



TestKingonline.com

No Pass No Pay!

**MCSE, CCNA, CCNP, OCP, CIW, JAVA, Sun Solaris, Checkpoint
World No 1 Cert Guides**

**Installing, Configuring and Administering Microsoft
SQL Server 2000 Enterprise Edition**

Exam 70-228

Congratulations!

You have purchased a TestKingonline. Study Guide.

This study guide is a complete collection of questions and answers that have been developed by our professional & certified team. You must study the contents of this guide properly in order to prepare for the actual certification test. The average time that we would suggest you for studying this study guide is approximately 10 to 20 hours and you will surely pass your exam. We guarantee it!

GOOD LUCK!

DISCLAIMER

This study guide and/or material is not sponsored by, endorsed by or affiliated with Microsoft, Cisco, Oracle, Citrix, CIW, Checkpoint, Novell, Sun/Solaris, CWNA, LPI, ISC, etc. All trademarks are properties of their respective owners.

Guarantee

If you use this study guide correctly and still fail the exam, send a scanned copy of your official score notice at: Sales@Testkingonline.com

We will gladly refund the cost of this study guide or give you an exchange of study guide of your choice of the Free.

This material is protected by copyright law and international treaties. Unauthorized reproduction or distribution of this material, or any portion thereof, may result in severe civil and criminal penalties, and will be prosecuted to the maximum extent possible u

Table of contents

PART I: Installing and Configuring SQL Server 2000 (9 questions).....	5
SECTION 1: Install SQL Server 2000. Considerations include clustering, default collation, file locations, number of instances, and service accounts. (4 questions)	5
SECTION 2: Upgrade to SQL Server 2000.....	10
SUBSECTION A: Perform a custom upgrade. (1 question)	10
SUBSECTION B: Upgrade to SQL Server 2000 from SQL Server 6.5. (0 questions).....	11
SUBSECTION C: Upgrade to SQL Server 2000 from SQL Server 7.0. (0 questions).....	11
SECTION 3: Create a linked server. (2 questions).....	11
SECTION 4: Configure SQL Mail and SQLAgentMail. (2 questions).....	14
SECTION 5: Configure network libraries. (0 questions)	16
SECTION 6: Troubleshoot failed installations. (0 questions)	16
PART II: Creating SQL Server 2000 Databases (23 questions).....	17
SECTION 1: Configure database options for performance. Considerations include capacity, network connectivity, physical drive configurations, and storage locations. (5 questions).....	17
SECTION 2: Attach and detach databases. (2 questions).....	25
SECTION 3: Create and alter databases.....	28
SUBSECTION A: Add filegroups. (1 question).....	28
SUBSECTION B: Configure filegroup usage. (2 questions)	30
SUBSECTION C: Expand and shrink a database. (2 questions).....	32
SUBSECTION D: Set database options by using the ALTER DATABASE or CREATE DATABASE statements. (1 question)	36
SUBSECTION E: Size and place the transaction log. (1 question)	37
SECTION 4: Create and manage objects. Objects include constraints, indexes, stored procedures, triggers, and views. (9 questions)	39
PART III: Managing, Monitoring, and Troubleshooting SQL Server 2000 Databases (38 questions) .55	55
SECTION 1: Optimize database performance. Considerations include indexing, locking, and recompiling. (13 questions)	55
SUBSECTION A: Optimize data storage. (1 question).....	78
SUBSECTION B: Optimize files and filegroups. (4 questions).....	80
SECTION 2: Manage database fragmentation. (0 questions).....	88
SECTION 3: Modify the database schema.....	88
SUBSECTION A: Perform disaster recovery operations. (7 questions)	88
SUBSECTION B: Perform backups. (4 questions)	98
SUBSECTION C: Recover the system state and restore data. (0 questions)	106
SECTION 4: Configure, maintain, and troubleshoot log shipping. (1 question)	106
SECTION 5: Perform integrity checks. Methods include configuring the Database Maintenance Plan Wizard and using the Database Consistency Checker (DBCC). (0 questions).....	108
SECTION 6: Troubleshoot transactions and locking by using SQL Profiler, SQL Server Enterprise Manager, or Transact-SQL. (8 questions).....	108
PART IV: Extracting and Transforming Data with SQL Server 2000 (20 questions)	122

SECTION 1: Set up IIS virtual directories to support XML. (1 question)	122
SECTION 2: Import and export data. Methods include the Bulk Insert task, the bulk copy program, Data Transformation Services (DTS), and heterogeneous queries. (9 questions)	123
SECTION 3: Develop and manage DTS packages. (2 questions).....	138
SECTION 4: Manage linked servers.	140
SUBSECTION A: Manage OLE DB Providers. (2 questions).....	140
SUBSECTION B: Configure security mapping. (0 questions)	143
SECTION 5: Convert data types. (1 questions).....	143
SECTION 6: Configure, maintain, and troubleshoot replication services. (5 questions).....	145
PART V: Managing and Monitoring SQL Server 2000 Security (38 questions)	154
SECTION 1: Configure mixed security modes or Windows Authentication. Considerations include client connectivity, client operating system, and security infrastructure. (9 questions)	154
SECTION 2: Create and manage log ons. (9 questions)	165
SECTION 3: Create and manage database users. (2 questions)	176
SECTION 4: Create and manage security roles. Roles include application, database, and server. (1 question).....	179
SUBSECTION A: Add and remove users from roles. (4 questions).....	180
SUBSECTION B: Create roles to manage database security. (1 question).....	187
SECTION 5: Enforce and manage security by using stored procedures, triggers, views, and userdefined functions. (6 questions).....	188
SECTION 6: Set permissions in a database. Considerations include object permissions, object ownership, and statement permissions. (3 questions).....	197
SECTION 7: Manage security auditing. Methods include SQL Profiler and C2 auditing. (3 questions)	202
PART VI: Managing, Monitoring, and Troubleshooting SQL Server 2000 (22 questions)	206
SECTION 1: Create, manage, and troubleshoot SQL Server Agent jobs. (6 questions).....	206
SECTION 2: Configure alerts and operators by using SQL Server Agent, System Monitor and Transact-SQL. (4 questions)	214
SECTION 3: Optimize hardware resource usage. Resources include CPU, disk I/O, and memory.	218
SUBSECTION A: Monitor hardware resource usage by using System Monitor. (2 questions).....	218
SUBSECTION B: Resolve system bottlenecks by using System Monitor. (2 questions)	221
SECTION 4: Optimize and troubleshoot SQL Server system activity. Activities include cache hits, connections, locks, memory allocation, recompilation, and transactional throughput.....	224
SUBSECTION A: Monitor SQL Server system activity by using traces. (5 questions).....	224
SUBSECTION B: Monitor SQL Server system activity by using System Monitor. (3 questions)	229

Total number of questions: 150

PART I: Installing and Configuring SQL Server 2000 (9 questions)

SECTION 1: Install SQL Server 2000. Considerations include clustering, default

collation, file locations, number of instances, and service accounts. (4 questions)

Question No: 1

You are the administrator of Microsoft Windows 2000 computer. You are preparing to install SQL Server 2000 on the computer.

Your company contains a variety of client computers that will connect to the SQL server 2000 computer by using a specific Net-Library, as shown in the following table.

Client computer Net-Library

Microsoft Windows 98 Named Pipes

Novell NetWare IPX/SPX

Apple Macintosh TCP/IP

You need to allow the client computers to connect to the SQL Server 2000 computer. You also want to

minimize the number of configuration changes required on the client computers.

Which three actions should you take? (Each correct answer presents part of the solution. Choose three)

- A. Install SQL Server 2000 as a named instance.
- B. Install SQL Server 2000 as the default instance.
- C. Configure the new instance for Mixed Mode Authentication.
- D. Configure the new instance for Windows Authentication.
- E. Configure the server to use the Named Pipes, IPX/SPX, and TCP/IP Net-Libraries.
- F. Configure the server to use the Multiprotocol Net-Library.

Answer: B, C, E.

Explanation:

B: We need to use a default instance of SQL Server in order to allow all the different clients to connect.

C: The client computers in this scenario will be running Windows 98, Novel NetWare and Apple Macintosh. SQL Server Authentication is required as these operating systems are not supported by

Windows Authentication, however, the use of SQL Server Authentication is offered only for backward compatibility, hence we need to implement Mixed Mode Authentication.

E: By using specific communication protocols, and not the Multiprotocol Net-Library, performance would be better. Every extraneous protocol adds overhead.

Note: Configuring Net-Libraries has been simplified for applications using the SQL Server 2000 client connectivity components. Selecting the client protocols to be installed on the application computer must be done during setup, and selecting the server protocols to be enabled on the instance of SQL Server 2000 must be done during the SQL Server 2000 setup. No other configuration is required. Once setup, a client can connect to any instance of SQL Server 2000 by specifying the network name of the database computer and the instance name.

Note: SQL Server 2000 can operate in either Windows Authentication mode, which allows a user to connect through a Windows NT 4.0 or Windows 2000 user account; Mixed mode, which allows users to connect to an instance of SQL Server 2000 by using either Windows Authentication or SQL Server Authentication. Furthermore, users who connect through a Windows NT 4.0 or Windows 2000 user account can make use of trusted connections in either Windows Authentication mode or Mixed mode. When a user connects to SQL Server 2000 with a nontrusted connection, SQL Server 2000 performs the authentication itself by checking if a SQL Server login account has been created and if the supplied password is correct. SQL Server Authentication is also used for backward compatibility because applications written for SQL Server 7.0 or earlier may require the use of SQL Server logins and passwords. In addition, SQL Server Authentication is required when an instance of SQL Server is running on Windows 98 because Windows Authentication mode is not supported on Windows 98.

Incorrect Answers:

A: The Multiprotocol Net-Library can only be used to connect to a default instance of SQL Server 2000 as it does not support named instances of SQL Server 2000 or server enumeration.

D: As client computers will be running Windows 98, Novel NetWare and Apple Macintosh, we cannot use Windows Authentication as these operating systems are not supported by Windows Authentication.

