol.
- D. Ensure that the server is not paused.

Answer: A

5. You are moving a SQL Server 2005 database named Timesheet to a new database server. The Timesheet database is used by the TimeEntry application. You need to minimize the amount of time that the TimeEntry application is unavailable while you move the database. What should you do?

- A. Move the database to the new server by using the SQL Management Object method in the Copy Database Wizard.
- B. Detach the current database. Copy the data files to the new server. Attach the files.
- C. Back up the database. Copy the backup file to the new server. Restore the database.
- D. Move the data files. Specify the new location by using ALTER DATABASE.

Answer: B

6. A support engineer reports that inserting new sales transactions in a SQL Server 2005 database results in an error. You investigate the error. You discover that in one of the databases, a developer has accidentally deleted some data in a table that is critical for transaction processing. The database uses the full recovery model. You need to restore the table. You need to achieve this goal without affecting the availability of other data in the database. What should you do?

- A. Back up the current transaction log. Restore the database with a different name and stop at the point just before the data loss. Copy the table back into the original database.
- B. Back up the current transaction log. Restore the database to the point just before the data loss.

- C. Restore the database from the existing backup files to a time just before the data loss.
- D. Restore the database to the point of the last full backup.

Answer: A

7. Your company uses SQL Server 2005. Users report that report execution is slow. You investigate and discover that some queries do not use optimal execution plans. You also notice that some optimizer statistics are missing and others are out of date. You need to correct the problem so that reports execute more quickly. Which two Transact-SQL statements should you use? (Each correct answer presents part of the solution. Choose two.)

- A. DBCC CHECKTABLE
- B. ALTER INDEX REORGANIZE
- C. UPDATE STATISTICS
- D. CREATE STATISTICS
- E. DBCC SHOW_STATISTICS
- F. DBCC UPDATEUSAGE

Answer: CD

8. You are responsible for implementing maintenance jobs on a SQL Server 2005 database server. Certain jobs run every Sunday and other jobs run at the beginning of every month. You need to schedule the jobs in the way that uses the least amount of administrative effort. What should you do?

- A. Create a job schedule that runs every Sunday. Assign weekly tasks to this schedule. Create a second schedule that runs on the first day of every month. Assign monthly tasks to this schedule.
- B. Create a job for each task that runs once a day. Use a Transact-SQL statement to check the date and day of the week. If the day is either a Sunday or the first day of the month, execute the code.
- C. Create a job schedule that runs once a day. Assign jobs to this job schedule. If the day is either a Sunday or the first day of the month, execute the jobs.
- D. Create a job for each task that runs once a week on Sunday. Add a second job schedule that runs the job on the first of the month.

Answer: A

9. You are responsible for managing the database schema for an order entry application in a SQL Server 2005 database. After a design review, a member of the development staff asks you to add a new column named `CommissionRate` to the `Product` table. Because the actual commission rates for all products are not known at this time, each product has a default commission rate of 12 percent. The rate can be modified as necessary. You add the new column. You need to configure the table to assign the default value as efficiently as possible. What should you do?

- A. Create an `INSERT` trigger to assign the default value to each item in the table.
- B. Create a `CHECK` constraint to validate the data and to assign the default value to each item in the table.
- C. Create an `UPDATE` trigger to update the default value for each new item in the table.
- D. Create a `DEFAULT` constraint to assign the default value specifying the `WITH VALUES` argument.

Answer: D

10. Your company Web site includes a page that customers use to send feedback about the company and its products. You use a SQL Server 2005 database to store the comments in the `Comments` column of a table named `Feedback`. You need to implement full-text searching so that you can run reports on the comments. Which two actions should you perform? (Each correct answer presents part of the solution. Choose two.)

- A. Create a nonclustered index on the `Comments` column.
- B. Execute the `USE Master Transact-SQL` statement.
- C. Create a full-text catalog.
- D. Create a full-text index on the `Comments` column.

Answer: CD

11. You are creating a SQL Server 2005 database for a mortgage company. The database will support a new Web-based application that will handle up to 1,000 simultaneous users. This application must quickly display the results of calculation-intensive operations, such as calculating mortgage payments and amortization schedules. You need to ensure that the database processes calculations as quickly and efficiently as possible. What should you do?

