



Exam : 070-431

**Title : Microsoft SQL Server 2005 Implementation
& Maintenance**

Ver : 11.29.07

QUESTION 1:

You work as the database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. All servers on the Certkiller .com network run Windows 2000 Server. The Certkiller .com network contains a SQL Server 2000 database server named Certkiller -DB01 that runs on Windows 2000 Server. Certkiller -DB01 contains a database named CK_Sales.

Certkiller .com intends migrating Certkiller -DB01 to SQL Server 2005 in the near future. To prepare for the migration, you decide to install SQL Server 2005 Express on a server named Certkiller -SR31. You verify that Certkiller -SR31 exceeds the minimum recommended hardware requirements for SQL Server 2005. What should you do next?

- A. Upgrade Certkiller -SR31 to Windows Sever 2003 and install Service Pack 1.
- B. Install and run the SQL Server Upgrade Advisor on Certkiller -SR31.
- C. Install SQL Server 2000 on Certkiller -SR31.
- D. Download and install the .NET Framework 2.0.

Answer: D

Explanation: SQL Server 2005 requires the .NET Framework 2.0. This is automatically installed when SQL Server 2005 is installed but is not automatically installed when SQL Server 2005 Express is installed.

Incorrect Answers:

A: You do not need to upgrade Certkiller -SR31 to Windows Sever 2003 Service Pack 1 as SQL Server 2005 Express can be installed on Windows 2000 Server. You will need to install Service Pack 4 for Windows 2000 Server though.

B: You should run the SQL Server Upgrade Advisor when you want to upgrade an instance of a SQL Server 2000 database to SQL Server 2005.

C: You do not need to install SQL Server 2000 before installing SQL Server 2005 Express.

Reference:

Microsoft SQL Server 2005 Books Online (2006), Index: installing SQL Server Express, before installing

QUESTION 2:

You work as a database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. All servers on the Certkiller .com network run Windows Server 2003 with the latest Service Packs and Hotfixes, while the client computers run Windows XP Professional, Windows 98 Second Edition, and Novell NetWare 6.5.

You have received instruction to prepare for a new installation of SQL Server 2005 on a member server named Certkiller -SR21. Certkiller -SR21 will host a

database named CK_Sales. All Certkiller .com users must be able to connect to Certkiller -SR21 by computer name. A Certkiller .com network administrator named Rory Allen must ensure that the client computers will be able to connect to Certkiller -DB01.

Which protocols can Certkiller .com users use to Certkiller -SR21?

- A. TCP/IP and Named Pipes
- B. TCP/IP and IPX/SPX
- C. Virtual Interface Adapter (VIA)
- D. Shared Memory

Answer: A

Explanation: By default, clients have TCP and Named Pipes as available protocols. You can manipulate the protocol ordering by using the SQL Server Client utility. The client application uses the protocols in the order specified on the client computer. If you are using SQL Server 2005, the protocol order is stored in the ProtocolOrder registry entry under the following registry subkey:

HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\MSSQLServer\Client\SN19.0

Incorrect Answers:

B: SQL Server 2005 does not support IPX/SPX. The newer versions of NetWare, such as NetWare 6.5, support TCP/IP which makes it the common protocol for all clients.

C: The Virtual Interface Adapter (VIA) can be used only by VIA hardware.

D: Shared Memory can only be used on the local computer and cannot be used as a network protocol.

Reference:

Microsoft SQL Server 2005 Books Online (2006), Index: client connections [SQL Server], about client network connections, choosing a network protocol

QUESTION 3:

You work as the database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. All servers on the Certkiller .com network run Windows 2000 Server. The Certkiller .com network contains a SQL Server 2000 database server named Certkiller -DB01 that runs on Windows 2000 Server. Certkiller -DB01 contains a database named CK_Sales.

Certkiller .com intends migrating Certkiller -DB01 to SQL Server 2005 in the near future. To prepare for the migration, you decide to install SQL Server 2005 Express on a server named Certkiller -SR23. Certkiller -SR23 was used for previously used by application developer and contains Microsoft Visual Studio 2003, .NET Framework 1.2 and Internet Explorer 5.5. You verify that Certkiller -SR23 exceeds the minimum recommended hardware requirements for SQL Server 2005. What should you do next? (Each correct answer provides part of the solution. Choose THREE.)

- A. Upgrade Certkiller -SR23 to Windows Sever 2003 and install Service Pack 1 for Windows Sever 2003.
- B. Install Service Pack 4 for Windows 2000 Sever on Certkiller -SR23.
- C. Upgrade the .NET Framework 1.2 to .NET Framework 2.0 on Certkiller -SR23.
- D. Uninstall the .NET Framework 1.2 and install the .NET Framework 2.0 on Certkiller -SR23.
- E. Upgrade Internet Explorer 5.5 to Internet Explorer 6.0 and install Service Pack 1 for Internet Explorer 6.0 on Certkiller -SR23.
- F. Upgrade Microsoft Visual Studio 2003 to Microsoft Visual Studio 2005 and install Service Pack 1 for Microsoft Visual Studio 2005 on Certkiller -SR23.

Answer: B, D, E

Explanation: SQL Server 2005 requires Service Pack 4 for Windows 2000 Sever, Internet Explorer 6.0 with Service Pack 1 and the .NET Framework 2.0. The .NET Framework 2.0 is automatically installed with other versions of SQL Server 2005 but is not automatically installed with SQL Server 2005 Express. You can upgrade .NET Framework 1.0 and upgrade .NET Framework 1.1 to .NET Framework 2.0 but .NET Framework 1.2 must be uninstalled first.

Incorrect Answers:

A: You do not need to upgrade Certkiller -SR31 to Windows Sever 2003 Service Pack 1 as SQL Server 2005 Express can be installed on a Windows 2000 Server once you install Service Pack 4 for Windows 2000 Server.

C: You can upgrade .NET Framework 1.0 and upgrade .NET Framework 1.1 to .NET Framework 2.0 but .NET Framework 1.2 must be uninstalled first.

F: SQL Server 2005 Express does not require any version of Microsoft Visual Studio.

Reference:

Microsoft SQL Server 2005 Books Online (2006), Index: installing SQL Server Express, before installing

QUESTION 4:

You work as the database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. All servers on the Certkiller .com network run Windows 2000 Server. The Certkiller .com network contains a SQL Server 2000 database server named Certkiller -DB01 that runs on Windows 2000 Server Service Pack 4. Certkiller -DB01 contains a database named CK_Trading.

Certkiller .com wants to migrate Certkiller -DB01 to SQL Server 2005 but wants to retain the CK_Trading database with the same instance name. You decide to perform an in-place upgrade to SQL Server 2005. You verify that Certkiller -DB01 exceeds the minimum recommended hardware requirements for SQL Server 2005.

What should you do next?

- A. Upgrade Certkiller -DB01 to Windows Sever 2003 Service Pack 1.

- B. Run the SQL Server Upgrade Advisor on Certkiller -DB01.
- C. Detach the CK_Trading database on Certkiller -DB01.
- D. Run the SQL Server 2005 Setup utility.

Answer: B

Explanation: Before you perform an in-place upgrade of an existing database, you should run the SQL Server Upgrade Advisor. The SQL Server Upgrade Advisor produces a list of items that must be addressed before and after performing the upgrade. This list of items is specific to the existing installation.

Incorrect Answers:

A: You do not need to upgrade Certkiller -DB01 to Windows Sever 2003 Service Pack 1 as SQL Server 2005 can be installed on Windows 2000 Server with Service Pack 4.

C: You need to detach a database when you need to move the database to another database instance or another database server. You do not need to detach the database before performing an in-place upgrade.

D: You should run the SQL Server Upgrade Advisor before you run the SQL Server 2005 Setup utility. The SQL Server Upgrade Advisor produces a list of items that must be addressed before and after performing the upgrade. This list of items is specific to the existing installation.

Reference:

Microsoft SQL Server 2005 Books Online (2006), Index: updating databases, detaching and attaching databases

QUESTION 5:

You work as the database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. All servers on the Certkiller .com network run Windows 2000 Server. The Certkiller .com network contains a SQL Server 7.0 database server named Certkiller -DB01 that runs on Windows 2000 Server. Service Pack 4 for Windows 2000 Server id installed on Certkiller -DB01.

Certkiller .com wants to migrate Certkiller -DB01 to SQL Server 2005. You decide to perform an in-place upgrade to SQL Server 2005. You verify that Certkiller -DB01 exceeds the recommended hardware requirements for SQL Server 2005.

What should you do next?

- A. Upgrade Certkiller -DB01 to Windows Sever 2003.
- B. Install Service Pack 4 for SQL Server 7.0 on Certkiller -DB01.
- C. Upgrade SQL Server 7.0 on Certkiller -DB01 to SQL Server 2000.
- D. Run the SQL Server 2005 Setup utility on Certkiller -DB01 and upgrade directly to SQL Server 2005.

Answer: B

Explanation: You can upgrade from SQL Server 7.0 with Service Pack 4 directly to SQL Server 2005. Thus you need only install Service Pack 4 for SQL Server 7.0.

Incorrect Answers:

A: You do not need to upgrade Certkiller -DB01 to Windows Sever 2003 as SQL Server 2005 can be installed on Windows 2000 Server with Service Pack 4.

C: You do not need to upgrade Certkiller -DB01 to SQL Server 2000 as SQL Server 7.0 with Service Pack 4 can be upgraded directly to SQL Server 2005. Besides, you will need to install Service Pack 3 for SQL Server 2000 before you can upgrade it to SQL Server 2005.

D: You need to install Service Pack 4 for SQL Server 7.0 before you can upgrade it directly to SQL Server 2005.

Reference:

Microsoft SQL Server 2005 Books Online (2006), Index: upgrading SQL Server, upgrading to SQL Server 2005

QUESTION 6:

You work as the database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. All servers on the Certkiller .com network run Windows Server 2003 and all client computers run Windows XP Professional. The Certkiller .com network contains two SQL Server 2000 database servers named Certkiller -DB01 and Certkiller -DB02 that are configured for log shipping. Certkiller -DB01 is the primary server and hosts a database named CK_Trading. Both Certkiller -DB01 and Certkiller -DB02 have Service Pack 3 for SQL Server 2000 installed. Certkiller .com wants to upgrade Certkiller -DB01 and Certkiller -DB02 to SQL Server 2005 Enterprise Edition but the CK_Trading database must remain available during the upgrade process. You verify that both Certkiller -DB01 and Certkiller -DB02 exceed the recommended hardware requirements for SQL Server 2005 Enterprise Edition. You then install Service Pack 1 for Windows Server 2003 on Certkiller -DB01 and Certkiller -DB02. What should you do next?

- A. Install Service Pack 4 for SQL Server 2000 on both database servers.
- B. Enable failover in the SQL Server 2000 log shipping configuration and upgrade Certkiller -DB02 to SQL Server 2005.
- C. Run the Database Maintenance Plan Wizard on Certkiller -DB01 to upgrade it to SQL Server 2005.
- D. Upgrade Certkiller -DB02 to SQL Server 2005 and then upgrade Certkiller -DB01 to SQL Server 2005.

Answer: B

Explanation: SQL Server 2005 does not support the Database Maintenance Plan Wizard which was integral to log shipping in SQL Server 2000. Hence, log shipping stops functioning when you upgrade a SQL Server 2000 log shipping configuration

directly to SQL Server 2005. To upgrade log shipping to SQL Server 2005, you must migrate the database servers to SQL Server 2005. If you want the database to remain available during the migration process, you must migrate with failover. To migrate with failover you must enable failover and upgrade Certkiller -DB02. Then you need to failover to Certkiller -DB02 and upgrade Certkiller -DB01. Then configure log shipping with the Certkiller -DB02 becoming the primary server and Certkiller -DB01 becoming the secondary server.

Incorrect Answers:

A: You do not need to install Service Pack 4 for SQL Server 2000 on the database servers as SQL Server 2000 with Service Pack 3 can be upgraded to SQL Server 2005. However, SQL Server 2005 does not support the Database Maintenance Plan Wizard which was integral to log shipping in SQL Server 2000. Hence, log shipping stops functioning when you upgrade a SQL Server 2000 log shipping configuration directly to SQL Server 2005.

C: SQL Server 2005 does not support the Database Maintenance Plan Wizard which was integral to log shipping in SQL Server 2000. Hence, log shipping stops functioning when you upgrade a SQL Server 2000 log shipping configuration directly to SQL Server 2005.

D: Removing log shipping from Certkiller -DB02 and upgrading Certkiller -DB01 to SQL Server 2005 will not ensure that the CK_Trading database remains available.

Reference:

Microsoft SQL Server 2005 Books Online (2006), Index: log shipping [SQL Server], upgrading

QUESTION 7:

You work as the database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. All servers on the Certkiller .com network run Windows Server 2003. The Certkiller .com network contains a Domain Controller named Certkiller -DC01 and a member server named Certkiller -SR21.

Certkiller .com wants to implement a SQL Server 2005 database that will store sales data. All Certkiller .com users will require access to the Sales_DB database. You have been instructed to install SQL Server 2005 on Certkiller -SR21 and to create a database named Sales_DB. You need to ensure that Certkiller .com users and applications can connect to the database instance by using only the computer name. What should you do during the installation of SQL Server 2005 on Certkiller-SR21?

- A. Select the Named Instance and use MSSQLSERVER as the instance name.
- B. Select the Named Instance and use Certkiller -DC01 as the instance name.
- C. Select the Default Instance and accept the default instance name.
- D. Select the Named Instance and use Certkiller -SR21 as the instance name.

Answer: C

Explanation: To ensure that users and applications can connect to the instance

using the computer name, you should select the Default Instance and accept the default instance name, MSSQLSERVER.

Incorrect Answers:

A, B, D: You cannot use a Named Instance as users and applications will need to specify the server name and the instance name when connection to the database.

QUESTION 8:

You work as the database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. All servers on the Certkiller .com network run on Windows Server2003 while the client computers run Windows XP Professional, Windows 98 Second Edition, and Novell NetWare 6.5. All computers are members of the Certkiller .com Active Directory domain. The Certkiller .com network contains a SQL Server 2005 database server named Certkiller -DB01 that hosts a database instance named MSSQLSERVER. Certkiller .com wants to implement another SQL Server 2005 instance for a database named Sales_DB. Sales_DB will store sales data for the company. All Certkiller .com users will require access to the Sales_DB database. You have been instructed to create the SQL Server 2005 instance for the Sales_DB database. You need to ensure that all users will be able to connect to the Sales_DB database. You cannot make configuration changes to the client computers. Your solution must optimize network traffic.

What should you do?

- A. Install the instance as a Named Instance and configure it to support Mixed Mode Authentication. The enable TCP/IP for the instance.
- B. Install the instance as a Default Instance and configure it to support Windows Authentication. The enable TCP/IP for the instance.
- C. Install the instance as a Named Instance and configure it to support SQL Authentication. The enable IPX/SPX for the instance.
- D. Install the instance as a Default Instance and configure it to support Windows Authentication. The enable IPX/SPX for the instance.

Answer: A

Explanation: The default instance is named MSSQLSERVER. This instance name is already in use; therefore you should use a Named Instance. You should use Mixed Mode Authentication to support NetWare users.

Incorrect Answers:

- B: You cannot use the Default Instance as the default instance is named MSSQLSERVER and is already in use. Furthermore, you need to use Mixed Mode Authentication as Windows Authentication does not support NetWare users.
- C: You can use SQL Authentication, but this will require that you create SQL Logins for all users and groups. It would be easier to use Windows Authentication for clients that user Windows XP Professional and Windows 98 SE, and create SQL Logins for the NetWare 6.5 users. Also, SQL Server 2005 does not support IPX/SPX. The newer versions of

NetWare, such as NetWare 6.5, support TCP/IP which makes it the common protocol for all clients.

D: You cannot use the Default Instance as the default instance is named MSSQLSERVER and is already in use. Furthermore, you need to use Mixed Mode Authentication as Windows Authentication does not support NetWare users. Also, SQL Server 2005 does not support IPX/SPX. The newer versions of NetWare, such as NetWare 6.5, support TCP/IP which makes it the common protocol for all clients.

QUESTION 9:

You work as the database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. All servers on the Certkiller .com network run on Windows Server2003 and all client computers run Windows XP Professional. The Certkiller .com network contains two SQL Server 2005 database servers that are configured in a cluster with two nodes named NodeA and NodeB. NodeA hosts a default database instance named MSSQLSERVER.

Certkiller .com wants to implement another SQL Server 2005 instance for a database named Sales_DB. Sales_DB will store sales data for the company. All Certkiller .com users will require access to the Sales_DB database. You have been instructed to create the SQL Server 2005 instance for the Sales_DB database on Node

A. You need to ensure that all users will be able to connect to the Sales_DB database. What should you do?

A. Run the SQL Server 2005 Setup utility on NodeA and select a Named Instance for the Sales_DB database.

B. Run the SQL Server 2005 Setup utility on NodeA and select the Default Instance for the Sales_DB database.

C. Move the existing instance from NodeA to NodeB. Run the SQL Server 2005 Setup utility on Node

A. Select the Default Instance and accept the default instance name

D. Run the SQL Server 2005 Setup utility on NodeB and select a Named Instance for the Sales_DB database.

Answer: A

Explanation: The default instance is named MSSQLSERVER. This instance name is already in use on NodeA; therefore you should use a Named Instance.

Incorrect Answers:

B: You cannot use the Default Instance as the default instance is named MSSQLSERVER and is already in use.

C: You do not need to move the existing instance as you can install up to 50 instances on a SQL Server 2005 database server.

D: You need to install the instance on NodeA not NodeB. You can install a new instance on NodeA if you use a Named Instance.

Reference:

QUESTION 10:

You work as the database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. All servers on the Certkiller .com network run on Windows Server2003 and all client computers run Windows XP Professional. The Certkiller .com network contains two SQL Server 2005 database servers that are configured in a cluster with two nodes named NodeA and NodeB.

NodeA contains two physical hard disk drives, one of which hosts the Quorum resource and contains a single partition while the other contains two logical partitions. NodeB also has two physical hard disk drives each of which contain a single logical partition. None of the hard disks in the cluster have any unallocated space left. The SQL Server 2005 cluster currently hosts four database instances. Certkiller .com wants to implement another SQL Server 2005 instance for a database named Manufacturing_DB. Manufacturing_DB will store data from the Certkiller .com Manufacturing department. You have been instructed to create the SQL Server 2005 instance for the Manufacturing_DB database. What should you do?

- A. Run the SQL Server 2005 Setup utility on NodeB and select a Named Instance for the Manufacturing_DB database.
- B. Create a second logical partition on one of the physical disk drives on NodeB. Run the SQL Server 2005 Setup utility on NodeB and select a Named Instance for the Manufacturing_DB database.
- C. Add a third physical disk drive on one of the nodes. Run the SQL Server 2005 Setup utility on that node and select Named Instance for the Manufacturing_DB database.
- D. Run the SQL Server 2005 Setup utility on NodeA and select a Named Instance for the Manufacturing_DB database.

Answer: C

Explanation: A single SQL Server 2005 database server can host up to 50 instances but a clustered SQL Server 2005 configuration can only host one instance per logical disk drive. There are currently five logical disk drives in the cluster but one is used for the Quorum resource this leaves you with four logical disk drives that can support a total of four instances. To add another instance, you would need to add another logical disk drive. As none of the hard disks have unallocated free space, you would need to add another hard disk to the cluster.

Incorrect Answers:

A, D: You cannot add another instance to any of the two nodes as a clustered SQL Server 2005 configuration can only host one instance per logical disk drive. There are currently five logical disk drives in the cluster but one is used for the Quorum resource this leaves you with four logical disk drives that can support a total of four instances. To add another instance, you would need to add another logical disk drive. As none of the hard disks have unallocated free space, you would need to add another hard disk to the cluster.

B: You cannot create another partition on any of the physical drives on NodeB as none of the hard disks in the cluster has unallocated free space that you can use to create another partition. You will need to add another hard disk to the cluster for the new instance.

QUESTION 11:

You work as the database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. All servers on the Certkiller .com network run Windows Server2003. You install SQL Server 2005 on a server named Certkiller -DB01 that runs on a Windows Server 2003 Service Pack 1. Certkiller -DB01 will contain a database named CK_Sales that will be used to store Internet sales transactions from the company's e-commerce Web site.

Certkiller -DB01 has seven 120 GB hard disk drives. Two of the hard disks are configured as a mirrored volume and contains the operating system and SQL Server 2005 is installation. Another two disks are configured as a RAID-1 volume and the rest are configured as a RAID-5 volume.

You to need to plan the placement of the database and the transaction logs. You want to implement fault tolerance and best performance for the CK_Sales database. What should you do?

- A. Place the database on the RAID-5 array and the transaction log on the RAID-1 array.
- B. Place the database and the transaction log on the RAID-5 array.
- C. Place the database on the mirrored volume and the transaction log on the RAID-5 array.
- D. Place the database on the RAID-1 array and the transaction log on the RAID-5 array.

Answer: A

Explanation: The RAID-5 array provides striping with parity. Striping improves read/write operations across multiple disks while parity provides fault tolerance. Placing the database on the RAID-5 array would thus provide best performance and fault tolerance. Performance can be improved by placing the transaction logs on a separate hard disk. The RAID-1 array provides striping which improves read/write operations across multiple disks.

Incorrect Answers:

B: RAID-5 array provides striping with parity. Striping improves read/write operations across multiple disks while parity provides fault tolerance. Performance can be improved even further by placing the transaction logs on a separate hard disk.

C: Mirrored volumes write the same data to two physical disks. This provides fault tolerance but does not provide improved performance. Fault tolerance and improved performance can be gained by placing the database on the RAID-5 volume.

D: The RAID-1 array provides striping which improves read/write operations across multiple disks but it does not provide fault tolerance.

QUESTION 12:

You work as the database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. All servers on the Certkiller .com network run Windows Server2003. The Certkiller .com network contains a SQL Server 2005 database server named Certkiller -DB01, and two Microsoft Exchange Server 2003 mail servers named Certkiller -EX01 and Certkiller -EX02. The Simple Mail Transfer Protocol (SMTP) is enabled on both mail servers.

Certkiller -DB01 hosts a database named CK_Sales that is used to store sales related data for the company. The CK_Sales database contains a table named Inventory. The manager of the Purchasing department has requested a database report containing a list of the 20 most popular products and the amount of stock the company has for each of the top 20 products. You have produced the report. You attempt to send the report by e-mail using the sp_send_dbmail procedure but receive a failed return code. You verify that the mail servers are online and that the SMTP service is running on them.

You need to send the report by e-mail from Certkiller -DB01.

What should you do?

- A. Enable Database Mail for the database instance.
- B. Install a Messaging Application Programming Interface (MAPI) client on Certkiller -DB01.
- C. Install Microsoft Office Outlook on Certkiller -DB01.
- D. Run the DatabaseMail90.exe utility on Certkiller -DB01.

Answer: A

Explanation:

You need to enable Database Mail before you can use the sp_send_dbmail procedure. Database Mail is not enabled by default but can be enabled by using the Database Mail Configuration Wizard, the SQL Server Surface Area Configuration tool, or sp_configure.

Incorrect Answers:

B, C: There is no need to install a mail client such as Microsoft Office Outlook or any other mail client on the database servers. SQL Server 2005 uses the Database Mail feature to send e-mail messages but this feature is not enabled by default.

D: The DatabaseMail90.exe utility is the utility that the sp_send_dbmail procedure uses to send the e-mail message. However, this utility will not run if Database Mail is not enabled.

Reference:

Microsoft SQL Server 2005 Books Online (2006), Index: sp_send_dbmail

QUESTION 13:

You work as the database administrator at Certkiller .com. Certkiller .com has five

departments named Accounting, Sales, Marketing, Research and Manufacturing. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. All servers on the Certkiller .com network run on Windows Server2003. The Certkiller .com network contains a SQL Server 2005 database server named Certkiller -DB01. Certkiller -DB01 hosts a database named CK_Sales that stores data for the Sales department. The CK_Sales database is accessed frequently throughout the day.

Mia Hamm, the manager of the Sales department is concerned that data may be lost should Certkiller -DB01 suffer hardware failure. Mia Hamm wants you to implement a recovery model and backup strategy for the CK_Sales database. You want to implement the highest level of recovery and you want to be able to restore the data for the CK_Sales database in the least possible time. What should you do? (Each correct answer provides part of the solution. Choose TWO.)

- A. Implement a Full Recovery model.
- B. Implement a Simple Recovery model.
- C. Perform a full backup and log backup every evening.
- D. Perform a full backup on Saturdays and differential backups Monday through Friday.
- E. Perform a full backup on Saturdays and inferential backups Monday through Friday.

Answer: A, D

Explanation: The Full Recovery model provides the highest level of recovery as it does not clear data in the log files until they are backed up. Your restore time is reduced by performing differential backups. With differential backups, all data since the last full backup is saved. This means that you only need the last full backup and the last differential backup when you need to restore data.

Incorrect Answers:

B: The Simple Recovery model clears the data in the log file as soon as the transaction is written to the database. This means that transactions that have been written to the database after the last backup could be lost in the event of hardware failure.

C: You should backup the log files throughout the day rather than at the end of the day as all transactions that occurred during the day could be lost should a hardware failure occur during the day.

D: Performing a full backup on Saturdays and inferential backups Monday through Friday would increase your restore time as each inferential backup performed after the last full backup must be restored in sequence. Thus, should a hardware failure occur on Friday, you would need to restore the last full backup from Saturday, and then the inferential backup from Monday, followed by the inferential backup from Tuesday until the last inferential backup.

QUESTION 14:

You work as the database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. All

servers on the Certkiller .com network run on Windows Server2003 while half the client computers run Windows XP Professional and the rest run UNIX. The Certkiller .com network contains a SQL Server 2000 database server named Certkiller -DB01 that runs on Windows 2000 Server. Certkiller -DB01 contains a database named CK_DB.

All Certkiller .com users will require access to the CK_DB database from the corporate network. You need to configure access to the CK_DB database to meet these requirements. You want to use the least amount of administrative effort in accomplishing this task.

What should you do?

- A. Configure the CK_DB database to support Mixed Mode Authentication. Create a SQL login for the Domain Users group. Create SQL logins with passwords for each UNIX user.
- B. Configure the CK_DB database to support Windows Authentication. Create a SQL login for each domain user.
- C. Configure the CK_DB database to support Mixed Mode Authentication. Create a SQL login for each domain user.
- D. Configure the CK_DB database to support Windows Authentication. Create a SQL login for the Domain Users group.

Answer: A

Explanation: You need to use Mixed Mode Authentication to allow access to the database for the UNIX users. Each UNIX user will require a SQL login with a password. The Windows XP Professional users can be granted access collectively by granting the Domain Users group access to the database.

Incorrect Answers:

B, D: You cannot use Windows Authentication as this will not allow the UNIX users to access the database.

C: You can create a SQL login for each user but it would require less administrative effort to create a SQL Login for the Domain Users group and separate SQL Logins for the UNIX users.

QUESTION 15:

You work as the database administrator at Certkiller .com. Certkiller .com has five departments named Accounting, Sales, Marketing, Research and Manufacturing. The Certkiller .com user accounts of each department are located in a domain group named after the department. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. All servers on the Certkiller .com network run on Windows Server2003. The Certkiller .com network contains a SQL Server 2005 database server named Certkiller -DB01.

You have been instructed to create a separate database on Certkiller -DB01 for each department. Each Certkiller .com department will use their database to store department-related data. Users in each department must be able to read and update

data in the database belonging to their department. No users must be able to access the database belonging to another department.

You need to configure access to the databases to meet these requirements. You want to use the least amount of administrative effort in accomplishing this task.

What should you do?

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

Answer: A

Explanation: The users of each department belong to a group named after the department. You can use these groups to configure access to the databases. Each database user should be added to the db_datareader and db_datawriter database roles for their respective databases. This will prevent users in one department from accessing another department's database.

Incorrect Answers:

B: You could create a Windows Authentication login for each domain user, configure the logins as database users for the appropriate database, and add each database user to the db_datareader and db_datawriter database roles but it would require less administrative effort to create a Windows Authentication login for each domain group.

C: You can use the domain groups to configure access to the databases but you should not add each database user to the db_ddladmins database role as this violates the principle of least privilege.

D: You could create a Windows Authentication login for each domain user but it would require less administrative effort to create a Windows Authentication login for each domain group. Furthermore, you should not add each database user to the db_ddladmins database role as this violates the principle of least privilege.

QUESTION 16:

You work as the database administrator at Certkiller .com. Certkiller .com has five departments named Accounting, Sales, Marketing, Research and Manufacturing. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. All servers on the Certkiller .com network run on Windows Server2003 and all client computers run Windows XP Professional. Certkiller .com has a

Certification Authority (CA) that has issued digital certificates to each client computer and each server on the Certkiller .com network.

The Certkiller .com network contains a SQL Server 2005 database server named Certkiller -DB01. Certkiller -DB01 hosts a database named CK_Accounts that stores data for the Accounting department. Users in the Accounting department use an in-house client application to connect to the CK_Accounts database.

Rory Allen is the manager of the Accounting department. Rory Allen is concerned about the security of the data transmitted between Certkiller -DB01 and the in-house client application. You are required to address this matter by implementing a security solution that will require authentication and encryption. What should you do?

- A. Implement a self-signed digital certificate on Certkiller-DB01, set the ForceEncryption option on Certkiller -DB01 to Yes and restart SQL Server service.
- B. Implement a self-signed digital certificate on Certkiller-DB01, set the ForceEncryption option on Certkiller -DB01 to Yes and restart Certkiller -DB01.
- C. Implement a self-signed digital certificate on Certkiller-DB01, set the ForceEncryption option on Certkiller -DB01 to No and restart the SQL Server service.
- D. Configure the Database Engine to use a digital certificate, set the ForceEncryption option on Certkiller -DB01 to No and restart Certkiller -DB01.
- E. Configure the Database Engine to use a digital certificate, set the ForceEncryption option on Certkiller -DB01 to Yes and restart Certkiller -DB01.

Answer: E

Explanation: To configure SQL Server 2005 to encrypt data, you must configure the Database Engine to use a digital certificate in the SQL Server Configuration Manager, set the ForceEncryption option to Yes and restart the server.

Incorrect Answers:

A: Certkiller .com has a Certification Authority (CA). You should use the CA to provide digital certificates for the Database Engine as the CA may reject the self-signed certificate. Also, once you make configuration changes to the SQL Server 2005 server, you must restart the server for the configuration changes to take effect. Restarting the SQL Server service will not ensure that the configuration changes will take effect.

B: Certkiller .com has a Certification Authority (CA). You should use the CA to provide digital certificates for the Database Engine as the CA may reject the self-signed certificate.

C: Certkiller .com has a Certification Authority (CA). You should use the CA to provide digital certificates for the Database Engine as the CA may reject the self-signed certificate. Also, if you want to ensure that data is encrypted you must set the ForceEncryption option to Yes. If the ForceEncryption option is set to No, then data encryption will only be used if the client requests it. Furthermore, once you make configuration changes to the SQL Server 2005 server, you must restart the server for the configuration changes to take effect. Restarting the SQL Server service will not ensure that the configuration changes will take effect.

D: If you want to ensure that data is encrypted you must set the ForceEncryption option to Yes. If the ForceEncryption option is set to No, then data encryption will only be used if the client requests it.

QUESTION 17:

You work as the database administrator at Certkiller .com. The Certkiller .com network consists of a single Windows Server 2003 domain named Certkiller .com. All servers on the Certkiller .com network run Windows Server 2003. The Certkiller .com departments include a Sales department, a Purchasing department, and an Accounts department. The Certkiller .com network contains a SQL Server 2000 database server named Certkiller -DB01 and a SQL Server 2005 database server named Certkiller -DB02. Certkiller -DB02 contains a database named CK_Sales that stores sales data for the company.

The Sales department makes extensive use of the CK_Sales database. Certkiller .com users in the Purchasing department make use of a Microsoft Access database that runs on an application server named Certkiller -SR04 while users in the Accounts departments make use of a SQL Server 2000 database that run Certkiller -DB01.

The manager of the Accounts department wants to be able to run distributed queries against the CK_Sales database and the databases in the Purchasing and Accounts departments. You must configure the Certkiller -SR04 and Certkiller -DB01 as linked servers on Certkiller -DB02. You need to select the required data sources to create the linked server definitions.

What should you do? (Choose all that apply.)

- A. Use the OLE DB Provider for Jet as a data source.
- B. Use the SQL Native Client OLE DB Provider as a data source.
- C. Use the OLE DB Provider for ODBC as a data source.
- D. Use the OLE DB Provider for DB2 as a data source.

Answer: A, B

Explanation: The Microsoft Access database requires the OLE DB Provider for Jet to communicate with the linked server while the SQL Server 2000 database requires the SQL Native Client OLE DB Provider.

Incorrect answers:

C: SQL Server 6.0 databases require the OLE DB Provider for ODBC but SQL Server 7.0 and later databases require the SQL Native Client OLE DB Provider.

D: The OLE DB Provider for DB2 is required for IBM DB2 databases.

Reference:

Microsoft SQL Server 2005 Books Online (2007), Index: linked servers [SQL Server], about linked servers

Microsoft SQL Server 2005 Books Online (2007), Index: linked servers [SQL Server], Access databases

Microsoft SQL Server 2005 Books Online (2007), Index: linked servers [SQL Server], DB2 databases

QUESTION 18:

You have just been hired as the database administrator at Certkiller .com. Certkiller .com currently uses several Microsoft Access databases to store its business information. The Microsoft Access databases are stored on two application servers named Certkiller -SR04 and Certkiller -SR05. You have been instructed to migrate some of the Microsoft Access databases on Certkiller -SR04 to SQL Server 2005. You must ensure that Certkiller .com users can access data in the Microsoft Access databases until all the data has been migrated to SQL Server 2005. During the migration process, Certkiller .com users must be able to run distributed queries against the data in both SQL Server 2005 and Microsoft Access. You install SQL Server 2005 on a member server named Certkiller -SR06. Certkiller -SR06 has limited disk space. What should you do?

- A. Configure Certkiller -SR04 as a linked server on Certkiller -SR06. Configure the OLE DB Provider for Jet as the data source.
- B. Configure Certkiller -SR04 as a linked server on Certkiller -SR06. Configure the OLE DB Provider for ODBC as a data source.
- C. Configure Certkiller -SR06 as a linked server on Certkiller -SR04. Configure the OLE DB Provider for ODBC as a data source.
- D. Configure Certkiller -SR06 as a linked server on Certkiller -SR04. Configure the OLE DB Provider for Jet as a data source.

Answer: A

Explanation: To support distributed queries, you must configure the Microsoft Access database server as a linked server on the SQL database server. Microsoft Access requires the OLE DB Provider for Jet to communicate with the linked server.

Incorrect answers:

B: Microsoft Access requires the OLE DB Provider for Jet to communicate with the linked server, not the OLE DB Provider for ODBC.

C, D: You cannot create linked servers in Microsoft Access. You must create them in SQL Server 2005.

Reference:

Microsoft SQL Server 2005 Books Online (2007), Index: linked servers [SQL Server], Access databases

QUESTION 19:

You work as a database administrator at Certkiller .com. The Certkiller .com network consists of a single Windows 2000 domain named Certkiller .com. All servers on the Certkiller .com network run Windows 2000 Server and all client computers run Windows XP Professional. Certkiller .com currently uses an ODBC-compliant

proprietary Relational Database Management System (RDBMS) to store business information.

You have been instructed to migrate some of the databases in the proprietary RDBMS to SQL Server 2005. You need to ensure that Certkiller .com users can access data in the RDBMS databases during the migration process. You deploy SQL Server client tools on all Certkiller .com client computers. You want to ensure that Certkiller .com users can access data in both database systems by running ad hoc queries against SQL Server.

What should you do?

- A. Create a linked server and a data source name (DSN) for the proprietary RDBMS system.
- B. Create a linked server and use the OLE DB Provider for SQL as a data source.
- C. Create a local partitioned view of the databases on the RDBMS system in SQL Server 2005.
- D. Create stored procedures in SQL Server 2005 to query the RDBMS system.

Answer: A

Explanation: To run distributed queries with a proprietary database system, you must create a linked server definition for the proprietary database system. As the proprietary database system is ODBC-compliant, you can create a data source name (DSN) for the proprietary database system that can be used to identify the proprietary database system.

Incorrect Answers:

B: The OLE DB Provider for SQL was used in SQL Server 2000. It is not used in SQL Server 2005 and has been replaced by the SQL Native Client OLE DB Provider.

C: A partitioned view must be derived from the same SQL Server instance. It cannot be created on remote databases.

D: Stored procedures can be run against linked server. However, the Certkiller .com users must be able to use ad hoc queries rather than stored procedures.

Reference:

Microsoft SQL Server 2005 Books Online (2007), Index: linked servers [SQL Server], about linked servers

QUESTION 20:

You work as the database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. All servers on the Certkiller .com network run Windows Server2003. The Certkiller .com network contains a single SQL Server 2005 database server named Certkiller -DB01. Certkiller -DB01 hosts a database named CK_Sales that stores sales data for the company. The CK_Sales database is configured to use the full recovery model.

You want to provide zero down time and high availability for the CK_Sales database. You decide to add a new database server which will be named

Certkiller -DB02 on the Certkiller .com network. You want to configure Certkiller -DB01 and Certkiller -DB02 as mirrored servers. You perform a full backup of the CK_Sales database on Certkiller -DB01. What should you do next on Certkiller -SR02?

- A. Install SQL Server 2005 and create a database instance.
- B. Restore the full backup of CK_Sales to Certkiller -DB02.
- C. Configure the database mirroring endpoint.
- D. Reproduce the user login from Certkiller -DB01 on Certkiller -DB02.

Answer: A

Explanation: Before you can implement mirrored servers, both servers must be running the same version of SQL Server 2005 (either Standard Edition, or Enterprise Edition) and each server must have a database instance.

Incorrect Answers:

B: You must first create a database instance on Certkiller -DB02 before you restore the full backup to Certkiller -DB02.

C: You must first create a database instance on Certkiller -DB02 and restore the full backup of the principle database to Certkiller -DB02 before you configure the database mirroring endpoint.

D: You must first create a database instance on Certkiller -DB02 and restore the full backup of the principal database to Certkiller -DB02 before you reproduce the user login from Certkiller -DB01 on Certkiller -DB02.

Reference:

Microsoft SQL Server 2005 Books Online (2006), Index: database mirroring [SQL Server], mirror database creation

QUESTION 21:

You work as the database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. All servers on the Certkiller .com network run Windows Server2003. The Certkiller .com network contains two member servers named Certkiller -SR15 and Certkiller -SR16, as well as a single SQL Server 2005 database server named Certkiller -DB01. Certkiller -DB01 hosts a database named CK_Projects that contains mission critical data.

You need to provide fault tolerance and improve the availability of the CK_Projects database. Your solution must provide zero down time and the highest level of recovery. You are unable to acquire any additional hardware. You verify the Certkiller -SR15 and Certkiller -SR16 exceeds the recommended hardware requirements for SQL Server 2005. You then install SQL Server 2005 on Certkiller -SR15 and Certkiller -SR16.

What should you do next?

- A. Configure the three SQL Server 2005 database servers in a failover cluster with two

active nodes and one passive node.

B. Configure the three SQL Server 2005 database servers in a failover cluster with three active nodes.

C. Configure the three SQL Server 2005 database servers in a failback cluster with three nodes.

D. Configure Certkiller -SR15 as a mirror server and Certkiller -SR16 as witness server.

Answer: D

Explanation: You can use mirrored servers to ensure the highest level of data protection. When you implement a witness server you ensure automatic failover when one of the mirrored servers becomes unavailable. This ensures zero-down time.

Incorrect Answers:

A, B, C: Clustering would require additional hardware in the form of a shared external hard disk drive.

Reference:

Microsoft SQL Server 2005 Books Online (2006), Index: witness [SQL Server]

QUESTION 22:

You work as the database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. All servers on the Certkiller .com network run Windows Server2003. The Certkiller .com network contains two SQL Server 2005 database servers named Certkiller -DB01 and Certkiller -DB02, and a member server named Certkiller -SR33. A database mirroring session operating in synchronous mode is in operation between Certkiller -DB01 and Certkiller -DB02, with Certkiller -DB01 being the principal database server. Certkiller -DB01 hosts a database named CK_Sales that stores sales data from Certkiller .com's e-Commerce Web site.

One Monday morning you discover that the CK_Sales database is unavailable. You discover that Certkiller -DB01 is offline while Certkiller -DB02 is online; however, failover to Certkiller -BD02 did not occur. This failure has adversely affected the profitability of Certkiller .com's e-Commerce division. You need to improve the availability of the CK_Sales database. You are unable to acquire any additional hardware.

What should you do?

A. Configure Certkiller -DB01 and Certkiller-DB02 in a failover cluster with two active nodes.

B. Configure the mirrored database session to run in asynchronous mode.

C. Install SQL Server 2005 on Certkiller -SR33 and configure it as a witness server.

D. Configure the mirrored database to operate in high-safety mode.

Answer: C

Explanation: High-safety mode supports synchronous operation in which all transactions must be committed to both partners. This provides zero-loss of data but requires a witness to provide automatic failover.

Incorrect Answers:

A: Failover clustering would require additional hardware in the form of a shared external hard disk drive. You, however, are unable to acquire additional hardware. Database mirroring is a software solution, but requires a witness server to provide failover.

B: In Asynchronous mode a transaction is committed to the principal database without requiring the principal database to wait for the mirror database to write the transaction log to disk. This maximizes performance for the mirrored database but at the expense of high-availability. When the principal server of a mirrored database in asynchronous operation fails, you must manually stop the mirror session, update the database and begin a new mirroring session.

D: High-safety mode supports synchronous operation in which all transactions must be committed to both partners. This provides zero-loss of data but requires a witness to provide automatic failover.

Reference:

Microsoft SQL Server 2005 Books Online (2006), Index: database mirroring [SQL Server], about database mirroring

Microsoft SQL Server 2005 Books Online (2006), Index: database mirroring [SQL Server], asynchronous

Microsoft SQL Server 2005 Books Online (2006), Index: database mirroring [SQL Server], synchronous

QUESTION 23:

You work as the database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. All servers on the Certkiller .com network run Windows Server2003. The Certkiller .com network contains three SQL Server 2005 database servers named Certkiller -DB01, Certkiller -DB02, and Certkiller -DB03. Database mirroring is configured between Certkiller -DB01 and Certkiller -DB02, with Certkiller -DB03 as the witness server. Certkiller -DB01 is the principal database server and hosts a database named CK_Sales that stores sales data from Certkiller .com's e-Commerce Web site.

The size of the CK_Sales database has increased over the last three month as Certkiller .com's customer base has increased. You are concerned about the performance of the CK_Sales database. You want to maximize database performance while retaining the mirrored configuration.

What should you do?