Note: Windows Authentication mode allows a user to connect through a Windows NT 4.0 or Windows 2000 user account and makes use of trusted connections. When a user connects to SQL Server 2000 a nontrusted connection, SQL Server 2000 performs the authentication itself by checking if a SQL Server login account has been created and if the supplied password is correct. Mixed mode allows users to connect to an instance of SQL Server 2000 by using either Windows Authentication or SQL Server Authentication.

F: The Multiprotocol Net-Library is used to support the different protocol that the clients use. This provides easy of administration. However, since unused protocols are installed, it also introduces overhead. This is not the best solution.

Note: The **Multiprotocol Net-Library** allows for the automatic selection of available network protocols to communicate with an instance of SQL Server 2000. This is convenient for connecting to multiple servers running different network protocols without having to reconfigure the client connection for each server. If the client and server Net-Libraries for TCP/IP Sockets, NWLink IPX/SPX, or Named Pipes are installed on the client and server, the Multiprotocol Net-Library will automatically choose the first available network protocol to establish a connection. The Multiprotocol Net-Library can only be used to connect to a default instance of SQL Server 2000 as it does not support named instances of SQL Server 2000 or server enumeration.

Question No: 2

You are the administrator of a new Microsoft Windows 2000 Server computer named CPSRV1. CPSRV1 is a member server in your company's Windows NT 4.0 domain, which is named CPDOMAIN. A separate Windows NT domain named CPUSERS contains all of your company's user accounts. A trust relationship exists between the CPDOMAIN domain and the CPUSERS domain. You are installing SQL Server 2000 on CPSRV1.

You specify a service account for each SQL Server service as shown in the following table:

Service	Service account	Account type
---------	-----------------	--------------

MSSQLServer	CPUSERS\sqlsvr	Domain administrator
-------------	----------------	----------------------

SQLServerAgent	CPUSERS\sqlagent	Domain administrator
----------------	------------------	----------------------

After the installation is complete, the MSSQLServer service and the SQLServerAgent service do not start. You need to configure the services to start without changing any security settings in either domain.

What should you do?

- A. Configure the services to run under the local system account
- B. Add CPSRV1 to the CPUSERS domain
- C. Add CPUSERS\sqlsvr and CPUSERS\sqlagent to the local administrators group on CPSRV1
- D. Add CPUSERS\sqlsvr and CPUSERS\sqlagent to the CPUSERS\Domain Admins groups.
- E. Add CPUSERS\sqlsvr and CPUSERS\sqlagent to the CPUSERS\Domain users groups.

Answer: C.

Explanation: MSSQLServer and MSServerAgent are running with accounts from the remote CPUSERS domain. These account does not have any specific local rights on the SQL Server. Local administrators rights and permissions are needed for these accounts are needed to make the SQL Server function. By adding the account sqlsvr and sqlagent from the other domain CPUSERS to the Local admin group on the SQL Server, it would be possible to start the SQL Server.

Incorrect Answers:

A: The **Local System** account does not require a password and does not have network access rights in a Windows NT 4.0 domain. This prevents the instance of SQL Server 2000 from interacting with other network servers in a Windows NT 4.0 environment. The SQL Server cannot run with a local system account in this scenario.

B: CPSRV1 is in the resource domain in the trust relationship. We cannot add a recourse from the resource domain to a group in the trusting domain.

D: The Domain Admins group represent the users who have broad administrative rights in a domain. These users are responsible for administration and delegation of authority.

E: By default, any domain user account is automatically added to the Domain Users group.

Question No: 3

You are the administrator of two Microsoft Windows 2000 advanced server computers. On these servers, you are configuring a database that will store accounting information for your company. The data must be available at all times. Interruptions in data connectivity should not last longer than five minutes. Any changes to the database should not require you to reconfigure the client computers. How should you configure the database?

A. Configure the database on the two servers as a SQL Server 2000 cluster.

B. Configure the database on one server, and then configure a standby database on the second server.

C. Configure the database on each server.

Use Microsoft Distributed Transaction Coordinator to keep the two servers perfectly synchronized.

D. Configure the database as a federated database, and then partition half the data on each server.

Answer: A.

Explanation: SQL Server 2000 failover clustering provides high availability support by allowing us to configure one failover cluster to automatically, with no manual configuration, fail over to any other node in the failover cluster configuration. In this way, we minimize system downtime and provide high server availability during an operating system failure or a planned upgrade. Before SQL Server 2000 failover clustering can be set up, we must install Windows NT 4.0, Enterprise Edition, Windows 2000 Advanced Server or Windows 2000 Datacenter Server, and the Microsoft Cluster Service (MSCS).

In this scenario a cluster would be preferred to a standby server because it would be available immediately. A standby server on the other would require manual configuration and more time to become online.

Incorrect Answers:

B: A standby server would require you to reconfigure clients to allow it to connect to the standby server when the main server goes down. It would require at least several minutes of manual administration before the standby server would be operational. But the scenario requires that no reconfiguration of the client computers should be necessary. Therefore we should use a SQL Server 2000 cluster not a standby server.

Note: A standby server is a second server that contains a copy of the databases on the primary server and that can be brought online in the event of a primary server failure or due to scheduled maintenance on the primary server. This allows users to continue working with databases when the primary server becomes unavailable. When the primary server becomes available again, any changes to the standby server's copies of databases must be restored back to the primary server.

C: The Microsoft Distributed Transaction Coordinator coordinates distributed transactions, but would not increase availability.

Note: The Microsoft Distributed Transaction Coordinator (MS DTC) is a transaction manager that allows client applications to include several different sources of data in one transaction. MS DTC coordinates committing the distributed transaction across all the servers enlisted in the transaction. An installation of SQL Server 2000 can participate in a distributed transaction by calling stored procedures on remote servers running SQL Server, automatically or explicitly promoting the local transaction to a distributed transaction and enlist remote servers in the transaction, and making distributed updates that update data on multiple OLE DB data sources. If these OLE DB data sources support the OLE DB distributed transaction interface, SQL Server 2000 can also enlist them in the distributed transaction. The MS DTC service coordinates the proper completion of the distributed transaction to ensure that either all of the updates on all the servers are made permanent, or, in the case of errors, all erased.

D: With Federated SQL Server 2000 servers, SQL Server 2000 databases can be spread across a group of autonomous database servers. These are capable of supporting the processing growth requirements of the largest Web sites and enterprise data-processing systems built with Microsoft Windows DNA, which divides the processing units of a data processing system into logical tiers. This however does not improve availability.

Question No: 4

You are the administrator of a new Microsoft Windows 2000 server computer named HQSQL5. HQSQL5 is a member server in your company's Windows NT 4.0 domain, which is named HQMAIN. After installing SQL Server 2000 on HQSQL5, you configure the MSSQLServer service account to use the HQMAIN\sqladmin user account, which is a member of the local administrators group on HQSQL5.

You can use the HQMAIN\sqladmin user account to log on the HQSQL5. However, the MSSQLServer service fails to start. You need to start the service. What should you do?

- A.** Revoke **Log on interactively** rights for the HQMAIN\sqladmin user account.
- B.** Add the HQMAIN\sqladmin user account to the HQMAIN\Domain Admins group.
- C.** Select the **Password Never Expires** check box for the HQMAIN\sqladmin user account.
- D.** Grant **Log on as a service** rights to the HQMAIN\sqladmin user account.

Answer: D.

Explanation: On Windows NT and Windows 2000 operating systems, SQL Server 2000 and SQL Server Agent are started and run as Windows services. Windows services require 'Log on as a service' rights, which allow a security principal to log on as a service, as a way of establishing a security context. The Local System account on the Windows computer always retains the right to log on as a service. Any service that runs under a separate account must be granted this right.

Incorrect Answers:

A: The log on interactively should only be applied to user accounts that runs services. Revoking it for the HQMAIN\sqladmin would make the MSSQLServer service to start.

B: The Domain Admins global group can represent the users who have broad administrative rights in a domain. Windows 2000 Server places all users with an Administrator account into this group automatically. This group is responsible for administration and delegation of authority. It is therefore not an appropriate group for Windows service.

C: The Password Never Expires check box specifies that the current Windows NT or Windows 2000 user account password does not need to be changed on a regular basis. However, it is only the MSSQLServer service that failed to start therefore the problem is not related to an expired password.

SECTION 2: Upgrade to SQL Server 2000.

SUBSECTION A: Perform a custom upgrade. (1 question)

Question No: 1

You are the administrator of a SQL server computer. The server is running SQL Server 6.5 and SQL Server 7.0.

You install a named instance of SQL Server 2000, and then run the SQL server upgrade wizard. On the database selection screen, some of the SQL Server 6.5 databases are not listed.

You cancel the SQL Server upgrade wizard.

You need to ensure that the SQL Server 6.5 databases are listed in the wizard. What should you do?

- A.** Uninstall SQL Server 7.0, and then rerun the SQL Server upgrade wizard.
- B.** Run the Microsoft SQL Server-switch application, and then rerun the SQL Server upgrade wizard.
- C.** Create a data transformation services package that imports the databases from SQL Server 6.5 to SQL Server 2000, and then execute the package.
- D.** Uninstall SQL Server 2000, and then reinstall SQL Server 2000 as the default instance.

Answer: D

Explanation: To run the SQL Server Upgrade Wizard, you must have a default instance of Microsoft SQL Server 2000 installed on your computer.

References:

BOL: Working with Instances and Versions of SQL Server

BOL: Upgrading Databases from SQL Server 6.5 (Upgrade Wizard)

Incorrect Answers:

A: It is not necessary to uninstall SQL Server 7.0. Just install SQL Server 2000 as the default instance. We cannot run the SQL Server upgrade wizard unless SQL Server 2000 is installed.

B: To run the SQL Server Upgrade Wizard, you must have a default instance of Microsoft SQL Server 2000 installed on your computer.