- A. Implement parameterized Transact-SQL queries in the application.

- B. Implement Transact-SQL stored procedures in the database.
- C. Implement CLR stored procedures in the database.
- D. Implement distributed Web services.

Answer: C

12. Your company uses a SQL Server 2005 database. This database contains a trigger named `trg_InsertOrders`, which fires when order data is inserted into the `Orders` table. The trigger is responsible for ensuring that a customer exists in the `Customers` table before data is inserted into the `Orders` table. You need to configure the trigger to prevent it from firing during the data import process. You must accomplish this goal while using the least amount of administrative effort. Which two Transact-SQL statements can you use to achieve this goal? (Each correct answer presents a complete solution. Choose two.)

- A. `ALTER TABLE OrdersDISABLE TRIGGER trg_InsertOrders`
- B. `DROP TRIGGER trg_InsertOrders`
- C. `DISABLE TRIGGER trg_InsertOrders ON Orders`
- D. `ALTER TRIGGER trg_InsertOrders ON Orders NOT FOR REPLICATION`
- E. `sp_settriggerorder@triggername= 'trg_InsertOrders', @order='None'`

Answer: AC

13. You work with a SQL Server 2005 database that provides banking information for customers. You want customers to see banking reports that combine data that is retrieved from the database with real-time investment information that comes from a Web service provided by a third party. The investment information must be current when the reports are executed. You need to create the appropriate objects that support the reports. What should you do?

- A. Publish the data in the database as an XML Web service by using the `FOR XML AUTO` clause.
- B. Create a table to store the banking information for each customer. Create a trigger that fires when data is inserted into the table that joins with the data coming from the Web service.
- C. Create a Transact-SQL stored procedure that uses a temporary table to store the banking information for each customer. Update the table with the values from the Web service.
- D. Have a developer in your company create an assembly that calls the remote Web service. Create a

CLR function by using the assembly. Call the CLR function and combine the results with banking information in the database.

Answer: D

14. Your company uses SQL Server 2005. You are implementing a series of views that are used in ad hoc queries. The views are used to enforce your companys security policy of abstracting data. Some of these views perform slowly. You create indexes on those views to increase performance, while still maintaining the companys security policy. One of the views returns the current date as one of the columns. The view returns the current date by using the GETDATE() function. This view does not allow you to create an index. You need to create an index on the view. Which two actions should you perform? (Each correct answer presents part of the solution. Choose two.)

- A. Remove all deterministic function calls from within the view.
- B. Remove all nondeterministic function calls from within the view.
- C. Schema-bind all functions that are called from within the view.
- D. Create the view and specify the WITH CHECK OPTION clause.

Answer: BC

15. Your data warehouse includes three separate SQL Server 2005 tables for storing customer data. Each table is used to horizontally partition data that is migrated from a SQL Server 2000 installation to SQL Server 2005. Data is partitioned as shown in the following table.

Table name	Stores customer surname beginning with letter
Customer1	A - H
Customer2	I - P
Customer3	Q - Z

Table nameStores customer surnamebeginning with letterCustomer1A C HCustomer2I C PCustomer3Q C Z The database does not allow customers to opt-out of e-mail marketing offers.

You need to add a new OptOutFlag column to each of the customer tables. You must ensure that this column is added to either all three customer tables or none of them. You must also ensure that the database is always in a consistent state.

Which Transact-SQL batch should you use?