- A. Configure the database mirroring session to run in synchronous operation.
- B. Configure the database mirroring session to run in asynchronous operation.
- C. Configure the database mirror to operate in high-safety mode.
- D. Configure the database mirror to operate in high-availability mode.

Answer: B

Explanation:

Asynchronous mode provides the best performance for mirrored databases at the expense of data recovery. During asynchronous mode, a transaction is committed to the principal database without requiring the principal database to wait for the mirror database to write the transaction log to disk. This means that the principal database can operate with minimum transaction lag.

Incorrect Answers:

A, C: High-safety mode supports synchronous operation. Transaction latency occurs in this mode as a transaction must be committed to both partners. In Asynchronous mode a transaction is committed to the principal database without requiring the principal database to wait for the mirror database to write the transaction log to disk. This maximizes performance for the mirrored database.

D: SQL Server 2005 does not have a high-availability mode for database mirroring. It only has high-safety mode, which supports synchronous operation; and high-performance mode, which supports asynchronous operation.

Reference:

Microsoft SQL Server 2005 Books Online (2006), Index: database mirroring [SQL Server], about database mirroring

Microsoft SQL Server 2005 Books Online (2006), Index: database mirroring [SQL Server], asynchronous

Microsoft SQL Server 2005 Books Online (2006), Index: database mirroring [SQL Server], synchronous

QUESTION 24:

You work as the database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. All servers on the Certkiller .com network run Windows Server2003. The Certkiller .com network contains a single database server named Certkiller -DB01, and a file server named Certkiller -SR22. Certkiller -DB01 runs SQL Server 2005

Workgroup Edition and hosts a database named CK_Sales that stores sales data for the company.

You want to implement a recovery model and backup strategy for the CK_Sales database. You decide to add a second database server named Certkiller -DB02 to the Certkiller .com network. You install SQL Server 2005 Workgroup Edition on Certkiller -DB02.

What should you do next?

- A. Configure Certkiller -DB01 and Certkiller -DB02 for log shipping.
- B. Configure database mirroring on Certkiller -DB01 and Certkiller -DB02.
- C. Configure Certkiller -DB01 and Certkiller -DB02 in a failover cluster with two active nodes.
- D. Configure Certkiller -DB01 and Certkiller -DB02 for snapshot replication.

Answer: A

Explanation: Of the available options, the only backup and recovery mechanism supported by SQL Server 2005 Workgroup Edition is log shipping.

Incorrect Answers:

B, D: Database mirroring and failover clustering requires SQL Server 2005 Standard Edition or SQL Server 2005 Enterprise Edition.

C: Snapshot replication requires SQL Server 2005 Enterprise Edition.

Reference:

Microsoft SQL Server 2005 Books Online (2006), Index: features [SQL Server]

QUESTION 25:

You work as the database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. All servers on the Certkiller .com network run Windows Server2003. The Certkiller .com network contains a single database server named Certkiller -DB01, and a file server named Certkiller -SR22. Certkiller -DB01 runs SQL Server 2005 Workgroup Edition and hosts a database named CK_Sales that stores sales data for the company. The CK_Sales database is configured to use the full recovery model. You want to implement a backup strategy for the CK_Sales database that includes log shipping. You successfully install SQL Server 2005 Workgroup Edition on a new database server named Certkiller -DB02.

What should you do next?

- A. Create a database instance on Certkiller -DB02.
- B. Perform a full backup of the CK_Sales database on Certkiller -DB01 and restore it to Certkiller -DB02.
- C. Configure Certkiller -DB01 and Certkiller -DB02 as active nodes.
- D. Reproduce the user login from Certkiller -DB01 on Certkiller -DB02.

Answer: B

Explanation: Log shipping requires that the primary database be backed up and restored to the secondary database server. This can be done automatically when you configure the secondary server for log shipping, or it can be done manually.

Incorrect Answers:

A: You have already installed SQL Server 2005 on Certkiller -DB02. You must create the initial database instance when you install SQL Server 2005.

C: Failover clustering uses nodes, not log shipping.

D: Database mirroring with failover may require that you reproduce the user login from Certkiller -DB01 on Certkiller -DB02. Log shipping does not have this requirement as it is a passive recovery mechanism that does not support automatic failover.

Reference:

Microsoft SQL Server 2005 Books Online (2006), Index: log shipping [SQL Server], configuring

QUESTION 26:

You work as the database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. All servers on the Certkiller .com network run Windows Server2003 and all client computers run Windows XP Professional. The Certkiller .com network contains a SQL Server 2005 database server named Certkiller -DB01. Certkiller -DB01 hosts a database named CK_Sales that stores sales data for the company. You install two new SQL Server 2005 database server named Certkiller -DB02 and Certkiller -DB03 on the Certkiller .com network. You want to configure log shipping between the three database servers with Certkiller -DB01 being the primary server. You want users to be able to use Certkiller -DB02 and Certkiller -DB03 to run queries and reports against the CK_Sales database. You backup the CK_Sales database and all transaction logs on Certkiller -DB01. What should you do next?

- A. Leave the database in RECOVERY mode when restoring the backups of the CK_Sales database and transaction logs to Certkiller -DB02 and Certkiller -DB03.
- B. Leave the database in NORECOVERY mode when restoring the backups of the CK_Sales database and transaction logs to Certkiller -DB02 and Certkiller -DB03.
- C. Leave the database in STANDBY mode when restoring the backups of the CK_Sales database and transaction logs to Certkiller -DB02 and Certkiller -DB03.
- D. Leave the database in READ-ONLY mode when restoring the backups of the CK_Sales database and transaction logs to Certkiller -DB02 and Certkiller -DB03.

Answer: C

Explanation: STANDBY mode allows read-only access to the secondary server that can be used for read-only query processing. This can be used to reduce the load on the primary server.

Incorrect Answers:

A, D: Restore operation in SQL Server 2005 does not support RECOVERY or READ-ONLY modes.

B: The execution of queries and reports are not supported on a database that is in NORECOVERY mode.

Reference:

Microsoft SQL Server 2005 Books Online (2006), Index: log shipping [SQL Server], secondary servers

QUESTION 27:

You work as the database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. All

servers on the Certkiller .com network run Windows Server2003 and all client computers run Windows XP Professional. The Certkiller .com network contains three SQL Server 2005 database servers named Certkiller -DB01, Certkiller -DB02, and Certkiller -DB03. Certkiller -DB01 hosts a database named CK_Sales that stores sales data for the company.

Log shipping is configured on Certkiller -DB01 and Certkiller -DB02 with Certkiller -DB01 being the primary server. Certkiller .com users run queries and reports against the CK_Sales database on Certkiller -DB02. You want to store all status information for Certkiller -DB01 and Certkiller -DB02 on Certkiller -DB03.

What should you do?

- A. Configure Certkiller -DB03 as a monitor server.
- B. Configure Certkiller -DB03 as a witness server.
- C. Run the stored procedure sp_add_log_shipping_alert_job on Certkiller -DB01 and Certkiller -DB02.
- D. Run the stored procedure sp_add_log_shipping_alert_job on Certkiller -DB03.

Answer: A

Explanation: To store status information for Certkiller -DB01 and Certkiller -DB02 on Certkiller -DB03 you must configure Certkiller -DB03 as a monitor server.

Incorrect Answers:

B: Witness servers are used with mirrored databases in high-safety mode to provide automatic failover. It is not used with log shipping.

C: The sp_add_log_shipping_alert_job stored procedure is used to configure a job for monitoring log shipping and must be run on the monitor server.

D: The sp_add_log_shipping_alert_job stored procedure is used to configure a job for monitoring log shipping and must be run on the monitor server. However, you must first configure Certkiller -DB03 as a monitor server.

Reference:

Microsoft SQL Server 2005 Books Online (2006), Index: log shipping [SQL Server], about log shipping

Microsoft SQL Server 2005 Books Online (2006), Index: log shipping [SQL Server], stored procedures, sp_add_log_shipping_alert_job

QUESTION 28:

You work as the database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. All servers on the Certkiller .com network run Windows Server2003 and all client computers run Windows XP Professional. The Certkiller .com network contains three SQL Server 2005 database servers named Certkiller -DB01, Certkiller -DB02, and Certkiller -DB03. Certkiller -DB01 hosts a database named CK_Sales that stores sales data for the company.

Log shipping is configured on Certkiller -DB01, Certkiller -DB02, and Certkiller -DB03 with Certkiller -DB01 being the primary server and Certkiller -DB03 being a monitor server. Certkiller .com wants to stop using Certkiller -DB03 as a monitor server and wants you to reconfigure Certkiller -DB03 to host a database named CK_Projects. What should you do?

- A. Remove log shipping on Certkiller -DB03.
- B. Remove log shipping on Certkiller -DB01.
- C. Run the stored procedure sp_delete_log_shipping_primary_database on Certkiller -DB01.
- D. Run the stored procedure sp_delete_log_shipping_primary_secondary on Certkiller -DB03.

Answer: B

Explanation: When a monitor server is added to a log shipping configuration, the monitor server cannot be changed. If you want to change the monitor server, you must first remove log shipping. This task must be performed on the primary server.

Incorrect Answers:

A: You must remove log shipping on the primary server, not on the monitor server.

C: The sp_delete_log_shipping_primary_database stored procedure is used to remove log shipping of the primary database including backup job as well as local and remote history. However, this stored procedure must be run after you have run the sp_delete_log_shipping_primary_secondary stored procedure to remove the secondary databases.

D: The sp_delete_log_shipping_primary_secondary stored procedure removes the secondary databases but it must be run on the primary server, not the monitor server.

Reference:

Microsoft SQL Server 2005 Books Online (2006), Index: log shipping [SQL Server], about log shipping

Microsoft SQL Server 2005 Books Online (2006), Index: log shipping [SQL Server], removing

Microsoft SQL Server 2005 Books Online (2006), Index: log shipping [SQL Server], stored procedures, sp_delete_log_shipping_primary_database

Microsoft SQL Server 2005 Books Online (2006), Index: log shipping [SQL Server], stored procedures, sp_delete_log_shipping_primary_secondary

QUESTION 29:

You work as the database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. All servers on the Certkiller .com network run Windows Server 2003 and all clients run Windows XP Professional.

The Certkiller .com network contains a database server named Certkiller -DB01

that is running SQL Server 2005, Standard Edition. Certkiller -DB01 contains a database named CK_Sales which is used to store data for the Sales department. Members of the Sales department usually access CK_Sales at the start of the day. You want to improve performance of the CK_Sales database by implementing database snapshots. However, the Snapshot Agent is not available in SQL Server Management Studio. What should you do?

- A. Create the snapshot in Replication Monitor.
- B. Install and start Replication services on Certkiller -DB01.
- C. Upgrade Certkiller -DB01 to SQL Server 2005, Enterprise Edition.
- D. Open a command prompt window on Certkiller -DB01 and run the snapshot.exe utility.

Answer: C

Explanation: You cannot create snapshots in SQL Server 2005, Standard Edition as this edition does not ship with the Snapshot Agent and does not support snapshots. Only SQL Server 2005, Enterprise Edition and SQL Server 2005, Developer Edition supports snapshots.

Incorrect Answers:

A, D: You can use Replication Monitor and the snapshot.exe command-line utility to create a snapshot in SQL Server 2005, Enterprise Edition and SQL Server 2005, Developer Edition. However, SQL Server 2005, Standard Edition does not ship with the Snapshot Agent and does not support snapshots.

B: You cannot create snapshots in SQL Server 2005, Standard Edition as this edition does not ship with the Snapshot Agent and does not support snapshots.

Reference:

Microsoft SQL Server 2005 Books Online (2006), Index: features [SQL Server]

QUESTION 30:

You work as the database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. All servers on the Certkiller .com network run Windows Server 2003 and all clients run Windows XP Professional.

The Certkiller .com network contains a SQL Server 2005 database server named Certkiller -DB01. Certkiller -DB01 contains a large database named CK_Catalog which is updated every Monday and every Thursday. An online version of the CK_Catalog database is available to customers on the Certkiller .com e-Commerce Web site. Clive Wilson, the Webmaster of the Certkiller .com e-Commerce Web site is responsible for updating the online version of the CK_Catalog database. To optimize performance of the CK_Catalog online database, you create database snapshots at 5:30 PM every Monday and every Thursday. Each time you create a new snapshot, you delete the snapshot from the previous week. The two snapshots are named ck_catss_monday and ck_catss_thursday. One

Tuesday morning Clive Wilson informs you that he imported the wrong data into the CK_Catalog online database. You need to revert the CK_Catalog online database to the ck_catss_thursday snapshot and allow Clive Wilson to import the correct data.

What should you do?

- A. Restore the ck_catss_thursday snapshot in Replication Monitor.
- B. Run the Transact-SQL statement: DROP DATABASE ck_catss_monday.
- C. Run the Transact-SQL statement: RESTORE DATABASE CK_Catalog FROM DATABASE_SNAPSHOT = 'ck_catss_thursday'.
- D. Set the Restrict Access option for the online CK_Catalog database to Single.

Answer: B

Explanation: When you want to revert a database back to a snapshot, you must drop any other snapshots of the database; therefore you need to drop the snapshot that you do not want to revert to.

Incorrect Answers:

A: You cannot revert a database to a snapshot in Replication Monitor. You can only revert a database to a snapshot using the RESTORE DATABASE Transact-SQL statement.

C: When you want to revert a database back to a snapshot, you must drop any other snapshots of the database; therefore you need to drop the ck_catss_monday snapshot first before you run the RESTORE DATABASE Transact-SQL statement.

D: When you run the RESTORE DATABASE Transact-SQL statement with the FROM DATABASE_SNAPSHOT clause, SQL marks the database as 'In restore'. When a database is marked 'In restore' you do not need to set the Restrict Access option to Single.

Reference:

Microsoft SQL Server 2005 Books Online (2006), Index: features [SQL Server]

QUESTION 31:

You work as the database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. All servers on the Certkiller .com network run Windows Server 2003 and all clients run Windows XP Professional.

The Certkiller .com network hosts a database that is used to store news stories. You need to run a query that returns the number of new stories in each category and subcategory that have been added in the last 7 days. Your stored procedure must also return the name of the subcategory and the category to which it belongs.

What query should you use?

A. SELECT c.CategoryName, c.SubcategoryName, COUNT(s.SubcategoryID)
FROM Samples AS s INNER JOIN Categories AS c
ON sSubcategoryID = c.SubcategoryID

```
WHERE DateDiff("day", s.Date, GETDATE()) <= 7
GROUP BY c.CategoryName, c.SubcategoryName
ORDER BY c.CategoryName, c.SubcategoryName
B. SELECT c.CategoryName, c.SubcategoryName, COUNT(s.SubcategoryID)
FROM Samples AS s INNER JOIN Categories AS c
ON s.SubcategoryID = c.SubcategoryID
WHERE DateDiff("week", s.Date, GETDATE()) <= 1
GROUP BY c.SubcategoryName
ORDER BY c.SubcategoryName, c.CategoryName
C. SELECT c.CategoryName, c.SubcategoryName, COUNT(s.SubcategoryID)
FROM Samples AS s INNER JOIN Categories AS c
ON s.SubcategoryID = c.SubcategoryID
WHERE DateDiff("week", s.Date, GETDATE()) <= 1
GROUP BY c.CategoryName, c.SubcategoryName
ORDER BY c.CategoryName, c.SubcategoryName
D. SELECT c.CategoryName, c.SubcategoryName, COUNT(s.SubcategoryID)
FROM Samples AS s INNER JOIN Categories AS c
ON s.SubcategoryID = c.SubcategoryID
WHERE DateDiff("day", s.Date, GETDATE()) <= 7
ORDER BY c.CategoryName, c.SubcategoryName
```

Answer: A

Explanation:

The COUNT aggregate function will return the number of non-null items within the result set based on the GROUP BY clause. To only retrieve the sample chapters that have been added in the last 7 days, you should use the DateDiff function to limit the result set to 7 days.

Incorrect Answers:

B: The DateDiff function will return sample chapters that have been added either this week or last week rather than the last 7 days. Also, when the COUNT aggregate function is used, the columns referenced in the SELECT statement must appear in the GROUP BY clause or a syntax error will occur.

C: The DateDiff function will return sample chapters that have been added either this week or last week rather than the last 7 days.

D: When the COUNT aggregate function is used, the columns referenced in the SELECT statement must appear in the GROUP BY clause or a syntax error will occur.

QUESTION 32:

You work as the database administrator at an investment company named Certkiller .com. Certkiller .com has its headquarters in New York and branch offices in Boston, Atlanta, Miami, and New Orleans. The Certkiller .com network contains a SQL Server 2005 database server named Certkiller -DB01 that is located at headquarters.

Certkiller -DB01 contains a large database named CK_Investments. The

CK_Investments database stores investment data for each office. You need to implement a stored procedure that will return the average length of an investment at each office.

What Transact SQL query should you use?

- A. `SELECT o.City, AverageInvestmentPeriod =
AVG(DateDiff(day,i.StartDate,i.EndDate))
FROM Investments AS i
INNER JOIN Offices AS o
ON o.OfficeID = i.OfficeID
GROUP BY o.OfficeID`
- B. `SELECT o.City, AverageInvestmentPeriod =
AVG(DateDiff(day,i.StartDate,i.EndDate))
FROM Investments AS i
INNER JOIN Offices AS o
ON o.OfficeID = i.OfficeID
GROUP BY o.City`
- C. `SELECT o.City, AverageInvestmentPeriod = AVG(i.StartDate - i.EndDate)
FROM Investments AS i
INNER JOIN Offices AS o
ON o.OfficeID = i.OfficeID
GROUP BY o.OfficeID`
- D. `SELECT o.City, AverageInvestmentPeriod = AVG(i.StartDate - i.EndDate)
FROM Investments AS i
INNER JOIN Offices AS o
ON o.OfficeID = i.OfficeID
GROUP BY o.City`

Answer: B

Explanation: The DateDiff function can be used to return the length of an investment period and the AVG aggregate function can be used to average the length of an investment period returned by the DateDiff function. When you use the AVG aggregate function, you can only include the aggregate and columns in the GROUP BY clause; therefore you need to group by city name.

Incorrect Answers:

A: When you use the AVG aggregate function, you can only include the aggregate and columns in the GROUP BY clause; therefore you cannot group by o.CityID as the o.CityID column is not retrieved by the SELECT statement.

C: You cannot use the minus operator on a datetime or smalldatetime column.

D: You cannot use the minus operator on a datetime or smalldatetime column. Also, when you use the AVG aggregate function, you can only include the aggregate and columns in the GROUP BY clause; therefore you cannot group by o.CityID as the o.CityID column is not retrieved by the SELECT statement.

QUESTION 33:

You work as the database administrator at an investment company named Certkiller .com. The Certkiller .com network contains a SQL Server 2005 database server named Certkiller -DB01 that contains a large database named CK_Sales. The Testing.com personnel in the Accounts department require the ability to view customer order information. You need to create a query that will allow Accounts department personnel the ability to view CustomerID values for all customers as well as: the orders that are placed by individual customers, and customers that has not yet placed any orders.

There is a foreign constraint on the CustomerID column of the CK_Sales.[Order Head] table that references the CustomerID column in the CK_Sales.Customer table. You want to create a single query that returns the CustomerID and OrderID. What query should you use?

- A. `SELECT c.CustomerID, o.OrderID
FROM CK_Sales.Customer c LEFT OUTER JOIN CK_Sales.[Order Head] c
ON c.CustomerID = o.CustomerID`
- B. `SELECT o.CustomerID, o.OrderID
FROM CK_Sales.Customer c LEFT OUTER JOIN CK_Sales.[Order Head] c
ON c.CustomerID = o.CustomerID`
- C. `SELECT c.CustomerID, o.OrderID
FROM CK_Sales.Customer c RIGHT OUTER JOIN CK_Sales.[Order Head] c
ON c.CustomerID = o.CustomerID`
- D. `SELECT c.CustomerID, o.OrderID
FROM CK_Sales.Customer c JOIN CK_Sales.[Order Head] c
ON c.CustomerID = o.CustomerID`

Answer: A

Explanation: This option shows a LEFT OUTER JOIN. All rows are thus returned from the left table, CK_Sales.Customer, whether there are matching rows in the right table or not. This means that all CustomerID values from CK_Sales.Customer will be returned. In the event of the customer having an order, those OrderID values are also returned.

Incorrect Answers:

B: Although this option represents a LEFT OUTER JOIN, the result will only return values from the CK_Sales.[Order Head] table because these are the only values that are specified in the SELECT clause.

C: This option shows that it is a RIGHT OUTER JOIN. This will result in all rows in CK_Sales.[Order Head] being returned regardless of whether there are matching rows in CK_Sales.Customer or not. Due to the foreign key relationship between CK_Sales.[Order Head] and CK_Sales.Customer, there shouldn't be CustomerID values in Sales.[Order Head] that does not appear in CK_Sales.Customer.

D: Although not explicitly identified, this option represents an INNER JOIN. Only rows

meeting the qualifying condition are returned. Thus if a CustomerID does not appear in the CK_Sales.[OrderHead] table, nothing will be returned for that particular customer.

QUESTION 34:

You work as the database administrator at an investment company named Certkiller .com. Certkiller .com has its headquarters in New York and branch offices in Boston, Atlanta, Miami, and New Orleans. The Certkiller .com network contains a SQL Server 2005 database server named Certkiller -DB01 that is located at headquarters. Certkiller -DB01 contains a large database named CK_Investments. The CK_Investments database stores investment data for each office. You need to implement a parameterized stored procedure that will decrement the AvailableInvestments table when a real estate investment is secured. The stored procedure must accept the OfficeID, InvestmentTypeID, StartDate and EndDate as input parameters.

What stored procedure should you use?

- A. UPDATE AvailableInvestments
SET NumRealEstate = NumRealEstate - 1
WHERE OfficeID = @officeid
AND InvestmentTypeID = @investmenttypeid
AND date BETWEEN @startdate and @enddate
- B. UPDATE AvailableInvestments
SET NumRealEstate = NumRealEstate - 1
WHERE OfficeID = officeid
AND InvestmentTypeID = investmenttypeid
AND date BETWEEN startdate and enddate
- C. UPDATE AvailableInvestments
SET NumRealEstate = NumRealEstate - 1
WHERE OfficeID = @officeid
AND InvestmentTypeID = @investmenttypeid
AND date >= @startdate
AND date <= @enddate
- D. DECLARE @date
WHILE @date < EndDate
UPDATE AvailableInvestments
SET NumRealEstate = NumRealEstate - 1
WHERE OfficeID = @officeid
AND InvestmentTypeID = @investmenttypeid
AND date = @date
BREAK

Answer: A

Explanation: You need to use the BETWEEN keyword to update records that fall between dates.

Incorrect Answers:

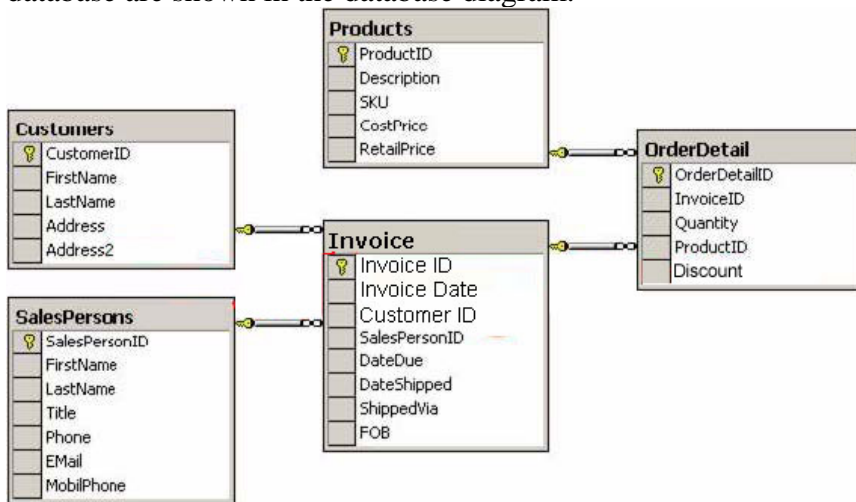
B: Input parameters require the 'at' symbol (@) as the first character.

C: You cannot use the compare operators (> < =) with a datetime or smalldatetime column.

D: Executing the change in a loop is not as efficient as executing that changes in a single statement.

QUESTION 35:

You work as the database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. All servers on the Certkiller .com network run Windows Server 2003. The Certkiller .com network contains a SQL Server 2005 database server named Certkiller -DB01. Certkiller -DB01 contains a database named CK_Sales. The tables in the CK_Sales database are shown in the database diagram.



Kara Lang is the manager of the Sales department. Kara Lang wants a report that indicates the performance of each sales person in the Sales department. She asks you to produce the report based on the SalesPersons and Invoices tables. The report must be presented in XML format.

What should you do?

- A. Run the query with the FOR XML clause.
- B. Run the query with the OPENXML statement.
- C. Use XQuery to query the CK_Sales database.
- D. Use XML DML to query the CK_Sales database.

Answer: A

Explanation: You can use the FOR XML clause in the SELECT statement to have the result set of a query returned in XML format rather than as table rows.

Incorrect Answers:

B, C: The OPENXML statement and XQuery can be used to query an XML document. It cannot be used to query a SQL database.

D: XML DML can be used to perform insert, delete and update operations on an XML document. It is not used to query a SQL database.

Reference:

Microsoft SQL Server 2005 Books Online (2006), Index: XML [SQL Server], FOR XML clause

Microsoft SQL Server 2005 Books Online (2006), Index: XML [SQL Server], OPENXML statement

Microsoft SQL Server 2005 Books Online (2006), Index: XML [SQL Server], XQuery

Microsoft SQL Server 2005 Books Online (2006), Index: XML [SQL Server], DML

QUESTION 36:

You work as the database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. Certkiller .com has its headquarters in Washington and branch offices in Detroit and Delaware. All servers on the Certkiller .com network run Windows Server 2003 and all client computers run Windows XP Professional.

The Certkiller .com network contains a SQL Server 2005 database server named Certkiller -DB01 that is located at headquarters. Certkiller -DB01 hosts a database named CK_Sales that stores sales data for the entire company. While Certkiller -DB01 is offline for maintenance purposes, all sales data are collected in XML documents. Once Certkiller -DB01 is back online, you must add the data from the XML documents to the CK_Sales database.

What should you do?

- A. Use the FOR XML clause to extract the data and store it in the database.
- B. Use the OPENXML statement to extract the data and store it in the database.
- C. Use XQuery to extract the data and store it in the database.
- D. Use XML DML to extract the data and store it in the database.

Answer: B

Explanation: The OPENXML statement to extract data from an XML document and stores in a SQL Server 2005 database.

Incorrect Answers:

A: The FOR XML clause is used in a SELECT statement to query a SQL database and format the result set in an XML document. It is not used to query a XML document.

C: XQuery can be used to query and process data in an XML document. It is not used to extract data from the XML document and store in a SQL database.

D: XML DML can be used to perform insert, delete and update operations on the XML document. It is not used to extract data from the XML document and store in a SQL database.

Reference:

Microsoft SQL Server 2005 Books Online (2006), Index: XML [SQL Server], FOR XML clause

Microsoft SQL Server 2005 Books Online (2006), Index: XML [SQL Server],

OPENXML statement

Microsoft SQL Server 2005 Books Online (2006), Index: XML [SQL Server], XQuery

Microsoft SQL Server 2005 Books Online (2006), Index: XML [SQL Server], DML

QUESTION 37:

You work as the database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. The network contains a database server named Certkiller -DB01 that hosts a database named CK_Books. The CK_Books database has a table named SampleChapters that is used to store sample chapters from books published by Certkiller .com. The schema for the SampleChapters table is shown in the following exhibit:

SampleChapters	
PK	SampleID
FK1	Book ID
FK2	AuthorID
	SampleText
FK3	SubcategoryID
	Date

The text for the sample chapters are stored in a column named SampleText. The SampleText column is defined as an xml data type column. The XML schema for the SampleText column is shown below:

```
<SampleChapter>
<BookTitle>title of book</BookTitle>
<Author>author of book</Author>
<ChapterText>
<ChapterTitle>title of chapter</ ChapterTitle >
<Intro>introduction or synopsis</Intro>
<Topic Title="topic title">topic text</Topic>
<Topic Title="next topic title">next topic text</Topic>
</ChapterText>
</SampleChapter>
```

You need to create indexes that will optimize performance for search queries that are run against the SampleText column. You need to ensure best performance for a query that return the book title, author, and introduction of a specific sample chapter. The search query is shown below:

```
SELECT SampleText.value('\SampleChapter\BookTitle[1]', 'varchar'),
SampleText.value('\SampleChapter\Author[1]', 'varchar'),
SampleText.value('\SampleChapter\ChapterText\ChapterTitle[1]',
'varchar'),
SampleText.value('\SampleChapter\ChapterText\Intro[1]', 'varchar')
FROM Samples
WHERE SampleID = @sampleID
```

You need to ensure that your solution meets Courseware Publisher's technical and business requirements.

What should you do? (Choose all that apply.)

- A. Create a PROPERTY secondary XML index on the SampleText column.
- B. Create a VALUE secondary XML index on the SampleText column.
- C. Create a nonclustered index on the AuthorID column.
- D. Create a clustered index on the SampleID column.
- E. Create a nonclustered index on the SampleID column.
- F. Create a primary XML index on the SampleText column.

Answer: A, D, F

Explanation: The SampleID column is the primary key column and will be used to locate chapter samples. You should create a clustered index on this column to improve search performance.

Keywords and concepts are stored in SampleText column which must hold XML data as the chapter samples must adhere to an XML schema. You need to use a PROPERTY secondary XML index which uses the SampleID and xml path and xml node values to create an index of the data. However, before you can create a secondary XML index, you must first create a primary XML index.

Incorrect Answers:

B: You need to use a VALUE secondary XML index as users will search for an attribute with a value such as <BookTitle>.

C: No searches are run against the AuthorID column; therefore you don't need to index the AuthorID column.

E: You should create a clustered index on the SampleID column as this provides better performance.

QUESTION 38:

You work as the database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. The network contains a database server named Certkiller -DB01 that hosts a database named CK_Books. The CK_Books database has a table named SampleChapters that is used to store sample chapters from books published by Certkiller .com. The schema for the SampleChapters table is shown in the following exhibit:

SampleChapters	
PK	SampleID
FK1	Book ID
FK2	AuthorID
FK3	SampleText
	SubcategoryID
	Date

The text for the sample chapters are stored in a column named SampleText. The SampleText column is defined as an xml data type column. The XML schema for the SampleText column is shown below:

```
<SampleChapter>
<BookTitle>title of book</BookTitle>
<Author>author of book</Author>
<ChapterText>
<ChapterTitle>title of chapter</ ChapterTitle >
<Intro>introduction or synopsis</Intro>
<Topic Title="topic title">topic text</Topic>
<Topic Title="next topic title">next topic text</Topic>
</ChapterText>
</SampleChapter>
```

You need to create a stored procedure that will return a list of sample chapters from the Samples table. The returned data must include the book title and the sample chapter title but must not return any XML tags.

What stored procedure you create?

- A. SELECT SampleText.nodes('\SampleChapter\BookTitle', 'varchar'),
SampleText.nodes('\SampleChapter\ChapterText\ChapterTitle', 'varchar')
FROM Samples
- B. SELECT SampleText.query('\SampleChapter\BookTitle'),
SampleText.query('\SampleChapter\ChapterText\ChapterTitle')
FROM Samples
- C. SELECT SampleText.value('\SampleChapter\BookTitle[1]', 'varchar'),
SampleText.value('\SampleChapter\ChapterText\ChapterTitle[1]', 'varchar')
FROM Samples
- D. SELECT SampleText.query ('child:BookTitle'),
SampleText.query ('child:child:ChapterTitle')
FROM Samples

Answer: C

Explanation: You must use the value method to query XML data in an xml data type column. The correct syntax for the value method is column_name.value('XQuery', 'data_type'). The '[1]' in the XQuery indicates that a single value is being returned.

Incorrect Answers:

A: The nodes method is used to create relational data from an xml column. It is not used to return scalar values.

B, D: The query method returns the XML tags together with the requested data.

QUESTION 39:

You work as the database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. The network contains a database server named Certkiller -DB01 that hosts a database named CK_Books. The CK_Books database has a table named SampleChapters that is used to store sample chapters from books published by Certkiller .com. The

schema for the SampleChapters table is shown in the following exhibit:

SampleChapters	
PK	SampleID
FK1	Book ID
FK2	AuthorID
FK3	SampleText
	SubcategoryID
	Date

The text for the sample chapters are stored in a column named SampleText. The SampleText column is defined as an xml data type column. Users are able to search sample chapter texts by using predefined keywords that are enclosed in XML <keyword> elements. The <keyword> elements can appear anywhere in the text. You need to create the indexes for the Sample table. What should you do? (Choose all that apply.)

- A. Create a VALUE secondary XML index on the SampleText column.
- B. Create a PATH secondary XML index on the SampleText column.
- C. Create a nonclustered index on the AuthorID column.
- D. Create a clustered index on the SampleID column.
- E. Create a nonclustered index on the SampleID column.
- F. Create a primary XML index on the SampleText column.

Answer: A, D, F

Explanation: The SampleID column is the primary key column and will be used to locate chapter samples. You should create a clustered index on this column to improve search performance.

Keywords are stored in SampleText column which must hold XML data. You need to use a VALUE secondary XML index as users will search for an attribute with a value such as <keyword>. However, before you can create a secondary XML index. You must first create a primary XML index before you can create a secondary XML index.

Incorrect Answers:

- B: You can use a PATH secondary XML index when you know the location of an element. In this scenario, keyword elements can appear anywhere in the sample chapter text.
- C: No searches are run against the AuthorID column; therefore you don't need to index the AuthorID column.
- E: You should create a clustered index on the SampleID column as this provides better performance.

QUESTION 40:

You work as the database developer for Certkiller .com. While busy designing queries for the retrieval of data from XML sources, you are writing a script that will generate an XML file to be imported into an application. This application assumed XML data includes both elements and attributes.

You need to make a choice of the most appropriate type of Transact-SQL statement to use to retrieve the data from the database.

What should you do?

- A. Make use of the OPENXML Transact-SQL statement.
- B. Make use of the sp_xml_preparedocument stored procedure.
- C. Make use of the SELECT ... FOR XML AUTO Transact-SQL statement.
- D. Make use of the SELECT ... FOR XML PATH Transact-SQL statement.

Answer: D

Explanation: The FOR XML PATH clause of the SELECT statement will allow you to use XPath to define the structure for the XML data that is returned. This structure can include both elements and attributes.

Incorrect Answers:

A: The OPENXML statement is used to insert XML data into relational tables. Not to retrieve a resultant set formatted as XML from relational tables.

B: The sp_xml_preparedocument system stored procedure is used prior to calling OPENXML.

C: The FOR XML AUTO clause of the SELECT statement will generate a result set of nested elements and as such do not make allowance for a mix of elements and attributes.

QUESTION 41:

You work as the database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. All servers on the Certkiller .com network run Windows Server 2003 and all client computers run Windows XP Professional. The Certkiller .com network contains a SQL Server 2005 database server named Certkiller -DB01 that hosts a database named CK_Projects. The CK_Projects database contains a table named Documents that contains an XML column named DocLocation. Certkiller .com users complain that queries against the DocLocation column are processed very slowly. You want to improve the performance of queries that are run against the DocLocation column by creating an index on the column.

What should you do?

- A. Create a primary index on the primary key of the Documents table and secondary index on the DocLocation column.
- B. Create a clustered index on the primary key of the Documents table and secondary index on the DocLocation column.
- C. Create a nonclustered index on the primary key of the Documents table and primary index on the DocLocation column.
- D. Create a clustered index on the primary key of the Documents table and primary index on the DocLocation column.

Answer: D

Explanation: You can create XML on XML data type columns to index all tags, values and paths in the column. This will improve query performance. The first XML index must be a primary index and the Documents table must have a clustered index on the primary key. A clustered index determines the physical ordering of the rows in a table and is required as the primary key is used to correlate XML index rows with the rows in the table.

Incorrect Answers:

A: You cannot create a unique index, clustered index or nonclustered index but not a primary index on a non-XML column. You must create a clustered index on the primary key as the primary key is used to correlate XML index rows with the rows in the table. Also, the first XML index must be a primary index. Secondary indexes can be created once the primary index has been created.

B: The first XML index must be a primary index. Secondary indexes can be created once the primary index has been created.

C: You must create a clustered index on the primary key as the primary key is used to correlate XML index rows with the rows in the table. A clustered index determines the physical ordering of the rows in a table while a nonclustered index creates a logical ordering.

Reference:

Microsoft SQL Server 2005 Books Online (2006), Index: indexes [SQL Server], creating

Microsoft SQL Server 2005 Books Online (2006), Index: indexes [SQL Server], XML

Microsoft SQL Server 2005 Books Online (2006), Index: indexes [SQL Server], clustered

Microsoft SQL Server 2005 Books Online (2006), Index: indexes [SQL Server], nonclustered

QUESTION 42:

You work as the database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. All servers on the Certkiller .com network run Windows Server 2003 and all client computers run Windows XP Professional.

The Certkiller .com network contains a SQL Server 2005 database server named Certkiller -DB01. The database instance running on Certkiller -DB01 uses the Active Directory account named ck_db01. You need to enable one of the databases to accept Simple Object Access protocol (SOAP) requests. You plan to accomplish this by creating a Hypertext Transfer Protocol (HTTP) endpoint that uses Kerberos authentication. You need to ensure that the Service Principal Name (SPN) is automatically associated with the database instance running on Certkiller -DB01. What should you do?

A. Configure the SQL Server service on Certkiller -DB01 to use the Local Admins account.

B. Configure the SQL Server service on Certkiller -DB01 to use the Local Service account.

- C. Configure the SQL Server service on Certkiller -DB01 to use the Local System account.
- D. Configure the SQL Server service on Certkiller -DB01 to use the Network Service account.

Answer: C

Explanation: To ensure that the Service Principal Name (SPN) is automatically associated with the instance running on Certkiller -DB01 you must configure the SQL Server service to use the Local Service account. A service that uses the Local System account can access all resources on the local computer as well as resources on the network.

Incorrect Answers:

- A: The SQL Server service should be run with the lowest possible privileges.
- B: The Local Service account has limited access to resources on the computer and uses a null session with anonymous authentication to access network resources.
- D: The Network Service account has limited access to resources on the computer but can access resources on the network.

Reference:

Microsoft SQL Server 2005 Books Online (2006), Index: SPNs [SQL Server]

Microsoft SQL Server 2005 Books Online (2006), Index: Service Principal Name

QUESTION 43:

You work as the database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. Some of the servers on the Certkiller .com network run Windows Server 2003 while the rest run Windows NT 4.0 Server. The client computers run Windows 98, Windows NT 4.0 Workstation and Windows XP Professional.

The Certkiller .com network contains a SQL Server 2005 database server named Certkiller -DB01. Certkiller -DB01 runs Windows Server 2003 and hosts a database named CK_Projects. You need to enable the CK_Projects database to accept Simple Object Access protocol (SOAP) requests. You plan to accomplish this by creating a Hypertext Transfer Protocol (HTTP) endpoint. You want to ensure that the most secure authentication method is used when client computers connect to the HTTP endpoint.

What should you do?

- A. Use the AUTHENTICATION = (INTEGRATED) clause in the CREATE ENDPOINT statement.
- B. Use the AUTHENTICATION = (KERBEROS) clause in the CREATE ENDPOINT statement.
- C. Use the AUTHENTICATION = (NTLM) clause in the CREATE ENDPOINT statement.
- D. Use the AUTHENTICATION = (DIGEST) clause in the CREATE ENDPOINT statement.

E. Use the AUTHENTICATION = (BASIC) clause in the CREATE ENDPOINT statement.

Answer: A

Explanation: When the authentication type of an endpoint is specified as INTEGRATED, the endpoint will respond with either Kerberos or NTLM, depending on which type the client uses in requesting authentication. No other authentication method will be used if authentication fails.

Incorrect Answers:

B: Kerberos authentication is supported by Windows 2000 and later. It is not supported by Windows 98 and Windows NT 4.0.

C: NTLM authentication is supported by Windows 98 and Windows NT 4.0. It is not as secure as Kerberos but Windows 98 and Windows NT 4.0 do not support Kerberos. Windows XP Professional can use NTLM but if integrated authentication is used, Windows 98 and Windows NT 4.0 will use NTLM and Windows XP Professional will use Kerberos.

D: Digest authentication is more secure than basic authentication but is not as secure as NTLM or Kerberos.

E: Basic authentication uses easily decoded base64-encoding and should only be used as a last resort unless the user that is granted permissions to the endpoint is a local user on the server computer itself.

Reference:

Microsoft SQL Server 2005 Books Online (2006), Index: CREATE ENDPOINT statement

Microsoft SQL Server 2005 Books Online (2006), Index: authentication [SQL Server], HTTP endpoints

QUESTION 44:

You work as the database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. All servers on the Certkiller .com network run Windows Server 2003 and all client computers run Windows XP Professional. The Certkiller .com network contains a SQL Server 2005 database server named Certkiller -DB01. Certkiller -DB01 hosts two database named CK_Staff and CK_Projects.

You create a Service Broker to enable the CK_Staff and CK_Projects databases to communicate. You create two message types named // Certkiller /Projects/Budget and // Certkiller .com/Projects/Expenses as well as the queue named // Certkiller .com/Expenses/ExpenseQueue.

You create a contract named // Certkiller .com/Expenses/ExpenseSubmission that uses the two message types and the queue. You want to create a new service named // Certkiller .com/Expenses that will listen for and accept messages from // Certkiller .com/Expenses/ExpenseSubmission by using // Certkiller .com/Expenses/ExpenseQueue.

Which CREATE SERVICE statement should you use?

A. CREATE SERVICE [// Certkiller .com/Expenses]
ON QUEUE [// Certkiller .com/Expenses/ExpenseQueue];
B. CREATE SERVICE [// Certkiller .com/Expenses]
ON QUEUE [// Certkiller .com/Expenses/ExpenseQueue]
([// Certkiller .com/Expenses/ExpenseSubmission]);
C. CREATE SERVICE [// Certkiller .com/Expenses]
([// Certkiller .com/Expenses/ExpenseSubmission]);
D. CREATE SERVICE [// Certkiller .com/Expenses]
ON QUEUE [// Certkiller .com/Expenses/ExpenseQueue]
([// Certkiller .com/Projects/Budget],
[// Certkiller .com/Projects/Expenses]);

Answer: B

Explanation: The CREATE SERVICE statement should have an ON QUEUE clause to specify which queue should be used, and must specify the contract which the service will listen for and accept messages from.

Incorrect Answers:

A: If the CREATE SERVICE statement does not specify the contract, then that service can only initiate a conversation and cannot listen and accept messages from any contract.

C: The CREATE SERVICE statement must have an ON QUEUE clause. The ON QUEUE clause specifies which queue the service can use.

D: The CREATE SERVICE statement should specify the contract that the service will listen for and accept messages from. It should not specify the messages.