C: We want to ensure that the SQL Server 6.5 databases are listed in the SQL Server 2000 installation wizard. To create DTS transformation packages that upgrade SQL Server 6.5 databases to SQL Server 2000 databases we must install SQL Server 2000 first.

SUBSECTION B: Upgrade to SQL Server 2000 from SQL Server 6.5. (0 questions)

SUBSECTION C: Upgrade to SQL Server 2000 from SQL Server 7.0. (0 questions)

SECTION 3: Create a linked server. (2 questions)

Question No: 1

You are the administrator of a SQL Server 2000 computer. Your company uses the server to store service contract information for its customers.

You are also the administrator of Oracle relational database management system (RDBMS) server. This server is used to store your company's financial information. The financial information is updated frequently throughout the day.

You need to create a series of reports that combine the service contract information and the financial information. These reports will be updated several times a day.

You want to create reports on the SQL Server computer by using the minimum amount of disk space. What should you do?

- A.** Set up SQL server replication to replicate the data from the oracle server to the SQL server computer.
- B.** Set up the oracle server as a linked server.
Create a view that joins the service contract information and the financial information.
- C.** Set up data transformation services (DTS) package that imports and transforms the database from the oracle server to the SQL server computer. Use SQL server agent to execute the DTS package throughout the day as needed.
- D.** Set up Microsoft ActiveX script that connects to the oracle server and imports the financial information into SQL server temporary table.
Create a view that joins the service contract information and the temporary table.

Answer: B.

Explanation: SQL Server 2000 permits the creation of links to OLE DB data sources called linked servers. After linking to an OLE DB data source, it is possible to reference rowsets from the OLE DB data sources as tables in Transact-SQL statements and to pass commands to the OLE DB data sources and include the resulting rowsets as tables in Transact-SQL statements. Each distributed query can reference multiple linked servers and can perform either update or read operations against each individual linked server. A single distributed query can perform read operations against some linked servers and update operations against other linked servers. The Microsoft OLE DB Provider for Oracle allows distributed queries to query data in Oracle databases.

Incorrect Answers:

A: Replication allows for the copying, distribution, and modification of data across an enterprise. This involves the copying of database data from one server to another. This is not the best solution in this scenario as this will require frequent replication because the reports must be updated several times a day.

C: Data Transformation Services is a set of graphical tools and programmable objects that allows for the extraction, transformation, and consolidation of data from disparate sources into single or multiple destinations. In this scenario, however, the Oracle database needs to be included in a query. Hence, Data Transformation Services are inappropriate.

D: Microsoft ActiveX scripts can be used to add functionality to Data Transformation Services packages, which is a set of graphical tools and programmable objects that allows for the extraction, transformation, and consolidation of data from disparate sources into single or multiple destinations. In this scenario, however, the Oracle database needs to be included in a query. Hence, Data Transformation Services are inappropriate.

Question No: 2

You are the administrator of a SQL Server 2000 computer named SQL1. You want to perform ad hoc distribution queries against a database that is stored on a SQL Server 2000 computer named SQL2. SQL2 contains several databases, and each of these databases uses a different collation. You want to ensure that comparisons in distributed queries are evaluated correctly. You also want to minimize administrative overhead. How should you configure SQL1?

- A.** Use the ALTER DATABASE statement to change the collation of the databases on SQL1.
- B.** Add SQL2 as remote server.
- C.** Add SQL2 as a linked server.

Select the Use Remote Collation check box, and do not specify a collation name.

D. Add SQL2 as a linked server.

Select the Use Remote Collation check box, and specify a collation name.

Repeat this process once for each database on SQL2.

Answer: C.

Explanation: We should set up SQL2 as a linked server. This would enable to issue distributed queries to it. We should also select Remote Collation option to use the collation on the remote server.

Note: SQL Server 2000 allows the creation of links to OLE DB data sources called linked servers. After a link has been created it is possible to reference rowsets from the OLE DB data sources as tables in Transact-SQL statements and pass commands to the OLE DB data sources and include the resulting rowsets as tables in Transact-SQL statements. Each distributed query can reference multiple linked servers and can perform either update or read operations against each individual linked server. A single distributed query can perform read operations against some linked servers and update operations against other linked servers.

The linked server options defined by using the sp_serveroption stored procedure control how SQL Server uses collations from linked servers. When the UseRemoteCollation option is set to TRUE, it specifies that the collation of remote columns is used for SQL Server data sources, and the collation specified in CollationName is used for data sources other than SQL Server. When FALSE, distributed queries always use the default collation of the local server instance, and CollationName and the collation of remote columns are ignored. If the UseRemoteCollation option is set to TRUE, we can specify the name of the collation used by the remote data source in the CollationName option. The collation must, however, be supported by SQL Server. This option should be used when accessing an OLE DB data source other than SQL Server, but whose collation matches one of the SQL Server collations

Incorrect Answers:

A: The COLLATE clause of the ALTER DATABASE statement can be used to change the collation used by the database. However, this solution is inappropriate as the SQL2 server contains several databases, each of which uses a different collation. We would thus not be able to specify all the collations.

B: Remote servers are set up in pairs with both servers configured to recognize each other as remote servers. The configuration options for both servers should be set so that each instance of SQL Server allows remote users to execute procedure calls. However, support for remote servers is provided for backward compatibility only. Applications that must execute procedures against remote instances of SQL Server 2000 should use linked servers instead.

D: In this scenario we cannot specify the collation name as SQL2 contains several databases, each of which use a different collation.

SECTION 4: Configure SQL Mail and SQLAgentMail. (2 questions)

Question No: 1

You are the administrator of a SQL Server 2000 computer. You create several jobs on the server and schedule them to run during the evening.

You configure an operator on the server and configure the jobs to send an e-mail to this operator if a job fails for any reason.

While testing the jobs, you discover that e-mail messages are not being sent to the operator when a job fails. You use SQL Server Enterprise Manager to confirm that SQL Mail is started.

You need to ensure that e-mail messages are sent to the operator whenever a job fails. What should you do?

- A. Configure SQLAgentMail to use a valid MAPI profile
 - B. Configure SQL Mail to use a mailbox that delegates permissions to the operator
 - C. Use the **net send** notification method to send the operator e-mail messages
 - D. Create a job that executes the **xp_startmail** stored procedure
- Configure the job to run whenever SQL Server 2000 starts

Answer: A.

Explanation: SQLAgentMail must be used when mail stored procedures are used to send notifications to operators by e-mail or pager. In order to use SQLAgentMail, a mail profile must be configured for SQL Server Agent to use to connect to the mail server and the mail server must be extended MAPI-compliant.

Incorrect Answers:

B: SQL Mail provides a way to receive e-mail messages generated by SQL Server 2000. Messages can be triggered to provide information about the status of a job or a warning caused by an alert. SQL Mail can include a result set in a reply to e-mail messages that contain queries and allows SQL Server to send and receive e-mail by establishing a client connection with a mail server. However, SQL Mail must have a mail profile created in the same user domain account or context that is used to start an instance of SQL Server 2000 as it looks for the defined mail profile in the domain account that triggered the e-mail message.

C: Net send is used to send messages to other users, computers, or messaging names across a network. The Messenger service must be running to receive messages. It is not used for e-mail or paging purposes.

D: The xp_startmail stored procedure is used to start a SQL Mail client session. If there is an existing mail session, xp_startmail does not start a new one. If mail is being used on the same computer on which SQL Server is also running, the mail client must be started either before xp_startmail is executed, or before SQL Server is started if SQL Mail is configured to automatically start when SQL Server starts.

Question No: 2

You are the administrator of several SQL Server 2000 computers. You configure these servers to send e-mail messages to SQL Server operators whenever a SQL server agent job fails or encounters an error.

Your company's software developers create a set of stored procedures that send query results in email messages about stored procedures that have not been successful.

By using a MAPI profile named MS Exchange Settings, you verify that SQLAgentMail is correctly configured and working properly. You need to configure the servers so that the stored procedures function correctly. Which two actions should you take? (Each correct answer presents part of the solution. Choose two.)

- A. Set the security permissions on your e-mail server to permit multiple connections to the SQL Server mailbox.
- B. Set the security permissions on your e-mail server to permit connections from the MSSQLServer service account.
- C. Configure the MSSQLServer service to use the same service account as the SQLServerAgent service.
- D. Configure the SQLServerAgent service to use the same service account as the MSSQLServer service.

- E. Configure SQL Mail to use the **MS Exchange Settings** MAPI profile.
- F. Configure SQL Mail to use the **Default Settings** MAPI profile.

Answer: C, E.

Explanation:

C: SQL Mail provides a way to receive e-mail messages generated by SQL Server 2000. Messages can be triggered to provide information about the status of a job or a warning caused by an alert. SQL Mail can include a result set in a reply to e-mail messages that contain queries and allows SQL Server to send and receive e-mail by establishing a client connection with a mail server. However, SQL Mail must have a mail profile created in the same user domain account or context that is used to start an instance of SQL Server 2000, i.e. it must use the same domain account that is used by the SQLServerAgent service, as it looks for the defined mail profile in the domain account that triggered the e-mail message.

E: It must also have a mail server that is extended MAPI-compliant and must be configured to use a valid MAPI profile. In this scenario we know that a MAPI profile named MS Exchange Settings works correctly with SQLAgentMail. The servers must also be configured to use the MS Exchange Settings MAPI profile.

Reference: INF: How to Configure SQL Mail (Q263556)

Incorrect Answers:

A, B: There is no need to set any security permissions on your e-mail server.

D: The MSSQLServer isn't running. MSSQLServer should be configured to use the same account as the SQLServerAgent service (which is running), not the opposite way around.

F: There is no guarantee that a default settings MAPI profile would work. We must use the MS Exchange Settings MAPI profile since we know that it works with SQLAgentMail.