A. BEGIN TRAN

ALTER TABLE CUSTOMER1

ADD OptOutFlag BIT NULL

ALTER TABLE CUSTOMER2

ADD OptOutFlag BIT NULL

ALTER TABLE CUSTOMER3

ADD OptOutFlag BIT NULL

IF @@error <> 0

ROLLBACK TRAN

ELSE

COMMIT TRAN

B. BEGIN TRAN

BEGIN TRY

ALTER TABLE CUSTOMER1

ADD OptOutFlag BIT NULL

ALTER TABLE CUSTOMER2

ADD OptOutFlag BIT NULL

ALTER TABLE CUSTOMER3

ADD OptOutFlag BIT NULL

END TRY

BEGIN CATCH

ROLLBACK TRAN

RETURN

END CATCH

COMMIT TRAN

C. ALTER TABLE CUSTOMER1

ADD OptOutFlag BIT NULL

ALTER TABLE CUSTOMER2

ADD OptOutFlag BIT NULL

ALTER TABLE CUSTOMER3

ADD OptOutFlag BIT NULL

D. BEGIN TRAN

ALTER TABLE CUSTOMER1

ADD OptOutFlag BIT NULL

ALTER TABLE CUSTOMER2

ADD OptOutFlag BIT NULL

ALTER TABLE CUSTOMER3

ADD OptOutFlag BIT NULL

COMMIT TRAN

Answer: B

16. You are creating a view to join the Customers and Orders tables in a SQL Server 2005 database. You need to ensure that the view cannot be affected by modifications to underlying table schemas. You want to accomplish this goal by using the least possible amount of overhead. What should you do?

A. Create CHECK constraints on the tables.

B. Create a DDL trigger to roll back any changes to the tables if the changes affect the columns in the view.

C. Create the view, specifying the WITH SCHEMABINDING option.

D. Create the view, specifying the WITH CHECK option.

Answer: C

17. Your company uses a SQL Server 2005 database. Customer data from your trading partners is imported into the Customers table every night. You need to ensure that the customer record is updated if it already exists in the Customers table. If the record does not exist, the data needs to be inserted into the Customers table. What should you do?

A. Create a FOR trigger.

B. Create an INSTEAD OF trigger.

C. Create an AFTER trigger.

D. Create a DDL trigger.

Answer: B

18. Your company uses a mission-critical database named DB1, which is located on a server named SQL1. You want to mirror DB1 on a second server named SQL2. Before you can establish database mirroring, you need to prepare DB1 and the servers for mirroring. Which three actions should you perform? (Each correct answer presents part of the solution. Choose three.)

- A. Set the recovery model of DB1 on SQL1 to bulk-logged recovery.
- B. Set the recovery model of DB1 on SQL1 to full recovery.
- C. Back up DB1 on SQL1. Restore the backup on SQL2; specify the NORECOVERY option.
- D. Back up DB1 on SQL1. Restore the backup on SQL2; specify the STANDBY option.
- E. Back up DB1 on SQL1. Restore the backup on SQL2; specify the RECOVERY option.
- F. Create endpoints on all participating servers.
- G. Set the AutoShrink property of DB1 on SQL1 to false.

Answer: BCF

19. A full backup of your database named DB1 is created automatically at midnight every day. Differential backups of DB1 occur twice each day at 10:00 and at 16:00. A database snapshot is created every day at noon. A developer reports that he accidentally dropped the Pricelist table in DB1 at 12:30. The last update to Pricelist occurred one week ago. You need to recover the Pricelist table. You want to achieve this goal by using the minimum amount of administrative effort. You must also minimize the amount of data that is lost. What should you do?

- A. Restore the most recent backup into a new database named DB1bak. Apply the most recent differential backup. Copy the Pricelist table from DB1bak to DB1.
- B. Delete all database snapshots except the most recent one. Restore DB1 from the most recent database snapshot.
- C. Recover DB1 from the most recent backup. Apply the most recent differential backup.
- D. Copy the Pricelist table from the most recent database snapshot into DB1.

Answer: D

20. Your company has two SQL Server 2005 computers named SQL1 and SQL2. Transaction log shipping occurs from SQL1 to SQL2 by using default SQL Server Agent schedule settings. You need to

reconfigure transaction log shipping to provide minimum latency on SQL2. What should you do?

A. On SQL1, reschedule the transaction log backup job so that it occurs every minute. On SQL2, maintain default schedule settings for both the log shipping copy and the restore jobs.

B. On SQL1, change the schedule type for the transaction log backup to Start automatically when SQL Server Agent starts. On SQL2, change the schedule types for both the log shipping copy and the restore jobs to Start automatically when SQL Server Agent starts.

C. On SQL1, maintain default schedule settings for the transaction log backup job. On SQL2, change the schedule types for both the log shipping copy and the restore jobs to Start automatically when SQL Server Agent starts.

D. On SQL1, reschedule the transaction log backup job so that it occurs every minute. On SQL2, reschedule both the log shipping copy and the restore jobs so that they occur every minute.

Answer: D