Reference:

Microsoft SQL Server 2005 Books Online (2006), Index: CREATE SERVICE statement

QUESTION 45:

You work as the database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. All servers on the Certkiller .com network run Windows Server 2003 and all client computers run Windows XP Professional. The Certkiller .com network contains a SQL Server 2005 database server named Certkiller -DB01. Certkiller -DB01 hosts two database named CK_Sales and CK_Inventory.

You create a Service Broker to enable the CK_Sales and CK_Inventory databases to communicate. You want to create a new service named // Certkiller .com/Income.

What should you do first?

- A. Create the queue.
- B. Create the contract.
- C. Create a database view.
- D. Create a new database instance.

Answer: A

Explanation: The CREATE SERVICE statement must have an ON QUEUE clause to specify which queue it should use. Therefore you must create the queue first.

Incorrect Answers:

B: The CREATE SERVICE statement can specify a contract if it must listen for a contract. If the CREATE SERVICE statement does not specify the contract, then that service can only initiate a conversation.

C: You do not need to create a database view based on the two databases but you do need to create the queue that the service will use.

D: You do not need to create a new database instance but you do need to create the queue that the service will use.

Reference:

Microsoft SQL Server 2005 Books Online (2006), Index: CREATE SERVICE statement

QUESTION 46:

You work as the database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. All servers on the Certkiller .com network run Windows Server 2003 and all client computers run Windows XP Professional. The Certkiller .com network contains a SQL Server 2005 database server named Certkiller -DB01. Certkiller -DB01 hosts two database named CK_Staff and CK_Projects.

You need to create a Service Broker to enable the CK_Staff and CK_Projects databases to communicate. You create three message types using the following Transact-SQL statements:

```
CREATE MESSAGE TYPE [// Certkiller .com/Staff/Name]
```

```
VALIDATION = WELL_FORMED_XML ;
```

```
CREATE MESSAGE TYPE [// Certkiller .com/Projects/Budget]
```

```
VALIDATION = WELL_FORMED_XML ;
```

```
CREATE MESSAGE TYPE [// Certkiller .com/Projects/Completion]
```

```
VALIDATION = EMPTY ;
```

You now want to create a contract named // Certkiller .com/Project/Progress between two services that include the two message types. You want to ensure that both the sender and the target is able to use // Certkiller .com/Staff/Name message type; but that only the sender is able to use the // Certkiller .com/Projects/Budget message type, and only the target is able to use the // Certkiller .com/Projects/Completion message type. Which CREATE CONTRACT statement should you use?

A. CREATE CONTRACT [// Certkiller .com/Project/Progress]

([// Certkiller .com/Staff/Name] SENT BY ANY

[// Certkiller .com/Projects/Budget] SENT BY ANY

[// Certkiller .com/ Projects/Completion] SENT BY ANY)

B. CREATE CONTRACT [// Certkiller .com/Project/Progress]

([// Certkiller .com/Staff/Name] SENT BY ANY

[// Certkiller .com/Projects/Budget] SENT BY TARGET

[// Certkiller .com/ Projects/Completion] SENT BY INITIATOR)

C. CREATE CONTRACT [// Certkiller .com/Project/Progress]
([// Certkiller .com/Staff/Name] SENT BY ANY
[// Certkiller .com/Projects/Budget] SENT BY INITIATOR
[// Certkiller .com/ Projects/Completion] SENT BY TARGET)
D. CREATE CONTRACT [// Certkiller .com/Project/Progress]
([// Certkiller .com/Staff/Name] SENT BY TARGET
[// Certkiller .com/Staff/Name] SENT BY INITIATOR
[// Certkiller .com/Projects/Budget] SENT BY INITIATOR
[// Certkiller .com/ Projects/Completion] SENT BY TARGET)
E. CREATE CONTRACT [// Certkiller .com/Project/Progress]
([// Certkiller .com/Staff/Name] SENT BY TARGET
[// Certkiller .com/Staff/Name] SENT BY INITIATOR
[// Certkiller .com/Projects/Budget] SENT BY TARGET
[// Certkiller .com/ Projects/Completion] SENT BY INITIATOR)

Answer: C

Explanation: The

// Certkiller .com/Staff/Name message must be used by both sender and target;
therefore your CREATE CONTRACT statement must include the
[// Certkiller .com/Staff/Name] SENT BY ANY line. Only the sender must be able to
use the // Certkiller .com/Projects/Budget message type; therefore your CREATE
CONTRACT statement must include the [// Certkiller .com/Projects/Budget] SENT BY
INITIATOR. Only the target must be able to use the
// Certkiller .com/Projects/Completion message type; therefore your CREATE
CONTRACT statement must include the [// Certkiller .com/Projects/Completion]
SENT BY TARGET.

Incorrect Answers:

A: The [// Certkiller .com/Projects/Budget] SENT BY ANY,
[// Certkiller .com/Projects/Completion] SENT BY ANY lines allows both sender and target
to use the // Certkiller .com/Projects/Budget and // Certkiller .com/Projects/Completion
message types. However, only the sender must be able to use the
// Certkiller .com/Projects/Budget message type, and only the target must be able to use the
// Certkiller .com/Projects/Completion message type.

B, E: The [// Certkiller .com/Projects/Budget] SENT BY TARGET,
[// Certkiller .com/Projects/Completion] SENT BY INITIATOR lines allow the target to use
the // Certkiller .com/Projects/Budget and the sender to use the
// Certkiller .com/Projects/Completion message types. However, the sender must be able to
use the // Certkiller .com/Projects/Budget message type, and the target must be able to use
the // Certkiller .com/Projects/Completion message type.

D: The [// Certkiller .com/Staff/Name] SENT BY TARGET,
[// Certkiller .com/Staff/Name] SENT BY INITIATOR, line allows both sender and
target to use the // Certkiller .com/Projects/Name message types. However, the correct
syntax to use here is [// Certkiller .com/Projects/Budget] SENT BY ANY.

Reference:

Microsoft SQL Server 2005 Books Online (2006), Index: CREATE CONTRACT statement

QUESTION 47:

You work as the database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. All servers on the Certkiller .com network run Windows Server 2003 and all client computers run Windows XP Professional. The Certkiller .com network contains a SQL Server 2005 database server named Certkiller -DB01 that hosts a database named CK_Data.

You install a new SQL Server 2005 database server named Certkiller -DB02. You create a new database named CK_Sales on Certkiller -DB02. You must populate the CK_Sales database with data from several tables in the CK_Data database.

You want to use the bcp command line utility to export the data to a text file. You want to accomplish this task as quickly as possible.

What should you do?

- A. Run the bcp command against each table that contains the required data.
- B. Run the bcp command against each column that contains the required data.
- C. Create a view based on the table columns that contain the required data and run the bcp command against resulting view.
- D. Run a query to return the required data and run the bcp command against the query's return set.

Answer: D

Explanation: The bcp command can be run against a table, a view, or a query.

Running the bcp utility against the result set of a query will require the least amount of time.

Incorrect Answers:

A: You can run the

bcp command against a table but you would need to run it against each table that contains the required data. You should first create a query that returns the required data and run the bcp command against the result set from the query.

B: The bcp command can be run against a table, a view, or a query. It cannot be run against a column.

C: You can run the bcp command against a table view but it would require less time to first create a query that returns the required data and run the bcp command against the result set from the query.

Reference:

Microsoft SQL Server 2005 Books Online (2006), Index: bcp utility [SQL Server], about bcp utility

QUESTION 48:

You work as the database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. All client computers run Windows XP Professional. The Certkiller .com network contains an Oracle 9i database server named Certkiller -DB01 that hosts a database named CK_Data. The CK_Data database consists of five comma-delimited files. Certkiller .com wants to migrate the CK_Data database to SQL Server 2005. You install a new SQL Server 2005 database server named Certkiller -DB02. You want to create a new database named CK_Data and all the required tables on Certkiller -DB02. You must import the data into the CK_Data database on Certkiller -DB02. You want to accomplish this task as quickly as possible. What should you do?

- A. Use the bcp command to import the required data.
- B. Use a BULK INSERT statement to import the required data.
- C. Use a SELECT statement with an INTO clause to import the required data.
- D. Use the Data Transformation Services (DTS) to import the required data.

Answer: B

Explanation: The bcp command and the BULK INSERT statement can be used to import data from a text file into a SQL Server 2005 database; however, the BULK INSERT statement provides the best performance.

Incorrect Answers:

A: The bcp command can be used to import data from a text file into a SQL Server 2005 database; however, the BULK INSERT statement provides the best performance.

C: The SELECT statement can be used to retrieve data from a SQL Server 2005 database. It cannot be used to retrieve data from a text file.

D: Data Transformation Services (DTS) can be used to import data from a single text file but BULK INSERT statement provides the best performance when data from multiple text files must be imported.

Reference:

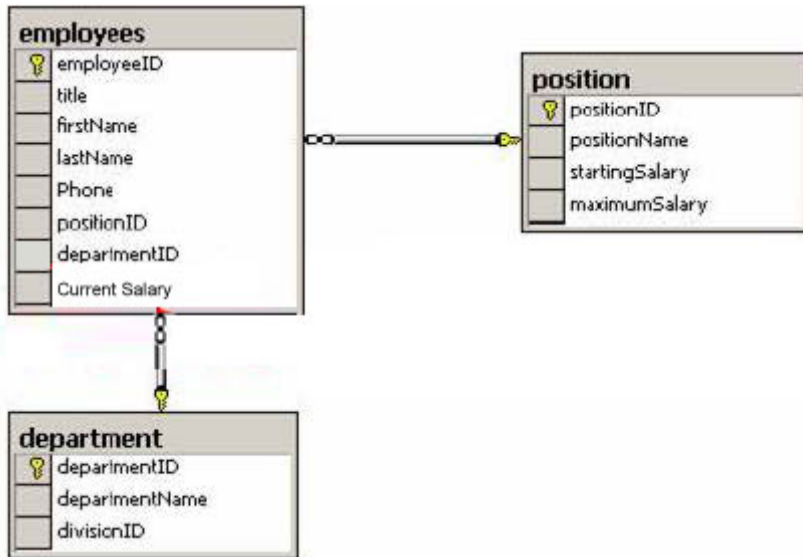
Microsoft SQL Server 2005 Books Online (2006), Index: bcp utility [SQL Server], about bcp utility

Microsoft SQL Server 2005 Books Online (2006), Index: BULK INSERT statement

Microsoft SQL Server 2005 Books Online (2006), Index: Bulk Insert task

QUESTION 49:

You work as the database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. All client computers run Windows XP Professional. The Certkiller .com network contains a SQL Server 2005 database server named Certkiller -DB01 that hosts a database named CK_Data. The tables in the CK_Data database are shown in the following database diagram.



You create a new database named CK_Staff on Certkiller -DB01. You want to transfer the list of employees from the CK_Data database to the CK_Staff database. You export the Employees table to a text file. The data must be imported as quickly as possible into the following table.



What should you do?

- A. Create an Integration Services task and specify a Copy Column Transformation.
- B. Create a Transform Data task and specify an ActiveX Script transformation.
- C. Create a Transform Data task and specify a Read File transformation.
- D. Create a Bulk Insert task and specify an XML format file.

Answer: D

Explanation: The best option to import the data as quickly as possible is to use a Bulk Insert task. However, bulk inserts do not support data transformations but a format file can be used with a bulk insert task. You can specify the columns that must be imported, and the order in which they must be imported in the format file.

Incorrect Answers:

A: The Copy Column Transformation is used to transform data from one table as it is being transferred to a new table. It does not copy data from a text file.

B, C: SQL Server 2005 does not support Transform Data tasks.

Reference:

Microsoft SQL Server 2005 Books Online (2006), Index: BULK INSERT statement, format files

QUESTION 50:

You work as the database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. You need to import data from several comma-delimited lists into a SQL Server 2005 database server named Certkiller -DB04. You need to use a SQL Server Integration Services (SSIS) package to accomplish this task. To this end you are designing the Data Flow elements that will define the importing of data from comma-delimited lists. You need to make a choice regarding the Data Flow elements that should be used in the Source - Data Flow step. What should you do?

- A. Use the DataReader Data Flow element.
- B. Use the Bulk Insert Data Flow element.
- C. Use the Flat File Data Flow element.
- D. Use the Raw File Data Flow element.
- E. Use the SQL Server Data Flow element.

Answer: C

Explanation: Data Flow defines sources, transformations and destinations. In this case you should make use of the Flat File source element. The Flat File element is used when importing data from a comma-delimited text file.

Incorrect Answers:

A: The DataReader source is used when importing data to an application that makes use of Active Data Objects (ADO.NET). In this scenario there is not need for ADO.NET connectivity.

B: Bulk Insert task can be added to a Data Flow, however, it does not allow for transformations.

D: Raw File is used when importing data to a column and the data does not require formatting. It can only be used as a destination in this case.

E: The SQL Server destination is used when you need to transform data and import it into SQL: Server on the computer where the package is run.

QUESTION 51:

You work as the database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. Certkiller .com hosts an online analytical processing (OLAP) system for a customer. The customer requested that Certkiller .com provide data transformation and load support. They prefer the data load to be schedules to occur at noon every day. You need to accommodate this customer. What should you do?

- A. Create an Integration Services package to provide data transformation and load support.

Run the package using the D T ExecUI command.

B. Create an Integration Services package to provide data transformation and load support.

Run the package using SQL Server Agent.

C. Create an Integration Services package to provide data transformation and load support.

Run the package using SQL Server Import and Export Wizard.

D. Create an Integration Services package to provide data transformation and load support.

Run the package using Business Intelligence Development Studio.

Answer: B

Explanation: The SQL Server Integration Services (SSIS) will meet the requirements. It is a platform for data integration and workflow solutions. It can be used to extract the data from the purchased lists, transform the data as required, validate against the Customers table, and load the data into the Contacts table. Further you need to make use of SQL Server Agent to run the package by creating a SQL Server Agent job and then executing the package as a job step. SQL Server Agent jobs can be scheduled for execution. This includes the ability to schedule recurring execution.

Incorrect Answers:

A: This is possible, but the DT ExecUI command does not make provision for scheduling recurring execution.

C: Though it is possible to make use of SQL Server Import and Export Wizard, it does not allow one to schedule recurring execution.

D: Though it is possible to make use of Business Intelligence Development Studio, it does not allow one to schedule recurring execution.

QUESTION 52:

You work as the database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. You need to import data from several comma-delimited lists into a SQL Server 2005 database server named Certkiller -DB02. You need to use a SQL Server Integration Services (SSIS) package to accomplish this task. To this end you are designing the Data Flow elements that will define the importing of data from comma-delimited lists. You need to make a choice regarding the Data Flow elements that should be used in the Destination - Data Flow step.

What should you do?

A. Use the Raw File Data Flow element.

B. Use the DataReader Data Flow element.

C. Use the Bulk Insert Data Flow element.

D. Use the Flat File Data Flow element.

E. Use the SQL Server Data Flow element.

Answer: E

Explanation: Data Flow defines sources, transformations and destinations. In this case you should make use of the SQL Server destination. The SQL Server destination is used when you need to transform data and import it into SQL: Server on the computer where the package is run.

Incorrect Answers:

A: Raw File is used when importing data to a column and the data does not require formatting.

B: The DataReader source is used when exporting data to an application that makes use of Active Data Objects (ADO.NET). In this scenario there is no need for ADO.NET connectivity.

C: Bulk Insert task can be added to a Data Flow, however, it does not allow for transformations.

D: The Flat File element is used when importing data from a comma-delimited text file.

QUESTION 53:

You work as the database administrator at Certkiller .com. Certkiller .com has its headquarters in New York, and branch offices in Boston, Chicago and Atlanta. Each office currently maintains its own database for customers in a single Microsoft Access table.

Certkiller .com wants to centralize the Microsoft Access databases in a single SQL Server 2005 database named Certkiller -DB01 that is located in the New York office. You enable the local administrator at each office to import data from the Microsoft Access databases into Certkiller -DB01. You need to ensure that the local administrator at each office has the minimum permissions necessary to perform the operation.

What should you do? (Each correct answer presents part of the solution. Choose TWO.)

- A. Use BULK INSERT to import the data to SQL Server 2005.
- B. Add the local administrators to the db_datawrite role.
- C. Add the local administrators to the db_dtsltduser role.
- D. Add the local administrators to the db_dtsoperator role.
- E. Use SQL Server Integration Service (SSIS) to create a package.

Answer: D, E

Explanation:

The legacy data must be transformed; SSIS allows you to transform and import data to SQL Server 2005. The local administrator needs permission to run the SSIS package; therefore you need to add them to the db_dtsoperator role.

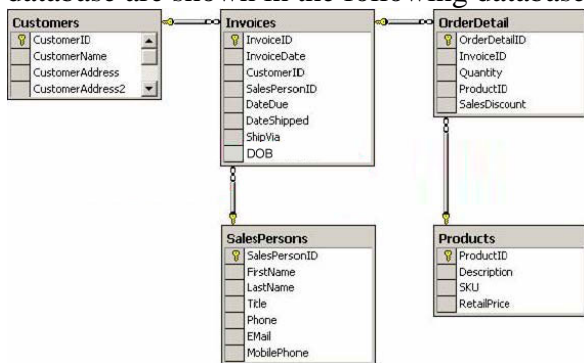
Incorrect Answers:

A: BULK INSERT does not transform data.

B, C: The db_dtsltduser role and db_datawrite role gives the administrator too much permissions.

QUESTION 54:

You work as the database administrator at Certkiller .com. All database servers run Windows Server 2003 and SQL Server 2005. All client computers run Windows XP Professional with Service Pack 2. The Certkiller .com network contains a database server named Certkiller -DB02 that runs SQL Server 2005 Enterprise Edition. Certkiller -DB02 hosts a database named CK_Sales. The tables in the CK_Sales database are shown in the following database diagram.



Certkiller .com has several Sales Representatives that use portable client computers. These portable client computers run SQL Server 2005 Express Edition. Certkiller .com has developed a new sales tracking application named salesTrack. The salesTrack application provides access to the Customers table in the CK_Sales database. Several Sales Representatives use the salesTrack application as part of Certkiller .com's Customer Relationship Management (CRM) strategy. You need to ensure that the Sales Representatives can interact with the CK_Sales database while offline. The Sales Representatives must be able to access information about their customers while offline; must be able to add new customers while offline; and must be able to synchronize with Certkiller -DC02 on demand using Web synchronization.

What should you do? (Each correct answer presents part of the solution. Choose two.)

- A. Implement snapshot replication.
- B. Implement merge replication.
- C. Implement transactional replication.
- D. Configure the portable client computers for pull subscriptions.
- E. Configure the portable client computers for push subscriptions.

Answer: B, D

Explanation: To allow offline access to the CK_Sales database and synchronization on demand using Web synchronization, you need to implement merge replication. To enable the portable client computers to synchronize with Certkiller -DC02 on demand using Web synchronization you must configure them for pull subscriptions.

During synchronization changes are propagated to the publisher and to other subscribers.

Incorrect Answers:

A: You cannot use snapshot replication to replicate transactions from multiple computers to a single database server. Also, with snapshot replication the entire database is sent during synchronization.

C: Transactional replication requires that data be entered at the Publisher. In this scenario, data will be entered at the subscribers.

E: Push replication does not support synchronization on demand from the subscriber. Only the publisher determines when synchronization occurs.

Reference:

Microsoft SQL Server 2005 Books Online (2006), Index: merge replication [SQL Server replication], about merge replication

Microsoft SQL Server 2005 Books Online (2006), Index: snapshot replication [SQL Server], about snapshot replication

Microsoft SQL Server 2005 Books Online (2006), Index: transactional replication, about transactional replication

Microsoft SQL Server 2005 Books Online (2006), Index: push subscriptions [SQL Server replication]

Microsoft SQL Server 2005 Books Online (2006), Index: pull subscriptions [SQL Server replication]

QUESTION 55:

You work as the database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. Certkiller .com has its headquarters in Washington and branch offices in Detroit and Delaware. All servers on the Certkiller .com network run Windows Server 2003 and all client computers run Windows XP Professional.

The Certkiller .com network contains four SQL Server 2005 database servers named Certkiller -DB01, Certkiller -DB02, Certkiller -DB03, and Certkiller -DB04.

Certkiller -DB01 and Certkiller -DB02 are located at headquarters while Certkiller -DB03 is located at the Detroit branch office and Certkiller -DB04 is located at the Delaware branch office. You need to implement a new database named OrdersDB to track sales orders from all three offices. You need to design replication for the OrdersDB database. You decide to implement updateable subscriptions with Certkiller -DB01 as a Publisher, Certkiller -DB02 as a Distributor, and Certkiller -DB03 and Certkiller -DB04 as Subscribers. You need to ensure that use the immediate update mode.

What should you do? (Each correct answer represents part of the solution. Choose two.)

- A. On Certkiller -DB01, create a Windows account for the Queue Reader Agent.
- B. On Certkiller -DB01, install and configure the Microsoft Distribution Coordinator (MS DTC).
- C. On Certkiller -DB02, install and configure the Microsoft Distribution Coordinator

(MS DTC).

D. On Certkiller -DB03 and Certkiller -DB04, install and configure the Microsoft Distribution Coordinator (MS DTC).

Answer: B, D

Explanation: To use immediate update mode, you need to install and configure the Microsoft Distribution Coordinator (MS DTC) on the Publisher and the Subscribers.

Incorrect Answers:

A: The Queue Reader Agent is not required when you use immediate update mode. It is only required when you use queued update mode.

C: The Microsoft Distribution Coordinator (MS DTC) should be installed and configured on the Publisher and the Subscribers. It is not required on the Distributor.

Reference:

Microsoft SQL Server 2005 Books Online (2006), Index: replication [SQL Server]

Microsoft SQL Server 2005 Books Online (2006), Index: immediate update mode [OLE DB]

Microsoft SQL Server 2005 Books Online (2006), Index: MS DTC

QUESTION 56:

You work as the database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. Certkiller .com has its headquarters in Washington and five field offices in the city. All servers on the Certkiller .com network run Windows Server 2003, all database servers run SQL Server 2005 and all client computers run Windows XP Professional. The Certkiller .com network contains a database server named Certkiller -DB01 that is located at headquarters. Certkiller -DB01 hosts a database named CK_Repairs.

Certkiller .com users in the Purchasing department order parts that technicians use to repair customer's computers at customer sites. The parts are shipped to the field offices. At present Certkiller .com tracks the usage of parts through data entered on the customer's invoice. Each field office has a client computer that runs SQL Server 2005 Standard Edition that is used to capture invoice data.

Certkiller .com wants to implement a new inventory tracking application for the parts. The new application will access a table named _Parts in the CK_Repairs database. You need to enable users at headquarters and at the various field offices to update the _Parts table. You decide to implement merge replication with CertKiller-DB01 as the Publisher and the SQL Server 2005 Standard Edition computers at the field offices as the subscribers.

To minimize the usage of disk space on the Subscribers, you need to ensure that only the Publisher is used to store conflict records.

What should you do?

A. Run the sp_addmergepublication stored procedure with the @conflict_logging

argument on Certkiller -DB01.

B. Run the sp_addmergesubscription stored procedure with the @publication argument on Certkiller -DB01.

C. Run the sp_addmergepublication stored procedure with the @centralized_conflicts argument on Certkiller -DB01.

D. Run the sp_addmergesubscription stored procedure with the @centralized_conflicts argument on Certkiller -DB01.

Answer: A

Explanation: The @conflict_logging = 'publisher' argument of the sp_addmergepublication stored procedure specifies that conflict records should be stored only on the Publisher.

Incorrect Answers:

A, D: When sp_addmergesubscription is run on the Publisher, it is used to register a subscriber. The sp_addmergesubscription stored procedure does not define the placement of the conflict records.

C: The centralized_conflicts = 'true' argument of the sp_addmergepublication stored procedure specifies that conflict records should be stored only on the Publisher. However, this argument is supported for backward compatibility with SQL Server 2000. The @conflict_logging = 'publisher' argument is the preferred argument for SQL Server 2005.

Reference:

Microsoft SQL Server 2005 Books Online (2006), Index: sp_addmergepublication

Microsoft SQL Server 2005 Books Online (2006), Index: sp_addmergesubscription

QUESTION 57:

You work as the database administrator at Certkiller .com. The Certkiller .com network contains a SQL Server 2005 database server named Certkiller -DB01 that runs on a Windows Server 2003 computer. Certkiller -DB01 hosts a database named CK_Sales that is used to store sales information for the company.

You create a maintenance job named DB_Job for Certkiller -DB01 and configure the job to run when the processor utilization on the server drops below 15%. A week later you notice that the scheduled maintenance job has not yet run. You open System Monitor on Certkiller -DB01 and notice that the processor utilization is currently at approximately 5%. You need to ensure that the maintenance job runs when the processor utilization drops below 15%.

What should you do? (Each correct answer represents part of the solution. Choose two.)

A. Open the DB_Job Properties dialog box in SQL Server Management Studio.

B. Open the SQL Server Agent Properties dialog box in SQL Server Management Studio.

C. Open the DB_Job Properties dialog box in SQL Server Configuration Manager.

D. Modify the advanced settings for each job step.

E. Modify the idle CPU condition settings.

Answer: B, E

Explanation: Maintenance jobs run under the SQL Server Agent. You must configure the CPU condition settings for the SQL Server Agent on the Advanced tab of the SQL Server Agent Properties dialog box. The CPU condition settings define the threshold and length of time that the SQL Server Agent uses to determine if the CPU is idle.

Incorrect Answers:

A, C: The SQL Server Agent determines when the CPU is idle. Therefore you should configure the threshold and length of time that the SQL Server Agent uses to determine if the CPU is idle. This can be done on the Advanced tab of the SQL Server Agent Properties dialog box, not the Properties dialog box of the job.

D: You must configure the CPU condition settings for the SQL Server Agent on the Advanced tab of the SQL Server Agent Properties dialog box. The CPU condition settings define the threshold and length of time that the SQL Server Agent uses to determine if the CPU is idle.

Reference:

Microsoft SQL Server 2005 Books Online (2006), Index: SQL Server Agent, CPU idle conditions

Microsoft SQL Server 2005 Books Online (2006), Index: indexes [SQL Server], reorganizing

QUESTION 58:

You work as the database administrator at Certkiller .com. The Certkiller .com network contains two SQL Server 2005 database servers named Certkiller -DB01 and Certkiller -DB02. Certkiller -DB01 hosts a database named CK_Sales that is used to store sales information for the company, while Certkiller -DB02 is used for reporting purposes only.

You want to be notified by e-mail when ever a syntax error event occurs on the CK_Sales database but you want to keep the number of e-mails to a minimum. You configure SQL Server Agent Mail and you define a SQL operator object on Certkiller -DB01 for your administrative user account.

What should you do? (Each correct answer represents part of the solution. Choose three.)

- A. Assign the alert to an operator.
- B. Create an alert for error number 170 on Certkiller -DB01.
- C. Create an alert for error number 170 on Certkiller -DB02.
- D. Configure the alert to Net Send the operators.
- E. Configure the alert to e-mail the operators.
- F. Create an alert for messages that contain the name ' Certkiller -DB01'.

Answer: A, B, E

Explanation: You need to create an alert for error number 170 on Certkiller -DB01. Error number 170 denotes a syntax error. The alert needs to be assigned to one or more operators and must be configured to notify the operator by e-mail.

Incorrect Answers:

C: Alerts should be created for the production database that is hosted on Certkiller -DB01. Certkiller -DB02 is used for reporting purposes only and does not require an alert.

D: You want the alert to notify by e-mail not by Net Send.

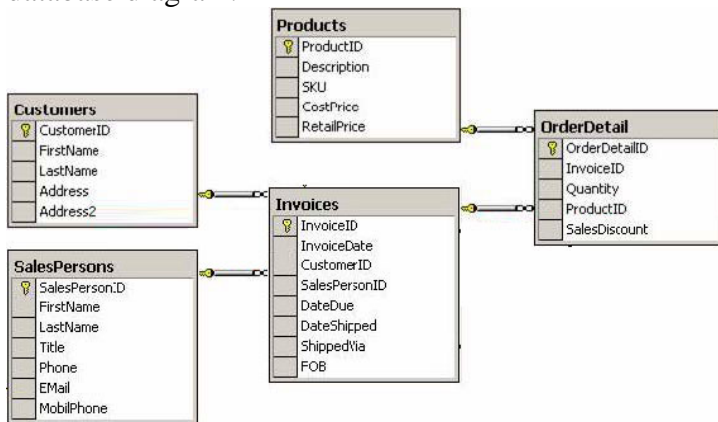
F: An alert for messages that contain the name ' Certkiller -DB01' would trigger for all events on Certkiller -DB01. You only want to be notified of syntax error events.

Reference:

Microsoft SQL Server 2005 Books Online (2006), Index: alerts [SQL Server], about alerts

QUESTION 59:

You work as the database administrator at Certkiller .com. The Certkiller .com network contains a SQL Server 2005 database server named Certkiller -DB01 that runs on a Windows 2000 Server computer. Certkiller -DB01 hosts a database named CK_Sales. The tables in the CK_Sales database are shown in the following database diagram.



An index named PK_Customers_CustomerID has been created on the Customers table. The index is used extensively throughout the day. Over the last few months, several new pages have been added to the index. Certkiller .com users complain that the performance queries against the index is very slow. You suspect that fragmentation is affecting query performance. You want to reduce fragmentation. You need to ensure that the index remain online and that no completed work is lost if the operation is interrupted.

What should you do?

- A. Run the ALTER INDEX PK_Customers_CustomerID ON CK_Sales.Customers REBUILD WITH (ONLINE = ON) Transact-SQL statement.
- B. Run the DBCC INDEXDEFRAG Transact-SQL statement.
- C. Run the ALTER INDEX PK_Customers_CustomerID ON CK_Sales.Customers

REORGANIZE Transact-SQL statement.

D. Drop and recreate the PK_Customers_CustomerID index.

Answer: C

Explanation: To ensure that no completed work is lost if the operation is interrupted, you must perform the operation on the index while it is online. The ALTER INDEX REORGANIZE statement is used to reduce fragmentation of an index while the index is online.

Incorrect Answers:

A: The ALTER INDEX REBUILD WITH ONLINE statement can be used to reduce fragmentation of an index while the index is online but data loss will occur should the operation be interrupted.

B: The DBCC INDEXDEFRAG can be used to defragment indexes on a table. However, this statement is supported for backward compatibility and is not recommended. The ALTER INDEX statement is preferred.

D: When you DROP an index, the index is offline. You need to ensure that the index remains online.

Reference:

Microsoft SQL Server 2005 Books Online (2006), Index: ALTER INDEX statement

Microsoft SQL Server 2005 Books Online (2006), Index: indexes [SQL Server], reorganizing

QUESTION 60:

You work as the database administrator at Certkiller .com. All servers on the Certkiller .com network run Windows Server 2003 and all database servers run SQL Server 2005. The Certkiller .com network contains a database server named Certkiller -DB01. Certkiller -DB01 hosts a database named CK_Sales that stores sales data for the company. The CK_Sales database is also used for online transaction processing (OLTP).

Several tables in the CK_Sales database have more than one million rows. You have created a clustered index on each table in the CK_Sales database as well as a nonclustered index on each table that is used frequently when the database was deployed. These indexes were based on projected query requirements.

Certkiller .com users complain that query performance has degraded over time. You want to optimize performance for the clustered and nonclustered indexes. You run the Index Tuning Wizard to generate index recommendations but no changes are recommended for the indexes. You suspect that the indexes are fragmented. You need to optimize the indexes and ensure that fragmentation of the indexes do not occur in the future.

What should you do?

A. Run ALTER INDEX statement with the REORGANIZE option and create a Database Consistency Check scheduled maintenance task.

B. Run ALTER INDEX statement with the REORGANIZE option and create a

Reorganize Index scheduled maintenance task.

C. Run ALTER INDEX statement with the REBUILD option and create a Rebuild Index scheduled maintenance task.

D. Run ALTER INDEX statement with the REBUILD option and create a Database Consistency Check scheduled maintenance task.

Answer: D

Explanation: The tables have clustered and nonclustered indexes; therefore you should run the ALTER INDEX statement with the REBUILD option to address the problem immediately. You should also create a Rebuild Index scheduled maintenance task to reduce the possibility of this problem reoccurring in the future. Clustered indexes affect the ordering of the table so, when clustered index is fragmented, the table is also fragmented and the index has to be rebuilt rather than reorganized.

Incorrect Answers:

A: The tables have clustered and nonclustered indexes. Clustered indexes affect the physical ordering of the data in the table. Thus, when clustered index is fragmented, the table is also fragmented and the index has to be rebuilt rather than reorganized. You should also create a Rebuild Index scheduled maintenance task to reduce the possibility of this problem reoccurring in the future. A Database Consistency Check scheduled maintenance task detects problems with the database but does not correct them.

B: The tables have clustered and nonclustered indexes. Clustered indexes affect the physical ordering of the data in the table. Thus, when clustered index is fragmented, the table is also fragmented and the index has to be rebuilt rather than reorganized.

D: You should run the ALTER INDEX statement with the REBUILD option to address the problem immediately but you need to create a Rebuild Index scheduled maintenance task to reduce the possibility of this problem reoccurring in the future. A Database Consistency Check scheduled maintenance task detects problems with the database but does not correct them.

QUESTION 61:

You work as the database administrator at Certkiller .com. The Certkiller .com network contains a SQL Server 2005 database server named Certkiller -DB01 that runs on a Windows Server 2003 computer. Certkiller -DB01 hosts a database named CK_Staff. The Full Recovery Model is implemented for the CK_Staff database.

Certkiller .com acquires another company named TestLabs.com. You import staff data from the new company into the CK_Staff database. Soon after ward you notice that Certkiller -DB01 is performing quite poorly. You discover that the transaction logs for the CK_Staff database has increased by almost five times in size and has consumed most of the free disk space. You need to recover disk space as soon as possible.

What should you do?

- A. Backup the transaction log files.
- B. Switch to the Simple Recovery Model.
- C. Run the DBCC SHRINKFILE Transact-SQL statement.
- D. Truncate the transaction log files.

Answer: C

Explanation: The DBCC SHRINKFILE statement is used to reduce the log files to a specified size.

Incorrect Answers:

A: Backing up the log file clears the log file but does not reduce the physical size of the transaction log file.

B: Switching to the Simple Recovery Model will not reduce the physical size of the transaction log file.

D: Truncation reduces the logical size of the transaction log file but it does not reduce the physical size of the transaction log file.

Reference:

Microsoft SQL Server 2005 Books Online (2006), Index: DBCC SHRINKFILE statement

Microsoft SQL Server 2005 Books Online (2006), Index: transaction log backups [SQL Server]

QUESTION 62:

You work as the database administrator at Certkiller .com. The Certkiller .com network contains a SQL Server 2005 database server named Certkiller -DB01 that runs on a Windows Server 2003 computer. Certkiller -DB01 hosts a database named CK_Orders. Several data capturers enter data into the CK_Orders database everyday.

Certkiller .com wants you to implement a disaster recovery plan for the CK_Orders database. You need to ensure that data can be easily recovered. You plan to perform a full backup every Saturday at 2:00 P.M.

What should you do?

- A. Open the Windows Backup utility on Certkiller -DB01. Configure a Full backup to run weekly on Saturdays at 2:00 P.M.
- B. Run the BACKUP DATABASE Transact-SQL statement on the CK_Orders database.
- C. Open the Database Backup dialog box from the CK_Orders database in the SQL Enterprise Manager console. Configure a Full backup to run weekly on Saturdays at 2:00 P.M.
- D. Open the Database Backup dialog box from the CK_Orders database in SQL Management Studio. Configure a Full backup to run weekly on Saturdays at 2:00 P.M.

Answer: D

Explanation: You configure a Full backup of a database in SQL Management

Studio. You can accomplish this by right-clicking the appropriate database and selecting Tasks and then Back Up from the context menu.

Incorrect Answers:

A: The Windows Utility can be used to backup data on a computer. It cannot be used to backup a database.

B: The BACKUP DATABASE statement can be used to backup a database but it cannot be used to schedule a recurring backup job.

C: SQL Server 2005 does not support the SQL Enterprise Manager console. SQL Server 2005 uses the SQL Management Studio.

Reference:

Microsoft SQL Server 2005 Books Online (2006), Index: database backups [SQL Server]

Microsoft SQL Server 2005 Books Online (2006), Index: BACKUP DATABASE statement

QUESTION 63:

You work as the database administrator at Certkiller .com. The Certkiller .com network contains a SQL Server 2005 database server named Certkiller -DB01 that runs on a Windows Server 2003 computer. Certkiller -DB01 hosts a database named CK_Orders. Several data capturers enter data into the CK_Orders database everyday. The Simple Recovery Model is implemented for the CK_Orders database. A Full backup of the CK_Orders database is performed every Saturday at 2:00 P.M.

You are concerned that data loss may occur should Certkiller -DB01 suffer a hard disk failure. You decide to implement additional backups of the CK_Orders database on a daily basis. You want to minimize the amount of time required for the backups as well as the time required to restore the database.

What should you do?

A. Perform differential backups of the CK_Orders database on every week day.

B. Switch the CK_Orders database to the Full Recovery Model.

C. Perform a Full backup of the CK_Orders database on every week day.

D. Perform incremental backups of the CK_Orders database on every week day.

Answer: A

Explanation: To minimize the amount of time required to perform the backups, you need to perform differential backups on the week days. Differential backup ups only backs up the data that has changed since the last Full backup.

Incorrect Answers:

B: The recovery model will not affect the time required for the backup or the restore.

C: Restoring data from transaction log backups will not minimize restore times.

D: SQL Server 2005 does not support incremental backups.

Reference:

Microsoft SQL Server 2005 Books Online (2006), Index: database backups [SQL Server]

QUESTION 64:

You work as the database administrator at Certkiller .com. The Certkiller .com network contains a SQL Server 2005 database server named Certkiller -DB01 that runs on a Windows Server 2003 computer. Certkiller -DB01 hosts a database named CK_Projects. The Simple Recovery Model is implemented for the CK_Projects database.

You need to implement a disaster recovery plan for the CK_Projects database. Your disaster recovery plan must ensure that data can be backed up as quickly as possible and that data can be recovered as quickly as possible following a disaster. You must also ensure that no more than three hours data is at risk at any one time. You plan to perform a full backup every Saturday at 12:00 P.M. You must now plan the backup schedule that will run during the week.

What should you do?

- A. Schedule a transaction log backup every three hours, Monday through Friday.
- B. Schedule a differential backup at 12:00 P.M., Monday through Friday and schedule a transaction log backup every three hours between differential backups.
- C. Schedule a differential backup every three hours, Monday through Friday.
- D. Schedule a full backup at 12:00 P.M., Monday through Friday and schedule a differential backup every three hours between differential backups.

Answer: C

Explanation: To ensure backups can be performed quickly, you need to implement differential backups every three hours. With differential backup ups, only the data that changed since the last Full backup will be saved. This means that you only need the last full backup and the last differential backup to restore the data.

Incorrect

Answer:

A, B: The Simple Recovery Model does not transaction log the backups as it truncates the transaction logs before the transaction logs are backed up. You must change the Recovery Model to Full or Bulk-logged in order to backup the transaction logs.

D: Performing full backups every day would increase backup times.

QUESTION 65:

You work as the database administrator at Certkiller .com. All servers on the Certkiller .com network run Windows Server 2003 and all database servers run SQL Server 2005. The Certkiller .com network contains a database server named Certkiller -DB01. Certkiller -DB01 hosts a database named CK_Sales that stores sales and product data. The Simple Recovery Model is implemented for the CK_Sales database. A Full backup of the CK_Sales database is performed every Saturday at 8:00 P.M.

The sales data is updated frequently throughout the day. Over the last six months the CK_Sales database has grown considerably. At present the Full backup requires

6 hours to complete. You are concerned that data loss may occur should Certkiller -DB01 suffer a hard disk failure. You decide to implement additional backups of the CK_Sales database on a daily basis. However, the backup may only run between midnight and 4:00 A.M. on week nights.

What should you do? (Each correct answer presents part of the solution. Choose two.)

- A. Schedule a Full backup of the CK_Sales database at 12:00 P.M. every week night.
- B. Schedule a differential backup of the CK_Sales database at 12:00 P.M. every week night.
- C. Create two filegroups. Backup a different filegroup every other week night.
- D. Backup the transaction log every week night.

Answer: C, D

Explanation:

Creating two filegroups and backing up a different filegroup every other week night will allow you to backup the database within the given time frame during the week. You should also backup the transaction logs every night to ensure full recoverability as only one filegroup is backed up each night.

Incorrect Answers:

A: A Full backup requires six hours to complete but the backup can only run from 12:00 P.M. until 4:00 P.M. during the week. Thus there is insufficient time to complete a Full backup during the week.

B: A differential backup will backup all data that has changed since the last Full backup. Towards the end of the week the differential backup may take quite some depending on the number of changes made since the last Full backup and may require more than 4 hours to complete.

Reference:

Microsoft SQL Server 2005 Books Online (2006), Index: database backups [SQL Server]

QUESTION 66:

You work as the database administrator at Certkiller .com. All servers on the Certkiller .com network run Windows Server 2003 and all database servers run SQL Server 2005. The Certkiller .com network contains a database server named Certkiller -DB01. Certkiller -DB01 hosts a database named CK_Sales that stores sales and product data. The CK_Sales database has three file groups named Primary, Secondary and Extra. Primary is the primary file group and contains the critical data from the database.

A Full backup of the CK_Sales database is performed every Saturday at 12:00 P.M., and a backup of the transaction log is performed every two hours during the working day. Certkiller -DB01 suffers a catastrophic failure. You replace Certkiller -DB01 with a new SQL Server 2005 database server named Certkiller -DB02. You now need to restore the Primary filegroup from a backup named bak_20060722 as quickly as possible. The Secondary and Extra file groups

will be restored later.
What should you do?

- A. Run RESTORE DATABASE CK_Sales FILEGROUP='Primary' FROM bak_20060722 WITH PARTIAL, NORECOVERY
- B. Run RESTORE DATABASE CK_Sales FILEGROUP='Primary' FROM bak_20060722 WITH PARTIAL
- C. Run RESTORE DATABASE CK_Sales FILEGROUP='Primary' FROM bak_20060722
- D. Run RESTORE DATABASE CK_Sales FILEGROUP='Primary' FROM bak_20060722 WITH NORECOVERY

Answer: A

Explanation: You need to perform a partial restore and then apply the transaction log backups. You must include the NORECOVERY option to be able to apply the transaction log backups.

Incorrect Answers:

B: You must include the NORECOVERY option to be able to apply the transaction log backups.

C: You need to perform a partial restore and then apply the transaction log backups. This means you must use the PARTIAL option and the NORECOVERY option to be able to apply the transaction log backups.

D: You must use the PARTIAL option to perform a partial backup.