SECTION 5: Configure network libraries. (0 questions)

SECTION 6: Troubleshoot failed installations. (0 questions)

PART II: Creating SQL Server 2000 Databases (23 questions)

SECTION 1: Configure database options for performance. Considerations include capacity, network connectivity, physical drive configurations, and storage locations. (5 questions)

Question No: 1

You are the administrator of a SQL Server 2000 computer. You have two new hard disks on which you will create a database named inventory. You want to insert, update, and delete data as quickly as possible.

Which two actions should you take? (Each correct answer presents parts of the solution. Choose two)

- A. Configure the hard disks as two mirrored NTFS volumes.
- B. Configure the hard disks as one mirrored NTFS volumes.
- C. Configure the hard disks as two independent NTFS volumes.
- D. Configure the hard disks as one extended NTFS volume.
- E. Place inventory_data.mdf on the first volume and inventory_log.ldf on the second volume.
- F. Place inventory_data.mdf on the first volume and inventory_data2.ndf and inventory_log.ldf on the second volume.

G. Place inventory_data.mdf and inventory_log.ldf on the same volume.

Answer: C, E.

Explanation: Placing log files and the data files on different physical disks or independent volumes, and placing the log file of a drive that does not hold the operating system can improve database performance as it allows the frequently written to log file to exist a disk with its own read/write heads and thus less seek time.

Note: A Windows 2000 supports basic disks, in which physical disks are divided into partitions, and dynamic disks, in which disks are divided into volumes that can comprise an entire disk, two or more entire disks, portions of a disk, or portions of up to 32 physical disks. A NTFS volume is volume that is formatted with the NTFS file system. Physical database design involves mapping the logical design onto physical media to allow for the data to be physically accessed and maintained as quickly as possible by taking advantage of the available hardware and software features. It is important to correctly design the database and to take advantage of hardware and software features early in the development cycle of a database application, because making changes to these components at a later stage it is difficult.

Incorrect Answers:

A: For redundancy and fault protection the log file can be placed on a mirrored volume. It would not be a good idea to put data files on a mirrored volume though. Performance when accessing the data would suffer.

Note: A mirrored volume is a fault-tolerant volume that duplicates data on two physical disks. It provides data redundancy as it copies the data on one volume to another volume on a different disk. When one of the physical disks fails, the data on the failed disk becomes unavailable, but the system can continue to operate by using the mirror disk. In such a set up, the data file and the log file will reside on the same drive. This will not improve database performance.

B: A mirrored volume cannot comprise of one volume only. It must comprise of two volumes on separate physical disks as a mirrored volume is a fault-tolerant volume that duplicates data on two physical disks. When one of the physical disks fails, the data on the failed disk becomes unavailable, but the system can continue to operate by using the mirror disk.


D: Data is written and read sequentially on extended volumes. As a result the data file and the log file will be placed on the same physical disk.

F: In this scenario the data file is separated into a primary data file, that uses the .mdf extension, and a secondary data file, that uses the .ndf extension. However, the secondary data file is placed on the same volume as the log file. This will lead to performance improvements only in cases where data is not read from the secondary data file.

G: Placing the data file and the log file on the same volume will not improve database performance as it will not allow parallel processing.

Question No: 2

You are the administrator of a SQL Server 2000 computer. The server contains a database named Inventory. The database is configured as shown in the exhibit.



Users frequently need details about parts. Users search for data by using the following query.

```

SELECT Parts.SKU, Locations.Description, Manufacturer.Name,
PartsLocations.Qty, PartsLocations.LastInventoried
FROM Parts
INNER JOIN Manufacturer ON Parts.ManufacturerID=
Manufacturer.ManufacturerID
INNER JOIN PartsLocation ON Parts.PartID = Parts.Location.PartsID
INNER JOIN Locations ON Parts.LocationID= Locations.LocationID
WHERE SKU?
  
```

You need to accelerate query response time. What should you do?

- A. Create a parameterised stored procedure to retrieve the data.
- B. Create a denormalized table that is maintained by triggers.
- C. Use ad hoc queries to retrieve the data.
- D. Create a nonclustered index on the primary key of each table.

Answer: D.

Explanation: A nonclustered index is similar to an index in a textbook. The data is stored in one place, the index in another, with pointers to the storage location of the data. The index contains entries describing the exact location in the table of the data values being searched for in the queries and should be used for columns that contain a large number of distinct values; queries that do not return large result sets; columns frequently involved in search conditions of a query that return exact matches; and decision-support-system

applications for which joins and grouping are frequently required. Create multiple nonclustered indexes on columns involved in join and grouping operations, and a clustered index on any foreign key columns.

Note: A better solution, not listed here, would be to create clustered indexes on the primary keys and nonclustered indexes on the foreign keys.

Incorrect Answers:

A: A parameterized query would not improve performance.

Parameterized queries are SQL queries written for reusability. They contain parameter markers as

placeholders for data that will change from execution to execution. In the Data Transformation Services (DTS) tasks that use parameterized queries.


B: Reasonable normalization often includes few regularly executed queries that use joins involving more than four tables. In this scenario four tables are being joined in the query. The normalization here is thus permissible.

Note: Database normalization often improves query performance. When useful indexes are available, the SQL Server 2000 query optimizer can select rapid, efficient joins between tables. Normalization is advantageous as it aids faster sorting and index creation; enables a larger number of clustered indexes; allows for narrower and more compact indexes; reduces the number of indexes per table, which improves the performance of INSERT, UPDATE, and DELETE statements; and reduces null values and the opportunity for inconsistency, which increase database compactness. However, as normalization increases, so do the number and complexity of joins required to retrieve data. Too many complex relational joins between too many tables can hinder performance.

C: An ad hoc query is a query that is composed for a particular task. The code in this scenario is an example of an ad hoc query.

Question No: 3

You are the administrator of a SQL Server 2000 computer. The server contains a database named Accounting. The data files are configured as shown in the exhibit.



The database has grown to 7 GB. The server has no more available disk space. You need to make more disk space available in the accounting database without affecting system performance. What should you do?

A. Compress drive E.

B. Add a new hard disk to the server.

Create a new file on the new hard disk as part of the PRIMARY filegroup.

C. Add a new hard disk to the server.

Create a new file on the new hard disk as part of a new filegroup.

D. Map a network drive to the server.

Create a new data file on the network drive as part of the PRIMARY filegroup.

E. Map a network drive on the server.

Create a new data file on the network drive as part of a new filegroup.

Answer: B.

Explanation: When a SQL Server 2000 server runs out of disk space, additional hard drives must be added to the server. A file can then be created on the new disk as part of an existing filegroup or as part of a new filegroup. In this scenario a new file should be placed on the new hard disk for best performance. This file can either be placed in the PRIMARY filegroup or a new filegroup could be created. The performance gain would be equal in both cases, but if we use a new filegroup we would have to configure it so that it would be used.

Incorrect Answers:

A: SQL Server data files cannot be installed on a system that uses compression.


C: When a SQL Server 2000 server runs out of disk space, additional hard drives must be added to the server. A file can then be created on the new disk as part of an existing filegroup or as part of a new filegroup. To use the new filegroup we would have to configure it (add data to it or make it the default filegroup). Without configuration the new filegroup will not be used. All data would still be added to the default PRIMARY filegroup.

D: Mapping a network drive to the server will increase the disk space available to the database but query time will be limited to the network environment and network transfer speeds. This will hamper query response time.

E: Mapping a network drive to the server will increase the disk space available to the database but query time will be limited to the network environment and network transfer speeds. This will hamper query response time.

Question No: 4

You are the administrator of a SQL Server 2000 computer. The server contains a database named Sales. The database will store sales transactions. Retail locations must be able to add transactions to the database 24 hours a day every day. The server is configured as shown in the exhibit.



You need to configure the data files and transaction log for the Sales database. What should you do?

A. Place the transaction log on physical disk 0 and the data file on the RAID-5 disk array.

- B. Place the transaction log on the RAID-5 disk array and the data file on physical disk 0.
- C. Place the transaction log and the data file on physical disk 0.
- D. Place the transaction log and the data file on the RAID-5 disk array.

Answer: D.


Explanation: Because retail locations must be able to add transactions to the database 24 hours a day every day, we must minimize any possible downtime and ensure that the server has high availability. Using a hardware-based RAID-5 disk array, we would increase availability as RAID-5 provides redundancy of all data on the array, allowing a single disk to fail and be replaced in most cases without system downtime.

Incorrect Answers:


- A:** By placing the transaction log on physical disk 0 we will not take advantage of the fault tolerance made possible by the RAID-5 disk array. The transaction log is thus not fault tolerant in this solution.
- B:** By placing the data file on physical disk 0 we will not take advantage of the fault tolerance made possible by the RAID-5 disk array. The data file is thus not fault tolerant in this solution.
- C:** By placing the transaction log and the data file on physical disk 0 we will not take advantage of the fault tolerance made possible by the RAID-5 disk array. In this solution the database is not fault tolerant and a disk failure would result in unnecessary downtime.

Question No: 5

You are the administrator of a SQL Server 2000 computer. The server is configured as shown in the Database Server Configuration exhibit.



You need to create a new database named Inventory. Employees in your company will use the Inventory database to track inventory data. Users will require immediate responses to queries that help them locate where parts are stored. The tables in the database will be configured as shown in the Database Schema exhibit.



The database will consume 14 GB of disk space. You must configure the data files and transaction log to accelerate query response time.

Which two courses of action should you take? (Each correct answer represents part of the solution. Choose two.)

- A. On drive C, create a transaction log.
On drive D, create a data file in the PRIMARY filegroup.
On drive E, create a data file in the SECONDARY filegroup.
- B. On each drive, create a transaction log.
On each drive, create a data file in the PRIMARY filegroup.
- C. On drive D, create a transaction log.
On drive E, create a data file in the PRIMARY filegroup.
- D. On the PRIMARY filegroup, create all tables and all indexes.
- E. On the PRIMARY filegroup, create all tables. On the SECONDARY filegroup, create all indexes.
- F. On the PRIMARY filegroup, create the Parts table and its indexes.
On the SECONDARY filegroup, create all other tables and their indexes.