QUESTION 67:

You work as the database administrator at Certkiller .com. All servers on the Certkiller .com network run Windows Server 2003 and all database servers run SQL Server 2005. The Certkiller .com network contains a database server named Certkiller -DB01. Certkiller -DB01 hosts a database named CK_Sales that stores sales and product data. Several data capturers enter data into the CK_Sales database everyday. The Full Recovery Model is implemented for the CK_Sales database. A Full backup of the CK_Sales database is performed every Saturday at 12:00 P.M., a differential backup is performed every week night at 12:00 P.M., and a backup of the transaction log is performed every two hours during the working day.

On Thursday morning you discover that Certkiller -DB01 suffered a hard disk failure sometime before the previous day's differential backup was performed. You replace the failed hard disk and restore the operating system and application on Certkiller -DB01. You now need to restore the CK_Sales database as quickly as possible.

What should you do?

- A. Restore the last Full backup.
Restore the last differential backup.

Restore all the transaction logs in order.

B. Restore the last Full backup.

Restore all the transaction logs in order.

C. Restore the last Full backup.

Restore the differential backup since the last Full backup.

Restore all the transaction logs since the last differential backup.

D. Restore the last Full backup.

Restore the last differential backup.

Restore all the transaction logs since the last differential backup.

Answer: D

Explanation: The Full Recovery Model is implemented for the CK_Sales database and you've performed a Full backup on Saturdays, daily differential backups and transaction log backups every two hours. In the Full Recovery Model, transactions in the transaction log are not cleared until they are backed up. Therefore you need to restore the last Full backup, restore the last differential backup, and restore all the transaction logs since the last differential backup.

Incorrect Answers:

A: You do not need to restore all the transaction log backups, only the transaction log backups since the last differential backup.

B: You could restore all the transaction log backups since the last Full backup but it would be quicker to restore the last differential backup performed since the last Full backup and then only the transaction log backups since the last differential backup.

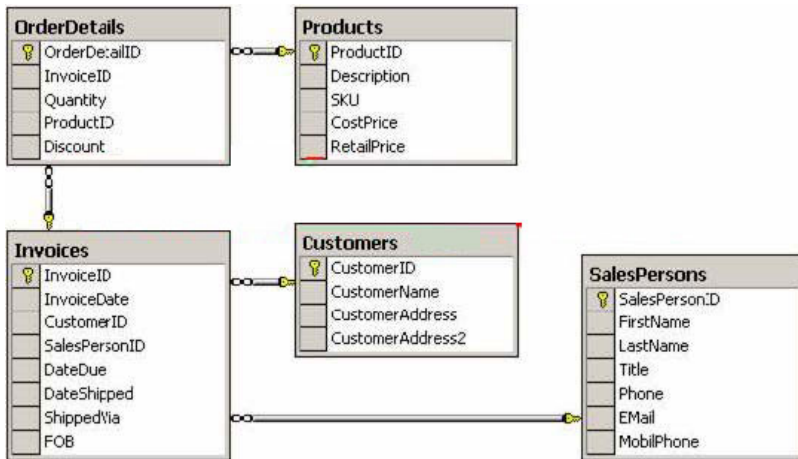
C: A differential backup will backup all data that has changed since the last Full backup. Therefore you only need to restore that last differential backup.

Reference:

Microsoft SQL Server 2005 Books Online (2006), Index: database backups [SQL Server]

QUESTION 68:

You work as the database administrator at Certkiller .com. All servers on the Certkiller .com network run Windows Server 2003 and all database servers run SQL Server 2005. The Certkiller .com network contains two database servers named Certkiller -DB01 and Certkiller -DB02. Certkiller -DB01 hosts a database named CK_Orders that stores sales data for the company. The tables in the CK_Orders database are shown in the following database diagram.



A full-text index has been created on the ProductID column of the Products table. You need to move the CK_Orders database from Certkiller -DB01 to Certkiller -DB02. You need to ensure that users can run queries against the CK_Orders database while it is being moved, and that the full-text index can be queried once the database has been moved to Certkiller -DB02. You plan to use the Copy Database Wizard to accomplish this task. What should you do? (Each correct answer presents part of the solution. Choose THREE.)

- A. Select the Use the SQL Management Object method option when using the Copy Database Wizard.
- B. Select the Use the detach and attach method option when using the Copy Database Wizard.
- C. Install SQL Server Integration Services (SSIS) on Certkiller -DB01 and Certkiller -DB02.
- D. Enable support for ad hoc distributed queries on Certkiller -DB02.
- E. Manually move the full-text catalog to Certkiller -DB02.
- F. Repopulate the full-text index on Certkiller -DB02 once the database has been moved.

Answer: A, C, F

Explanation: The SQL Server Integration Services (SSIS) is required for the Copy Database Wizard and must be installed on both the source and destination servers. To ensure that users can query the database while it is being moved, you must select the Use the SQL Management Object method option when using the Copy Database Wizard. The Use the SQL Management Object method allows active connections to the database while it is being moved. Once the database is moved, you must repopulate the full-text index. The Copy Database Wizard will move the full-text catalog to Certkiller -DB02 if the Use the SQL Management Object method is selected but it does not repopulate the index. Only once the index is repopulated can it be queried.

Incorrect Answers:

B: The Use the detach and attach method option does not allow active connections to

the database while it is being moved and thus will not allow the database to be queried while it is being moved.

D: Ad hoc distributed queries are not required to ensure that users can query the database while it is being moved.

E: When the Use the SQL Management Object method option is used, the Copy Database Wizard will move the full-text catalog to Certkiller -DB02 therefore you do not need to move the full-text catalog manually, but you will need to repopulate the index.

Reference:

Microsoft SQL Server 2005 Books Online (2006), Index: Copy Database Wizard

QUESTION 69:

You work as the database administrator at Certkiller .com. All servers on the Certkiller .com network run Windows Server 2003 and all database servers run SQL Server 2005. The Certkiller .com network contains a database server named Certkiller -DB01 B01. Certkiller -DB01 hosts a database named CK_Manufacturing that stores data from the Manufacturing department and a database named CK_Sales that stores sales data for the company.

Due to the growth of the CK_Sales database, Certkiller -DB01 is running low on free disk space. It is anticipated that the CK_Sales database will grow even further over the next few months. You add a second SQL Server 2005 database server named Certkiller -DB02 to the Certkiller .com network. You want to move the CK_Sales database to Certkiller -DB02. You need to accomplish task in minimum time.

What should you do?

A. Perform a Full backup of the CK_Sales database and its transaction logs.

Restore the Full backup to Certkiller -DB02.

Restore the transaction logs to Certkiller -DB02.

B. Use the sp_detach_db stored procedure to detach the CK_Sales database on Certkiller -DB01.

Copy the data and log files for CK_Sales to Certkiller -DB02.

Use the sp_attach_db stored procedure to attach the CK_Sales database to Certkiller -DB02.

C. Stop the SQL Server service on Certkiller -DB01.

Use the sp_detach_db stored procedure to detach the CK_Sales database on Certkiller -DB01.

Copy the data and log files for CK_Sales to Certkiller -DB02.

Use the sp_attach_db stored procedure to attach the CK_Sales database to Certkiller -DB02.

D. Create a new database named CK_Sales on Certkiller -DB02.

Use the SQL Server Import and Export Wizard to copy data from the CK_Sales database on Certkiller -DB01 to the new CK_Sales database on Certkiller -DB02.

Answer: B

Explanation: The fastest method of moving a database from one server to another is to detach the database and attach it to the destination server. This can be accomplished by running the sp_detach_db stored procedure to detach the database on source server, copying the data and log files for the database to the destination server, and then running the sp_attach_db stored procedure to attach the database to the destination server.

Incorrect Answers:

A: The Use the detach and attach method option does not allow active connections to the database while it is being moved and thus will not allow the database to be queried while it is being moved.

C: You can move the database by restoring a backup of the database to the destination server but it is quicker to detach and attach the database.

D: The SQL Server Import and Export Wizard can be used to move a database but it is quicker to detach and attach the database.

Reference:

Microsoft SQL Server 2005 Books Online (2006), Index: sp_attach_db

Microsoft SQL Server 2005 Books Online (2006), Index: sp_detach_db

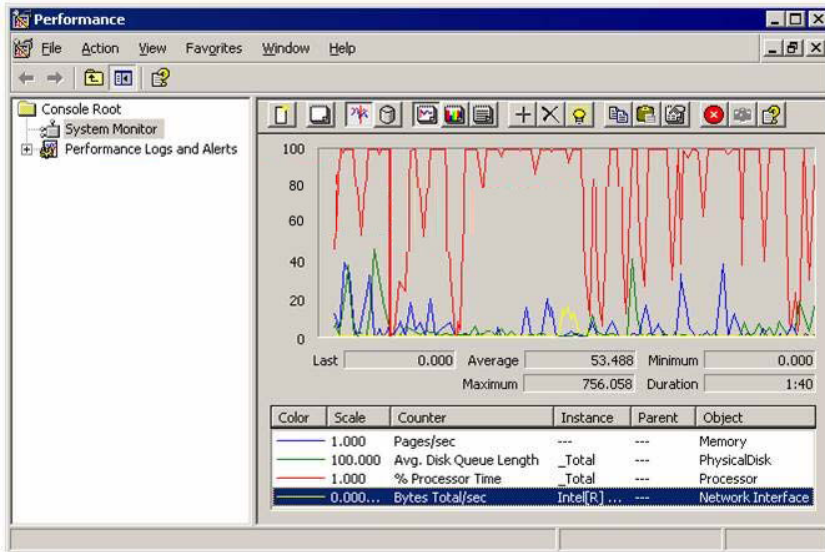
Microsoft SQL Server 2005 Books Online (2006), Index: attaching databases [SQL Server], onto another server instance

Microsoft SQL Server 2005 Books Online (2006), Index: detaching databases [SQL Server]

Microsoft SQL Server 2005 Books Online (2006), Index: SQL Server Import and Export Wizard

QUESTION 70:

You work as the database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. All servers on the Certkiller .com network run Windows Server 2003 and all client computers run Windows XP Professional. The Certkiller .com network contains a SQL Server 2005 database server named Certkiller -DB01 that hosts a database named CK_Sales. The CK_Sales database stores sales data for the company. Certkiller .com users complain that the performance of the CK_Sales database has deteriorated over the last few weeks. You use System Monitor to monitor the performance of Certkiller -DB01 and receive the output as shown in the exhibit.



You also notice that the ratio of SQL Recompilations/sec to Batch Requests/sec is unusually high. You suspect that the number of recompiles is causing the high processor time value. You need to improve the performance of the CK_Sales database.

What should you do?

- A. Turn off automatic updates of statistics for all tables in the CK_Sales database.
- B. Install an additional processor on Certkiller -DB01.
- C. Run the Database Engine Tuning Advisor.
- D. Use SQL Server Profiler to identify the stored procedures being recompiled.

Answer: D

Explanation: You can use SQL Server Profiler to identify the stored procedures that are being recompiled. SQL Profiler will indicate which stored procedures are being recompiled and why each recompilation is occurring.

Incorrect Answers:

A: Turning off automatic updates of statistics for the tables may improve performance but it does not address the problem caused by the recompiles.

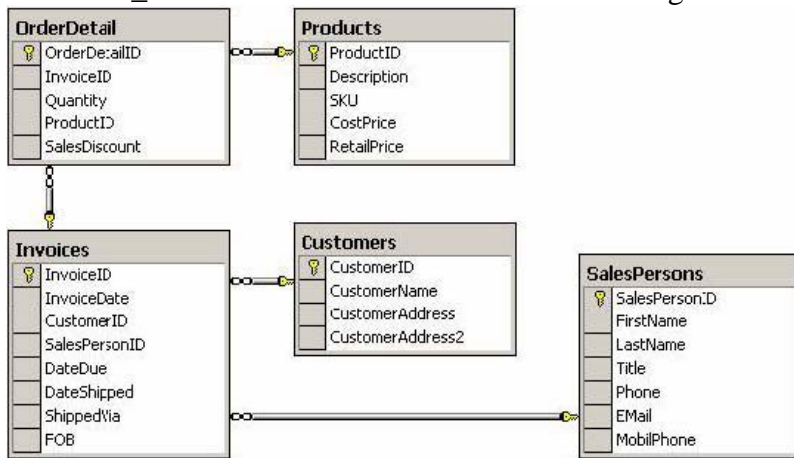
B: The high processor usage could be caused by the high number of recompiles. You should first reduce the number of recompiles before installing extra hardware on the server.

C: You can use the Database Engine Tuning Advisor to determine whether indexing and partitioning of a table would improve database performance but you cannot use it to identify which stored procedures are being recompiled.

QUESTION 71:

You work as the database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. All servers on the Certkiller .com network run Windows Server 2003 and all client

computers run Windows XP Professional. The Certkiller .com network contains a SQL Server 2005 database server named Certkiller -DB01. Certkiller -DB01 hosts a database named CK_Sales that stores sales data for the company. The tables in the CK_Sales database are shown in the following database diagram.



The recent increase in database usage at the company has resulted in several Certkiller .com users complaining of timeouts when they try to retrieve sales orders from the CK_Sales database. You need to determine whether partitioning the OrderDetails table would improve database performance.

What should you do? (Each correct answer presents part of the solution. Choose two.)

- A. Run Database Engine Tuning Advisor.
- B. Create a performance log file to monitor SQL Server:Buffer Manager:Page reads/sec.
- C. Run SQL Server Profiler to replay the trace file and the log file.
- D. Use the SQL Server Profiler Tuning template to create a trace file.
- E. Create a performance log file to monitor Logical Disk: Disk Read Bytes/sec

Answer: A, D

Explanation:

You can use the Database Engine Tuning Advisor to determine whether indexing and partitioning of a table would improve database performance. The Database Engine Tuning Advisor analyzes a workload file that you can create by running SQL Server Profiler and creating a trace based on the SQL Server Profiler Tuning template.

Incorrect Answers:

B: The SQL Server:Buffer Manager:Page reads/sec counter is used to monitor the read activity on an instance of SQL Server. It does not help you determine whether partitioning a table will improve database performance.

C: You can use the SQL Server Profiler to analyze a trace file and a log file to determine which queries are causing excessive resource utilization. However, this information does not help you determine whether partitioning a table will improve database performance.

E: The Logical Disk: Disk Read Bytes/sec counter is used to monitor the read activity on a logical disk. It does not help you determine whether partitioning a table will improve

database performance.

Reference:

Microsoft SQL Server 2005 Books Online (2007), Index: Database Engine Tuning Advisor [SQL Server]

Microsoft SQL Server 2005 Books Online (2007), Index: SQL Server Profiler

Microsoft SQL Server 2005 Books Online (2007), Index: performance counters [SQL Server]

Microsoft SQL Server 2005 Books Online (2007), Index: max server memory option

QUESTION 72:

DRAG DROP

You work as the database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. All servers on the Certkiller .com network run Windows Server 2003 and all client computers run Windows XP Professional. The Certkiller .com network contains a SQL Server 2005 database server named Certkiller -DB01 that hosts a database named CK_Sales that stores sales data for the company.

A sharp increase in sales has resulted in several Certkiller .com users complaining of timeouts when they try to retrieve sales orders from the CK_Sales database. You need to determine whether partitioning the Orders table would improve database performance.

What should you do? (To answers, select the appropriate actions from the pane on the left and place them in the pane on the right.)

Steps, Select from these	Steps, place here
Run the Database Engine tuning advisor	Place here
Run SQL Server Profiler to replay the trace file and the log file.	Place here
Monitor the SQL Server:BufferManager: Page read/sec performance counter	Place here,if any
Use SQL Server Profiler Tuning template to create a trace file.	Place here,if any

Answer:

Steps, Select from these	Steps, place here
	Run the Database Engine Tuning Advisor.
Run SQL Server Profiler to replay the trace file and the log file.	Use SQL Server Profiler Tuning template to create a trace file.
Monitor the SQL Server:BufferManager: Page read/sec Performance Counter	Place here, if any.
	Place here, if any.

Explanation:

You can use the Database Engine Tuning Advisor to determine whether indexing and partitioning of a table would improve database performance. The Database Engine Tuning Advisor analyzes a workload file that you can create by running SQL Server Profiler and creating a trace based on the SQL Server Profiler Tuning template.

Incorrect Answers:

The SQL Server:Buffer Manager:Page reads/sec counter is used to monitor the read activity on an instance of SQL Server. It does not help you determine whether partitioning a table will improve database performance.

You can use the SQL Server Profiler to analyze a trace file and a log file to determine which queries are causing excessive resource utilization. However, this information does not help you determine whether partitioning a table will improve database performance.

Reference:

Microsoft SQL Server 2005 Books Online (2007), Index: Database Engine Tuning Advisor [SQL Server]

Microsoft SQL Server 2005 Books Online (2007), Index: SQL Server Profiler

Microsoft SQL Server 2005 Books Online (2007), Index: performance counters [SQL Server]

Microsoft SQL Server 2005 Books Online (2007), Index: max server memory option

QUESTION 73:

You work as the database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. All servers on the Certkiller .com network run Windows Server 2003 and all client computers run Windows XP Professional. The Certkiller .com network contains a file server named Certkiller -SR10 and a SQL Server 2005 database server named Certkiller -DB01. Certkiller -DB01 hosts a database named CK_Sales that stores sales data for the company.

Certkiller .com users in the Accounts department use a proprietary application to access the CK_Sales database. After a few months Certkiller .com users in the Accounts department complain that the application takes a considerable amount of time to perform certain operations. You need to determine which queries that are used by the application are performing poorly. You intend to use SQL Server Profiler to replay the queries. You need to ensure that your attempt to resolve the problem has the least possible impact on network performance and the performance of Certkiller -DB01.

What should you do?

- A. Create a trace and save the data to a table in the tempdb.
- B. Create a trace and save the data to file on a shared folder on Certkiller -SR10.
- C. Create a trace and save the data to a file on Certkiller -DB01.
- D. Create a trace and save the data to a table in the user database.

Answer: C

Explanation: You should save the data as a file to minimize the performance impact on the database server.

Incorrect Answers:

A: To minimize the impact of saving the data, you should save the data as a file rather than as a table. Furthermore, saving the data to the tempdb will not ensure that you can replay the data at a later stage as the tempdb is rebuilt every time the database server is

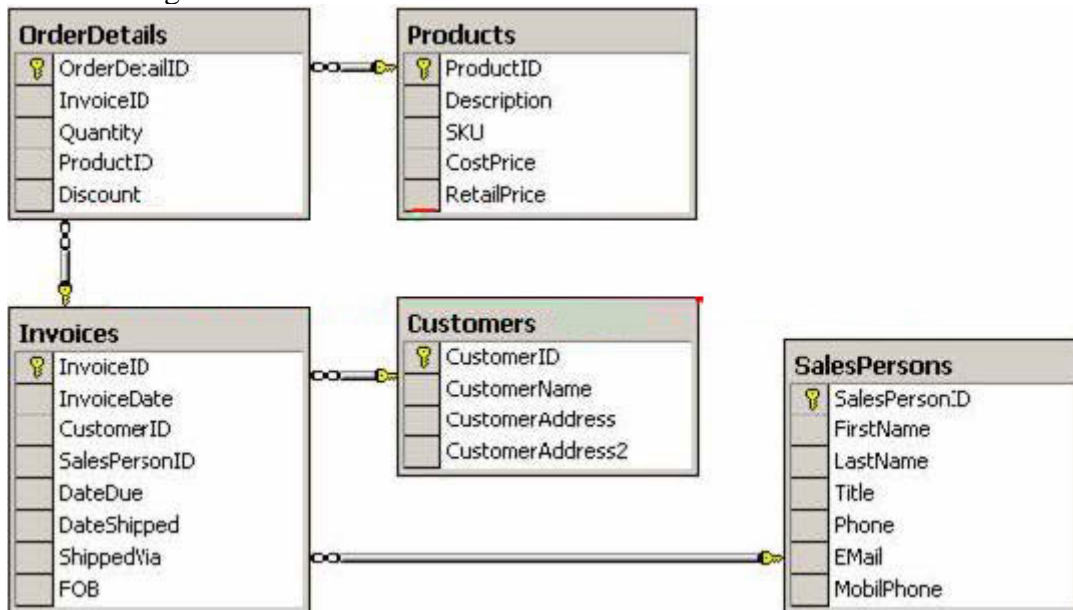
rebooted.

B: Saving the data to a network server and replaying the file at a later stage will increase network traffic and will impact on performance for users.

D: To minimize the impact of saving the data, you should save the data as a file rather than as a table.

QUESTION 74:

You work as the database administrator at Certkiller .com. All servers on the Certkiller .com network run Windows Server 2003 and all database servers run SQL Server 2005. The Certkiller .com network contains a database server named Certkiller -DB01 B01. Certkiller -DB01 hosts a database named CK_Sales that stores sales data for the company. There is at least one index on each table in the CK_Sales database. The tables in the CK_Sales database are shown in the following database diagram.



Certkiller .com users in the Finance department run several stored procedures against the CK_Sales database to create month end reports. You have recently altered some of the stored procedures. How Finance department users complain that the month-end reports take longer to create. You need to improve the performance of the stored procedures.

What should you do? (Each correct answer presents part of the solution. Choose two.)

- A. Use System Monitor to monitor the execution of the stored procedures and save the results to a trace file.
- B. Use SQL Server Profiler with the Tuning template to generate a trace file for the stored procedures.
- C. Use the Index Tuning Wizard to analyze the trace file.
- D. Run the stored procedures in the Query Editor and save the results to a trace file.
- E. Use the Database Engine Tuning Advisor to analyze the trace file.

Answer: B, E

Explanation: You should use SQL Server Profiler with the Tuning template to generate a trace file for the stored procedures and then use the Database Engine Tuning Advisor to analyze the trace file and tune the database.

Incorrect Answers:

A: The Windows Server 2003 System Monitor cannot be used to monitor the execution of a stored procedure. System Monitor is used to monitor system resources such as processor, hard disk, and services, etc.

C: The Index Tuning Wizard has been replaced in SQL Server 2005 by the Database Engine Tuning Advisor.

D: The results from the Query Editor cannot be saved to a trace file and cannot be used by the Database Engine Tuning Advisor to analyze the trace file and tune the database.

Reference:

Microsoft SQL Server 2005 Books Online (2006), Index: optimizing databases [SQL Server]

Microsoft SQL Server 2005 Books Online (2006), Index: performance [SQL Server], Database Engine Tuning Advisor

Microsoft SQL Server 2005 Books Online (2006), Index: performance [SQL Server], queries

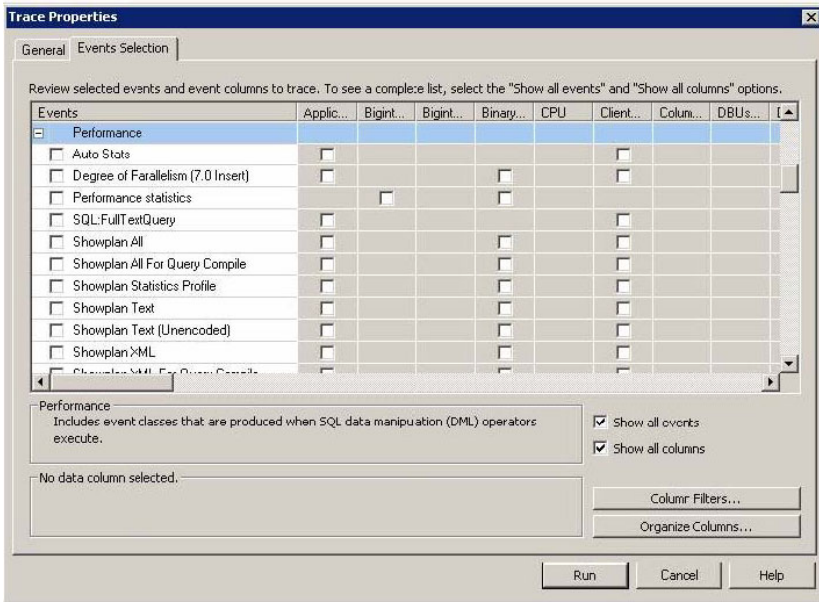
QUESTION 75:

HOTSPOT

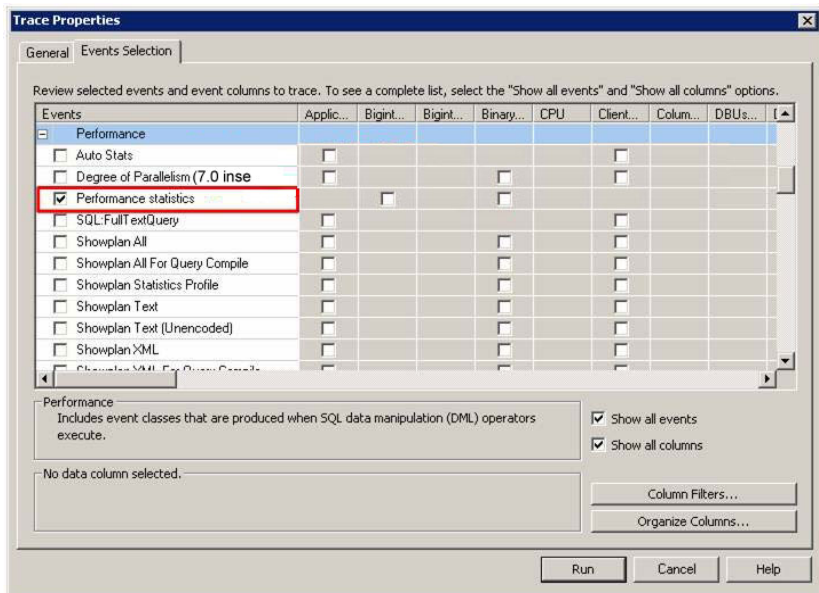
You work as the database administrator at Certkiller .com. All servers on the Certkiller .com network run Windows Server 2003 and all database servers run SQL Server 2005. The Certkiller .com network contains a database server named Certkiller -DB01 B01. Certkiller -DB01 hosts four databases that are named CK_Staff, CK_Sales, CK_Manufacturing and CK_Finance. Certkiller .com users in the Human Resources, Sales, Manufacturing and Finance departments use various applications to access the databases on Certkiller -DB01.

Certkiller .com recently acquired another company named Testlabs.com. Users in TestLabs.com also access the databases on Certkiller -DB01. You need to determine if Certkiller -DB01 has sufficient RAM for its plan cache. You decide to use SQL Server Profiler to run a trace against Certkiller -DB01. You use a blank template to create a new trace definition. You want to collect information on when each compiled plan is first cached, removed from the cache, or recompiled. You need to minimize any impact of the trace on database performance.

What should you do? To answer, select the event class that you should include in your trace.



Answer:



Explanation:

The Performance statistics event class capture information related to when each compiled query plan is first cached, removed from the cache, or recompiled.

Reference:

Microsoft SQL Server 2005 Books Online (2006), Index: query plans [SQL Server], caching

QUESTION 76:

You work as the database administrator at Certkiller .com. All servers on the Certkiller .com network run Windows Server 2003 and all database servers run SQL Server 2005. The Certkiller .com network contains a database server named Certkiller -DB01 B01. Certkiller -DB01 hosts various databases that

Certkiller .com users access on a daily basis.

Certkiller .com recently acquired another company named Testlabs.com. Users in TestLabs.com also access the databases on Certkiller -DB01. You need to determine if Certkiller -DB01 has sufficient RAM for its plan cache. You decide to use SQL Server Profiler to run a trace against Certkiller -DB01. You use a blank template to create a new trace definition. You want to collect information on when each compiled plan is first cached, removed from the cache, or recompiled. You need to minimize any impact of the trace on database performance. What should you do?

- A. Include the Auto Stats event class in your trace.
- B. Include the Showplan All event class in your trace.
- C. Include the Performance Statistics event class in your trace.
- D. Include the Showplan Statistics Profile event class in your trace.

Answer: C

Explanation: The Performance statistics event class capture information related to when each compiled query plan is first cached, removed from the cache, or recompiled.

Incorrect Answers:

A: The Auto Stats event class captures specified information related to when SQL automatically updates statistics for indexes and columns.

B: The Showplan All event class captures detailed information related to when each plan is compiled and has a serious impact on database performance.

D: The Showplan Statistics Profile event class captures information each time a SQL statement is executed and has a serious impact on database performance.

Reference:

Microsoft SQL Server 2005 Books Online (2006), Index: query plans [SQL Server], caching

QUESTION 77:

You work as the database administrator at Certkiller .com. All servers on the Certkiller .com network run Windows Server 2003 and all database servers run SQL Server 2005. The Certkiller .com network contains a database server named Certkiller -DB01 B01. Certkiller -DB01 hosts a database named CK_Manufacturing that stores data from the Manufacturing department. Several tables in the CK_Manufacturing database have more than one million rows. You have created a clustered index on each table in the CK_Manufacturing database as well as a nonclustered index on each table that is used frequently. You want to optimize performance for the CK_Manufacturing database. You create a trace file for a workload in SQL Server Profiler. You now want to tune the workload file. What should you do?

- A. Run the Index Tuning Wizard.
- B. Run the dta utility at a command prompt.
- C. Run the sqlmaint utility at a command prompt.
- D. Run the profiler90 utility at a command prompt.

Answer: B

Explanation: The dta utility is the command line version of the Database Engine Tuning Advisor. It can be used to analyze the workload file and tune the database.

Incorrect Answers:

A: The Index Tuning Wizard has been replaced in SQL Server 2005 by the Database Engine Tuning Advisor. The dta utility is the command line version of the Database Engine Tuning Advisor

C: The sqlmaint utility is used to run database maintenance plans that have been created in previous versions of SQL Server.

D: The profiler90 utility is used to launch the SQL Server Profiler from the command line. It is not used to tune the workload file.

Reference:

Microsoft SQL Server 2005 Books Online (2006), Index: dta utility

Microsoft SQL Server 2005 Books Online (2006), Index: performance [SQL Server], Database Engine Tuning Advisor

Microsoft SQL Server 2005 Books Online (2006), Index: sqlmaint utility

Microsoft SQL Server 2005 Books Online (2006), Index: profiler90 utility

QUESTION 78:

You work as the database administrator at Certkiller .com. All servers on the Certkiller .com network run Windows Server 2003 and all database servers run SQL Server 2005. The Certkiller .com network contains a database server named Certkiller -DB01 B01. Certkiller -DB01 hosts a database named CK_Sales that stores sales data for the company. CK_Sales is also used for online transaction processing (OLTP).

Several tables in the CK_Sales database have more than one million rows. You have created a clustered index on each table in the CK_Sales database as well as a nonclustered index on each table that is used frequently. You want to optimize performance for the clustered and nonclustered indexes. You want to generate index recommendations.

What should you do?

- A. Run the Index Tuning Wizard and save the recommended changes to a file.
- B. Run the Database Engine Tuning Advisor and save the data to a SQL script file.
- C. Run the SQL Server Profiler with the Tuning template and save the result to a file.
- D. Run the sys.dm_db_index_operational_stats statement and save the results to a file.

Answer: C

Explanation: You should use SQL Server Profiler with the Tuning template to generate a workload file and then use the Database Engine Tuning Advisor to analyze the workload file.

Incorrect Answers:

A: The Index Tuning Wizard has been replaced in SQL Server 2005 by the Database Engine Tuning Advisor.

B: you should first use SQL Server Profiler with the Tuning template to generate a workload file and then use the Database Engine Tuning Advisor to analyze the workload file.

D: you can't use the sys.dm_db_index_operational_stats statement to generate index recommendations.

Reference:

Microsoft SQL Server 2005 Books Online (2006), Index: optimizing databases [SQL Server]

Microsoft SQL Server 2005 Books Online (2006), Index: performance [SQL Server], Database Engine Tuning Advisor

Microsoft SQL Server 2005 Books Online (2006), Index: sys.dm_db_index_operational_stats

QUESTION 79:

You work as the database administrator at Certkiller .com. All servers on the Certkiller .com network run Windows Server 2003 and all database servers run SQL Server 2005. The Certkiller .com network contains a database server named Certkiller -DB01 B01. Certkiller -DB01 hosts a database named CK_Sales that stores sales data for the company.

Certkiller .com recently acquired another company named Testlabs.com. Users in TestLabs.com also access the CK_Sales database on Certkiller -DB01. A

Certkiller .com user reports that she often receives an error message when she accesses the CK_Sales database. You discover that an errant process is causing the error message. You need to terminate the errant process while minimizing the impact on users that are currently connected to the database. You identify the session ID (SPID) of the errant process as SPID 34.

What should you do next?

- A. Run the KILL SPID 34 Transact-SQL statement.
- B. Stop and restart the SQL Server Agent service.
- C. Locate the errant process in task manager and end the process tree.
- D. Run the KILL 34 Transact-SQL statement.

Answer: D

Explanation: The errant process can be terminated by using the KILL statement and specifying the session ID as in KILL 34.

Incorrect Answers:

A: The KILL statement does not use the SPID clause, only the SPID number.

B: Stopping and restarting the SQL Server Agent service will affect all users that are connected to the database.

C: You cannot use Task Manager to terminate a SQL Server process. Task Manager terminates all processes associated with an application.

Reference:

Microsoft SQL Server 2005 Books Online (2006), Index: KILL statement

QUESTION 80:

You work as the database administrator at Certkiller .com. All servers on the Certkiller .com network run Windows Server 2003 and all database servers run SQL Server 2005. Six month after implementing a Web application that access a Certkiller .com database; Web site users complain about poor response times from the Web site. You need to determine if the problem is caused by deadlocks. If the problem is caused by deadlocks, you want to determine which objects are involved in the deadlocks.

What should you do?

- A. Use SQL Server Profiler to run a trace.
- B. Run the sp_lock stored procedure.
- C. Query the sys.dm_tran_database_transactions dynamic management view.
- D. Query the sys.dm_tranlocks dynamic management view.

Answer: A

Explanation: You can use SQL Server Profiler to monitor locks and deadlocks.

Incorrect Answers:

B, D: The sp_lock stored procedure has been deprecated in SQL Server 2005 in favor of the sys.dm_tranlocks dynamic management view. However, the sys.dm_tranlocks dynamic management view allows you to view current locks and deadlocks. It does not allow you to view a history of deadlocks

C: The sys.dm_tran_database_transactions dynamic management view provides information about transactions rather than locks and deadlocks.

QUESTION 81:

You work as the database administrator at Certkiller .com. All servers on the Certkiller .com network run Windows Server 2003 and all database servers run SQL Server 2005. The Certkiller .com network contains a database server named Certkiller -DB01 B01. Certkiller -DB01 hosts a database named CK_Manufacturing that stores data from the Manufacturing department.

One day Certkiller .com users complain that they cannot access the CK_Manufacturing database on Certkiller -DB01. When you attempt to open the SQL Server management Studio, you discover that Certkiller -DB01 is not responding. You are unable to connect using sqlcmd with a standard connection. You need to be able to connect to Certkiller -DB01 to be able to diagnose the

problem.

What should you do?

- A. Run sqlcmd with the -A option from the command prompt.
- B. Run the sp_configure 'remote admin connections', 1; statement.
- C. Run the DBCC CHECKDB statement.
- D. Determine the session ID of the errant process and terminate that process.

Answer: A

Explanation: The sqlcmd command with the -A option to establish a dedicated administration connection (DAC) to Certkiller -DB01. You will be able to run queries to diagnose the problem once a DAC is established.

Incorrect Answers:

B: By default, only local dedicated administration connections (DACs) are allowed as DAC only listens on the loop-back IP address (127.0.0.1), port 1434. The sp_configure 'remote admin connections', 1; statement allows a remote DAC to be established. It does not establish the DAC.

C: You will not be able to run the DBCC CHECKDB statement on a nonresponsive server until you establish a dedicated administration connection (DAC). Furthermore, the DBCC CHECKDB statement is used to check the structural and logical integrity of a database. It is not used to diagnose a nonresponsive server.

D: You will not be able to identify the session ID of any process on a nonresponsive server until you establish a dedicated administration connection (DAC).

Reference:

Microsoft SQL Server 2005 Books Online (2006), Index: sqlcmd commands

Microsoft SQL Server 2005 Books Online (2006), Index: remote admin connections option

Microsoft SQL Server 2005 Books Online (2006), Index: DBCC CHECKDB statement

QUESTION 82:

You work as the database administrator at Certkiller .com. All servers on the Certkiller .com network run Windows Server 2003 and all database servers run SQL Server 2005 with the default settings. The Certkiller .com network contains a database server named Certkiller -DB01 that hosts a database named CK_Sales. The CK_Sales database stores sales data for the company.

One morning Certkiller .com users complain that they are unable to connect to Certkiller -DB01. You log on to your client computer named Certkiller -WS444. You attempt to connect to Certkiller -DB01 using the Dedicated Administrator Connection (DAC) but the connection fails. You need to be able to connect to Certkiller -DB01 from Certkiller -WS444.

What should you do? (Each correct answer presents part of the solution. Choose two.)

- A. Start the SQL Server Agent service on Certkiller -DB01.

- B. Start the SQL Browser service on Certkiller -DB01.
- C. Run the sp_configure 'remote admin connections', 0; statement on Certkiller -DB01.
- D. Enable the remote admin connection option on Certkiller -DB01.
- E. Enable named pipes on Certkiller -DB01.

Answer: B, D

Explanation: By default, only local dedicated administration connections (DACs) are allowed as DAC only listens on the loop-back IP address (127.0.0.1). You need to enable the remote admin connection option on Certkiller -DB01 by running the sp_configure 'remote admin connections', 1; statement. You should also start the SQL Browser service on Certkiller -DB01 as the SQL Server client uses this service to determine the port number that DAC is listening on. Alternatively, you can configure DAC to use a specific port.

Incorrect Answers:

B: The Dedicated Administrator Connection (DAC) does not require the SQL Server Agent service.

C: The sp_configure 'remote admin connections', 0; statement allows SQL Server to accept only local dedicated administration connections (DACs) using port 1434 on the loop-back IP address (127.0.0.1). This is the default setting.

E: The Dedicated Administrator Connection (DAC) uses TCP/IP and does not require named pipes.

QUESTION 83:

You work as the database administrator at Certkiller .com. All servers on the Certkiller .com network run Windows Server 2003 and all database servers run SQL Server 2005. Certkiller .com has its headquarters in Chicago and branch offices in Dallas and Miami. You work in the Miami office. Head office contains a database server named Certkiller -DB01 B01. Certkiller -DB01 hosts a database named CK_Sales that stores sales data for the company.

Several Certkiller .com users in the Miami office start complaining about difficulty in connecting to the CK_Sales database. You discover that connection between the clients and Certkiller -DB01 is being terminated. You need to determine the cause of the network connectivity problem as quickly as possible. Unfortunately, the administrator at the Chicago office is on sick leave and you are unable to travel to the Chicago office. You need to access Certkiller -DB01 from your administrative computer.

What should you do?

- A.
Open the Event Viewer and connect to Certkiller -DB01 and review the System and Application logs.
- B. Open the SQL Server Management console and review the SQL Server Agent log.
- C. Run the Network Monitor and capture network traffic to and from Certkiller -DB01.
- D. Open the SQL Server Management console run the SQL Server Profiler.

Answer: A

Explanation: you can use the Event Viewer to connect to a remote server and review the logs on the remote server. Client connectivity problems will appear in the Application log and the System log.

Incorrect Answers:

B: The SQL Server Agent log records events related to the SQL Server Agent service. The SQL Server Agent service is not responsible for network connectivity to the database server.

C: The Network Monitor is used to capture and analyze network traffic. It does not provide information related to the cause of network connectivity problems.

D: The SQL Server Profiler is used to monitor server and database activity but not network connectivity issues.

QUESTION 84:

You work as the database administrator at Certkiller .com. All servers on the Certkiller .com network run Windows Server 2003 and all database servers run SQL Server 2005. The Certkiller .com network contains a database server named Certkiller -DB01 B01. Certkiller -DB01 hosts a database named CK_Sales that stores sales data for the company.

You create a maintenance job that runs two statements against the CK_Sales database and schedule the job to run every evening at 7:30 P.M. The following morning you confirm that the job ran on schedule. A month later you discover that the job has failed and no longer runs on schedule. You need to determine the cause of the problem.

What should you do?

A. Manually run the job and use the SQL Server Profiler to monitor the execution of the job. Save the results to a trace file.

B. Run the sp_trace_setstatus stored procedure.

C. Search the Application log in Event Viewer for SQL-related error messages.

D. Run the sp_help_jobhistory stored procedure.

Answer: D

Explanation: The sp_help_jobhistory stored procedure returns information about the cause of a job failure.

Incorrect Answers:

A: SQL Server Profiler does not return information about the cause of a job failure.

B: The sp_trace_setstatus stored procedure is used to start, stop or close a trace. It does not return information about the cause of a job failure.

C: You can use the Application log to determine the cause of a job failure as the Application log does not record such events.

Reference:

Microsoft SQL Server 2005 Books Online (2006), Index: sp_help_jobhistory

Microsoft SQL Server 2005 Books Online (2006), Index: sp_trace_setstatus

Microsoft SQL Server 2005 Books Online (2006), Index: error messages [SQL Server]

QUESTION 85:

You work as the database administrator at Certkiller .com. All servers on the Certkiller .com network run Windows Server 2003 and all database servers run SQL Server 2005. The Certkiller .com network contains a database server named Certkiller -DB01 B01. Certkiller -DB01 hosts a database named CK_Sales that stores sales data for the company.

You create a maintenance named JobA that is run manually by Certkiller .com users in the Sales department. A month later you become concerned that Sales department users may neglect running JobA and want determine when JobA was last run.

What should you do? (Each correct answer presents a complete solution. Choose two.)

- A. Run the sp_monitor stored procedure.
- B. Run the sp_help_jobhistory stored procedure.
- C. Run the sp_help_jobactivity stored procedure.
- D. Run the sp_help_job stored procedure.

Answer: B, C

Explanation: The sp_help_jobactivity and the sp_help_jobhistory stored procedures returns information about a job, including when the job was last run.

Incorrect Answers:

A: The sp_monitor stored procedure returns server activity statistics. It does not return information related to when a job was last run.

D: The sp_help_job stored procedure returns information about jobs used by SQL Server to perform automated activities. It does not return information related to when a manual job was last run.

Reference:

Microsoft SQL Server 2005 Books Online (2006), Index: sp_monitor

Microsoft SQL Server 2005 Books Online (2006), Index: sp_help_jobhistory

Microsoft SQL Server 2005 Books Online (2006), Index: sp_help_jobactivity

Microsoft SQL Server 2005 Books Online (2006), Index: sp_help_job

QUESTION 86:

You work as the database administrator at Certkiller .com. You are responsible for writing procedure documents for the company. All servers on the Certkiller .com network run Windows Server 2003. The Certkiller .com network contains a database server named Certkiller -DB01 B01. Certkiller -DB01 was recently upgraded from SQL Server 2000 to SQL Server 2005.

You need to review and update the procedure for reducing fragmentation on

indexes. The current procedure document recommends the use of the DBCC SHOWCONTIG and DBCC INDEXDEFRAG statements. You need to alter the procedure document to recommend the use of commands and statements that will be supported in future versions of SQL Server.

What should you do?

- A. Recommend the use for the sys.dm_db_partition_stats and ALTER INDEX REORGANIZE statements.
- B. Recommend the use for the sys.dm_db_index_physical_stats and ALTER INDEX REORGANIZE statements.
- C. Recommend the use for the sys.dm_db_partition_stats and ALTER INDEX REBUILD statements.
- D. Recommend the use for the sys.dm_db_index_physical_stats and ALTER INDEX REBUILD statements.

Answer: B

Explanation: The sys.dm_db_index_physical_stats and ALTER INDEX REORGANIZE statements are used to reduce fragmentation of indexes and will be supported in future versions of SQL Server and replaces the DBCC SHOWCONTIG and DBCC INDEXDEFRAG statements respectively.

Incorrect Answers:

- A: The sys.dm_db_partition_stats does not replace either the DBCC SHOWCONTIG and DBCC INDEXDEFRAG statements.
- C: The sys.dm_db_partition_stats and the ALTER INDEX REBUILD statements do not replace either the DBCC SHOWCONTIG and DBCC INDEXDEFRAG statements.
- D: The ALTER INDEX REBUILD statement does not replace either the DBCC SHOWCONTIG and DBCC INDEXDEFRAG statements.

Reference:

Microsoft SQL Server 2005 Books Online (2006), Index: DBCC SHOWCONTIG
Microsoft SQL Server 2005 Books Online (2006), Index: DBCC INDEXDEFRAG
Microsoft SQL Server 2005 Books Online (2006), Index: ALTER INDEX REBUILD
Microsoft SQL Server 2005 Books Online (2006), Index:
sys.dm_db_index_physical_stats
Microsoft SQL Server 2005 Books Online (2006), Index: sys.dm_db_partition_stats

QUESTION 87:

You work as the database administrator at Certkiller .com. All servers on the Certkiller .com network run Windows Server 2003 and all database servers run SQL Server 2005. The Certkiller .com network contains a database server named Certkiller -DB01 B01.

Certkiller -DB01 hosts a database named CK_Products that contains 60 lookup tables. The lookup tables should remain static at all times and users should not be able to alter the static data. You discover that some of the data in the lookup tables have been altered. You need to ensure that users cannot change any of the data in

the lookup tables.
What should you do?

- A. Create a new filegroup and move the lookup tables to the new filegroup.
Enable the Read-Only option on the filegroup.
- B. Create a view of the lookup tables.
Allow users to access the lookup tables through the view only.
- C. Create stored procedures for modifying data in the lookup tables.
Allow users to modify data through the stored procedures only.
- D. Create a new database role and add all users to the new role.
Grant SELECT permissions to the new role.

Answer: A

Explanation: You can prevent users from modifying data in the lookup tables by moving the tables to a separate filegroup and making the file group read-only.

Incorrect Answers:

B: Views are used to limit the part of the database that the users can see. It does not prevent them from modifying the data in the underlying table.

C: Users should not be able to modify the data in the lookup tables. Allowing them to modify data through the stored procedures only would not meet this requirement.

D: This option will require considerable administrative effort as there are 60 look up tables. You will need to grant SELECT permissions for each of these tables. Using a read-only filegroup would require less effort.

Reference:

Microsoft SQL Server 2005 Books Online (2007), Index: filegroups [SQL Server], read-only

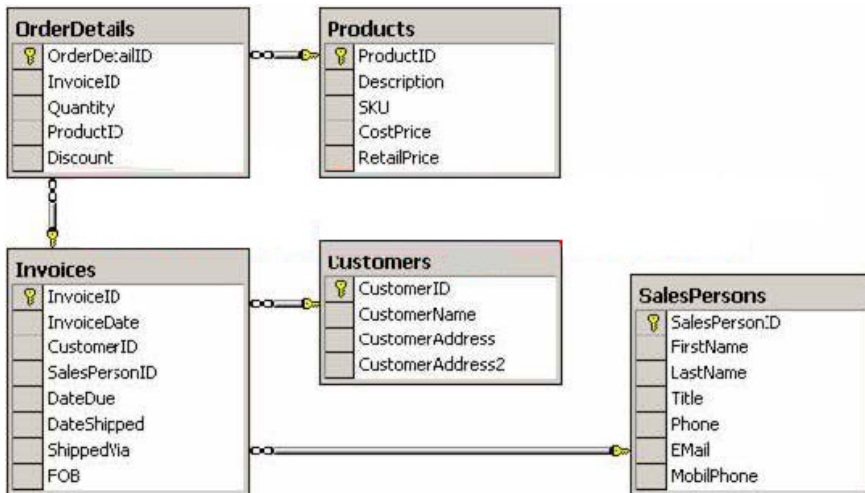
Microsoft SQL Server 2005 Books Online (2007), Index: views [SQL Server]

Microsoft SQL Server 2005 Books Online (2007), Index: database roles [SQL Server]

Microsoft SQL Server 2005 Books Online (2007), Index: stored procedures [SQL Server]

QUESTION 88:

You work as the database administrator at Certkiller .com. All servers on the Certkiller .com network run Windows Server 2003 and all database servers run SQL Server 2005. The Certkiller .com network contains a database server named Certkiller -DB01 B01. Certkiller -DB01 hosts a database named CK_Sales that stores sales data for the company. The tables in the CK_Sales database are shown in the following database diagram.



Several products listed in the Products table have been discontinued. You want to delete the data from the Products table that refer to the discontinued products. You also want to copy data from the Invoices table that are related to the discontinued products to a separate table named DiscontinuedProductInvoices. You decide to use triggers to accomplish this task. The first trigger will delete rows from the Products table while the second trigger will copy data from the Invoices table to the DiscontinuedProductInvoices table. You want the second trigger to execute once the first trigger is executed.

What should you do?

- A. Run the `sp_configure 'nested triggers', 0` statement.
- B. Run the `sp_configure 'nested triggers', 1` statement.
- C. Run the `CREATE TRIGGER` statement and create nested triggers.
- D. Run the `TRIGGER_NESTLEVEL` statement.

Answer: C

Explanation: You can use the `CREATE TRIGGER` statement to create nested triggers. A nested trigger is initiated by another trigger.

Incorrect Answers:

A, B: The nested trigger option of the `sp_configure` command specifies whether nested triggers are permitted. Setting this option to 0 disables nested triggers while 1 enables nested triggers. By default nested triggers are enabled.

D: The `TRIGGER_NESTLEVEL` returns the number of triggers executed for the statement that fired the trigger.

Reference:

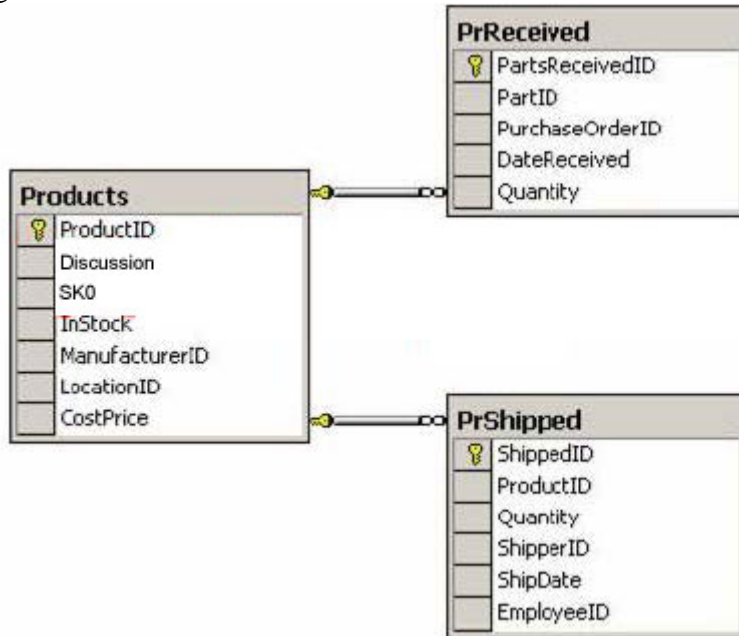
Microsoft SQL Server 2005 Books Online (2007), Index: nested DML triggers

Microsoft SQL Server 2005 Books Online (2007), Index: nested triggers option

Microsoft SQL Server 2005 Books Online (2007), Index: `TRIGGER_NESTLEVEL` function

QUESTION 89:

You work as the database administrator at Certkiller .com. All servers on the Certkiller .com network run Windows Server 2003 and all database servers run SQL Server 2005. The Certkiller .com network contains a database server named Certkiller -DB01 B01. Certkiller -DB01 hosts a database named CK_Products. The tables in the CK_Products database are shown in the following database diagram.



When products are shipped, the PrShipped table is updated and when products are received, the PrReceived table is updated. You want the InStock column in the Products table to be updated automatically when the PrShipped and PrReceived tables are updated.

What should you do?

- A. Create nested triggers on the PrShipped and PrReceived tables.
- B. Create UPDATE triggers on the PrShipped and PrReceived tables.
- C. Create nested triggers on the Products table.
- D. Create UPDATE triggers on the Products table.

Answer: B

Explanation: You should create UPDATE triggers on the PrShipped and PrReceived tables. Whenever these tables are updated, the trigger will fire. The trigger should be configured to update the InStock field in the Products table.

Incorrect Answers:

A, C: A nested trigger allows one trigger to initiate another trigger. This is not what is required. What are required are UPDATE triggers on the PrShipped and PrReceived tables that are configured to update the InStock column in the Products table.

D: The UPDATE triggers should be created on the PrShipped and PrReceived tables and should be configured to update the InStock field in the Products table.

Reference:

Microsoft SQL Server 2005 Books Online (2007), Index: nested triggers option
Microsoft SQL Server 2005 Books Online (2007), Index: TRIGGER_NESTLEVEL
function
Microsoft SQL Server 2005 Books Online (2007), Index: triggers [SQL Server]

QUESTION 90:

You work as the database administrator for an IT training company named Certkiller .com. All servers on the Certkiller .com network run Windows Server 2003 and all database servers run SQL Server 2005. The Certkiller .com network contains a database server named Certkiller -DB01 B01. Certkiller -DB01 hosts a database named CK_Training that stores training data for the company. You are instructed to create checklists that lecturers can use to prepare training courses. You must use information from text files stored on Certkiller -DB01's hard disk drive to create the checklists for the different certifications that Certkiller .com offers training for. You want to use Visual Studio .NET 2005 to build the checklists by implementing a common language runtime (CLR) function. What should you do?

- A. Use the CREATE FUNCTION Transact-SQL statement to reference the assembly.
- B. Use the CREATE ASSEMBLY Transact-SQL statement to register the assembly.
- C. Register the assembly in the global assembly cache (GAC).
- D. Sign the assembly with a strong name.

Answer: B

Explanation: You need to register the assembly by using the CREATE ASSEMBLY statement.

Incorrect Answers:

- A: You need to register the assembly before you can reference it.
- C: You do not need to register an assembly in the global assembly cache (GAC). The GAC is used to allow multiple applications to share an assembly.
- D: You need to sign the assembly with a strong name if it is to be registered in the global assembly cache (GAC). However, you do not need to register an assembly in the global assembly cache (GAC). The GAC is used to allow multiple applications to share an assembly.

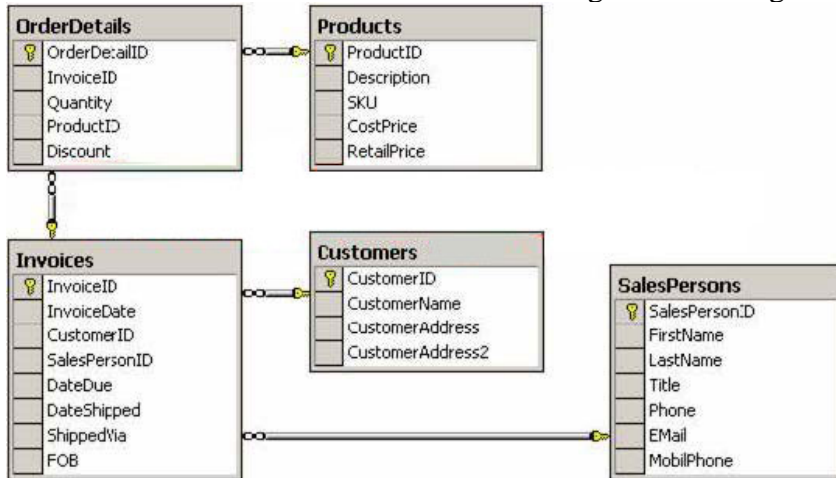
Reference:

Microsoft SQL Server 2005 Books Online (2007), Index: CREATE ASSEMBLY
statement
Microsoft SQL Server 2005 Books Online (2007), Index: CREATE FUNCTION
statement
Microsoft SQL Server 2005 Books Online (2007), Index: assemblies

QUESTION 91:

You work as the database administrator for an IT training company named

Certkiller .com. All servers on the Certkiller .com network run Windows Server 2003 and all database servers run SQL Server 2005. The Certkiller .com network contains a database server named Certkiller -DB01 B01. Certkiller -DB01 hosts a database named CK_Sales that stores sales data for the company. The tables in the CK_Sales database are shown in the following database diagram.



Certkiller .com decides to switch to commission-based salaries for its Sales staff. You need to generate information that Rory Allen, the manager of the Sales department, will use to calculate the commission for each sales person. Your solution must meet the following criteria:

1. You must be able to use multiple SELECT statements to produce the result set.
2. The result set must be returned as a table for which the PRIMARY KEY constraint is defined.
3. Rory Allen must be able to reference the result set in the FROM clause of a SELECT statement.

What should you do?

- A. Create a stored procedure.
- B. Create a view.
- C. Create a user-defined function.
- D. Create a Common Runtime Language (CLR) function.

Answer: C

Explanation: You need a user-defined function to meet the requirements for your solution.

Incorrect Answers:

A: Result sets from a stored procedure cannot be in the FROM clause of a SELECT statement.

B: A view only supports a single SELECT statement. You need to be able to use multiple SELECT statements to produce the result set.

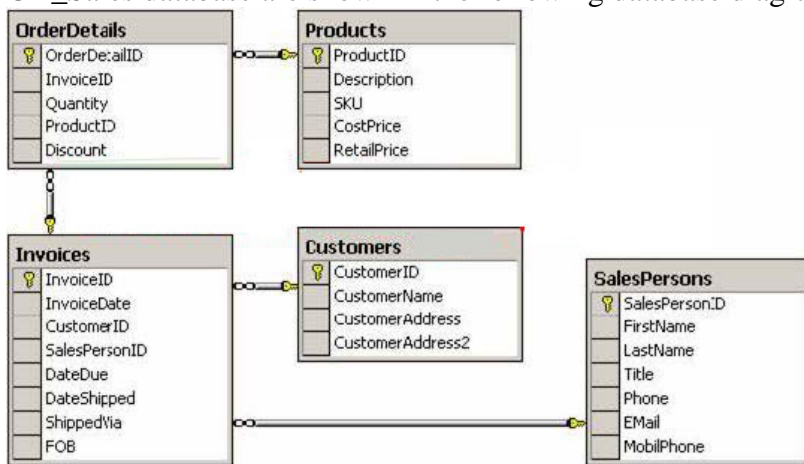
D: You cannot define a PRIMARY KEY constraint on the result set returned from a Common Runtime Language (CLR) function.

Reference:

Microsoft SQL Server 2005 Books Online (2007), Index: stored procedures [SQL Server]
Microsoft SQL Server 2005 Books Online (2007), Index: stored procedures [SQL Server], creating
Microsoft SQL Server 2005 Books Online (2007), Index: views [SQL Server]
Microsoft SQL Server 2005 Books Online (2007), Index: functions [SQL Server],

QUESTION 92:

You work as the database administrator for an IT training company named Certkiller .com. All servers on the Certkiller .com network run Windows Server 2003 and all database servers run SQL Server 2005. The Certkiller .com network contains a database server named Certkiller -DB01 B01. Certkiller -DB01 hosts a database named CK_Sales that stores sales data for the company. The tables in the CK_Sales database are shown in the following database diagram.



A PRIMARY KEY column has been declared on the ProductID column of the Products table and the InvoiceID column of the Invoices table. You are creating a table named Inventory to track stock levels. You need to ensure that combination of the DateShipped, ProductID and InvoiceID is unique for each row in the Inventory table.

What should you do?

- A. Create a rule.
- B. Create a table-level constraint.
- C. Create a UNIQUE constraint for each of the three columns.
- D. Create a CHECK constraint for each of the three columns.

Answer: B

Explanation: A table-level constraint allows you to ensure that a combination of columns is unique for each row in the table.

Incorrect Answers:

A, D: A rule and CHECK constraints are used to ensure FOREIGN KEY integrity. They are not used to ensure that data in a column or combination of columns is unique.
C: UNIQUE constraints do not allow a value to be repeated in the column. This means

that every value in the DateShipped column must be unique; each value in the ProductID column must be unique; and each value in the InvoiceID column must be unique.

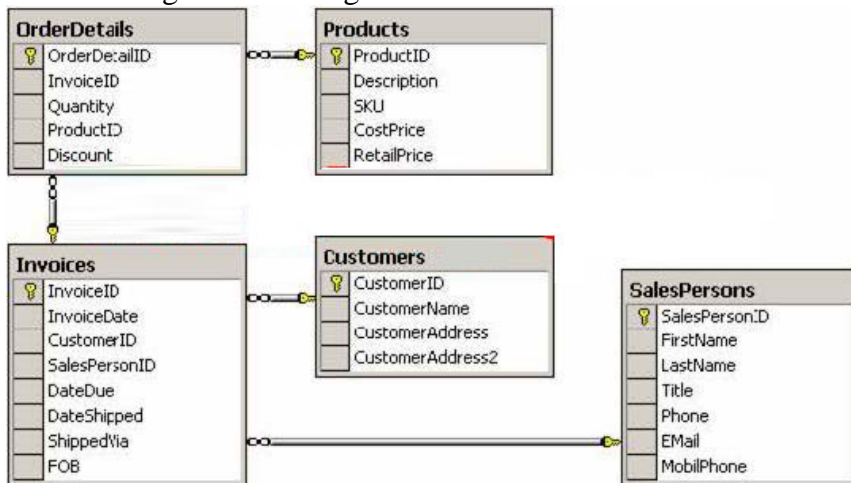
Reference:

Microsoft SQL Server 2005 Books Online (2007), Index: constraints [SQL Server]

Microsoft SQL Server 2005 Books Online (2007), Index: rules [SQL Server]

QUESTION 93:

You work as the database administrator at Certkiller .com. All servers on the Certkiller .com network run Windows Server 2003 and all database servers run SQL Server 2005. The Certkiller .com network contains a database server named Certkiller -DB01 B01. Certkiller -DB01 hosts a database named CK_Sales that stores sales data for the company. The tables in the CK_Sales database are shown in the following database diagram.



Certkiller .com decides to switch to commission-based salaries for its Sales staff. The manager of the Sales department will run queries against the Invoices table to calculate the commission for each sales person. To ensure the success of the queries, you want to enforce the values that the SalesPersonID column in the Invoices table will accept.

What should you do?

- A. Create a foreign key constraint that references the CK_Sales.SalesPersons table.
- B. Create a check key constraint that references the CK_Sales.SalesPersons table.
- C. Create a DML trigger that references the CK_Sales.SalesPersons table.
- D. Create a unique constraint on the SalesPersonID column of the CK_Sales.Invoices table.

Answer: A

Explanation: A foreign key constraint that references the CK_Sales.SalesPersons table will ensure that only SalesPersonIDs that exist in the SalesPersons table may be used in the Invoices table.

Incorrect Answers:

B: A check constraint defines the value that can be entered into a column.

C: A DML trigger fires when a UPDATE, INSERT or DELETE statement is run against a table. This is not what is required.

D: A unique constraint ensures that a value in a column only appears once in that column. This is not what is required.

Reference:

Microsoft SQL Server 2005 Books Online (2007), Index: constraints [SQL Server]

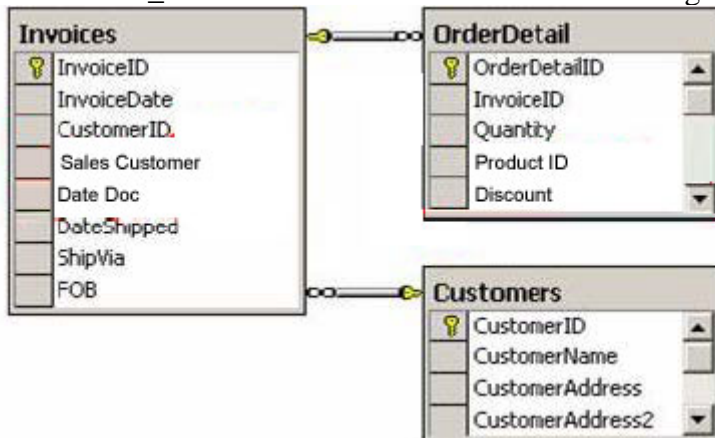
Microsoft SQL Server 2005 Books Online (2007), Index: constraints [SQL Server], vs.

DML triggers

Microsoft SQL Server 2005 Books Online (2007), Index: DML triggers

QUESTION 94:

You work as the database administrator at Certkiller .com. All servers on the Certkiller .com network run Windows Server 2003 and all database servers run SQL Server 2005. The Certkiller .com network contains a database server named Certkiller -DB01 B01 that has two RAID-5 arrays for storing data files. Certkiller -DB01 hosts a database named CK_Sales that stores sales data for the company. CK_Sales is also used for online transaction processing (OLTP). The tables in the CK_Sales database are shown in the following database diagram.



Each of the tables in the CK_Sales database contains more than one million rows. You have created a clustered index on each table. You have also created several nonclustered indexes on the Invoices and OrderDetail tables as these two tables are used very frequently. The OrderDetail and Customers tables are used in joins. You want to optimize performance for the database. You decide to create two filegroups named Filegroup1 and Filegroup2. You will place each filegroup on a separate RAID-5 array. You need to decide where to place the tables and indexes. What should you do? (Each correct answer presents part of the solution. Choose three.)

- A. Place the nonclustered indexes for the Invoices and OrderDetail tables on Filegroup1.
- B. Place the Invoices and OrderDetail tables on FileGroup1.
- C. Place the nonclustered indexes for the Invoices and OrderDetail tables on Filegroup2.
- D. Place the Customers table on Filegroup2.
- E. Place the OrderDetail and Customers tables on Filegroup2.

Answer: B, C, D

Explanation:

You can improve database performance by placing the frequently used tables on one filegroup, and their nonclustered indexes on another. In addition, you should place the two tables that are used in joins on separate disks. The Invoices and OrderDetail tables are frequently used so they should be placed on one filegroup while their nonclustered indexes should be placed on another filegroup. The OrderDetail and Customers tables are used in joins. Therefore these two tables should be on different file groups. As the OrderDetail table is on Filegroup1, the Customers table should be placed on Filegroup2.

Incorrect Answers:

A: You can improve database performance by placing the frequently used tables on one filegroup, and their nonclustered indexes on another filegroup. You only have the option of placing the Invoices and OrderDetail tables on FileGroup1; therefore the nonclustered indexes for the Invoices and OrderDetail tables should be placed on Filegroup2.

E: You can improve database performance by placing the two tables that are used in joins on separate filegroups. The OrderDetail and Customers tables are used in joins; therefore these two tables should be on different file groups.

Reference:

Microsoft SQL Server 2005 Books Online (2006), Index: filegroups [SQL Server]

Microsoft SQL Server 2005 Books Online (2006), Index: filegroups [SQL Server], index placement

QUESTION 95:

You work as the database developer for Certkiller .com. The Certkiller .com network contains a database server named Certkiller -DB01. Certkiller -DB01 hosts a database named CK_Sales that stores sales data for the company. Certkiller .com Sales Representatives should be allowed to check the current commissions due to them.

You need to optimize the indexing strategies for the CK_Sales database. You need to design the indexes for the Orders table. The following query is frequently executed, though it is not the most commonly executed query.

```
SELECT Salesrepresentative, SUM(Commission)
FROM Orders
WHERE Date BETWEEN @startDate AND @endDate
```

```
GROUP BY Salesrepresentative
```

```
ORDER BY Salesrepresentative
```

You need to use the appropriate statement to create the best index to accommodate this query.

What statement should you use?

A. CREATE INDEX ix_Commission
ON Orders(Salesrepresentative, Date, Commission)

B. CREATE CLUSTERED INDEX ix_Commission
ON Orders(Salesrepresentative, Date)
C. CREATE INDEX ix_Commission
ON Orders(Date)
INCLUDE (Salesrepresentative);
D. CREATE INDEX ix_Commission
ON Orders(Date, Salesrepresentative)
INCLUDE (Commission);

Answer: D

Explanation: The Date column is used to select the records and the Salesrepresentative column is used to group and order the records. This means that both these columns have to be key columns. The Date column is used in a BETWEEN comparison, thus it should also be the first in the query. Furthermore, it has higher selectivity than the Salesrepresentative column.

Incorrect Answers:

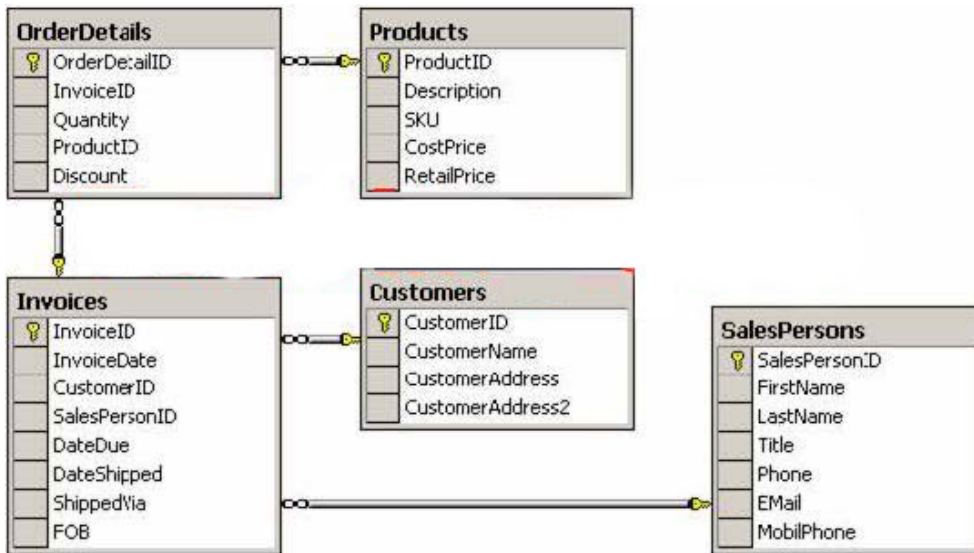
A: The column used for equality or BETWEEN comparisons should be listed first. This must then be followed by the most selective column, then the rest of the predicate columns in order of decreasing selectivity. And, although it is possible that you can create an index by using a computed column as the key column, it is recommended that key columns be kept as narrow as possible. This means that making use of an included column for Commission is a better option.

B: A Clustered index should have high selectivity. The Salesrepresentative column does not have high selectivity. This means that this clustered index will not be appropriate for any other queries done against the table like retrieving order information for instance. A table can only have one unique index.

C: The Salesrepresentative column is used in the GROUP BY clause. Thus it would be better suited as the key column rather than an included column. Furthermore, this is not a covering index as it does not include Commission as an included column.

QUESTION 96:

You work as the database administrator at Certkiller .com. All servers on the Certkiller .com network run Windows Server 2003 and all database servers run SQL Server 2005. The Certkiller .com network contains a database server named Certkiller -DB01. Certkiller -DB01 hosts a database named CK_Sales that stores sales data for the company. The tables in the CK_Sales database are shown in the following database diagram.



You need to optimize the CK_Sales database to support a redesigned stored procedure. You need add a key column to an existing clustered index on the OrderDetails table but you must ensure that the table remains available to users while the index is built. What should you do?

- A. Run the CREATE INDEX statement with the DROP EXISTING and ONLINE options set to ON.
- B. Disable the index. Then run the CREATE INDEX statement with the DROP EXISTING and ONLINE options set to ON
- C. Run the ALTER INDEX statement with the REORGANIZE keyword and the ONLINE option set to ON.
- D. Run the ALTER INDEX statement with the REBUILD keyword and the ONLINE option set to ON.

Answer: A

Explanation:

You must drop the clustered index and recreate it if you want to add a column to an existing clustered index. You can drop and recreate an existing index using the CREATE INDEX statement and setting the DROP EXISTING option to ON. To ensure that users can access the table while the index is being built, you must also set the ONLINE option to ON.

Incorrect Answers:

- B: A clustered index cannot be rebuilt online if the index is disabled.
- C, D: You cannot use the ALTER INDEX statement to add a column to a clustered index. You must drop the clustered index and recreate it if you want to add a column to an existing clustered index.

QUESTION 97:

You work as the database administrator at Certkiller .com. All servers on the Certkiller .com network run Windows Server 2003 and all database servers run SQL Server 2005. The Certkiller .com network contains a database server named Certkiller -DB01.

Certkiller -DB01 hosts a database named CK_News that stores news items from the company. The CK_News database has a table named NewsStories that stores the news story text. These news stories are written in Microsoft Word. A data capturer uses an application to copy the news story to the CK_News database. You need to redesign the NewsStories table so that you can implement full-text searches on it. You backup the CK_News database and drop the NewsStories table. What Transact-SQL statement should you use to recreate the table?

- A. CREATE TABLE NewsStories
(
NewsID bigint UNIQUE,
ReporterID int REFERENCES Employees.StaffID,
NewsText varchar(max),
Date smalldatetime NOT NULL
)
- B. CREATE TABLE NewsStories
(
NewsID bigint PRIMARY KEY CLUSTERED,
ReporterID int REFERENCES Employees.StaffID,
NewsText xml,
Date smalldatetime NOT NULL
)
- C. CREATE TABLE NewsStories
(
NewsID bigint PRIMARY KEY CLUSTERED,
ReporterID int REFERENCES Employees.StaffID,
NewsText varbinary(max),
FileType varchar(8),
Date smalldatetime NOT NULL
)
- D. CREATE TABLE NewsStories
(
NewsID bigint UNIQUE,
ReporterID int REFERENCES Employees.StaffID,
NewsText image,
FileExtension varchar(8),
Date smalldatetime NOT NULL
)

Answer: C

Explanation:

You should create the NewsID column as the primary key so that it can be used as the key column for the full-text search. Because the data capturer uses an application to copy the news story from Microsoft Word to the NewsStories table, you need to define the NewsText column as either varchar(max) or image. To enable Full-text indexing on a Microsoft Word document, you also need a column that stores the file extension so that the full-text search uses the correct filter to read the document.

Incorrect Answers:

A: Full-text search requires a key column that contains unique, non-null values. A UNIQUE constraint does not ensure that the column has no null values; it limits null values to one row. Because the data capturer uses an application to copy the news story from Microsoft Word to the NewsText table, you need to define the NewsText column as either varchar(max) or image and not varchar(max). Also, to enable Full-text indexing on a Microsoft Word document, you need a column that stores the file extension so that the full-text search uses the correct filter to read the document.

B: Because the Web master uses an application to copy the news story from Microsoft Word to the NewsStories table, you need to define the NewsText column as either varchar(max) or image and not xml. Also, to enable Full-text indexing on a Microsoft Word document, you need a column that stores the file extension so that the full-text search uses the correct filter to read the document.

D: Full-text search requires a key column that contains unique, non-null values. A UNIQUE constraint does not ensure that the column has no null values; it limits null values to one row.

QUESTION 98:

You work as the database administrator at Certkiller .com. All servers on the Certkiller .com network run Windows Server 2003 and all database servers run SQL Server 2005. The Certkiller .com network contains a database server named Certkiller -DB01 B01 that has two RAID-5 arrays for storing data files. Certkiller -DB01 hosts a database named CK_Products that stores data on the products sold by Certkiller .com.

Certkiller .com wants to allow prospective customers search for information in the CK_Products database via a Web-based interface. The prospective customers will use key words and phrases to search the CK_Products database. You decide to implement a Full-Text Search on the database. You install the Full-Text Search component on Certkiller -DB01 and start the Microsoft Search service. You then create the required catalog and full-text index. You modify the full-text index by running the following Transact-SQL statement:

```
ALTER FULLTEXT INDEX ON CK_Products.Description  
ADD (shortName)  
WITH NO POPULATION
```

Later you discover that you cannot perform full-text queries against the shortName column. You need to ensure that full-text queries can be performed against the shortName column.

What should you do?

- A. Drop and recreate the full-text index.
- B. Run the following statement:
`ALTER FULLTEXT INDEX ON CK_Products.Description
ENABLE`
- C. Stop and restart the SQL Server Agent.
- D. Run the following statement:
`ALTER FULLTEXT INDEX ON CK_Products.Description
START FULL POPULATION`

Answer: D

Explanation: If the WITH NO POPULATION clause is used you add or drop a column from a full-text index, the full-text index will not be populated after the ADD or DROP column operation. The full-text index will only be populated if you run the START...POPULATION command.

Incorrect Answers:

A: The full-text queries against the shortName column failed because the index was not populated after the column was added. You can drop and recreate the full-text index but it would be easier to repopulate the index. You can run the START...POPULATION command to repopulate the index.

B: This statement enables the full-text index; however, the full-text index is enabled by default.

C: The full-text queries against the shortName column failed because the index was not populated after the column was added. Stopping and restarting the SQL Server Agent will not repopulate the full-text index.

Reference:

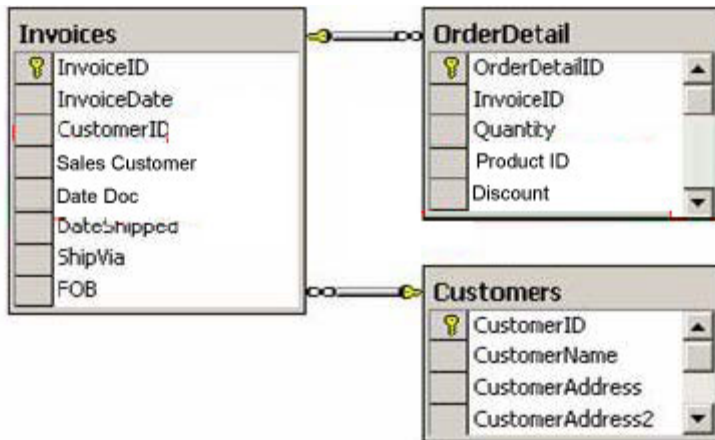
Microsoft SQL Server 2005 Books Online (2007), Index: ALTER FULLTEXT INDEX statement

Microsoft SQL Server 2005 Books Online (2007), Index: full-text indexes [SQL Server]

QUESTION 99:

You work as the database administrator at Certkiller .com. All servers on the Certkiller .com network run Windows Server 2003 and all database servers run SQL Server 2005. Certkiller .com has its headquarters in Chicago and branch offices in Miami, and San Francisco. The Certkiller .com network contains three database servers named Certkiller -DB01, Certkiller -DB02, and Certkiller -DB03.

Certkiller -DB01 hosts a database named CK_Sales that stores sales data for the company. Transactional replication of the CK_Sales database is configured with Certkiller -DB01 as the Publisher and Certkiller -DB02 and Certkiller -DB03 are configured as Subscribers. Certkiller -DB01 is located at headquarters while Certkiller -DB02 is located in Miami and Certkiller -DB03 is located in San Francisco. The tables in the CK_Sales database are shown in the following database diagram.



The Customers table contains more than 200,000 rows. Certkiller .com uses a custom application to access data in the Customers table. Certkiller .com users in the Sales department at each office use the custom application to access and update data for the customers assigned to his or her office.

You want to improve database performance by implementing views on the Customers table.

What type of view should you implement?

- A. Partitioned views.
- B. Replication views.
- C. Indexed views.
- D. Distributed partitioned views.

Answer: D

Explanation: Distributed partitioned views allow you to partition data in the Customers table into smaller tables based on the different offices, and you can distribute the tables across the database servers in each office. This also allows the servers to share the query processing load.

Incorrect Answers:

- A: Partition views reside on a single server while distributed partitioned views can be distributed across the database servers in each office. This allows the servers to share the query processing load and greatly improves performance.
- B: Replication views do not improve performance.
- C: Indexed views are used to improve performance of queries run against a view but you can achieve even better performance by using distributed partitioned views. Distributed partitioned views allow the three database servers to share the query processing load.

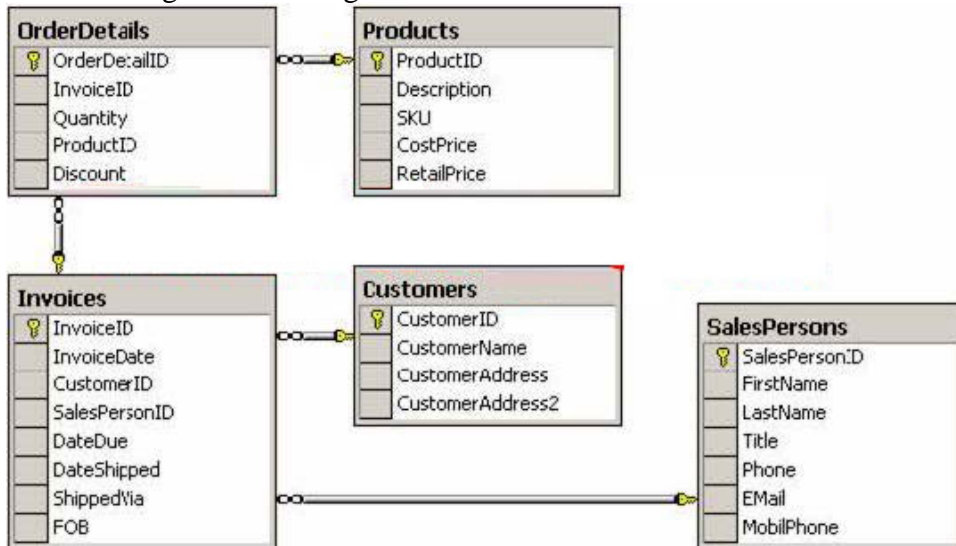
Reference:

Microsoft SQL Server 2005 Books Online (2007), Index: views [SQL Server]

QUESTION 100:

You work as the database administrator at Certkiller .com. All servers on the Certkiller .com network run Windows Server 2003 and all database servers run SQL Server 2005. The Certkiller .com network contains a database server named

Certkiller -DB01 B01. Certkiller -DB01 hosts a database named CK_Sales that stores sales data for the company. The tables in the CK_Sales database are shown in the following database diagram.



The CK_Sales database currently uses a single filegroup. Certkiller .com decides to switch to commission-based salaries for its Sales staff. The manager of the Sales department will run queries against the Invoices table to calculate the commission for each sales person. To improved database performance, you decide to partition the Invoices table by SalesPersonID and store each partition in a separate filegroup. You also want to create a partitioned index for the Invoices table.

What should you do before you can create the partition scheme? (Each correct answer presents part of the solution. Choose two.)

- A. Create the partition function.
- B. Create the index.
- C. Create the filegroups.
- D. Create the table.

Answer: A, C

Explanation: You need to create the partition function and the filegroups that you want to use in the partition scheme before you can create the partition scheme.

Incorrect Answers:

B, D: You can create the partition scheme before you create the tables and indexes.

Reference:

Microsoft SQL Server 2005 Books Online (2007), Index: partitions [SQL Server]

Microsoft SQL Server 2005 Books Online (2007), Index: partitioned views [SQL Server]

QUESTION 101:

SIMULATION

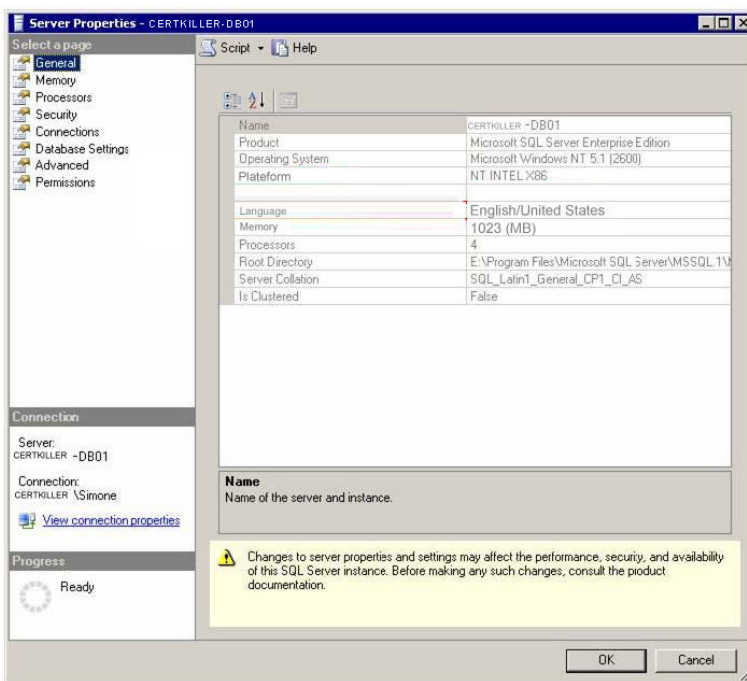
You work as the database administrator at Certkiller .com. All servers on the Certkiller .com network run Windows Server 2003 and all database servers run SQL

Server 2005. The Certkiller .com network contains a database server named Certkiller -DB01 B01. Certkiller -DB01 hosts a database named CK_Sales that stores sales data for the company. CK_Sales is the only non-system database on Certkiller -DB01 and is accessed through Certkiller .com's e-Commerce Web site. You need to prevent users from executing nested trigger against the CK_Sales database.

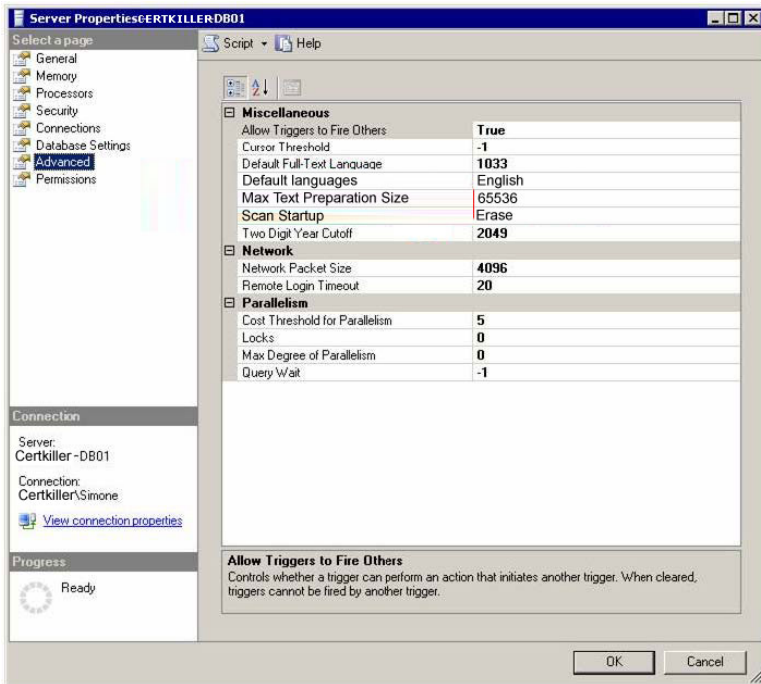
What should you do? (To answer, modify the appropriate server and database settings in SQL Server Management Studio (SSMS)).