Answer: A, E.

Explanation:

A: The transaction log should be placed on a separate disk. A transaction log would not benefit of two or more disks since it is sequential. Create two data files on two separate physical disks would improve performance.

E: Separating the tables and the indexes improve performance for in particular joins.

Note: With SQL Server 2000 it is possible to create tables or indexes on specific filegroups. This allows us to control where the database's tables and indexes are physically located as filegroups can be placed on different hard drives. Placing tables in one filegroup and the table's nonclustered indexes in another filegroup on different physical disk can improve database performance because it will allow separate threads to access the tables and indexes. However, a table and its clustered index cannot be separated into different filegroups as the clustered index determines the physical order of the data in the table. Furthermore, the transaction log should be placed on a physically separate disk. The transaction log file is written serially; therefore, using a separate, dedicated disk allows the disk heads to stay in place for the next write operation. This will further improve performance.

Incorrect Answers:

B: To improve performance, the transaction log should not be placed on the same physical disk that the

tables or indexes are placed on. The transaction log file is written serially; therefore, placing it on a separate, dedicated disk allows the disk heads to stay in place for the next write operation and thus improves performance. Furthermore, placing tables in one filegroup and the table's nonclustered indexes in another filegroup on different physical disk can improve database performance because it will allow separate threads to access the tables and indexes. However, a table and its clustered index cannot be separated into different filegroups as the clustered index determines the physical order of the data in the table.

C: The transaction log should be placed on a physically separate disk. The transaction log file is written serially; therefore, using a separate, dedicated disk allows the disk heads to stay in place for the next write operation. This will further improve performance. In this solution, the tables and their indexes are placed on the same disk. This solution thus does not take advantage of the performance improvements that can be realised through the correct utilization of the available hardware.

D: In this solution, the tables and their indexes are placed on the same disk. This solution thus does not take advantage of the performance improvements that can be realised through the correct utilization of the available hardware.

F: Greater performance gains can be realized by placing tables in one filegroup and the table's nonclustered indexes in another filegroup on different physical disk. This will allow separate threads to access the tables and indexes.

SECTION 2: Attach and detach databases. (2 questions)

Question No: 1

You are the administrator of a SQL Server 2000 computer. Each night you need to send a copy of the server's msdb database in an e-mail message to the administrator of another SQL Server computer.

You create a job that contains the following steps:

- **Detach the msdb database**
- **Send the database to the administrator in an e-mail message**
- **Reattach the msdb database**

You test the job, and it fails on the first step. You need to ensure that the msdb database is mailed to the administrator every night.

What should you do?

A. Ensure that the SQLServerAgent service is running under a user account that has database owner access to the **msdb** database

B. Delete the first and last steps in the job

C. Configure the job to back up the msdb database to a temporary file
Send the file to the administrator in an e-mail message

D. Insert a new job step before the first step
Configure the new step to terminate all processes that use the **msdb** database

Answer: C.

Explanation: The current job fails since the detach msdb statement fails. It fails because detaching a database requires that no one is using the database. The msdb database is used by the SQL Server Agent service, which performs scheduled activities such as backups and replication tasks.

We want to send the msdb database. The msdb database could be detached if all processes that use it were terminated. Then we could send the database file and reattach the msdb database. This would stop scheduled tasks from being run during this period. A better approach would be to backup the database and send the temporary file instead.

Note: A new SQL Server 2000 installation automatically includes six databases: master, model, tempdb, pubs, Northwind, and msdb. The master, model, and tempdb databases cannot be detached.

Incorrect answers:

A: The problem that the SQL Service Agent is using the msdb database. This problem would still be there even user account running the SQL Server Agent service is changed.

B: The msdb database is used by the SQL Server Agent service and the database file would be locked. It cannot be sent directly in an e-mail message.

D: The msdb database could be detached if all processes that use it were terminated. Then we could send the database file and reattach the msdb database. This would stop scheduled tasks from being run during this period.

Question No: 2

You are the administrator of two SQL Server 2000 computers. One of these servers contains a 4-GB database named marketing. You want to remove the marketing database from one server and add it to the other as quickly as possible.

What should you do?

A. Detach the database from the original server by using the `sp_detach_db` stored procedure. Copy the database and the transaction log files to the new server, and attach them by using the `sp_attach_db` stored procedure.

B. Use the DTS export wizard to transfer all database objects from the original server to the new server. Drop the database from the original server.

C. Run a full backup of the database on the original server. Create a new database named marketing on the new server. Restore the backup in the new marketing database. Drop the database from the original server.

D. Shut down the original server. Copy the database and the transaction log files to the new server. Use the DISK INIT and DISK REFIT statements to attach the data file to the new server. Drop the database from the original server.

Answer: A.

Explanation: Detaching and attaching databases is sometimes a useful technique for quickly moving a database from one computer to another. To move a database, or database file, to another server or disk the database must first be detached from the original server or disk; moved to the other server or disk and then attached on the server or disk. In addition the new location of the moved file(s) must be specified. The `sp_detach_db` stored procedure is used to detach the database and can run UPDATE STATISTICS on all tables before detaching the database while the `sp_attach_db` stored procedure is used to attach the database in its new location.

Incorrect Answers:

B: DTS can import data from a text file or from an OLE DB data source such as a Microsoft Access 2000 database into SQL Server. It can also export data from SQL Server to an OLE DB data destination. It provides for high-speed data loading from text files into SQL Server tables. However, it is not the fastest way to move a whole database from one server to another.

C: Before a database can be restored from a backup to a new location; the database must first be recreated in the destination as the backup holds only the data files that belong in the various tables. It does not hold the table definitions. This is not the fastest method to move a database.

D: The DISK INIT and DISK REFIT commands cannot be used to move a database from one server to another.


Note: DISK INIT and DISK REINIT are features that are used in a SQL Server 6.x. In SQL Server 6.x, DISK INIT is used to create database or transaction log devices to store the specified database or transaction log when DISK INIT followed either a CREATE DATABASE or ALTER DATABASE statement while DISK REINIT is used to restore device entries to appropriate system tables when the device entry was missing from sysdevices. In SQL Server 2000 the CREATE DATABASE and ALTER DATABASE statements allow for the creation of separate data and log files. Both CREATE DATABASE and ALTER DATABASE create operating system files and databases in a single step. It is thus recommended that all references to DISK INIT should be removed in SQL Server 2000 and replaced with references to either CREATE DATABASE or ALTER DATABASE.

SECTION 3: Create and alter databases.


SUBSECTION A: Add filegroups. (1 question)

Question No: 1

You are the administrator of a SQL Server 2000 computer. The server contains a database named Contracts. The server is configured as shown in the server configuration exhibit.



The database files are configured as shown in the database properties exhibit.



The database developers have been creating new tables and indexes without specifying a filegroup. The primary filegroup is reaching its maximum capacity.

You want the developers to continue adding new objects, but you do not want then to change the way they create objects. You do not want to run out of disk space. You also want to minimize the time it takes to administer the database.

What should you do?

- A. Back up the existing objects on the PRIMARY filegroup. Drop them from the database. Re-create them on the SECONDARY filegroup.
- B. Set the file growth on the PRIMARY filegroup to **UNLIMITED**.
- C. Set the PRIMARY filegroup so that it is read-only.
- D. Set the SECONDARY filegroup as the default filegroup.

Answer: D.

Explanation: By setting the secondary filegroup to the default filegroup, user data will be added to this filegroup and the database would be allowed to grow.

Note: In SQL Server 2000 a database is mapped to a set of operating-system files, which stores all data and objects in the database, such as tables, stored procedures, triggers, and views. In addition to the log file, there are two operating-system files: the primary data file, which contains the startup information for the database and is used to store data, and the secondary data file, which contains all of the data that is not part of the primary data file. If the primary file can hold all of the data in the database, databases do not need to have secondary data files but every database must have one primary data file. Some databases may be large enough to need multiple secondary data files or to use secondary files on separate disk drives to spread data

across multiple disks. These files can be placed into different filegroups. This allows files to be grouped together for administrative and data allocation/placement purposes. The filegroup that contains the primary file is the PRIMARY filegroup. This filegroup contains the primary data file and any other files that are not put into another filegroup. User-defined filegroups are called SECONDARY filegroups and are any filegroups that are specifically created by the user when first creating or later altering the database. One filegroup is designated as the DEFAULT filegroup. When objects are created in the database without specifying to which filegroup they belong, they are assigned to the DEFAULT filegroup. By default, the PRIMARY filegroup is designated the DEFAULT filegroup. The DEFAULT filegroup can be changed using the ALTER DATABASE statement. This allows any objects that do not have a filegroup specified when they are created are allocated to the data files in the new DEFAULT filegroup. However, allocation for the system objects and tables remains within the PRIMARY filegroup, not the new DEFAULT filegroup.

Incorrect Answers:

A: The database startup information and the system objects and tables must remain in the PRIMARY filegroup. These database objects cannot be moved to a SECONDARY filegroup.

B: The FILEGROWTH property can only be specified for operating-system files, i.e., primary or secondary data files. The FILEGROWTH property cannot be specified for a filegroup.

C: SQL Server 2000 allows filegroups to be marked as read-only. A filegroup marked read-only cannot be modified. Any user-defined filegroup, i.e., SECONDARY filegroup can be marked as read-only. However, the PRIMARY filegroup cannot be marked as read-only as it contains the system objects and tables. These include the tempdb and deleted database tables.

SUBSECTION B: Configure filegroup usage. (2 questions)

Question No: 1

You are the administrator of a SQL Server 2000 database. You import a table of geographic information from a Microsoft access database into a SQL Server 2000 database. The table has 12,000 rows. Each row averages 5,000 bytes. The table contains lookup data that does not change. You want to minimize the size of the data file and the time required to back up the data. Which two actions should you take? (Each correct answer presents part of the solution. Choose two)

- A.** Create a 60-MB data file named geography.ndf
- B.** Create a 95-MB data file named geography.ndf
- C.** Create a 60-MB data file named geography.mdf
- D.** Create a 95-MB data file named geography.mdf
- E.** Place the table in the PRIMARY filegroup.
- F.** Place the table in a new filegroup named LOCATION.

Answer: B, F.

Explanation: As SQL Server 2000 supports backing up or restoring individual files or file groups within a database, it is possible to reduce back up time by creating a separate data file table so that only that table will be backed up. This separate data file will have to be a secondary data file because a database can only have one primary data file, which is identified by the .mdf extension, it is the starting point of the database and locates all the other files in the database. Secondary data files are identified by the .ndf extension and comprise all of the data files except the primary data file. A database need not have a secondary data file and it is also possible to create multiple secondary data files on database.

Rows are stored in data pages. Rows cannot cross different data pages. The size of data page in SQL Server 2000 (or SQL 7.0) is 8,192 bytes. We can only store one row in every data page in this scenario. Thus, to

estimate the size of the table we would need to multiply the number of rows, 12,000 in this scenario, by 8,192 bytes. This will give us a table size of 98,304,000 bytes. As computer data storage is calculated using the binary rather than the decimal system, there is 1,024 bytes per Kb and 1,024 Kb per MB, rather than 1,000 bytes per Kb and 1,000 Kb per MB, we thus require a database size of at least 93.75 MB.

Incorrect Answers:

A: The data size of this table is insufficient as a minimum size of 93.75 MB is required.

C: The data size of this table is insufficient as a minimum size of 93.75 MB is required. Furthermore, the .mdf file extension indicate in this solution implies that the data file will be placed in the primary filegroup. It would thus not be possible to backup the data file pertaining to the table only.

D: The .mdf file extension indicate in this solution implies that the data file will be placed in the primary filegroup. It would thus not be possible to backup the data file pertaining to the table only.

E: As SQL Server 2000 supports backing up or restoring individual files or file groups within a database, it is possible to reduce back up time by creating a separate data file table so that only that table will be backed up. This separate data file will have to be a secondary data file because a database can only have one primary data file, which is identified by the .mdf extension, it is the starting point of the database and locates all the other files in the database.

Question No: 2

You are the administrator of a database that contains 64 lookup tables. These tables store static data that should not change. However, users report that some of this data is being changed. You need to prevent users from modifying the data.

You want to minimize changes to your security model and to your database applications. How should you modify the database?

A. Create a filegroup named LOOKUP. Move the lookup tables to this filegroup. Select the read only check box for the filegroup.

B. Create a database role named datamodifier. Grant SELECT permissions to the datamodifier role. Add all users to the role.

C. Deny INSERT, UPDATE, and DELETE permissions for all users. Create stored procedures that modify data in all tables except lookup tables. Require users to modify data through these stored procedures.

D. Create a view of the lookup tables. Use the view to allow users access to the lookup tables.

Answer: A.

Explanation: SQL Server 2000 allows filegroups other than the primary filegroup to be marked as readonly. A filegroup that is marked read-only cannot be modified. Tables that must not be modified can be placed on a filegroup that can then be mark as read-only. This will prevents accidental updates.

Incorrect Answers:

B: This solution effectively permits all users to modify any table in the database, as they would all have SELECT permissions to the database.

C: This solution will effectively prevent all users from being able to performing these tasks on all tables in the database as they are denied permissions to run INSERT, UPDATE, or DELETE statements against the database. A better solution would be to apply the deny permissions at the table level of the lookup tables. This would, however, require unnecessary administrative time.

D: Views are virtual tables that are used to filter tables or rows that the users should not have access to, however, if the users have access to the base tables, they can still run INSERT, UPDATE, and

DELETE statements against the base table.

SUBSECTION C: Expand and shrink a database. (2 questions)

Question No: 1

You are the administrator of a SQL Server 2000 computer. The server contains a database named Inventory. The Inventory database has a table named StorageLocations that stores the location of

parts contained in your company's warehouses. The StorageLocations table is configured as shown in the exhibit.

StorageLocations				
	Column Name	Data Type	Length	Allow Nulls
🔑	StorageLocationID	int	4	
	LocationDescription	varchar	50	
	PartID	int	4	
	UnitsStored	bigint	8	
	DateStored	datetime	8	
	TagID	int	4	

The LocationDescription field is usually described with a name 10 to 25 characters in length. The locations never store more than 100,000 units of any given part.

You want to modify the table's schema to save space. You cannot lose any existing data. You want to do this by using the minimum amount of administrative time and server resources.

Which Transact-SQL statement should you execute?

- A.** ALTER TABLE [dbo].[StorageLocations]
ALTER COLUMN [UnitsStored] [int] NOT NULL
- B.** ALTER TABLE [dbo].[StorageLocations]
[LocationDescription] [char] (25) COLLATE
SQL_Latin1_General_CP1_CI_AS NOT NULL
- C.** ALTER TABLE [dbo].[StorageLocations]
ALTER COLUMN [UnitsStored] [smallint] NOT NULL
- D.** ALTER TABLE [dbo].[StorageLocations]
[LocationDescription] [nvarchar] (25) COLLATE
SQL_Latin1_General_CP1_CI_AS NOT NULL

Answer: A.

Explanation: In this scenario there are two considerations that must go into the table schema modification. The space allocated to the LocationDescription and the use of data types. The scenario states that the LocationDescription field is usually described with a name, 10 to 25 characters in length. The locations never store more than 100,000 units of any given part. In the table the parts is represented by the PartID column; the amount of units by UnitsStored Storage column; the locations by the LocationID column and the LocationDescription field by the LocationDescription column. Because the LocationDescription field is usually described with a name, 10 to 25 characters in length, does not mean that the character length of this field is never longer than 25 characters. We thus cannot reduce the character length of this column. We must thus turn our attention to the UnitsStored column. SQL Server 2000 has four integer data types. These are bigint, which has a storage size of 8 bytes and is whole numbers that range from -9223372036854775808 through 9223372036854775807; int, which has a storage size of 4 bytes and ranges from -2,147,483,648 through 2,147,483,647; smallint, which has a storage size of 2 bytes and ranges from -32,768 through 32,767; and tinyint, which has a storage size of 1 byte and ranges from 0 through 255. As no more than

100,000 units of any given part will be held in storage and int ranges from -2,147,483,648 through 2,147,483,647, this data field can be represented by the int data type. ALTER COLUMN is used in conjunction with ALTER TABLE. The latter modifies a table definition by altering, adding, or dropping columns. Through ALTER COLUMN a specified column's data type length can be altered.

Incorrect Answers:

B: The scenario states that the LocationDescription field is usually described with a name, 10 to 25 characters in length. This does not mean that the character length of this field is never longer than 25 characters. If we ALTER COLUMN to a char data type field with a field length of only 25, truncation of the LocationDescription will occur if the description exceeds 25 characters. We thus cannot reduce the character length of this column.

C: As no more than 100,000 units of any given part will be held in storage and int ranges from -2,147,483,648 through 2,147,483,647, this data field can be represented by the int data type. smallint, however, does not provide a sufficient range as it only ranges up to 32,767 while we require a range of up to at least 100,000.

Note: SQL Server 2000 uses four integer data types, namely, bigint, which has a storage size of 8 bytes and is whole numbers that range from -9223372036854775808 through 9223372036854775807; int, which has a storage size of 4 bytes and ranges from -2,147,483,648 through 2,147,483,647; smallint, which has a storage size of 2 bytes and ranges from -32,768 through 32,767; and tinyint, which has a storage size of 1 byte and ranges from 0 through 255

D: While the char data type is a fixed-length character data with a specified length has a storage size in bytes that is equal to the specified data length, varchar is a variable-length character data with a specified length but has a storage size in bytes that is equal to the actual length of the entered data and not the specified length of the data. The varchar data type should be used when it is expected that the data values in a column will vary considerably in size. The use of this data type can reduce database size but this solution restricts the data length to 25 characters. The scenario states that the LocationDescription field is usually described with a name, 10 to 25 characters in length. This does not mean that the character length of this field is never longer than 25 characters. If we ALTER COLUMN to a varchar data type field with a field length of only 25, truncation of the LocationDescription will occur if the description exceeds 25 characters. We thus cannot reduce the character length of this column.

Question No: 2

You are the administrator of a SQL Server 2000 computer. The server contains a database that stores inventory data. Another database administrator has created a number of scheduled jobs to maintain the inventory database. Each weeknight the following jobs occur.

- **A BULK INSERT job imports data at 10:00 P.M.**
- **Indexes are rebuilt at 10:15 P.M.**
- **Data integrity checks are performed at 10:30 P.M.**
- **A differential backup is performed at 10:45 P.M.**
- **A DBCC SHRINKDATABASE job runs at 11:00 P.M.**

You notice that the final job often fails and returns the following error message: "Server Msg 3140, Level 16, State 3. Could not adjust the space allocation for file 'inventory_data'."

You need to ensure that the final job runs without errors. What should you do?

A. Create a new DBCC SHRINKFILE job to shrink the inventory_data file. Schedule the new job to run at the same time as the DBCC SHRINKDATABASE job.

- B.** Modify the DBCC SHRINKDATABASE job so that it uses DBCC SHRINKFILE statement to shrink each file individually.
- C.** Increase the time between the data integrity checks and the differential backup.
- D.** Increase the time between the differential backup and the DBCC SHRINKDATABASE job.

Answer: D.

Explanation: The DBCC SHRINKDATABASE cannot be executed until the previous job step, the differential backup, has been completed. We should increase the time between these two job steps, or even better configure the last job step to run only after the differential backup has been completed.

Note: The DBCC SHRINKDATABASE statement shrinks data files on a per-file basis but shrinks log files as if all the log files existed in one contiguous log pool. The target size for the data and log files cannot be smaller than the minimum size of a file that was specified when the file was originally created, or the last explicit size set with a file size changing operation such as the ALTER DATABASE statement with the MODIFY FILE option or the DBCC SHRINKFILE statement.