Answer:

In SQL Server Management Studio (SSMS), right click the database server named Certkiller -DB01 in the Object Explorer pane and select Properties from the context menu.



In the left-hand pane, under Select a page, click on Advanced.

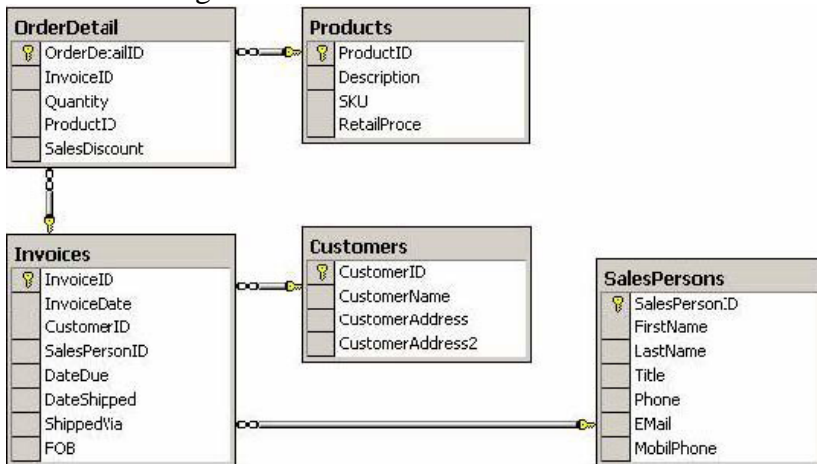


In the right-hand pane, click on the Allow Triggers to Fire Others option and select False from the drop down list.

QUESTION 102:

SIMULATION

You work as the database administrator at Certkiller .com. All servers on the Certkiller .com network run Windows Server 2003 and all database servers run SQL Server 2005. The Certkiller .com network contains a database server named Certkiller -DB01 B01. Certkiller -DB01 hosts a database named CK_Sales that stores sales data for the company. The tables for the CK_Sales database are shown in the following exhibit.



There is a unique clustered index on the identity column of each table.

Certkiller .com users complain that a query they run against the OrderDetails table in the CK_Sales database takes a long time to complete. The query is shown in the

following exhibit:

```
SELECT Products.ProductID, Products.RetailPrice,  
OrderDetails.ProductID, OrderDetails.Quantity  
FROM OrderDetails INNER JOIN Products  
ON OrderDetails.ProductID = Products.ProductID  
WHERE OrderDetails.Quantity > 10  
GROUP BY OrderDetails.ProductID
```

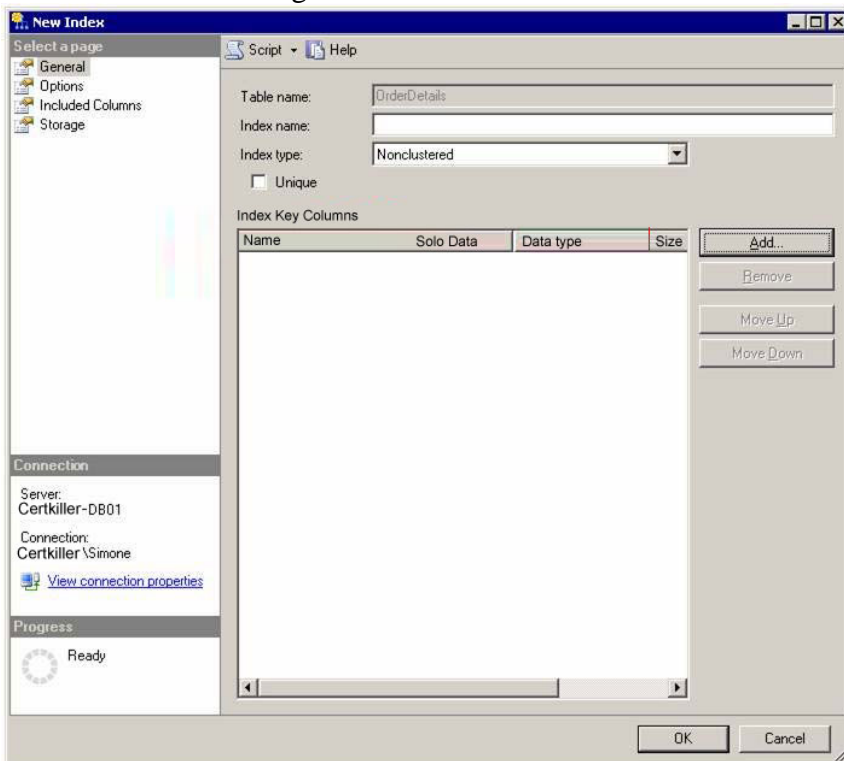
You need to improve the performance of the query. You want to accomplish this by creating an index on the OrderDetails table. However, you cannot drop the unique clustered index on the identity column. You need to ensure that your index consumes as little disk space as possible. You must also use as few columns as possible in the index.

What should you do? (To answer, create the required indexes in the New Index dialog box.)

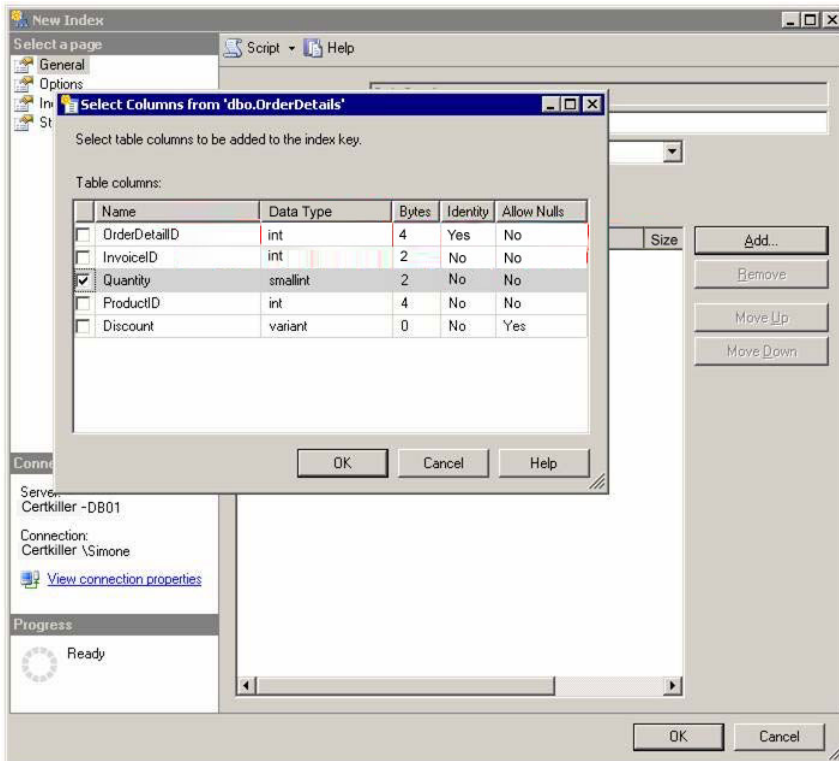
Answer:

In SQL Server Management Studio (SSMS), navigate to and expand the CK_Sales database. Expand the Tables node and then the OrderDetails table. Right click the Indexes node under the OrderDetails node and select New Index from the context menu to open the New Index dialog box.

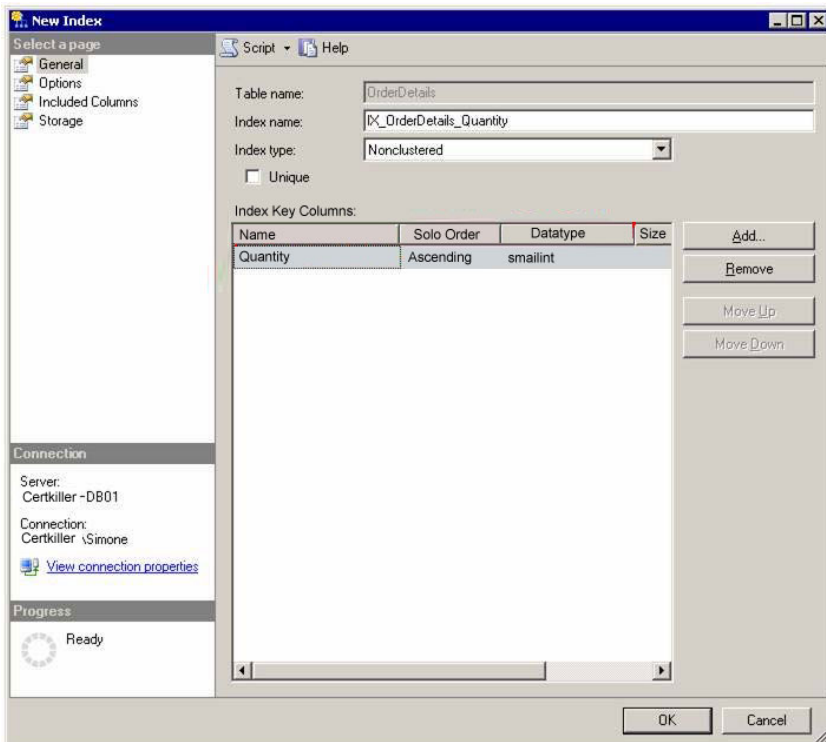
In the General page of the New Index dialog box, click on the Add ... button to open the column selection dialog box.



Select Quantity in the Select Columns for 'dbo.OrderDetails' dialog box and then click OK.



Enter a name for the nonclustered index and click OK on the New Index dialog box.



Explanation:

The slow query uses a field called Quantity in the WHERE clause, you need to create a nonclustered index on the Quantity field to improve performance.

QUESTION 103:

SIMULATION

You work as the database developer for Certkiller .com. The Certkiller .com network contains a database server named Certkiller -DB01. Certkiller -DB01 hosts a database named CK_Sales that stores sales data for the company. Certkiller .com Sales Representatives should be allowed to check the current commissions due to them.

You need to optimize the indexing strategies for the CK_Sales database. You need to design the indexes for the Orders table. The following query is frequently executed, though it is not the most commonly executed query.

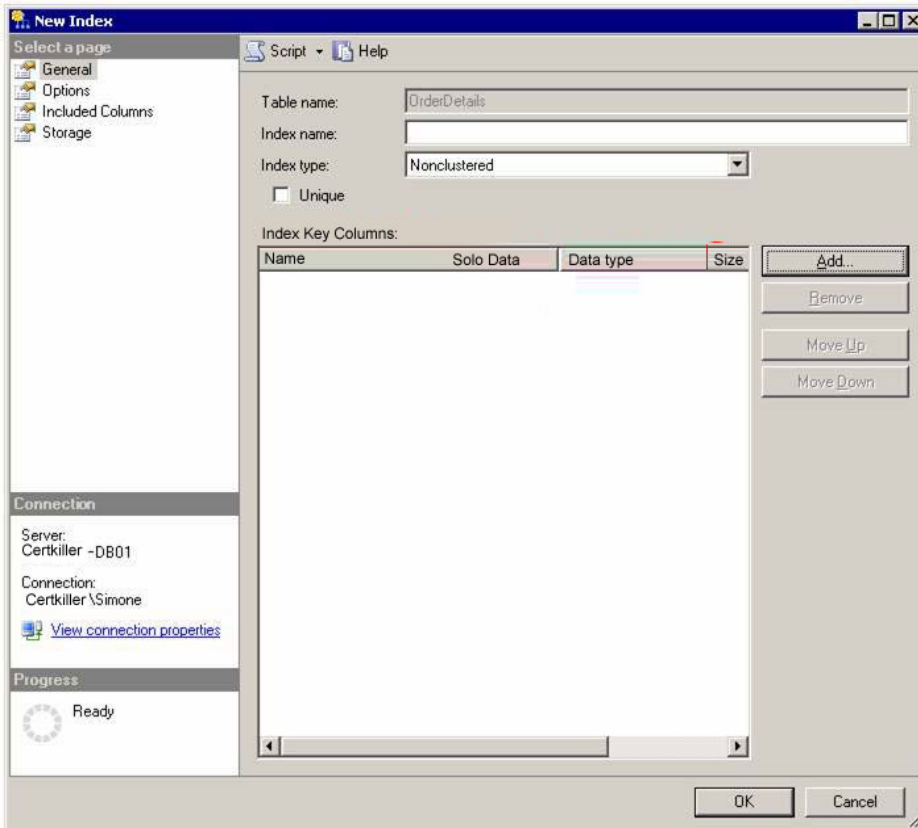
```
SELECT SalesPerson, SUM(Commission)
FROM Orders
WHERE InvoiceDate BETWEEN @startDate AND @endDate
GROUP BY SalesPerson
ORDER BY SalesPerson
```

What should you do? (To answer, create the required indexes in the New Index dialog box.)

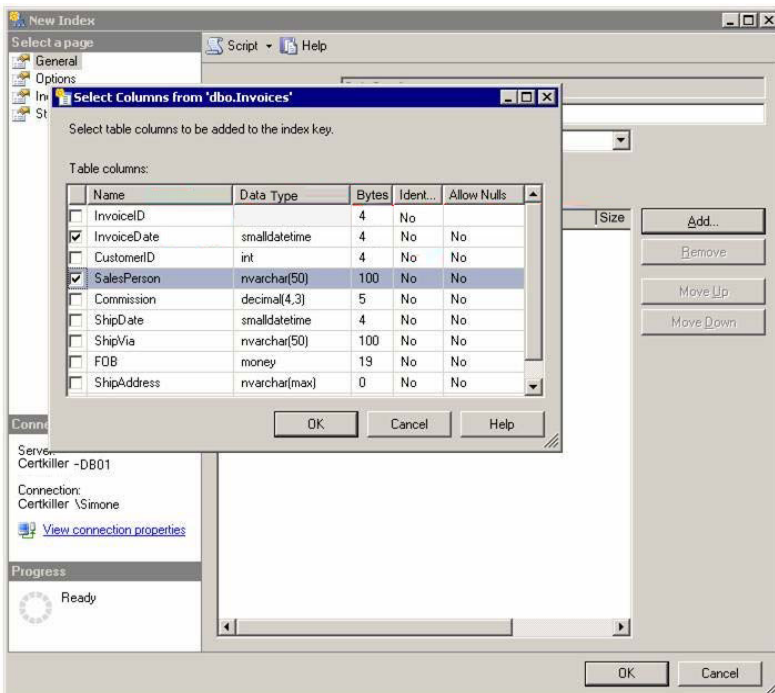
Answer:

In SQL Server Management Studio (SSMS), navigate to and expand the CK_Sales database. Expand the Tables node and then the Invoices table. Right click the Indexes node under the Invoices node and select New Index from the context menu to open the New Index dialog box.

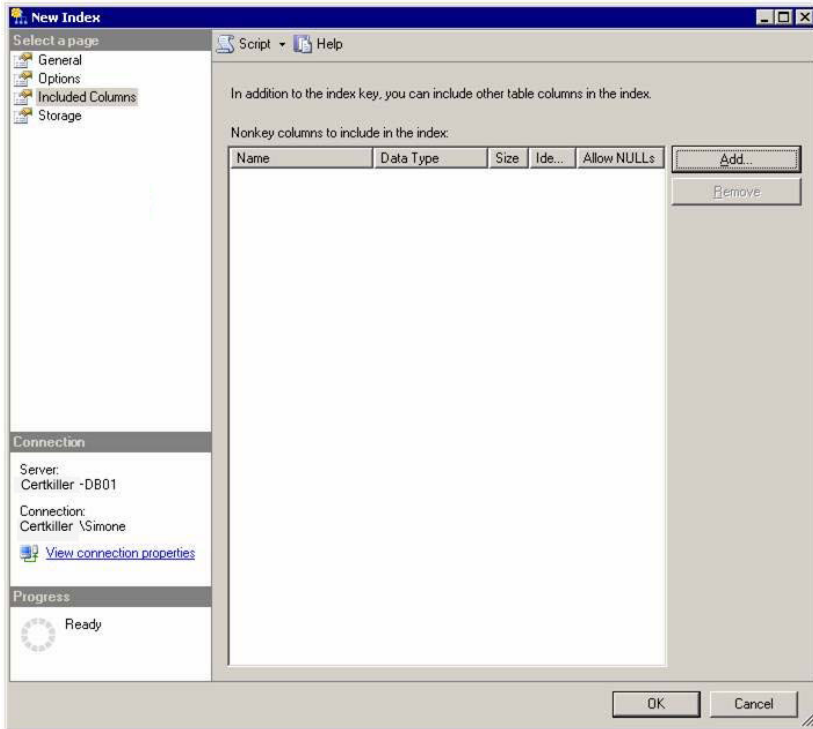
In the General page of the New Index dialog box, click on the Add ... button to open the column selection dialog box.



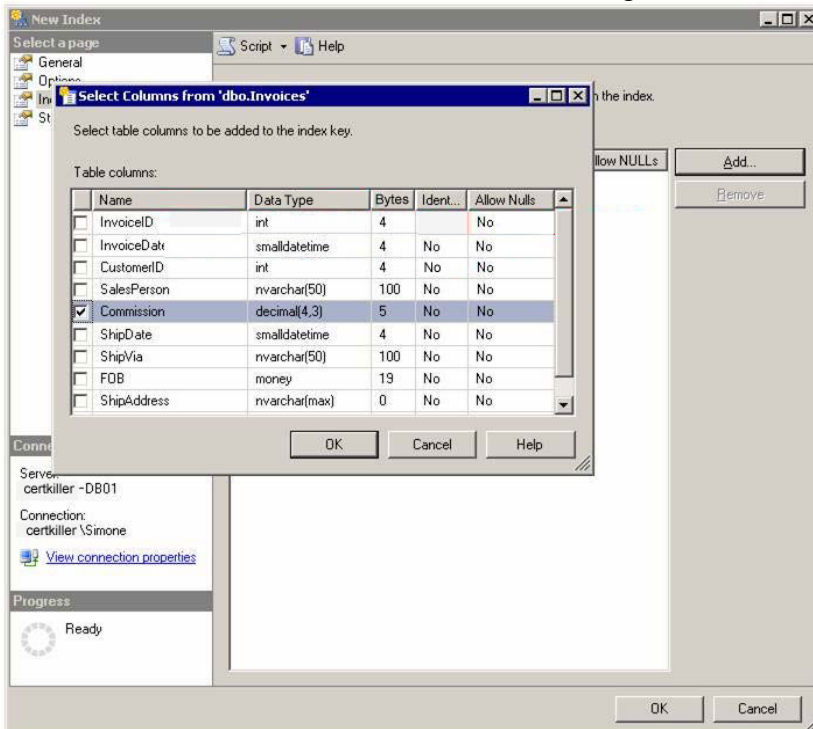
Select InvoiceDate and SalesPerson from the Select Columns for 'dbo.Invoices' dialog box and then click OK.



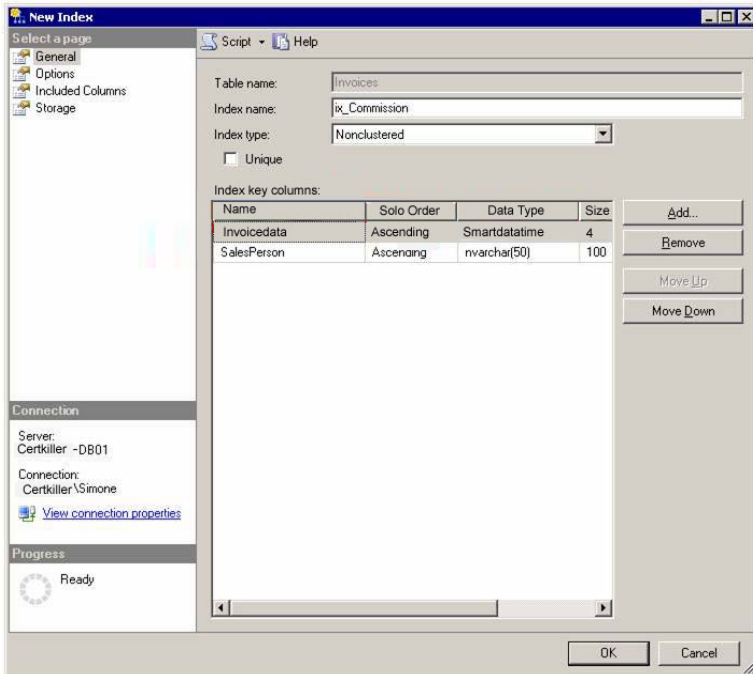
In the New Index dialog box, click on Include Columns in the left-hand pane, under Select a page.



Then click the Add ... button to open the column selection dialog box, select Commission from the Select Columns for 'dbo.Invoices' dialog box and then click OK.



Return to the General page of the New Index dialog box by clicking on General in the left-hand pane, enter a name for the new index, and click OK.



Explanation: The InvoiceDate column is used to select the records and the SalesPerson column is used to group and order the records. This means that both these columns have to be key columns. The InvoiceDate column is used in a BETWEEN comparison, thus it should also be the first in the query. Furthermore, it has higher selectivity than the SalesPerson column.

QUESTION 104:

SIMULATION

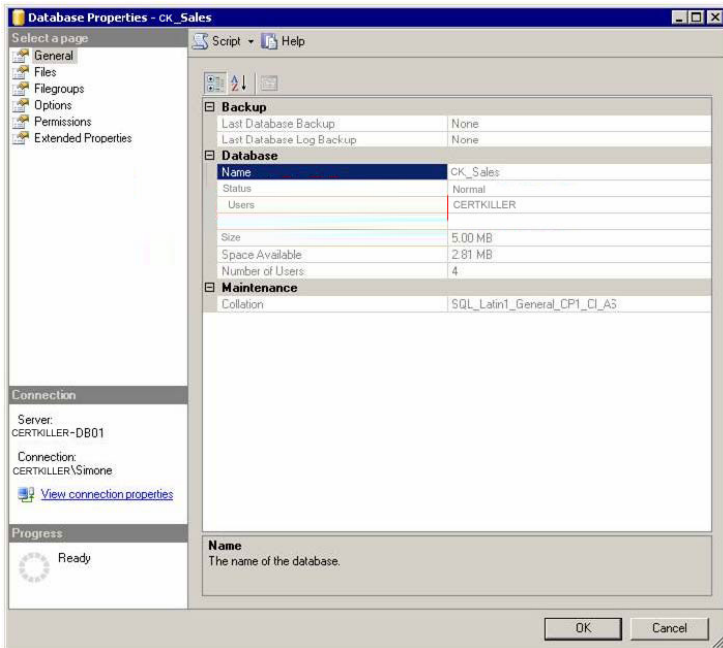
You work as the database administrator at Certkiller .com. All servers on the Certkiller .com network run Windows Server 2003 and all database servers run SQL Server 2005. The Certkiller .com network contains a database server named Certkiller -DB01 B01. Certkiller -DB01 hosts a database named CK_Sales that stores sales data for the company.

The backup routine for the CK_Sales database includes transaction log backups. At the close of every business day, two SQL Server Integration Services (SSIS) packages import thousands of rows of data into the CK_Sales database. You need to minimize the log space that is used by these import operations.

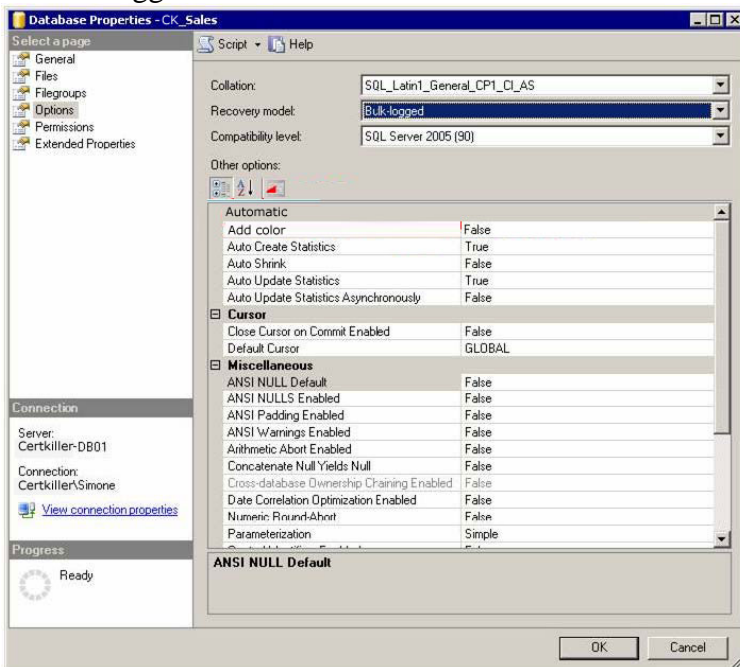
What should you do? (To answer, modify the appropriate settings by using SQL Server Management Studio (SSMS))

Answer:

In SQL Server Management Studio (SSMS), expand the database server named Certkiller -DB01 in the Object Explorer pane and right-click on the CK_Sales database. Then select Properties from the context menu to open the Properties dialog box for the CK_Sales database.



In the left-hand pane, under Select a page, click on Options and set the Recovery Model to Bulk-logged and click OK.



QUESTION 105:

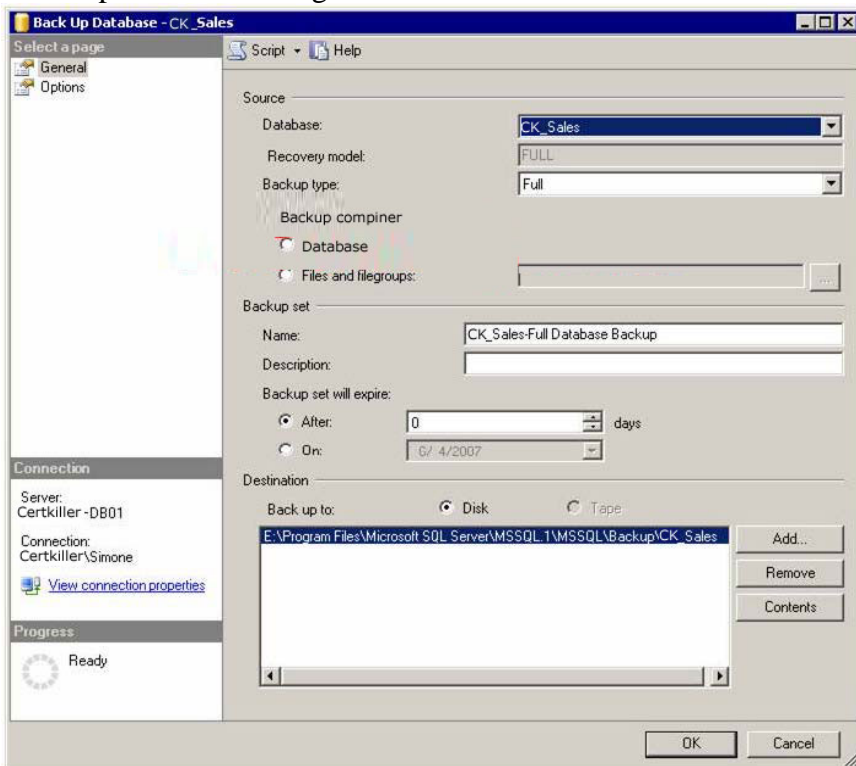
SIMULATION

You work as the database administrator at Certkiller .com. All database servers on the Certkiller .com network run SQL Server 2005. The Certkiller .com network contains a database server named Certkiller -DB01 B01 that runs on Windows 2000 Server. Certkiller -DB01 hosts a database named CK_Sales that stores sales data for the company.

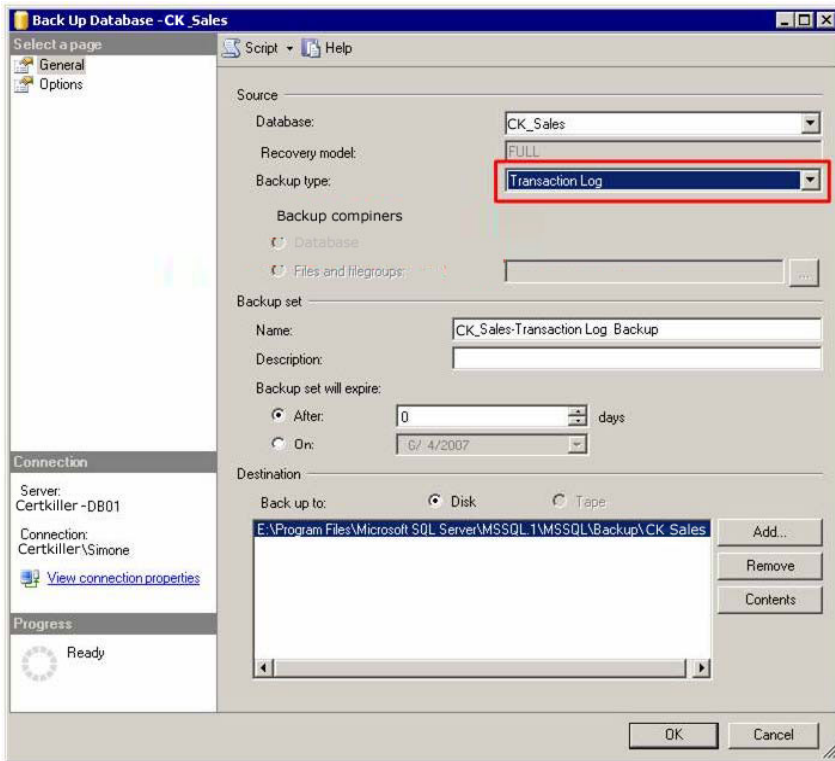
You are required to move the CK_Sales database to a different instance of SQL Server 2005. Before you move the database, you must perform a backup that contains all of the changes since the last transaction log backup. You must ensure that no changes can be applied to the database once the backup is completed. What should you do? (To answer, to perform the backup by using the Back Up Database dialog box.)

Answer:

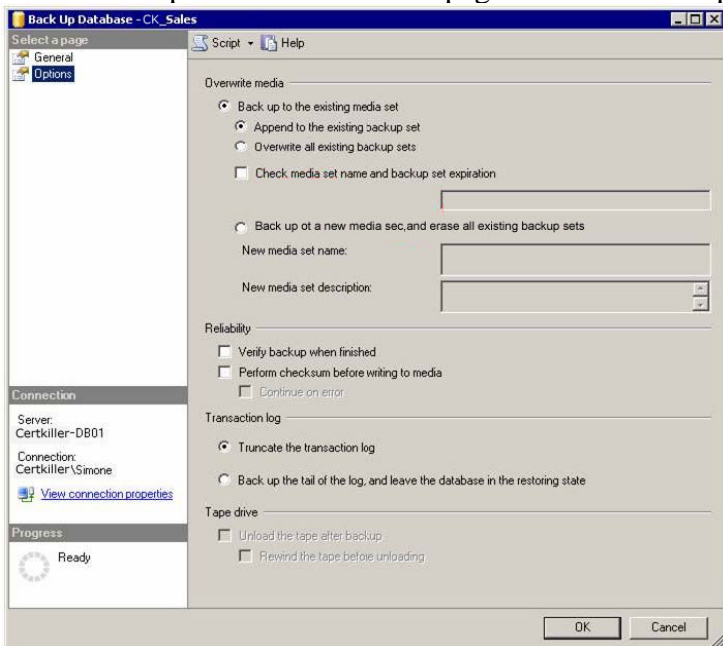
In SQL Server Management Studio (SSMS), expand the database server named Certkiller -DB01 and the Database node in the Object Explorer pane. Right-click on the CK_Sales database, point to Tasks on the context menu and select Back Up... to open the Back Up Database dialog box.



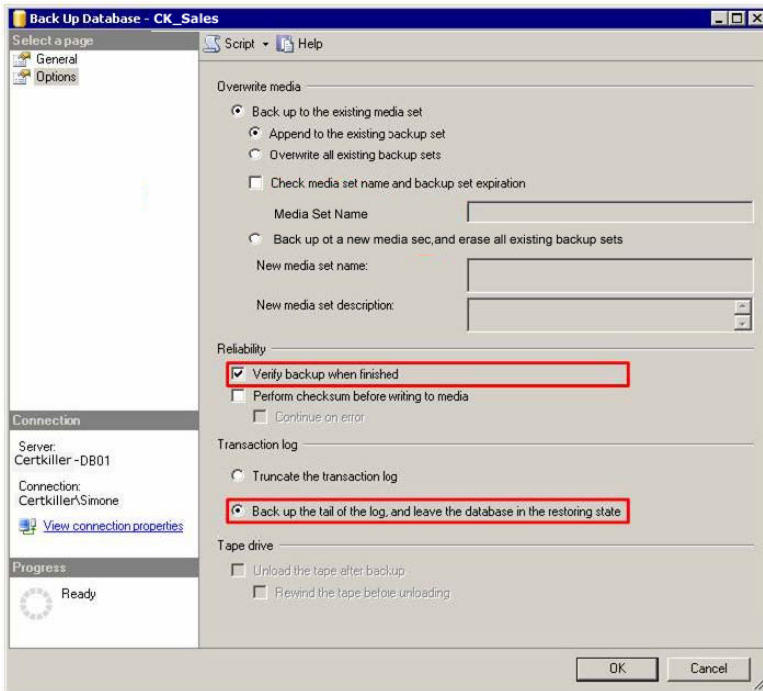
On the General page of the Back Up Database dialog box, change the Backup type to Transaction Log.



Then click Options under Select a page in the left-hand pane.



In the Options page, select the Verify backup when finished check box under Reliability, and Back up the tail of the log, and leave the database in the restoring state radio button under Transaction log.



Then click OK.

QUESTION 106:

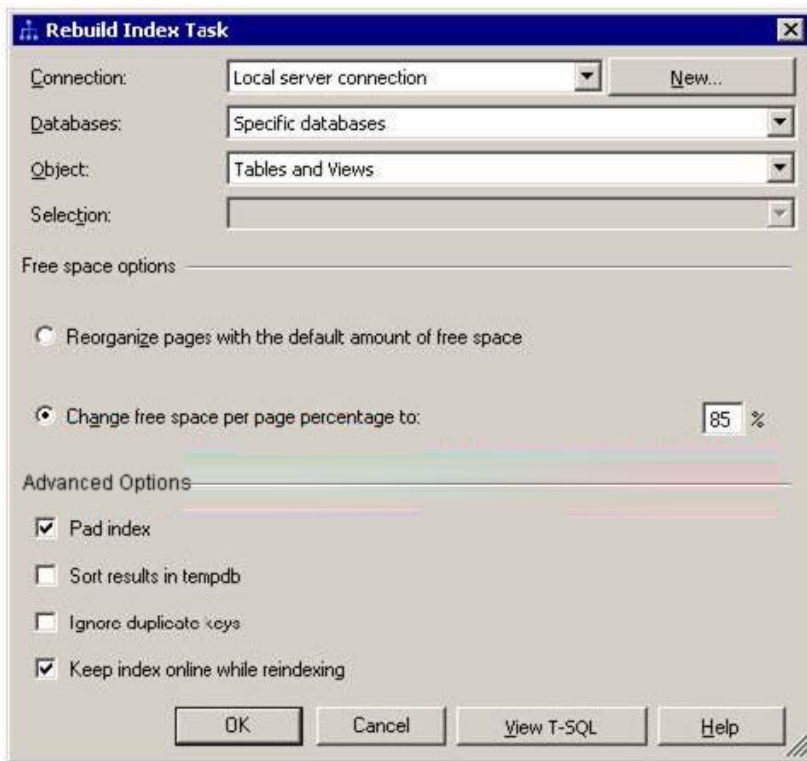
SIMULATION

You work as a database administrator at Certkiller .com. The Certkiller .com network contains a database server named Certkiller -DB01. Certkiller -DB01 hosts a database named CK_Sales that stores sales data for the company. Certkiller .com Sales Representatives should be allowed to check the current commissions due to them.

You have optimized the indexing strategies for the CK_Sales database but you are concerned about performance degradation of the indexes. You want to design a maintenance plan to rebuild the indexes in the CK_Sales database. You want the indexes to be rebuilt with a fill factor of 85 at both the leaf level and the balanced tree of the index. You also need to ensure that all tables in the database remain accessible by users while the maintenance plan is running.

What should you do? (To answer, configure the appropriate settings on the Define Rebuild Index Task page in the Maintenance Plan Wizard.)

Answer:



QUESTION 107:

SIMULATION

You work as the administrator at Certkiller .com. Certkiller .com has purchased a new server named Certkiller -DB01. The hard disk subsystem on Certkiller -DB01 is shown in the following table.

Volume	RAID level	Disk size
C:	RAID-1	2 × 10GB
D:	RAID-1	2 × 10GB
E:	RAID-0	2 × 20GB

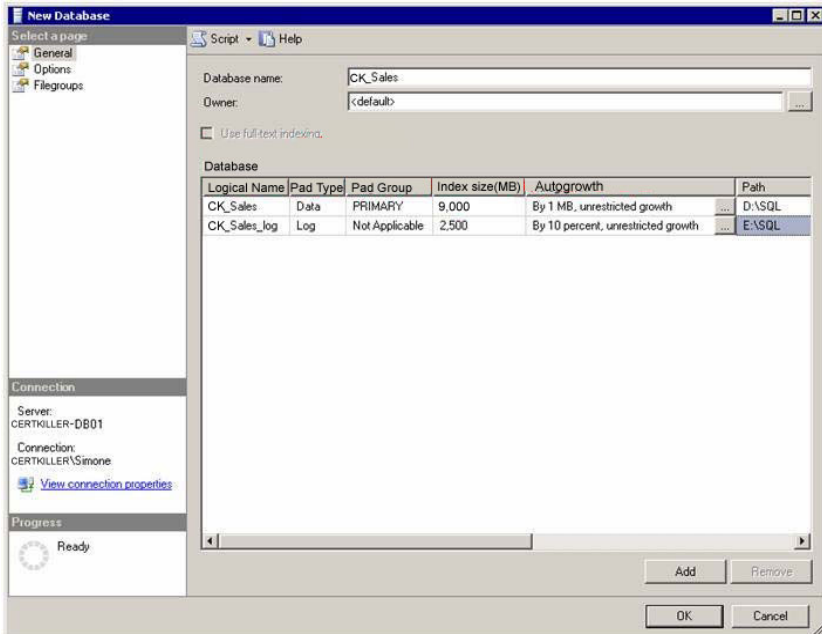
You are required to install Windows Server 2003 and SQL Server 2005 on separate volumes on Certkiller -DB01. You must then create a database named CK_Sales that will store sales data for Certkiller .com. The data file and log file should be stored in a folder named SQL in the root of different volumes and should have the default settings for automatic file growth and maximum file size. The data file and log file must be 9,000 MB and 2,500 MB in size respectively. The transaction log must be stored on a fault-tolerant volume and must be set to be automatically truncated. No database files should be stored on the system drive. What should you do? (To answer, create the CK_Sales database using the SQL Server Management Studio (SSMS).)

Answer:

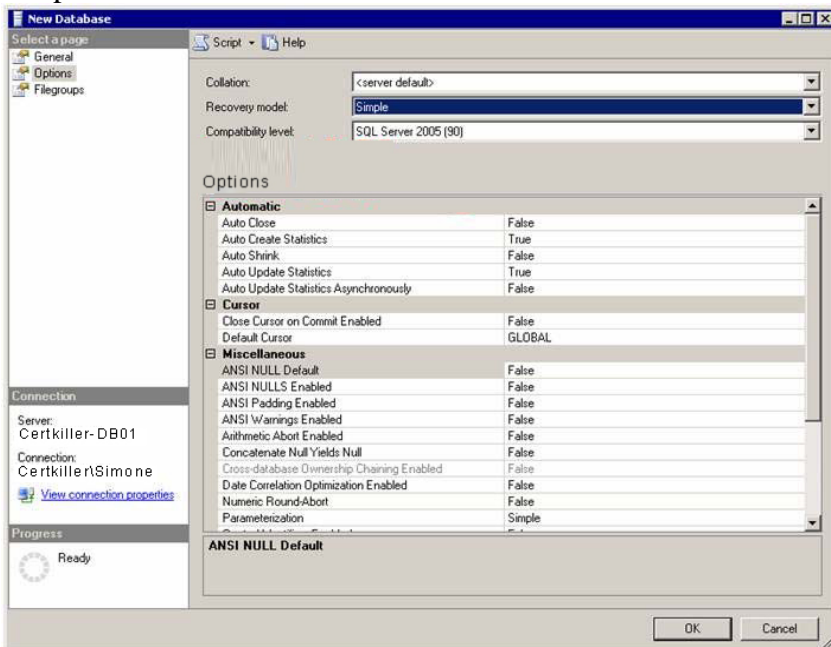
Right-click the Database node in the Object Explorer of SQL Server Management Studio

(SSMS) and select New Database ... from the context menu to open the New Database dialog box.

In the General page of the New Database dialog box, enter CK_Sales in the Database name text box and set the Initial Size for the data file to 9,000 and the log file to 2,500. Then set the Path for the data file to D:\SQL and the path for the log file to [E:\SQL](#).



Then select Options under Select a page and ensure that the Recovery Model is set to Simple and click OK.



QUESTION 108:

SIMULATION

You work as a database administrator at Certkiller .com. The Certkiller .com network contains two SQL Server 2005 Servers named Certkiller -DB01 and Certkiller -DB02. Certkiller -DB01 contains a database named CK_Staff that stores personnel information for the company.

The Sales department makes extensive use of the CK_Sales database while Certkiller .com users in the Accounts departments make use of an in-house application to access a database named CK_Personnel that runs Certkiller -DB02. Certkiller .com has upgraded and reconfigured the Accounting application. The new application will run on Certkiller -DB01 and must be able to run distributed queries against the CK_Sales database and the CK_Personnel database.

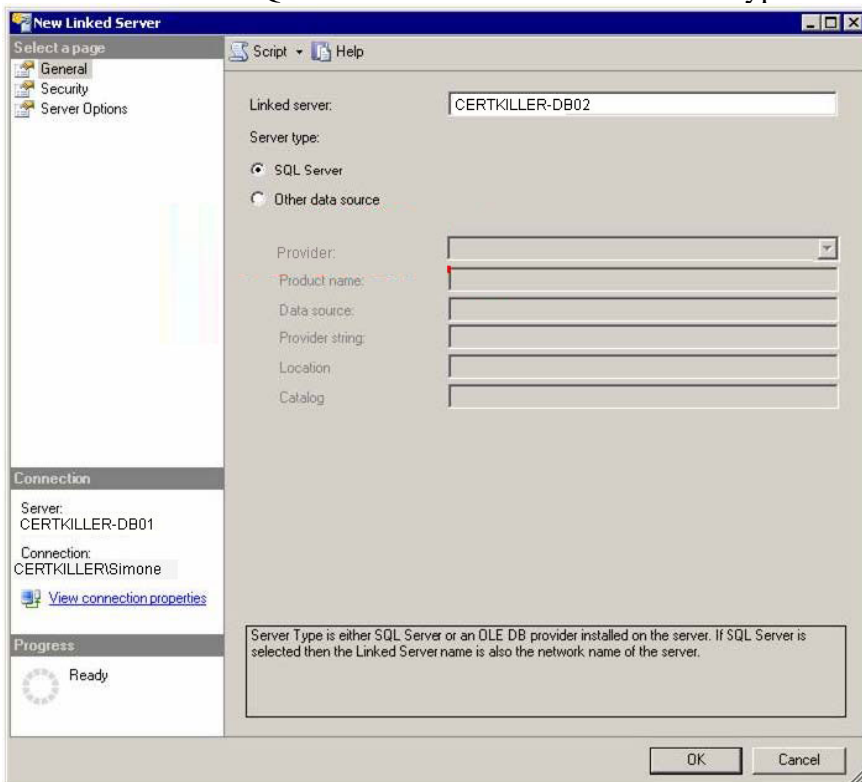
You have created a login named CK_App1 that the application will use to access Certkiller -DB02 through Certkiller -DB01. You need to provide the application with a login named CK_App2 and the password Certkiller that it will use for the linked server connection to Certkiller -DB02. You also need to ensure that the application can perform distributed queries. You must apply the fewest possible settings.

What should you do? (To answer, configure the appropriate settings in the SQL Server Management Studio (SSMS).)

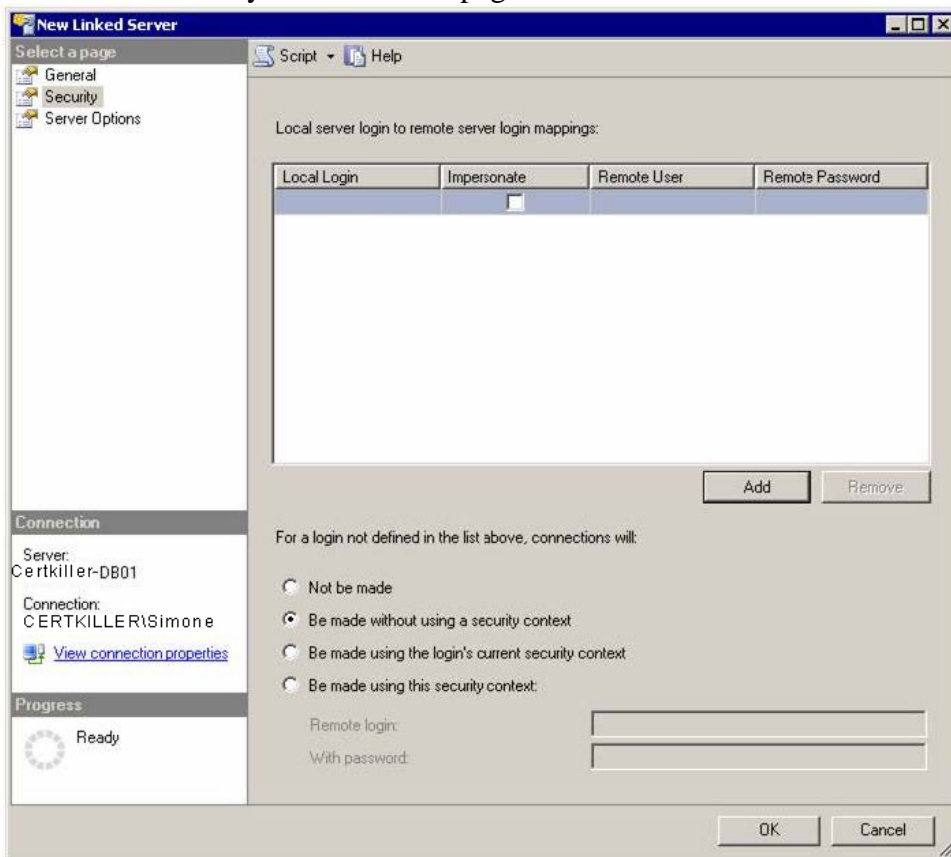
Answer:

In the Object Explorer of SQL Server Management Studio (SSMS), expand the Server Objects node and right-click on the Linked Servers node. Then select New Linked Server ... from the context menu to open the New Linked Sever dialog box.

In the New Linked Server dialog box, enter Certkiller -DB02 in the Linked server text box and select the SQL Server radio button under Server type.

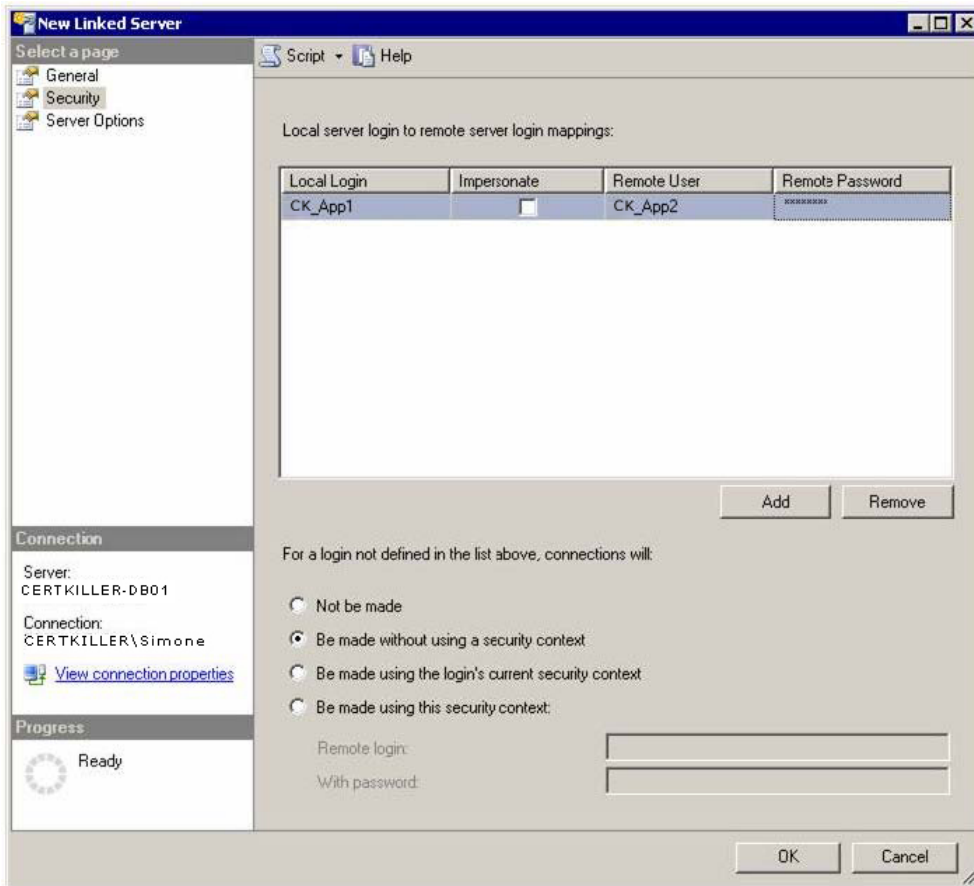


Then select Security under Select a page and click the Add button.



Select CK_App1 from the drop down list under Local Login, CK_App2 under Remote

User, and Certkiller under Remote Password. Then click OK.



QUESTION 109:

You are employed as a database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. You have received instruction from the CIO to prepare a new installation of SQL Server 2005. The new SQL Server 2005 database server is named Certkiller -DB01. A Certkiller .com user named Mia Hamm is a member of the Marketing department. You instruct her to identify the protocols that the client computers might use to connect to the server. Mia Hamm needs to identify the protocols that are needed to accomplish this goal. What should she identify? (Choose all that apply)

- A. She should use the Multiprotocol.
- B. Mia Hamm should use the Named Pipes.
- C. She should use the Shared Memory.
- D. She should use the TCP/IP
- E. Mia Hamm should use the Virtual Interface Adapter (VIA).

Answer: B, D

Explanation: The clients have by default TCP and Named Pipes as available protocols. You can make use of the SQL Server Client utility to manipulate the protocol ordering. The client application uses the protocols in the order specified on the client computer. When you make use of the SQL Server 2005, the protocol order is stored in the ProtocolOrder registry entry under the following registry subkey:

HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\MSSQLServer\Client\SN19.0

Incorrect Answers:

C: Shared Memory can only be used on the local computer and cannot be used as a network protocol.

E: The Virtual Interface Adapter (VIA) can be used only by VIA hardware.

Reference:

Microsoft SQL Server 2005 Books Online (2006), Index: client connections [SQL Server], about client network connections, choosing a network protocol

QUESTION 110:

Certkiller .com has employed you as a database administrator. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. You have been informed by the CIO that your application should access data that is located on two SQL Server 2005 computers. The two SQL Servers are named Certkiller -DB02 and Certkiller -DB03. The CIO informs you that you have the permissions to create a stored procedure on Certkiller -DB02 to support the application. However, on Certkiller -DB03, you only have permissions to select data.

A Certkiller .com user named Andy Reid is a member of the Research and Development department. You instruct Andy Reid to write the stored procedure on Certkiller -DB02. The stored procedure accesses Certkiller -DB03 by using the OPENQUERY Transact-SQL statement. Andy Reid discovers that the query fails when it is executed. Andy Reid needs to troubleshoot the origin of the error. What should he do?

- A. Andy Reid needs to add Certkiller -DB03 as a remote server to Certkiller -DB02.
- B. By making use of the four-part syntax of server database schema table Andy Reid should join the two servers.
- C. Andy Reid should add Certkiller -DB03 as a linked server to Certkiller -DB02.
- D. By using an alias Andy Reid should reference Certkiller -DB03.

Answer: C

Explanation: A linked server needs to be defined for every external data source you want to access. The security context needs to be configured under which the distributed queries will run. After creating a linked server you can make use of the Transact-SQL OPENQUERY function to execute the distributed queries.

QUESTION 111:

You are employed as a database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. The Certkiller .com network contains a SQL Server 2005 computer named Certkiller -DB01.

You have received instruction from the CIO to create an HTTP endpoint that will be used to provide customer data to external applications. A Certkiller .com user named Andy Booth is a junior technician appointed in the Research and Development department. You instruct him to create a stored procedure named dbo.usp_GetResData to retrieve the information in the CK_R&D database. Andy Booth creates the endpoint by using the following code.

```
CREATE ENDPOINT SQLEP_R&DData
AS HTTP
(PATH = '/R&DData',
AUTHENTICATION = (INTEGRATED),
PORTS = (CLEAR),
SITE = ' Certkiller -DB01')
FOR SOAP
(WEBMETHOD 'PersonData'\
(NAME=' Certkiller .dbo.usp_GetResData'),
BATCHES = DISABLED,
WSDL = DEFAULT,
DATABASE = 'CK_R&D',
NAMESPACE = 'http:// Certkiller .com/R&D')
```

You receive various complaints by users stating that they are unable to get any data when they connect to the endpoint. As the administrator you connect to the endpoint and discover that it is not responding. You instruct Andy Booth to modify the endpoint in order that data is returned as expected.

What should he do?

- A. In order to receive data Andy Booth must specify WSDL = 'pr_GetResData'.
- B. It is important that Andy Booth specify STATE = Started in order to rectify the data.
- C. In order to rectify the error Andy Booth should specify BATCHES = ENABLED.
- D. It is important that Andy Booth change the AUTHENTICATION property to KERBEROS.

Answer: B

Explanation: The possible states for an endpoint are STARTED, STOPPED, and DISABLED. For an endpoint to respond to requests, the state must be set to STARTED. To comply with the SQL Server 2005 "off by default" approach to security, the default state is STOPPED.

QUESTION 112:

You work as a database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. The

Certkiller .com network contains two SQL Server 2005 computers named Certkiller -DB01 and Certkiller -DB02

You have been informed by the CIO that the transaction log shipping occurs from Certkiller -DB01 to Certkiller -DB02 by using default SQL Server Agent schedule settings. A Certkiller .com user named Kara Lang is a junior technician appointed in the Research and Development department. You instruct Kara Lang to reconfigure the transaction log shipping to provide minimum latency on Certkiller -DB02.

What should Kara Lang do?

A. Kara Lang should reschedule the transaction log backup job on Certkiller -DB01 in order for it to occur every minute.

It is vital that the reschedule the transaction log backup job on Certkiller -DB01 in order for it occur every minute.

Thereafter she should reschedule the log shipping copy as well as the restore jobs on Certkiller -DB01 in order for it to occur every minute.

B. Kara Lang should maintain default schedule settings for the transaction log backup job on Certkiller -DB01.

It is important that she changes the schedule types for the log shipping copy on Certkiller -DB02.

Then she should restore jobs on Certkiller -DB02 to start automatically when SQL Server Agent commences.

C. Thereafter she should maintain default schedule settings on Certkiller -DB02 for both the log shipping copy and the restore jobs.

D. It is important that Kara Lang change the schedule type for the transaction log backup on Certkiller -DB01 to start automatically when SQL Server Agent starts.

Then she should change the schedule types for the log shipping copy on Certkiller -DB02.

Thereafter Kara Lang should change the schedule types on Certkiller -DB02 to restore jobs to start automatically, as soon as the SQL Server Agent commences.

Answer: A

Explanation: In order to reduce latency on Certkiller -DB02 Kara Lang needs to get a copy from Certkiller -DB01 as soon as possible after every write. One minute is the shortest possible interval for the transaction log backup. As soon as the backup is complete Kara Lang should copy it and restore it on Certkiller -DB02 as soon as possible.

QUESTION 113:

You are employed as a database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. Certkiller .com makes use of multiple servers in a distributed environment. The Certkiller .com network contains two SQL Server 2005 database server named Certkiller -DB01 and Certkiller -DB02. Both Certkiller -DB01 and

Certkiller -DB02 makes use of SQL Server Authentication and make use of different logins. A Certkiller .com employee named Kara Lang works in the Research and Development department. You instruct her to write a distributed query that joins the data on Certkiller -DB01 with the data on Certkiller -DB02. What should she do?

- A. She should run SQL Server Profiler over a Remote Desktop connection to Certkiller -DB01. Store the trace in a file on a separate server.
- B. Certkiller -DB02 needs to be configured as a distributed server. Thereafter Kara Lang should use pass-through authentication.
- C. Both database servers should use the same login name as the security context.
- D. Kara Lang should configure Certkiller -DB02 as a linked server in order to imitate the remote login.
- E. Kara Lang should configure Certkiller -DB02 as a remote server. Thereafter she should write the query on Certkiller -DB01.

Answer: D

Explanation

: Special attention needs to be paid on the security context for the external connection when you make use of linked servers to access external data sources. You can configure the linked server to use one of the following three security modes:

Self-mapping - When a linked server is created, this mode is added for all local logins so SQL Server tries to connect to the external data source using the current user's login credentials. The same login and password must exist on the remote server. This is the default behaviour.

Delegation - This mode impersonates the Windows local credentials; the connection forwards the credentials of an authenticated Windows user to the linked server. The Windows user account and password must exist on the linked server.

Remote Credentials - This mode lets you map local logins to remote logins on the external data source. Delegation of operating system logins is the securest mechanism.

QUESTION 114:

You are employed as a database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. Your duties include administrating a SQL Server 2005 database server named Certkiller -DB01.

Your institute is well known for the sale of their textbooks. You work for a company that sells books. You are creating a report for a SQL Server 2005 database. This report will indicate the all the sales representatives and their total sales for the present month. A Certkiller .com user named Dean Austin is a junior technician in the Sales department. You inform him that the report should only include those sales representatives who met their sales quota for the present month.

The monthly sales quota for Certkiller .com is \$2,000. Certkiller .com makes use of date parameters that are passed in variables named @StartDate and @EndDate. You instruct Dean Austin to create a report that will meet those requirements.

Which SQL query should Dean Austin use?

- A. SELECT s.SalesRep,
SUM(ISNULL (o.OrderTotal,0.00))AS SumOrderTotal
FROM SalesAgent s
JOIN OrderHeader o ON s.SalesRepID = o.SalesRepID
WHERE o.OrderDate BETWEEN @StartDate AND @EndDate AND o.OrderTotal
>= 2000
GROUP BY s.SalesRep
- B. SELECT s.SalesRep,
SUM(ISNULL (o.OrderTotal,0.00)) AS SumOrderTotal
FROM SalesAgent s
JOIN OrderHeader o ON s.SalesRepID = o.SalesRepID
WHERE o.OrderDate BETWEEN @StartDate AND @EndDate
GROUP BY s.SalesRep
HAVING SUM(o.OrderTotal) >= 2000
- C. SELECT s.SalesRep,
SUM(ISNULL(o.OrderTotal,0.00)) AS SumOrderTotal
FROM SalesAgent s
JOIN OrderHeader o ON s.SalesRepID = o.SalesRepID
WHERE o.ordertotal = 2000 AND o.OrderDate BETWEEN @StartDate AND
@EndDate
GROUP BY s.SalesRep
HAVING SUM(o.OrderTotal) >= 2000
- D. SELECT s.SalesRep,
SUM(ISNULL(o.OrderTotal,0.00))AS SumOrderTotal
FROM SalesAgent s
JOIN OrderHeader o ON s.SalesRepID = o.SalesRepID
WHERE o.OrderDate BETWEEN @StartDate AND @EndDate
GROUP BY s.SalesRep

Answer: B

QUESTION 115:

You are employed as a database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. Certkiller .com has its headquarters in London and a branch office in Paris. You are based at headquarters. Your duties include administrating a SQL Server 2005 database server named Certkiller -DB01

You have been informed by the CIO that the Research and Development department wants a report that evaluates the customer activity in the previous quarter between the office in London and the branch office in Paris. Certkiller .com wants the information to be sorted by surname and first name. A Certkiller .com user named Amy Walsh is a junior technician in the IT department. You inform her to restore a recent backup of the Paris database onto your server.

Amy Walsh writes the queries to build the report, ordering the information by the LastName and FirstName columns. Amy Walsh reviews the information and detects that the client lists from the Paris database is sorted different to the one in the headquarters. You receive a notification from the Research and Development department that they require the revised data within 15 minutes for a presentation. Amy Walsh needs to implement the fastest possible solution that will guarantee that the information from both databases is sorted identically. What should Amy Walsh do?

- A. Amy Walsh needs to use the SQL Server Import and Export Wizard to copy the data from the Paris database into new tables whilst specifying the same collation as the London database.
- B. Amy Walsh needs to copy the data in the Paris database to a new database with the same collation as the London database by using the Copy Database Wizard.
- C. The query needs to be modified by Amy Walsh on the Paris database to use the COLLATE setting in the ORDER BY clause. Thereafter she should specify the same collation as the London database in the query.
- D. The format file should be modified to specify the same collation as the London database. Thereafter she can import the table again.

Answer: C

Explanation:

COLLATE {collation_name}

Specifies that the ORDER BY operation should be performed according to the collation specified in collation_name, and not according to the collation of the column as defined in the table or view. collation_name can be either a Windows collation name or a SQL collation name. COLLATE is applicable only for columns of the char, varchar, nchar, and nvarchar data types.

QUESTION 116:

You are employed as a database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. Your duties include administrating a SQL Server 2005 database server named Certkiller -DB01.

A Certkiller .com user named Ally Wagner is a member of the Research and Development department. You instruct her to create a stored procedure that will delete data from the Contact table in a Certkiller -DB01. The stored procedure encompasses the subsequent Transact-SQL statement to handle any errors that occur.

```
BEGIN TRY
BEGIN TRANSACTION
DELETE FROM CK_Sales.Products
WHERE ProductID = @ProductID
```

```
COMMIT TRANSACTION
END TRY
BEGIN CATCH
DECLARE @ErrorMessage nvarchar(2000)
DECLARE @ErrorSeverity int
DECLARE @ErrorState int
SELECT @ErrorMessage = ERROR_MESSAGE(),
@ErrorSeverity = ERROR_SEVERITY(),
@ErrorState = ERROR_STATE()
RAISERROR (@ErrorMessage, @ErrorSeverity, @ErrorState)
END CATCH
```

Ally Wagner tests the stored procedure and detects that it leaves open transactions. You inform her to modify the stored procedure in order that it properly handles the open transactions. What should she do?

- A. Ally Wagner should add a ROLLBACK TRANSACTION command to the CATCH block.
- B. A COMMIT TRANSACTION command should be added to the CATCH block.
- C. In order for the stored procedures to properly handle the open transactions she should add a ROLLBACK TRANSACTION command to the TRY block.
- D. It is important that Ally Wagner remove the COMMIT TRANSACTION command from the TRY block to properly handle the open transactions.

Answer: A

Explanation: If an exception occurs from one of the statements in the TRY block, control is branched to the CATCH block, where the exception can be handled, logged, and so on. In this scenario you want to ROLLBACK the transaction if there is an error.

QUESTION 117:

You are employed as a database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. You have received instruction from the CIO to create an online catalog application that will display the product information on the Certkiller .com Web site. The product information is stored in a SQL Server 2005 database server named Certkiller -DB01. The CIO informs you that the data is stored as relational data. However, the data should be passed to the application as an XML document by using FOR XML.

A Certkiller .com employee named Clive Wilson works in the Research and Development department. You instruct him, to test the application. Clive Wilson detects that not all the items matching the query appear in the XML document. He determines that only the products that have values for all elements in the schema appear in the XML document. You instruct him to modify the Transact-SQL statement in order that all products matching the query appear in the XML

document.

What should Clive Wilson do?

- A. Clive Wilson should modify the replace value of clause in the query.
- B. The HAVING clause needs to be added to the query.
- C. In order for all matching queries to appear in the XML document Clive Wilson needs to add an XML index to the table that contains the product information.
- D. It is important that Clive Wilson add the XSINIL argument to the ELEMENTS directive in the query.

Answer: D

Explanation: If you add the XSINIL instruction to the ELEMENTS clause in the FOR XML construction, Certkiller -DB01 will generate an empty XML element for NULL values.

QUESTION 118:

You are employed as a database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. You have received instruction from the network CIO to create a Web-based application to manage data aggregation for reports. The Application that you create will connect to a SQL 2005 database named CK_Data and another page of your application has controls which are used to execute stored procedures in a database named CK_Sales.

Certkiller .com has an existing Service Broker connection between the CK_Data and CK_Sales databases. You have received additional instruction to add two new message types to the existing service. In each database, you create message types named CKReport and CKResult. You are required to add the two new message types to the existing service.

What should you do first?

- A. You have to create a contract between the services. This should be accomplished by using the subsequent statement:
`CREATE CONTRACT ProcessReport
(CKReport SENT BY INITIATOR,
CKResult SENT BY TARGET)`
- B. A queue should be created on every database with the ACTIVATION argument set to CK_Data.dbo. CKReport.
- C. Services for every database needs to be created by using the subsequent statement:
`CREATE SERVICE CK_Data
ON QUEUE CKReport`
- D. You have to create a conversation between the databases. This can be done by making use of the subsequent statement:
`BEGIN DIALOG FROM SERVICE ' CKReport '
TO SERVICE ' CKResult '`

Answer: A

Explanation: A contract contains a list of message types and the services that are allowed to send them.

Syntax:

```
CREATE CONTRACT contract [AUTHORIZATION owner ]  
( { {message_type_name | [DEFAULT ] }  
SENT BY { INITIATOR | TARGET | ANY }  
} [ ,...n ] ) [;]
```

Key:

contract Name of the contract

owner Owner of the contract(database user or role)

message_type_name Name of a message type to be included

QUESTION 119:

You work as the database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. Your duties at Certkiller .com include importing data into SQL Server 2005 databases. A Certkiller .com user named Amy Wilson is a member of the Sales department. You have received numerous text files that contain the sales transactions from various stores across the country. The data are separated by semicolons in columns. You instruct Amy Wilson to import the files into the sales database. What should she do?

- A. In order to import the files into the sales database Amy Wilson needs to use the BULK INSERT statement with the default arguments.
- B. Amy Wilson needs to create a custom format file.
She also needs to specify a semicolon as the row terminator.
- C. The bcp command needs to be used and Amy Wilson should specify a semicolon as the field terminator.
- D. In order to import the files into the sales database Amy Wilson needs to use the bcp command with the default arguments.

Answer: C

Explanation: The bcp utility bulk copies data between an instance of Microsoft SQL Server 2005 and a data file in a user-specified format. The bcp utility can be used to import large numbers of new rows into SQL Server tables or to export data out of tables into data files. To import data into a table, you must either use a format file created for that table or understand the structure of the table and the types of data that are valid for its columns.

QUESTION 120:

You are employed as a database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. You are employed at the regional Marketing offices of Certkiller .com.

Your duties at Certkiller .com encompass the importing and exporting data in the SQL Server 2005 databases. You have received instruction from the CIO to send a text file to the main office asks that contains the updated contact data for the clients in your area.

The CIO requests that the data be sorted by the State, LastName, and FirstName columns. A Certkiller .com employee named Mia Hamm works in the Sales department. You instruct Mia Hamm to forward the necessary information to the database administrator at the main office. Mia Hamm needs to accomplish this by using the minimum of effort.

What should she do?

- A. A format file needs to be created for the export operation.
- B. Mia Hamm needs to identify State, LastName as well as FirstName in the ORDER BY clause in the bcp queryout command.
- C. The data needs to be copied into a new table that has a clustered index on State, LastName and FirstName. Thereafter the data can be exported.
- D. Mia Hamm needs to identify State, LastName as well as FirstName in the ORDER hint in the bcp out command.

Answer: B

Explanation: The bcp utility bulk copies data between an instance of Microsoft SQL Server 2005 and a data file in a user-specified format. The bcp utility can be used to import large numbers of new rows into SQL Server tables or to export data out of tables into data files. Except when used with the queryout option, the utility requires no knowledge of Transact-SQL. Queryout copies from a query and must be specified only when bulk copying data from a query.

Syntax:

```
bcp
{[[database_name.][owner.]]{table_name | view_name} | "query"}
{in | out | queryout | format} data_file
[-mmax_errors] [-fformat_file] [-x] [-eerr_file]
[-Ffirst_row] [-Llast_row] [-bbatch_size]
[-n] [-c] [-N] [-w] [-V (60 | 65 | 70 | 80)] [-6]
[-q] [-C { ACP | OEM | RAW | code_page } ] [-tfield_term]
[-rrow_term] [-iinput_file] [-ooutput_file] [-apacket_size]
[-Sserver_name[\instance_name]] [-Ulogin_id] [-Ppassword]
[-T] [-v] [-R] [-k] [-E] [-h"hint [...n]" ]
```

QUESTION 121:

You are employed as a database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. The

Certkiller .com network contains two SQL Server 2005 computers named Certkiller -DB01 and Certkiller -DB02.

Certkiller -DB01 and Certkiller -DB02 both take part in replication.

Certkiller -DB01 is the Publisher and Distributor of a publication named CK_Pub1. CK_Pub1 is the only publication on Certkiller -DB01 while Certkiller -DB02 is the only Subscriber. The CIO requests a status report regarding the replication latencies.

A Certkiller .com user named Rory Allen is a member of the Research and Development department. You instruct him to determine the current latencies between the Publisher and Distributor as well as between the Distributor and Subscriber using the Replication Monitor on Certkiller -DB01.

What should Rory Allen do?

- A. Rory Allen should select the Subscription Watch List tab for Certkiller -DB01. Then he should double-click the Certkiller -DB02 subscription. Thereafter he can view the duration details on the Publisher to Distributor History tab as well as the Distributor to Subscriber History tab.
- B. It is important that Rory Allen select the Tracer Tokens tab for the Pub1 publication. Thereafter he can select the Insert Tracer option and wait for the requested latency values for the Certkiller -DB02 subscription to appear.
- C. The Subscription Watch List tab for Certkiller -DB01 should be selected. Thereafter Rory Allen can view the Latency column for the Certkiller -DB02 subscription.
- D. Rory Allen needs to select the All Subscriptions tab for the Pub1 publication. Thereafter Rory Allen can view the Latency column for the Certkiller -DB02 subscription.

Answer: B

Explanation: Tracer tokens are a new feature of the transactional replication engine in SQL Server 2005. They are specialized transactions that are used to gain timing information. Since tracer tokens are picked up by the replication engine and transit with every other transaction, they can be used to determine the amount of time it takes a transaction to move from the publisher to distributor and then from the distributor to each subscriber. With this piece of information, you can now answer the two most predominant questions for replication: How far behind am I and how long will it take to catch up?

QUESTION 122:

You work as a database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. Your duties include implementing maintenance on a SQL Server 2005 database server named Certkiller -DB01.

At present there are certain jobs that run every Sunday whilst other jobs run at the start of every month. A Certkiller .com user named Amy Walsh is a member of the

Finance department. You instruct her to schedule the jobs in the way that uses the minimum administrative effort.

What should Amy Walsh do?

- A. Amy Walsh needs to create a job schedule that runs once a day.
Then she can assign jobs to this job schedule.
Thereafter the jobs can be executed if the day is either a Sunday or the first day of the month.
- B. A job needs to be created for every task that runs daily.
Then Amy Walsh can use a Transact-SQL statement to check the date and day of the week.
Thereafter a code needs to be executed if the day is either a Sunday or the first day of the month.
- C. Amy Walsh needs to create a job schedule that runs every Sunday.
Then she can assign weekly tasks to this schedule.
She can then create a second schedule that runs on the first day of every month.
Thereafter she can assign monthly tasks to this schedule.
- D. Amy Walsh needs to create a job for every task that runs weekly on Sundays.
Thereafter she should add a second job schedule that runs the job on the first of the month.

Answer: C

Explanation: You need 2 schedules or both tasks will be executed on Sundays and on the first day of the month. Creating a job schedule for every task that runs once a week instead of creating one schedule containing all the tasks would use more administrative effort than needed.

QUESTION 123:

You are employed as a database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. The Certkiller .com network contains a SQL Server 2005 database server named Certkiller -DB01.

You discover that a power failure occurred on the storage area network (SAN) where Certkiller -DB01 is located. A Certkiller .com employee named Clive Wilson is a junior technician in the Research and Development department. You inform Clive Wilson to check the allocation as well as the structural and logical integrity of all databases, which encompass the system catalogs.

What should he do?

- A. DBCC CHECKTABLE needs to be executed for every table.
- B. Clive Wilson needs to execute DBCC CHECKDB.
- C. It is important that Clive Wilson is able to execute DBCC CHECKCATALOG.
- D. Clive Wilson needs to execute DBCC CHECKFILEGROUP for every filegroup.

Answer: B

Explanation: DBCC CHECKDB will check the logical and physical integrity of all the objects in the specified database by:-

1. Running DBCC CHECKALLOC on the database.
2. Running DBCC CHECKTABLE on every table and view in the database.
3. Running DBCC CHECKCATALOG on the database.
4. Validating the contents of every indexed view in the database.
5. Validating the Service Broker data in the database

QUESTION 124:

You work as a database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. You have received a report from the CIO stating that as soon as new sales transactions is inserted in a SQL Server 2005 database it results in an error

A Certkiller .com user named Ally Wagner is a junior technician in the Sales department. You instruct her to investigate the error. Ally Wagner determines that a developer has accidentally deleted some data in a table that is critical for transaction processing in one of the databases.

The database makes use of the full recovery model. Ally Wagner needs to restore the table. You inform Ally Wagner to accomplish this goal without affecting the accessibility of other information in the database.

What should she do?

- A. The database needs to be restored to the point of the last full backup.
- B. Ally Wagner needs to restore the database from the existing backup files to a time prior to the data loss.
- C. Ally Wagner needs to back up the current transaction log.
Then she can restore the database with a different name and stop at the point just before the data loss.
Thereafter the table can be copied back into the original database.
- D. Back up The current transaction log needs to be backed up.
Thereafter Ally Wagner can restore the database to the point prior to the data loss.

Answer: C

Explanation: Ally Wagner needs to restore the table that has been accidentally altered. All other suggestions to solution will restore the complete database.

Therefore, you will lose data that may have been written to other tables after the point in time where you stop the recovery.

QUESTION 125:

You are employed as a database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com.

You have been informed by the CIO that the msdb database on a SQL Server 2005 computer is corrupted. This database should therefore be restored. A Certkiller .com employee named Dean Austin is employed as a junior technician in the Finance department. You inform him that the databases should be backed up daily. The backup files of the database are written to a network share. However, the file names do not clearly specify which databases are in every file. Dean Austin needs to locate the right backup file as quick as possible. You inform him that the first file in the list is named DB_Backup.bak. What Transact-SQL statement should Dean Austin use?

- A. Dean Austin should use the RESTORE DATABASE MSDBFROM DISK = N\\Server1\Backup\DB_Backup.bak Transact-SQL statement.
- B. The Transact-SQL statement RESTORE LABELONLYFROM DISK = N\\Server1\Backup\DB_Backup.bak should be used.
- C. Dean Austin should use the RESTORE VERIFYONLYFROM DISK = N\\Server1\Backup\DB_Backup.bak Transact-SQL statement.
- D. Dean Austin should use the RESTORE HEADERONLYFROM DISK = N\\Server1\Backup\DB_Backup.bak Transact-SQL statement.

Answer: D

QUESTION 126:

You work as the database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. You have received instruction from the CIO to implementing transaction log shipping for a database named CK_Data from a server named Certkiller -DB01 to a server named Certkiller -DB02. A Certkiller .com user named Rory Allen is appointed as a junior technician in the Research and Development department. You inform Rory Allen that CK_Data is 100 GB in size. It is therefore too big to transfer over the network in a reasonable amount of time. You instruct Rory Allen to minimize the impact on the network. Rory Allen needs to accomplish this by initializing the secondary database. What actions should Rory Allen perform? (Choose all that apply)

- A. It is vital that a complete backup of CK_Data to portable media is performed. Then he needs to restore the secondary database from that backup. Thereafter the STANDBY option needs to be specified.
- B. Rory Allen needs to identify the simple recovery model for CK_Data.
- C. Prior to activating the transaction log shipping to the secondary database Rory Allen needs to perform the subsequent statement on the primary server. The statement that Rory Allen needs to perform is:-BACKUP LOG TESTDATA WITH TRUNCATE_ONLY
- D. A partial backup of CK_Data needs to be performed to portable media. Thereafter a secondary database should be restored from that backup.

Then the RECOVERY option should be identified.

E. It is important that Rory Allen identifies either the full or the bulk-logged recovery model for CK_Data.

Answer: A, E

Explanation: Standby mode is the RESTORE command option Rory Allen needs to use to configure standby servers. In this mode, the secondary database is available for read-only access by users and applications. Log shipping requires that the primary database is configured with the Full or Bulk-Logged recovery model.

QUESTION 127:

You are employed as a database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com.

A full backup of your database named CK_Data is created automatically at midnight every day. You have been informed by the CIO that differential backups of CK_Data occur twice a day at 10:00 and at 16:00. A database snapshot is thus created daily at noon. A Certkiller .com user named Dean Austin is appointed as the technician in the Research and Development department.

The CIO informs you that a developer accidentally dropped the NewProjects table in CK_Data at 12:30. You check and notice that the NewProjects was last updated one week ago. You instruct Dean Austin to recover the NewProjects table. He needs to accomplish this by using the minimum amount of administrative effort. It is important that he reduces the amount of information that is lost.

What should he do?

- A. All database snapshots need to be deleted except the latest one. Thereafter CK_Data can be restored from the latest database snapshot.
- B. It is important that Dean Austin restores the latest backup into a new database named CK_Data_Temp. The latest differential backup can then be applied. Thereafter the NewProjects table from CK_Data_Temp should be copied to CK_Data.
- C. Dean Austin should copy the NewProjects table from the latest database snapshot into CK_Data.
- D. From the latest backup Dean Austin can recover the CK_Data and apply the latest differential backup on it.

Answer: C

Explanation: SQL Server 2005 introduces the capacity for administrators to generate and use a read-only, stable view of a database. The database snapshot offers this capability without the overhead of creating a complete copy of the database or storage requirement. As the primary database diverges from the snapshot, the snapshot gets its own copy of original pages when they are modified. The snapshot may be used to recover an accidental change to a database by simply reapplying the pages from the snapshot to the primary database. As the last

snapshot was created at noon the same day that the Pricelist table was dropped and the fact that there has been no changes to the table makes it possible to just copy the table from the snapshot.

QUESTION 128:

You are employed as a database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. Certkiller .com has a server named Certkiller -DB01 that is configured to run SQL Server 2005 Enterprise Edition.

Certkiller -DB01 has 2 GB of RAM. However, 1.6 GB is used by the default SQL Server database engine instance. 100 MB is the average data growth of all databases combined per month. You have received several reports from clients stating that the execution times are increasing.

A Certkiller .com employee named Kara Lang works in the Research and Development department. You instruct her to assess whether more RAM is required. Kara Lang needs to make use of the System Monitor to create a counter log. This counter log will aid her in deciding whether to add more RAM.

Which performance object should she add to the counter log?

- A. SQLServer:General Statistics should be added by Kara Lang to the counter log.
- B. Kara Lang should add SQLServer:SQL Statistics to the counter log.
- C. She should add MSAS 2005:Cache to the counter log.
- D. Kara Lang should add SQLServer:Buffer Manager to the counter log.
- E. MSAS 2005:Memory should be added to the counter log.

Answer: D

Explanation: The SQL Server:Buffer Manager object will show you

- * Low Buffer cache hit ration
- * Low Page life expectancy
- * High number of Checkpoint pages/sec
- * High number Lazy writes/sec

Insufficient memory and I/O overhead are usually related bottlenecks. SQLServer performance depends heavily on the I/O subsystem. Unless your database fits into physical memory, SQLServer constantly brings database pages in and out of the buffer pool. This generates substantial I/O traffic. Similarly, the log records need to be flushed to the disk before a transaction can be declared committed. And finally, SQLServer uses tempdb for various purposes such as to store intermediate results, to sort, to keep row versions and so on. So a good I/O subsystem is critical to the performance of SQLServer.

QUESTION 129:

You are employed as a database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. Your duties at Certkiller .com are the administrating of a database named CK_Data that is

situated on a SQL Server 2005 computer named Certkiller -DB03.

As the administrator you receive a warning that the drive on which the CK_Data log file is situated is near capacity. The transaction log is backed up every five minutes. You notice that this is steadily growing. You suspect that an uncommitted transaction may be the reason and you want to investigate it. You want to identify both the server process ID and the start time of the oldest active transaction in CK_Data.

What should you do?

A. The first thing you need to do is open a query window.

Then you should connect to the master database.

Thereafter you should perform the following statement:

```
SELECT TOP 1 spid, last_batch
```

```
FROM sys.sysprocesses
```

```
WHERE dbid = db_id('CK_Data') AND open_tran > 0
```

```
ORDER BY last_batch
```

B. You need to connect to the CK_Data database.

Then you should execute DBCC OPENTRAN.

Thereafter you should view the SPID and Start time rows.

C. You need to connect to the master database and execute DBCC OPENTRAN.

Then you must view the SPID and Start time rows.

D. You need to open the Activity Monitor in the SQL Server Management Studio.

Then you can select the Process Info page and apply the following filter settings:

Database = CK_Data

Open Transactions = Yes

View the Process ID and Last Batch columns.

Answer: B

Explanation: The DBCC OPENTRAN command gives you the information you need by displaying information about the oldest active transaction and the oldest distributed and nondistributed replicated transactions.

Syntax

DBCC OPENTRAN

[(['database' | database_id | 0])]

[WITH TABLERESULTS] [, [NO_INFOMSGS]]

]

QUESTION 130:

You are employed as a database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. Your duties at Certkiller .com include the administration of the SQL Server computer named Certkiller -DB03. Certkiller -DB03 was installed using default settings. A Certkiller .com employee named Clive Wilson is appointed as the junior technician in the Research and Development department. A power failure occurs. You detect

that the SQL Server (MSSQLSERVER) service on your database server does not start. You instruct Clive Wilson to troubleshoot the error.

What actions should Clive Wilson perform? (Choose all that apply)

- A. The C:\Program Files\Microsoft SQL Server\MSSQL.1\MSSQL\LOG\SQLAgent.out file should be viewed in Notepad.
- B. The system log needs to be viewed in Event Viewer.
- C. Clive Wilson should view in Notepad the C:\Program Files\Microsoft SQL Server\MSSQL.1\MSSQL\LOG\ErrorLog file.
- D. Clive Wilson should view in Notepad the C:\Program Files\Microsoft SQL Server\MSSQL.1\MSSQL\LOG\ErrorLog.1 file.
- E. Clive Wilson should view the application log in Event Viewer.

Answer: B, C, E

Explanation: To resolve database errors, you should review the SQL Server error logs and the Windows Application and the Windows System Event Log for errors. Even though a variety of informational messages is logged to each of these locations, any error with a severity level of 16 or higher automatically gets logged to the SQL Server error log and the Windows Application Event log. Errors with a severity of 16 or higher are critical errors that you need to investigate immediately.

QUESTION 131:

You work as the database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. Certkiller .com makes use of a SQL Server 2005.

A Certkiller .com employee named Andy Booth is appointed as a junior technician in the Research and Development department. You have received a query from various users stating that the report execution is slow. You instruct Andy Booth to troubleshoot the error. Andy Booth discovers that some queries do not use optimal execution plans.

It has come to your attention that some optimizer statistics are missing while others are out of date. Andy Booth needs to solve this problem in order for the reports to be executed faster. Andy Booth needs to identify the Transact-SQL statements that need to be used.

What Transact-SQL statements should he identify? (Choose all that apply)

- A. Andy Booth needs to identify both UPDATE STATISTICS and CREATE STATISTICS as the Transact-SQL statements.
- B. ALTER INDEX REORGANIZE and CREATE STATISTICS needs to be identified as the Transact-SQL statements that should be used.
- C. Andy Booth needs to identify both UPDATE STATISTICS and DBCC SHOW_STATISTICS as the Transact-SQL statements.
- D. Andy Booth needs to identify the DBCC CHECKTABLE.

Answer: A

Explanation: When a DBA creates an index, the query optimizer stores statistical information about the indexed columns. CREATE STATISTICS will create statistics on columns that are not contained in indexes but that are used in query predicates. UPDATE STATISTICS will optimize query performance.

QUESTION 132:

You work as the database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. Your duties at Certkiller .com are the administrating of the SQL Server 2005 computer named Certkiller -DB01. You need to configure Certkiller -DB01 to use TCP/IP with all default settings. A Certkiller .com employee named Rory Allen works in the Finance department. Certkiller .com policy requires that every server make use of a firewall. You inform Rory Allen that you are able to connect to the SQL Server instance from the local workstation. However, the client computers are unable to connect to the SQL Server instance. Rory Allen needs to identify the most likely cause of the connection issues. What should he do?