Incorrect Answers:

A: The DBCC SHRINKDATABASE command shrinks the whole database. Shrinking a single data file would not be the optimal solution to shrink the database.

Note: The DBCC SHRINKFILE statement shrinks the size of the specified data or log file in the related database and is used to reduce the size of a file to smaller than its originally created size. The minimum file size for the file is then reset to the newly specified size. By using the EMPTYFILE argument in conjunction with the DBCC SHRINKFILE statement before executing the ALTER DATABASE statements, we can remove the data in a file. Other users can be working in the database when the file is shrunk, thus the database being shrunk does not have to be in single-user mode when using the DBCC SHRINKFILE statement. The DBCC SHRINKDATABASE statement shrinks data files on a per-file basis but shrinks log files as if all the log files existed in one contiguous log pool. The target size for the data and log files cannot be smaller than the minimum size of a file that was specified when the file was originally created, or the last explicit size set with a file size changing operation such as the ALTER DATABASE statement with the MODIFY FILE option or the DBCC SHRINKFILE statement. These two functions thus adjust the space allocation on a per-file basis and cannot run at the same time.

B: The DBCC SHRINKDATABASE statement shrinks data files on a per-file basis already. There is no need to use the DBCC SHRINKFILE statement to shrink each file individually.

Note: The DBCC SHRINKDATABASE statement shrinks log files as if all the log files existed in one contiguous log pool. The target size for the data and log files cannot be smaller than the minimum size of a file that was specified when the file was originally created, or the last explicit size set with a file size changing operation such as the ALTER DATABASE statement with the MODIFY FILE option or the DBCC SHRINKFILE statement.

C: Database integrity checks such as the DBCC CHECKDB statement should be run before any major system changes, such as a hardware or software upgrade, or after a hardware failure, or if a system problem is suspected. It is not necessary to run database integrity checks as part of the normal backup or maintenance procedures.

SUBSECTION D: Set database options by using the ALTER DATABASE or CREATE DATABASE statements. (1 question)

Question No: 1

You are the database administrator for an accounting company. You are installing SQL Server 2000 on a new Microsoft Windows 2000 Server computer. The computer will run two client/server database applications that your company has purchased. Each application uses a separate database. During the installation, you specify SQL_Latin1_General_CP1_CI_AI as the default collation for the SQL server computer. After completing the installation, you discover that one of the new applications, named Financials, is sorting information incorrectly.

You contact the application vendor and discover that the Financials application requires the SQL_Latin1_CP437_BIN collation. You need to configure the correct collation for the applications without diminishing the performance of the SQL Server computer.

What should you do?

- A.** Rerun the SQL Server Setup and specify the SQL_Latin1_General_CP437_BIN collation.
- B.** Run the **rebuildm** utility and specify the SQL_Latin1_General_CP437_BIN collation.
- C.** Use the ALTER DATABASE statement to specify the SQL_Latin1_General_CP437_BIN collation for the **Financials** database.
- D.** Use the ALTER TABLE statement to specify the SQL_Latin1_General_CP437_BIN collation for each table in the **Financials** database.

Answer: C.

Explanation: The collation of a database can be specified in the COLLATE clause of the ALTER DATABASE statement. This will change the default collation of the database that was specified when the database was created, or to alter the collation for a database where the collation was not specified when the database was created and thus the default collation of the server was used.

Reference: BOL: ALTER DATABASE

Incorrect Answers:

A: The collation of a database can be altered using the COLLATE clause of the ALTER DATABASE statement. It is therefore not necessary to reinstall the instance of SQL Server 2000 or to rerun the setup program. Furthermore, rerunning the SQL Server 2000 setup program would disrupt the performance of the SQL Server computer.

B: The rebuildm utility alters the master database and can be used to change the collation settings for an instance of SQL Server 2000. However, it is only the Financial database that needs to be altered as it is financial application is the only application that uses the SQL_Latin1_General_CP437_BIN collation.

D: The COLLATE clause can be specified at different levels. It can be specified for the database, a table, a column or an expression. In this scenario the collation can be altered at the database level as this will reduce administrative effort when new tables are added to the database.

SUBSECTION E: Size and place the transaction log. (1 question)

Question No: 1

You are the administrator of a SQL Server 2000 computer. You are configuring a database for an inventory application. The hard disks on your server are configured as shown in the exhibit.

The operating system files are located on drive C. Your database will store a maximum of 6 GB of data and requires a maximum of 2 GB of disk space for the transaction log.

You want to optimize database performance. What should you do?

- A. Add a 2-GB transaction log to drive D, a 3-GB data file to drive E, and a 3-GB data file to drive F
- B. Add a 1-GB transaction log to drive C, a 1-GB transaction log to drive D, a 3-GB data file to drive E, and a 3-GB data file to drive F
- C. Add a 1-GB transaction log to drive E, a 1-GB transaction log to drive F, a 3-GB data file to drive E, and a 3-GB data file to drive F
- D. Add a 2-GB transaction log to drive F, a 3-GB data file to drive D, and a 3-GB data file to drive E

Answer: D.

Explanation: Basically there are two methods to improve performance in this scenario:

Method 1: put the transaction log file on a separate physical disk.

One way to improve performance and recoverability of your database is to place the data files and the transaction log files on separate disk subsystems, and to use dedicated disks for your data files so that data read and write (I/O) requests are not competing with other I/O requests on those disks. This is (A).

Method 2: data files split across two physical disks

Splitting your data files among a number of disks results in increased I/O performance in particular when multiple controllers are used like in this scenario. This is (D).

To decide which is better we must think about the application. This is an inventory system and we can assume that the queries would be more frequent than additions and update of rows. So we should choose method 2 since it gives the best performance for reading data.

Note: Physical database design involves mapping the logical design onto physical media to allow for the data to be physically accessed and maintained as quickly as possible by taking advantage of the available hardware and software features. It is important to correctly design the database and to take advantage of hardware and software features early in the development cycle of a database application, because making changes to these components at a later stage it is difficult. Placing log files and the data files on different physical disks, and placing the log file of a drive that does not hold the operating system can improve database performance as it allows the frequently written to log file to exist a disk with its own read/write heads.

Incorrect Answers:

A: In an inventory system data is read more frequently and not so many changes are made to the database. We should optimize the data files not the log file.

B: Placing log files and the data files on different physical disks and placing the log file of a drive that does not hold the operating system can improve database performance as it allows the frequently written to log file to exist a disk with its own read/write heads. In this part of the log file is placed on drive C, which would also hold the operating system.


C: Placing log files and the data files on different physical disks can improve database performance as it allows the frequently written to log file to exist a disk with its own read/write heads. In this solution the log file and the data file are placed on the same physical drive.

SECTION 4: Create and manage objects. Objects include constraints, indexes, stored procedures, triggers, and views. (9 questions)

Question No: 1

You are the administrator of a SQL Server 2000 computer. The server contains a database named sales. You need to change the way customer Ids appear in the Customers table. The database schema

is shown in the exhibit.



You need to automate the process of updating the primary key tools. You also want to minimize records locks and administration within the database during the update process.

What should you do?

- A. Add an ON UPDATE CASCADE constraint to the CustomerID field in the Customers table. Modify the values in the CustomerID field in the Customers table.
- B. Create a duplicate record that has a new CustomerID value. Update the foreign key fields in the invoices, contacts, and quotes tables with the new value.
- C. Disable the FOREIGN KEY constraints. Within a transaction, modify the values in the CustomerID field in the Customers table and all related foreign key values in the invoices, contacts, and quotes tables. Re-enable the FOREIGN KEY constraints after the keys are changed.
- D. Create a Data Transformation Services package. Use the package to transform the CustomerID value and the values of the related foreign keys in the invoices, contacts, and quotes tables.

Answer: A.

Explanation: A cascading referential integrity constraints allows for the automatic updating of primary key tools by defining the actions SQL Server 2000 takes when a user deletes or updates a key to which existing foreign keys point.

The ON UPDATE CASCADE option specifies that if an attempt is made to update a key value in a row, where the key value is referenced by foreign keys in existing rows in other tables, all of the foreign key values are also updated to the new value specified for the key. If cascading referential actions have also been defined on the target tables, the specified cascading actions are also taken for the key values updated in those tables.

Reference: BOL: Cascading Referential Integrity Constraints

Incorrect Answers:

B: In this solution the FOREIGN KEY fields are updated manually.

C: Constraints allow you to define the way SQL Server 2000 enforces the integrity of a database. Constraints define rules regarding the values allowed in columns and are the standard mechanism for enforcing integrity. Using constraints is preferred to using triggers, rules, and defaults. A FOREIGN KEY is a column or a combination of columns used to establish and enforce a link between the data in two tables and are used in conjunction with the primary keys on a referenced table. It can be used together with cascading referential constraints to automate the process of the updating primary keys. Disabling the foreign keys and re-enabling them after the data has been changed will not allow for the automation of the primary key tools. Running the whole update process within a transaction would lock the table.

D: Data Transformation Services is a set of graphical tools and programmable objects that allows for the extraction, transformation, and consolidation of data from disparate sources into single or multiple destinations.

Question No: 2

You are the administrator of a SQL Server 2000 computer. One of the databases on the server contains a table named complaints. This table is used to store information about customer complaints. The customer service representatives in your company add and edit the information in the complaints table. They work with this table by using a number of client applications, including a web-based application and a Microsoft Windows 32-bit application.

The customer service manager discovers that some customer complaints are marked as closed before they are resolved to the customer's satisfaction.

You need to notify the customer service manager whenever a complaint is marked as closed. You do not want to make any changes to the client applications.

What should you do?

- A.** Create an UPDATE trigger that sends an e-mail message.
- B.** Create a stored procedure that sends an e-mail message. Use the stored procedures to make changes to the data.
- C.** Create a user-defined function that sends an e-mail message. Use the function to make changes to the data.
- D.** Create a rule that validates data entry. Bind the rule to the **CompliantClosed** column.

Answer: A.