- A. Rory Allen needs to make sure that all client computers connect by using Shared Memory protocol.
- B. It is important that the server is not stopped.
- C. It is important that Rory Allen make sure that port 1433 is open in your firewall.
- D. Rory Allen needs to make sure that port 443 is open in your firewall.
- E. Rory Allen needs to make sure that all client computers connect by using the Virtual Interface Protocol.

Answer: C

Explanation: The default instance of the Database Engine makes use of port 1433, but that can be changed. Instances of SQL Server 2005 Express Edition, SQL Server Mobile and named instances of the Database Engine make use of dynamic ports. The SQL Server Browser service lets users connect to instances of the Database Engine that are not listening on port 1433, without knowing the port number. In order to make use of the SQL Server Browser you need to open UDP port 1434.

QUESTION 133:

You work as a database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. Your duties include administrating two SQL Server 2005 database named Certkiller -DB01 and Certkiller -DB02.

A Certkiller .com employee named Clive Wilson writes various ad hoc queries against the company databases. Clive Wilson has access to the Client database on

Certkiller -DB01. Clive Wilson does not have any access to the Marketing database on Certkiller -DB02. You need to allow Clive Wilson to write queries that join data from both Certkiller -DB01 and Certkiller -DB02.

What should you do?

A. Instruct Clive Wilson to write the queries on Certkiller -DB02 by using the OPENQUERY statement.

Thereafter she can specify Certkiller -DB01 as the server name.

B. A linked server should be created on Certkiller -DB01 to Certkiller -DB02.

Then configure the linked server to use impersonation.

C. Create a linked server on Certkiller -DB01 to Certkiller -DB02.

Then configure the linked server to use mapped logins.

D. Instruct Clive Wilson to indicate the SQL Server object names by using four-part notation.

Answer: C

Explanation: You need to define a linked server for every external data source you want to access. Thereafter you can configure the security context under which your distributed queries will run. You can configure the linked server to use one of the following security modes:

Self-mapping - When a linked server is created, this mode is added for all local logins, so SQL Server tries to connect to the external data source using the current user's login credentials. The same login and password must exist on the remote server. This is the default behaviour.

Delegation - This mode impersonates the Windows local credentials; the connection forwards the credentials of an authenticated Windows user to the linked server. The Windows user account and password must exist on the linked server.

Remote Credentials - This mode lets you map local logins to remote logins on the external data source.

QUESTION 134:

You work as a database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. Certkiller .com has multiple servers in a distributed environment.

Your duties at Certkiller .com encompass administrating two SQL Server 2005 computers named Certkiller -DB01 and Certkiller -DB02. Every server makes use of SQL Server Authentication and they use different logins. You need to write a distributed query that joins the data on Certkiller -DB01 with the data on Certkiller -DB02.

What should you do?

A. You should configure Certkiller -DB02 as a linked server in order to impersonate the remote login.

B. You should configure Certkiller -DB02 as a distributed server. Then you will be

able to use pass-through authentication.

C. You should make sure that both Certkiller -DB01 and Certkiller -DB02 makes use of the same login name as the security context for each server.

D. You should configure Certkiller -DB02 as a remote server and then write the query on Certkiller -DB01.

Answer: A

Explanation: You need to define a linked server for each external data source you want to access and then configure the security context under which your distributed queries will run. You can configure the linked server to use one of the following security modes:

Self-mapping - When a linked server is created, this mode is added for all local logins, so SQL Server tries to connect to the external data source using the current user's login credentials. The same login and password must exist on the remote server. This is the default behavior.

Delegation - This mode impersonates the Windows local credentials; the connection forwards the credentials of an authenticated Windows user to the linked server. The Windows user account and password must exist on the linked server.

Remote Credentials - This mode lets you map local logins to remote logins on the external data source.

QUESTION 135:

You work as a database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. Your duties at Certkiller .com are the administrating of the SQL Server computers.

The Certkiller .com network contains a SQL Server 2005 database server named Certkiller -DB01. Certkiller -DB01 contains a table that contains of quite a few indexes. During routine monitoring you notice that data modification performance has degraded. You also believe that some of the indexes are not used.

You want to identify which indexes have not been used by any queries since the last time Certkiller -DB01 has started

What should you do?

A. Use the sys.dm_fts_index_population dynamic management view.

B. Use the sys.dm_exec_requests dynamic management view.

C. Use the sys.dm_db_index_usage_stats dynamic management view.

D. Use the sys.dm_db_index_physical_stats dynamic management view.

Answer: C

Explanation: The sys.dm_db_index_usage_stats dynamic management view contains core statistics about each index within a database. If you want to find the number of seeks, scans, lookups, or updates that have occurred with an index, you

should use this view. You can determine the last time an index was used by examining the last_user_seek, last_user_scan and last_user_lookup columns.

QUESTION 136:

You are employed as a database administrator at Certkiller .com. Your duties at Certkiller .com include administrating a SQL Server 2005 database computer named Certkiller -DB04. None of the databases on Certkiller -DB04 contain any views. You want to tune a database named CK_Stats that is hosted on Certkiller -DB04. You to accomplish this task using the Database Engine Tuning Advisor (DTA). Certkiller -DB04 contains a workload file that is appropriate for DT

A. You need

to identify any missing nonclustered indexes. You want the existing structures remain intact, and that newly recommend structures are partitioned for best performance.

What should you do? (Each correct answer presents part of the solution. Choose THREE.)

- A. Specify the use of the Nonclustered indexes physical design structure.
- B. Specify the use of the Indexes physical design structure.
- C. Specify Full partitioning as the partitioning strategy.
- D. Specify No partitioning as the partitioning strategy.
- E. Specify Aligned partitioning as the partitioning strategy.
- F. Specify Keep all existing physical design structures (PDS) in database option.
- G. Specify Do not keep all existing physical design structures (PDS) in database option.

Answer: A, C, F

Explanation: You want to identify missing nonclustered indexes; therefore you must specify the Nonclustered indexes physical design structure. You also want recommended structures to be partitioned for best performance. This means that you need the Full partitioning strategy. Finally, you want existing structures to remain intact. Therefore you must specify Keep all existing PDS.

Incorrect Answers:

B: Specifying the Indexes physical design structure will cause the DTA to consider adding indexes only. We just want nonclustered indexes.

D: When No partitioning is specified, the DTA does not consider partitioning in its recommendations.

E: When Aligned partitioning is specified, the DTA only considers partitioning that is aligned to existing partitions.

G: The Do not keep all existing PDS option will cause the DTA to consider dropping all existing physical design structures.

QUESTION 137:

You work as a database administrator at Certkiller .com. The Certkiller .com network

consists of a single Active Directory domain named Certkiller .com. The Certkiller .com network contains a SQL Server 2005 database server named Certkiller -DB01.

Certkiller -DB01 contains a database named CK_Sales that is used to process sales orders via an order processing application. You received a report from a Certkiller .com user named Rory Allen informing you that the order processing application stopped responding in the middle of an order transaction. Rory Allen's SQL Server session ID is 54.

You need to determine if session 54 is blocked by another connection. If the result is positive, you need to find the blocking session ID.

What should you do? (Each correct answer presents a complete solution. Choose TWO.)

- A. In Activity Monitor, view the BlockedBy column for session 54 on the Process Info page.
- B. In Activity Monitor, view the Request Mode column for session 54 on the Locks by Process page.
- C. Open a new query window and run the statement: `SELECT * FROM sys.dm_exec_requests WHERE session_id = 54`. Then view the blocking_session_id column.
- D. Open a new query window and run the statement: `SELECT * FROM sys.dm_exec_sessions WHERE session_id = 54`. Then view the status column.

Answer: A, C

Explanation: The Process Info page shows the properties of current processes. Blocked By tells you the Process ID (SPID) of a blocking process. Blocking indicates whether this process is blocking others. 1=yes; 0=no.

The sys.dm_exec_request DMV can be used to determine whether a process is being blocked and to identify the process that is creating the blocking. If a value greater than 0 exists in the blocking_process_id column, the process is being blocked by the SPID logged in the column.

QUESTION 138:

You are employed as a database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. The Certkiller .com network contains a SQL Server 2005 database server named Certkiller -DB03.

You have received a report from the Certkiller .com users that they receive deadlock error messages in an order processing application.

You need to monitor which objects and SQL Server session IDs are involved when deadlock conditions occur. You also want information about each participant in the deadlock.

What should you do?

- A. On Certkiller -DB03, use the SQL Server Profile and trace the Lock:Timeout event.
- B. On Certkiller -DB03, use the System Monitor and observe the SQLServer:Locks - Number of Deadlocks/sec counter.
- C. On Certkiller -DB03, use the SQL Server Profile and trace the Lock:Deadlock event.
- D. On Certkiller -DB03, use the SQL Server Profile and trace the Lock:Deadlock Chain event.

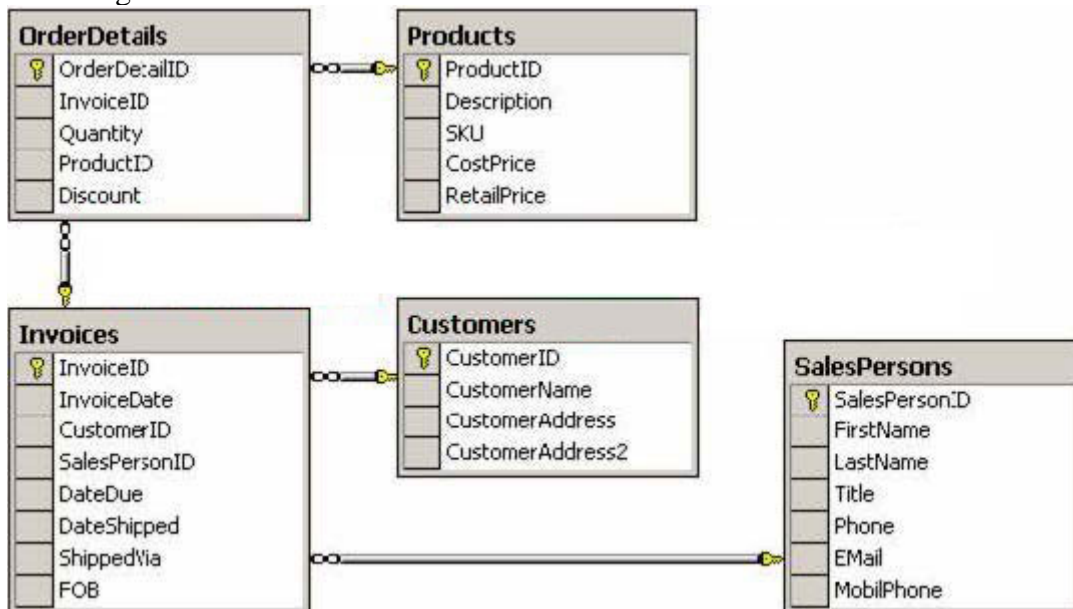
Answer: D

Explanation: Use SQL Server Profiler to identify the cause of a deadlock. A deadlock occurs when there is a cyclic dependency between two or more threads, or processes, for some set of resources within SQL Server. Using SQL Server Profiler, you can create a trace that records, replays, and displays deadlock events for analysis.

QUESTION 139:

You are work as a database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. The Certkiller .com network contains a SQL Server 2005 database server named Certkiller -DB08. Your duties at Certkiller .com include the administrating of Certkiller -DB08.

Certkiller -DB08 hosts a database named CK_Sales that stores sales transactions for the company. The table structure for the CK_Sales database is shown in the following exhibit:



You need to modify the Products table to add a new column named ProductName. You want a friendly name for each product will be stored in this column. The table currently contains data. The Marketing department has not yet created a friendly name for each product. You need to add this new column by using the least amount of effort.

What should you do?

- A. You need to define the new column as NULL. Update the ProductName column to the same value as the Description column. Alter the ProductName column to be NOT NULL.
- B. You need to define the new column as NOT NULL with a default value of 'Undefined.'
- C. You need to define the new column as NULL. Use application logic to enforce the data constraint.
- D. You need to define the new column as NULL with a default value of 'Undefined.'

Answer: B

Explanation:

You don't want to create the new column to accept NULL values and therefore you should create it as NOT NULL and by default enter 'Undefined' as the value for each product until the Sales Department changes the value to the correct friendly name.

QUESTION 140:

You work as a database administrator at Certkiller .com. Certkiller .com stores client information in a data warehouse. The data warehouse contains three separate SQL Server 2005 tables for storing client information.

The information stored in three separate tables named CK_Staff1, CK_Staff 2 and CK_Staff 3. The tables are used to horizontally partition information that is migrated from a SQL Server 2000 installation to a SQL Server 2005.

The information is partitioned as shown in the table displayed below:

Table	Stores employee surnames beginning in letters
CK_Staff1	A - J
CK_Staff2	K - P
CK_Staff3	R - Z

The database is configured in order that Certkiller .com employees are unable to opt-out of e-mail notifications to their personal email addresses. You have received instruction from the CIO to add a new OptOutFlag column to every CK_Staff table. You need to ensure the following:

1. The column is added to either all three CK_Staff tables or none of them.
2. The database is always in a consistent state.

Which Transact-SQL batch should you use?

A. BEGIN TRAN
ALTER TABLE CK_Staff1
ADD OptOutFlag BIT NULL
ALTER TABLE CK_Staff2
ADD OptOutFlag BIT NULL
ALTER TABLE CK_Staff3

```
ADD OptOutFlag BIT NULL
COMMIT TRAN
B. ALTER TABLE CK_Staff1
ADD OptOutFlag BIT NULL
ALTER TABLE CK_Staff2
ADD OptOutFlag BIT NULL
ALTER TABLE CK_Staff3
ADD OptOutFlag BIT NULL
C. BEGIN TRAN
ALTER TABLE CK_Staff1
ADD OptOutFlag BIT NULL
ALTER TABLE CK_Staff2
ADD OptOutFlag BIT NULL
ALTER TABLE CK_Staff3
ADD OptOutFlag BIT NULL
IF @@error <> 0
ROLLBACK TRAN
ELSE
COMMIT TRAN
D. BEGIN TRAN
BEGIN TRY
ALTER TABLE CK_Staff1
ADD OptOutFlag BIT NULL
ALTER TABLE CK_Staff2
ADD OptOutFlag BIT NULL
ALTER TABLE CK_Staff3
ADD OptOutFlag BIT NULL
END TRY
BEGIN CATCH
ROLLBACK TRAN
RETURN
END CATCH
COMMIT TRAN
```

Answer: D

Explanation: You should do this using the BEGIN CATCH. If there is an error reported in any of the tables when the OptOutFlag column is inserted will result that the whole transaction will rollback.

QUESTION 141:

You work as a database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. Your duties at Certkiller .com encompass administrating SQL Server 2005 database named Certkiller -DB08.

You have received instruction from the CIO to implementing a series of views that are used in ad hoc queries. The views are used to put into effect the Certkiller .com security policy of abstracting data. You notice that some of these views perform slowly therefore you create indexes on those views to increase the performance. You do this whilst maintaining the company's security policy.

The current date is returned by one of the views as one of the columns. The current date returned by the view is returned by using the GETDATE() function. This view does not permit you to create an index. You need to create an index on the view.

What should you do? (Choose all that apply)

- A. All nondeterministic function calls should be removed from within the view.
- B. Create the view with the WITH CHECK OPTION phrase.
- C. Remove all the deterministic function calls from within the view.
- D. Schema-bind all the functions that are called from within the view.

Answer: A, D

Explanation: The GETDATE() is a nondeterministic function, meaning that it does NOT have to return the same value every times it is called, and therefore the results of the function can not be indexed. When you specify the SCHEMABINDING option, you cannot drop any tables, views, or functions referenced by the view without first dropping the view.

QUESTION 142:

You are employed as a database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. Your duties at Certkiller .com include maintaining a SQL Server 2005 database named Certkiller -DB06.

Analysts in the Sales department of Certkiller .com usually make use of a view named v_ClientSales to join the Clients and the Sales tables in the Certkiller -DB06. The analysts make use of this view to combine the total sales by clients per month. You want to increase the performance of the view.

What should you do?

- A. You should create two separate views that do not contain any joints. Then name the one view named v_Clients for the Clients table and the other v_Sales for the Sales table.
- B. You should update the view to use an outer join between the Clients and Sales tables.
- C. You should create a stored procedure for the business analysts that make use of the v_ClientSales view.
- D. You should create a clustered index on the v_ClientSales view.

Answer: D

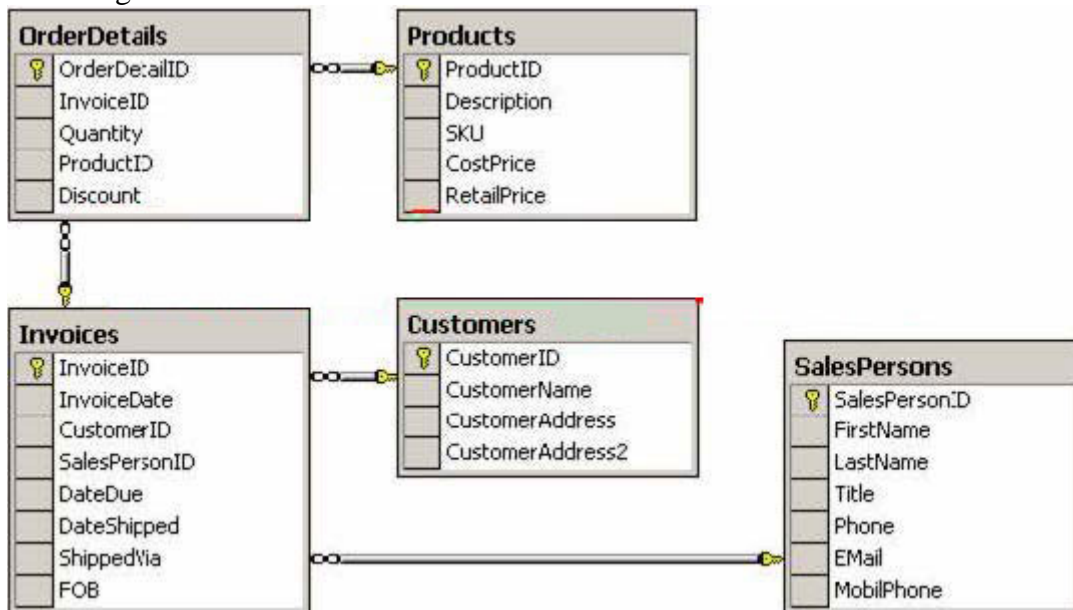
Explanation: As a view is nothing more than a virtual table containing data from

one or more tables you have the option to create a clustered index on the view in order to obtain an indexed view and obtain performance benefits.

QUESTION 143:

You are work as a database administrator at Certkiller .com. Your duties at Certkiller .com include the administrating of the SQL Server 2005 computer named Certkiller -DB07.

Certkiller -DB07 hosts a database named CK_Sales that stores sales transactions for the company. The table structure for the CK_Sales database is shown in the following exhibit:



You are busy creating a view to join the Customers and Invoices tables. You want to ensure that the view cannot be affected by modifications to underlying table schemas.

What should you do?

- A. You need to create CHECK constraints on the tables.
- B. You need to create a DDL trigger to roll back any changes to the tables if the changes affect the columns in the view.
- C. You need to create the view, specifying the WITH SCHEMABINDING option.
- D. You need to create the view, specifying the WITH CHECK option.

Answer: C

Explanation: When you create a view and specify the SCHEMABINDING option, you cannot drop any tables, views, or functions referenced by the view without first dropping the view.

QUESTION 144:

You are work as a database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. The Certkiller .com network contains a SQL Server 2005 database server named Certkiller -DB06.

Certkiller -DB06 hosts a database that contains a trigger named trg_InsertOrders. The trigger, trg_InsertOrders is configured to 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 want to use the minimum amount of administrative effort to configure the trigger to prevent it from firing during the data import process?

Which of the following Transact-SQL statements can be used? (Each 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: A, C

Explanation:

You do not want to alter or delete the trigger, so that you can use the trigger again with the least administrative effort. You should use the DISABLE TRIGGER Transact-SQL command.

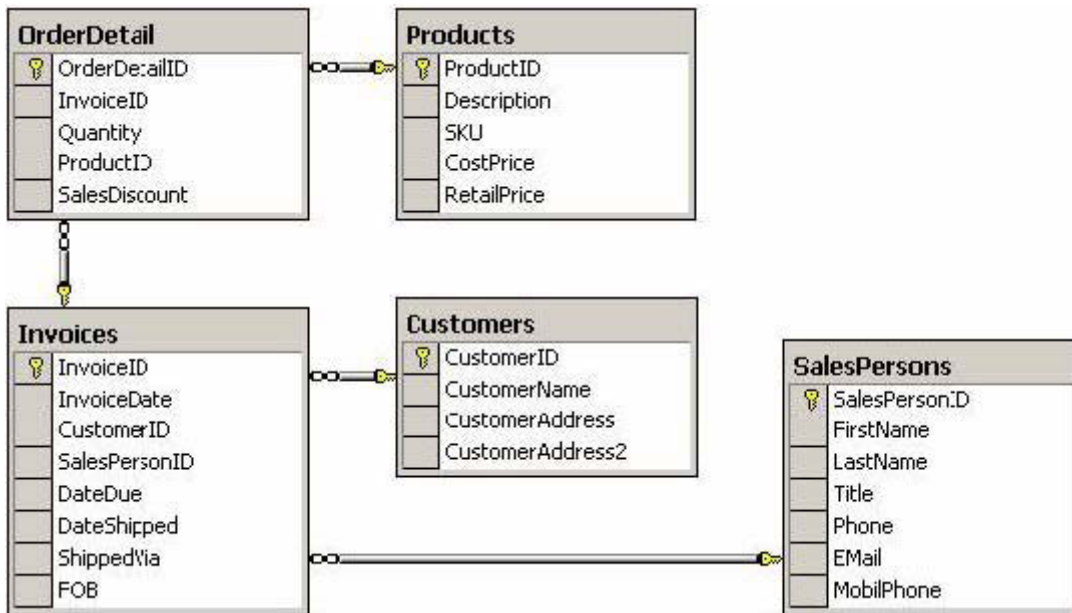
Syntax:

```
DISABLE TRIGGER { [ schema . ] trigger_name [ ,...n ] | ALL } ON {  
object_name | DATABASE | ALL SERVER } [ ; ]
```

QUESTION 145:

You are work as a database administrator at Certkiller .com. All servers on the Certkiller .com network run Windows 2000 Server and all client computers run Windows 2000 Professional. The Certkiller .com network contains a SQL Server 2005 computer named Certkiller -DB02.

CERTKILLER-DB02 hosts a database named CK_Sales that stores the sales information for the company. The table structure for the CK_Sales database is shown in the following exhibit:



The Description column in the Products table is defined as an nchar column. Several values in the Description column contain preceding or trailing spaces. You want to implement a mechanism that will ignore preceding or trailing spaces when data from the Products table is retrieved. You want this solution to be available for reuse in Transact-SQL statements and views. What should you do?

- A. You need to create DML triggers that query the inserted and deleted tables.
- B. You need to create a stored procedure that calls the LTRIM and RTRIM built-in functions.
- C. You need to create a Transact-SQL function that calls the LTRIM and RTRIM built-in functions.
- D. You need to call the TRIM built-in function.

Answer: C

Explanation:

RTRIM is a Transact-SQL command that returns a character string after truncating all trailing blanks. LTRIM is a Transact-SQL command that returns a character expression after it removes leading blanks.

QUESTION 146:

You work as a database administrator at Certkiller .com. Your duties at Certkiller .com include administering a SQL Server 2005 database computer named Certkiller -DB07.

The application developers in the Research and Development department have created an assembly that contains a CLR function. The purpose of the CLR function is to read data from a spreadsheet, perform several calculations and returning the information to Certkiller -DB07. You are required to register the

assembly with Certkiller -DB07 by using the CREATE ASSEMBLY statement. You must ensure that your solution uses the least privileged security permission set. What should you do?

- A. You should make use of the UNSAFE permission set.
- B. You should make use of the Default permission set.
- C. You should make use of the SAFE permission set.
- D. You should make use of the EXTERNAL_ACCESS permission set.

Answer: D

Explanation: To keep CLR code from running amok, Microsoft created a three-tier security model for how CLR code can be invoked: SAFE, EXTERNAL_ACCESS and UNSAFE. The SAFE permissions set is essentially the same as what a conventional stored procedure would be able to do. It can't modify anything outside of SQL Server itself. EXTERNAL_ACCESS allows access through .NET to the Registry and the file system. UNSAFE is aptly named. Code marked as UNSAFE can't do anything, and it should really not be used outside of a debugging or experimental context. Most programmers should never need to use anything higher than EXTERNAL_ACCESS.

QUESTION 147:

You work as a database administrator at Certkiller .com. Your duties at Certkiller .com encompass administrating a SQL Server computer named Certkiller -DB03.

You have received instruction from the CIO to retrieve data concerning a user who is currently logged in on Certkiller -DB03. You create a function that returns scalar information about the activity time for a particular user.

What should you do? (Choose all that apply)

- A. You should create a function that will return a numeric value which will signify the number of hours that particular user has logged for the current day.
- B. You should create a function that returns a numeric value that represents the number of hours a user has logged for the current month.
- C. You should create a function that returns a list of values that represent the login times for the given user.
- D. You should create a function that returns a list of values that represent the people who have logged more hours than the current user has logged.

Answer: A, B

Explanation: The values representing login times will only show at what time a user logs in not for how long and values representing users who have logged more hours than the current user has logged does not have to show activity time for all users.

QUESTION 148:

You work as a database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. Your duties at Certkiller .com include the administration of a SQL Server 2005 database named Certkiller -DB02.

Certkiller -DB02 hosts a database named CK_Sales that stores sales transactions for the company. Certkiller .com wants to provide customers real-time delivery tracking information using data retrieved from Certkiller -DB02 and data from a parcel tracking Web service provided by a third party. You need to ensure that the real-time tracking information is current when queries are executed. You need to create the appropriate objects that support the queries. What should you do?

- A. You should create a table to store the tracking information for every shipped item. Then you should create a trigger that fires when information is inserted into the table that joins with the data coming from the parcel tracking Web service.
- B. You should publish the data in the database as an XML Web service by using the FOR XML AUTO clause.
- C. You should create an assembly that calls the parcel tracking Web service and a CLR function using the assembly. Then you should call the CLR function and combine the results with customer information in the database.
- D. You should create a Transact-SQL stored procedure that uses a temporary table to store the tracking information for each customer. Then you should update the table with the values from the parcel tracking Web service.

Answer: C

QUESTION 149:

You work as a database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. The Certkiller .com network contains a SQL Server 2005 database server named Certkiller -DB01. Your duties include administrating Certkiller -DB01. Certkiller -DB01 is hosting Certkiller .com's commercial Web sites. You are creating a database on Certkiller -DB01 to support a new Web-based application that will handle up to 2,000 simultaneous users. The application that will access this database must quickly display the results of calculation-intensive operations. You need to ensure that the database processes calculations as quickly and efficiently as possible. What should you do?

- A. You need to implement parameterized Transact-SQL queries in the application.
- B. You need to implement Transact-SQL stored procedures in the database.
- C. You need to implement CLR stored procedures in the database.

D. You need to implement distributed Web services.

Answer: C

Explanation: A Common Language Runtime (CLR) will generally perform better than T-SQL and CRUD (Create, Read, Update, Delete) operations for complex math, string manipulation and other tasks that go beyond data access.

QUESTION 150:

You work as a database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. Your duties include administrating a SQL Server 2005 database named Certkiller -DB04.

It has come to your attention that schema changes were made on Certkiller -DB04. Due to these changes on Certkiller -DB04 your Web site stopped functioning. It is however unclear who made the changes. You have received mandate from the CIO that all changes to the database schema be tracked.

What should you do?

- A. You need to create a DDL trigger to roll back any changes to the tables if the changes affect the columns in the view.
- B. You need to implement a DML AFTER trigger that writes data about schema changes to a log table.
- C. You need to implement a DML INSTEAD OF trigger that writes data about schema changes to a log table.
- D. You need to implement DDL AFTER triggers that write user and schema information to a log table.
- E. You need to create a stored procedure that writes data about schema changes to a log.

Answer: D

Explanation:

A trigger is a specialized implementation of a Transact-SQL or CLR batch that automatically runs in response to an event within the database. You can create two types of triggers in Certkiller -DB04: data manipulation language (DML) triggers and data definition language (DDL) triggers. DDL triggers, which run in response to DDL events that occur on the server such as creating, altering, or dropping an object are used for database administration tasks such as auditing and controlling object access.

QUESTION 151:

You work as a database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. Your duties at Certkiller .com include managing the database schema for an order entry application in a SQL Server 2005 database named Certkiller -DB04.

A Certkiller .com employee named Mia Hamm works as part of the design review committee. Mia Hamm requests that a new column named Commission be added to the CK_Product table. The actual commission rates for the products are not known at present. Therefore, every product has a default commission rate of 10 percent.

The table must allow you to modify the rate if required.

You need to configure the table to assign the default value as efficiently as possible.

What should you do?

A. You need to create an UPDATE trigger to update the default value for every new item in the CK_Products table.

B. You need to create an INSERT trigger to assigns the default value to every item in the CK_Products table.

C. You need to create a DEFAULT constraint and assign the default value by specifying the value in the WITH VALUES argument.

D. You need to create a CHECK constraint to validate the data and then an assign the default value to every item in the CK_Products table.

Answer: C

Explanation:

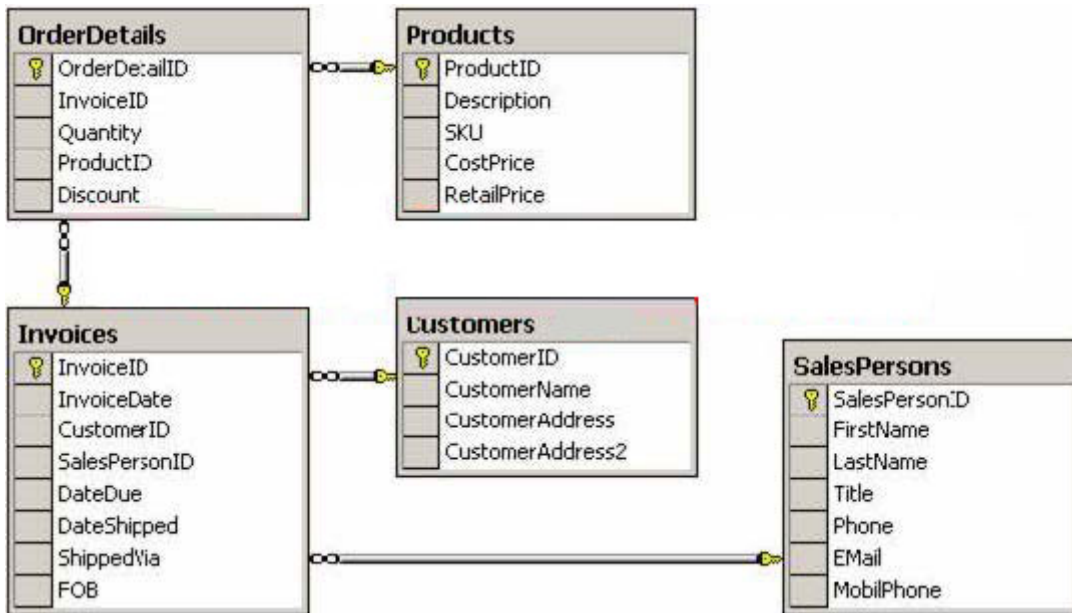
[CONSTRAINT constraint_name] - Specifies the start of a DEFAULT definition. To maintain compatibility with earlier versions of SQL Server, a constraint name can be assigned to a DEFAULT. constraint_name must follow the rules for identifiers, except that the name cannot start with a number sign (#). If constraint_name is not specified, a system-generated name is assigned to the DEFAULT definition.

WITH VALUES -Specifies that the value given in DEFAULT constant_expression is stored in a new column added to existing rows. If the added column allows null values and WITH VALUES is specified, the default value is stored in the new column, added to existing rows. If WITH VALUES is not specified for columns that allow nulls, the value NULL is stored in the new column in existing rows. If the new column does not allow nulls, the default value is stored in new rows regardless of whether WITH VALUES is specified.

QUESTION 152:

You work as a database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. Certkiller .com makes use of a database named CK_Sales that is hosted on a database server named Certkiller -DB01.

The table structure for the CK_Sales database is shown in the following exhibit:



Certkiller .com users report that table scans on the Invoices table causes a slow query. The Invoices table is a large table that is used regularly. You investigate and discover that the query contains the following statement:

SELECT InvoiceID, InvoiceDate FROM Invoices WHERE DateDue = <value>

You need to provide maximum query performance. However, the Invoice table should at all times remain available to users.

What should you do?

A. In order to accomplish this you should execute the following statement:-

```
USE CK_Sales
```

```
GO
```

```
CREATE INDEX index1
```

```
ON Invoices(DateDue)
```

```
INCLUDE(InvoiceID, InvoiceDate)
```

```
WITH (ONLINE = ON)
```

```
GO
```

B. In order to accomplish this you should update all statistics on the Invoices table.

C. In order to accomplish this you should use the CREATE STATISTICS statement to create the missing statistics on the Invoices.DateDue column.

D. In order to accomplish this you should execute the following statement:-

```
USE CK_Sales
```

```
GO
```

```
CREATE INDEX index1
```

```
ON Invoices(DateDue, InvoiceID, InvoiceDate)
```

E. In order to accomplish this you should set the priority boost server option to 1.

Answer: A

Explanation: You want to create an index on the columns 1 to 3 in order to have quicker response when querying data in columns 1 and 2 based on a value in

column 3.

Syntax

Relational Index

CREATE [UNIQUE] [CLUSTERED | NONCLUSTERED] INDEX index

ON object (column [ASC | DESC] [,...n])

[INCLUDE (column [,...n])]

[WITH (option [,...n])]

[ON { partition_scheme (column) | filegroup | default }]

[;]

QUESTION 153:

You are hired as a database consultant at Certkiller .com. The Certkiller .com network contains a SQL Server database server named Certkiller -DB01.

Certkiller -DB01 host a database named CK_Sales that stores line items from Certkiller .com's online sales transactions. Certkiller .com processes about 40,000 transactions every day. An application that processes the sales transactions requires a clustered index on a column named TransID.

You need to create a table that supports an efficient reporting solution that queries the transactions by date.

What should you do? (Each correct answer presents part of the solution. Choose TWO.)

- A. You should place a nonclustered index on the date column.
- B. You need to add a unique clustered index on the date column.
- C. You must map each partition to a filegroup, with each filegroup accessing a different physical drive.
- D. You should create a partitioning scheme that partitions the data by date.

Answer: A, D

Explanation: After you build a clustered index, you can create nonclustered indexes on the table. In contrast with a clustered index, a nonclustered index does not force a sort order on the data in a table. In addition, you can create multiple nonclustered indexes to most efficiently return results based on the most common queries you execute against the table. By using partitions, you can place a subset of a table or index on a designated filegroup. This capability lets you separate specific pieces of a table or index onto individual filegroups and effectively manage file input/output (I/O) for volatile tables. Additionally, as organizations collect more and more data and keep it longer and longer for analysis purposes, tables continue to grow larger and larger. Managing such massive tables can be difficult. With partitioning, however, you can segregate data within a table based on age.

QUESTION 154:

You are employed as a database administrator at Certkiller .com. All servers on the

Certkiller .com network run Windows 2000 Server and all client computers run Windows 2000 Professional.

Certkiller .com contains a SQL Server 2005 database server named Certkiller -DB01 that hosts a database named CK_Data. Certkiller .com also contains another SQL Server 2005 database server named Certkiller -DB02, which has the latest copy of CK_Data. The database CK_Data on Certkiller _DB01 runs 24/7.

During routine monitoring you detect a high number of full scans on Certkiller -DB01. You then come to the conclusion that extra indexes are needed on CK_Data. A workload file that is suitable for Database Engine Tuning Advisor (DTA) already exists.

You need to analyze the workload file by using DTA without compromising the performance on Certkiller -DB01 during the analysis. You also need to make sure that availability during the implementation of any recommendations suggested by the DTA is maintained.

What actions should you take?

A. On Certkiller -DB01, you need to store the workload file.

Start DTA on Certkiller -DB02 and connect to Certkiller -DB01.

Specify all workload and tuning options as necessary.

In the Advanced Tuning Options dialog box, select the Generate only online recommendations check box.

B. On Certkiller -DB02, you need to store the workload file.

Start DTA on Certkiller -DB02 and connect to Certkiller -DB02.

Specify all workload and tuning options as necessary.

In the Advanced Tuning Options dialog box, select the Generate only online recommendations check box.

C. On Certkiller -DB01, you need to store the workload file.

Start DTA on Certkiller -DB01 and connect to Certkiller -DB01.

Specify all workload and tuning options as necessary.

In the Advanced Tuning Options dialog box, select the All recommendations are offline check box.

D. On Certkiller -DB02, you need to store the workload file.

Start DTA on Certkiller -DB02 and connect to Certkiller -DB02.

Specify all workload and tuning options as necessary.

In the Advanced Tuning Options dialog box, select the All recommendations are offline check box.

Answer: B

Explanation: One of the requirements is to ensure maximum performance on Certkiller -DB01 during analysis, therefore storing the workload file on Certkiller -DB01 when running DTA is out of the question. You only want suggestions from the DTA that can be implemented while the database is online and therefore you should choose the Generate only online recommendations.

QUESTION 155:

You are work as a database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. The Certkiller .com network contains a SQL Server 2005 database server named Certkiller -DB01. Your duties at Certkiller .com include administrating Certkiller -DB01.

Certkiller .com has a Web site that has a page for customers to send feedback about Certkiller .com's products. You use a database CK_Products to store the comments in the Comments column of a table named Feedback.

You want to implement full-text searching so that you can run reports on the Comments column.

What should you do? (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.

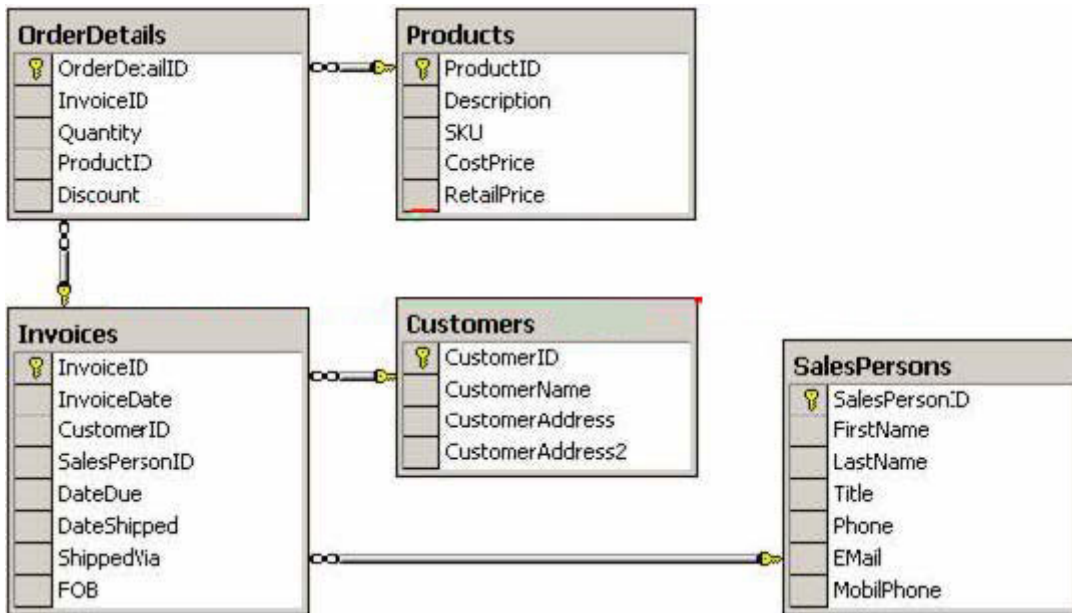
Answer: C, D

Explanation: Although you create full-text indexes on columns within tables in SQL databases, the full-text indexes are maintained in a structure outside of SQL Server called a full-text catalog. A full-text catalog can contain one or more full-text indexes.

QUESTION 156:

You work as a database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. Your duties at Certkiller .com include administrating a SQL Server 2005 database named Certkiller -DB09.

Certkiller -DB09 contains a database named CK_Sales that stores sales transactions for the company. The table structure for the CK_Sales database is shown in the following exhibit:



The Invoices table contains more than 300 million rows of information. In this table some information is historical and some is current. You want to partition the information in order to increase performance and optimize maintenance of the CK_Sales database.

What should you do?

- A. You should implement vertical partitioning in order to increase the performance and optimize maintenance.
- B. You should implement horizontal partitioning in order to increase the performance and optimize maintenance.
- C. You need to implement Distributed partitioning in order to increase the performance and optimize maintenance.
- D. You should implement a raw partition in order to increase the performance and optimize maintenance.

Answer: B

Explanation:

Table and index partitioning - intended to improve performance of operations performed on large tables. The basic concept is straightforward and involves splitting a table into several units (called partitions), which can be accessed independently of each other, limiting impact of I/O intensive activities performed on the table's data (queries, data loads, backups and restores, maintenance tasks - such as index rebuilds and defragmentations, as well as operations that would result in lock escalation to the table level). The most common method of splitting data is horizontal partitioning, in which rows of a table matching mutually exclusive criteria (such as, range of dates or letters in alphabet, for datetime and character data, respectively) are placed in designated partitions.

QUESTION 157:

You work as a database administrator at Certkiller .com. The Certkiller .com network consists of a single Active Directory domain named Certkiller .com. Certkiller .com makes use of a SQL Server 2005 database named CK_Sales. CK_Sales is located on a database server named Certkiller -DB01.

A new Certkiller .com employee named Mia Hamm works in the Finance department. You have received instruction from the CIO to give a new employee, Mia Hamm, access to CK_Sales. You create a login named mhamm by using the following Transact-SQL statement. CREATE LOGIN mhamm WITH PASSWORD = Certkiller -DB01'

Mia Hamm reports that as soon as she logs in, she receives the following error message: "Login failed. You are not associated with a trusted SQL Server connection." You need to resolve the error and allow the Mia Hamm to gain access to Certkiller -DB01.

What should he do?

- A. The login name needs be created with square brackets ([]).
- B. The login needs to be given access to a specific database by using the CREATE USER Transact-SQL statement.
- C. The SQL Server security mode should be changed from Windows Authentication mode to SQL Server and Windows Authentication mode.
- D. The SQL Server security mode should be changed from SQL Server and Windows Authentication mode to Windows Authentication mode.

Answer: C

Explanation: The syntax to create a Windows login is:

CREATE LOGIN [Domain\User] FROM WINDOWS

The Syntax to create a SQL Server login is

CREATE LOGIN login_name WITH PASSWORD='password'

You have created a SQL Server login on a SQL Server that is configured to only accept Windows Authentication mode.