Explanation: Adding an appropriate UPDATE trigger that notifies the service manager whenever a complaint is marked is closed would accomplish the requirements. The client applications application would not have to be changed.

Note: In SQL Server 2000, triggers are a special type of stored procedure that fire automatically when an UPDATE, INSERT, or DELETE statement is issued against a table or view. Triggers are powerful tools that can be used to enforce business rules automatically when data is modified and can automate the processing for a company. Like stored procedures, triggers contain Transact-SQL statements which can include instructions to send an e-mail message when a particular column is updated

Incorrect Answers:

B: If a stored procedure is created and used the client application would have to be adjusted accordingly. But we don't want to make any changes in the client application.

Note: A stored procedure is a set of Transact-SQL statements compiled into a single execution plan and can return data as output parameters, which can return either data or a cursor variable; return

codes, which are always an integer value; a result set for each SELECT statement contained in the stored procedure or any other stored procedures called by the stored procedure; or as a global cursor that can be referenced outside the stored procedure. Stored procedures assist in achieving a consistent implementation of logic across applications as the SQL statements and logic needed to perform a task can be coded and tested once in a stored procedure. Thereafter each application or user that needs to perform that task can execute the stored procedure. Coding business logic into a single stored procedure also offers a single point of control for ensuring that business rules are correctly enforced. Stored procedures can also be used to customize security permissions in much the same way as views. Stored procedure permissions can be used to prevent users from modifying data in certain columns and then grant users permissions to execute the procedure.

C: If a user defined function is created and used the client application would have to be adjusted accordingly. But we don't want to make any changes in the client application.


Note: In programming languages, functions are subroutines used to encapsulate frequently performed logic. Any code that must perform the logic incorporated in a function can call the function rather than having to repeat all of the function logic. SQL Server 2000 supports built-in functions and user-defined functions. Built-in functions operate as defined in the Transact-SQL Reference and cannot be modified while user-defined functions allow the user to define their own Transact-SQL functions using the CREATE FUNCTION statement. User-defined functions take zero or more input parameters, and return a single value. Some user-defined functions return a single, scalar data value, such as an int, char, or decimal value.

D: Just creating a rule would only check the input, it would not notify the service manager as required.

Note: Rules are a backward-compatibility feature of SQL Server 2000 and performs the same functions as CHECK constraints. CHECK constraints are the preferred, standard way to restrict the values in a column as they are more concise than rules.

Question No: 3

You are the administrator of a SQL Server 2000 computer. The server contains a database named inventory. The database has a Parts table that has a field named InStock. When parts are shipped, a table named PartsShipped is updated. When Parts are received, a table names PartsReceived is updated. The relationship of these tables is shown in the exhibit.



You want the database to update the InStock field automatically. What should you do?

- A. Add triggers to the PartsShipped and the PartsReceived tables that update the InStock field in the Parts table.
- B. Create a user-defined function those calculations that calculates current inventory by running aggregate queries on the PartsShipped and PartsReceived tables.
- C. Use a view that creates an InStock as a part of an aggregate query.
- D. Create stored procedures for modifying the PartsReceived and the PartsShipped tables that also modify the InStock field in the parts table. Use these procedures exclusively when modifying data in the PartsReceived and the PartsShipped tables.

Answer: A.

Explanation: Triggers are a special class of stored procedure defined to fire automatically when an UPDATE, INSERT, or DELETE statement is issued against a table or view. They are powerful tools that can be used to enforce business rules automatically when data is modified. Triggers can extend the integrity checking logic of SQL Server constraints, defaults, and rules, although constraints and defaults should be used instead whenever they provide all the needed functionality. In this scenario an AFTER UPDATE trigger can be used to update the tables to update the PartsShipped and the PartsReceived tables that update the InStock column in the parts table.

Incorrect Answers:

B: A user-defined function is a subroutine that is made up of one or more Transact-SQL statements and can be used to encapsulate code for reuse. It can be used to run aggregate calculations. However, a simple update of the InStock column, and not an aggregate function, is required when changes are made to the PartsShipped and PartsReceived tables. This can be accomplished through the use of triggers.

C: A view would not be able to detect changes in any of the columns.

Note: A view is a virtual table that allows data to be accessible through it rather than through the base table. A view can be referenced through Transact-SQL statements in the same way a table is. In

addition a table can be used to restrict a user to specific rows and/or columns in a table, join columns from multiple tables so that they appear as a single virtual table and can aggregate information instead of supplying details to the user.

D: A simple update of the InStock column, and not a stored procedure, is required when changes are made to the PartsShipped and PartsReceived tables. This can be accomplished through the use of triggers.

Note: A stored procedure is a group of Transact-SQL statements compiled into a single execution plan. Stored procedures can return data as output parameters; return codes; a result set for each SELECT statement contained in the stored procedure or any other stored procedures called by the stored procedure; or a global cursor that can be referenced outside the stored procedure. Stored procedures assist in achieving a consistent implementation of logic across applications and can improve performance.

Question No: 4

You are the administrator of a SQL Server 2000 computer. The server contains a database named sales. Your company uses the database to store the sales department's responses to requests for price quotations. The database is configured as shown in the exhibit.

Developers in your company create an application used for saving quotations. The application executes the following transact-SQL statement:

```
UPDATE QuotationSummary  
SET CustomerName = 'Jean-Paul Deloria'  
WHERE QuotationID = 12
```

When the application executes the statement, the developers receive the following error:

Server: Mrh4403, Level 16, Stats 1, Line 1

View or function 'QuotationSummary' is not updateable because it contains aggregates:

You want developers to be able to use this UPDATE statement without generating errors. What should you do?

- A.** Create an INSTEAD OF trigger on the view to update the composite tables.
- B.** Grant the developers UPDATE permissions on each base table.
- C.** Add a cascading update trigger on each base table.
- D.** Create a QuotationSummary table and populate the table with data from the base tables.

Answer: A.

Explanation: SQL Server 2000 uses INSTEAD OF triggers and partitioned views to enhance the class of updateable views. INSTEAD OF triggers can be created on a view in order to make a view updateable. The INSTEAD OF trigger is executed instead of the data modification statement on which the trigger is defined. This trigger allows the user to specify the set of actions that need to take place in order to process the data modification statement. Thus, if an INSTEAD OF trigger exists for a view on a given data modification statement, the corresponding view is updateable through that statement. If the view is a partitioned view, the view is updateable, subject to certain restrictions.

Incorrect Answers:

B: Before a user can use a database; they must be given certain permissions that are determined by the type of activity they will be permitted on the database. To run UPDATE statements against a database the user must be granted UPDATE permissions. This can be specified at the table or


database level. In this scenario, granting the permission at the database level would require less administrative effort. However, the error message is not related to permissions but to the updateability of the table.

C: Cascading referential integrity constraints makes it possible to define the actions SQL Server 2000 takes when a user attempts to delete or update a key to which existing foreign keys point. Cascading referential integrity constraints fire AFTER triggers after all of the cascading referential actions directly caused by the original DELETE or UPDATE are performed.

D: Creating a table that combines the composite tables that the view is based on would deny us the ability to use the view to return aggregated column values.

Question No: 5

You are the administrator of a SQL Server 2000 computer. The server contains a database named Inventory. Users report that several storage locations in the UnitsStored field contain negative numbers. You examine the database's table structure. The table properties are configured as shown in the exhibit.



You correct all the negative numbers in the table. You must prevent the database from storing negative numbers. You also want to minimize use of server resources and physical I/O.

Which transact-SQL statement should you execute?

- A. ALTER TABLE dbo.storageLocations ADD CONSTRAINTS
CK_storageLocations_UnitsStored
CHECK (UnitsStored >= 0)
- B. CREATE TRIGGER CK_UnitsStored On StorageLocations
FOR INSERT, UPDATE AS
IF INSERTED, UnitsStored < 0 ROLLBACK TRAN
- C. CREATE TABLE ck_unitsstored As @Units >= 0

```
GO
Sp_bindrule 'OK_UnitsStored',
'StorageLocations, UnitsStored'
GO
D. CREATE PROC UpdateUnitsStored
(StorageLocationID int, @UnitsStored bigint)AS
IF @UnitsStored < 0
DATSERERROR (50099, 17)
ELSE
UPDATE storahelocations
SET UnitsStored = @UnitsStored
WHERE StorageLocationID = @storageLocationID
```

Answer: A.

Explanation: CHECK constraints are used to enforce domain integrity by restricting the values that are accepted by a column. They determine which values are valid by testing a logical expression that is not based on data in another column. The problem posed by this scenario can be solved through the implementation of such a constraint. In this solution the CHECK constraint is added in line 1 and the CHECK expression is specified in line 3. The expression specifies that the values contained in the UnitsStored column must be more than or equal to 0.

Incorrect Answers:

B: This solution uses a trigger to perform the CHECK constraint, which a trigger can do. However, it is recommended that triggers only be used when the CHECK constraint cannot meet the functionality required by the application. This is thus not the best answer.

C: In this solution creates a new table is created with the CREATE TABLE statement. This is inappropriate to the problem, as it does not address the insertion of negative values in the existing table.

D: This solution creates a stored procedure to perform the CHECK constraint.

Question No: 6

You are the administrator of a SQL Server 2000 computer. The server contains a database named MedicalRecords.

The database tables are configured as shown in the exhibit.

Patients				
Column Name	Data Type	Length	Allow Nulls	
PatientID	int	4		
Title	char	4	✓	
FirstName	char	25		
LastName	char	35		
Address	varchar	50		
Address2	varchar	50	✓	
City	char	25		
State	char	2		
Zip	char	10		
Phone	char	15	✓	

PatientsDoctors				
Column Name	Data Type	Length	Allow Nulls	
PatientID	int	4		
DoctorID	int	4		
CaseID	int	4		
DateOpened	datetime	8		
DateClosed	char	10	✓	

PatientsInsurance				
Column Name	Data Type	Length	Allow Nulls	
PatientID	int	4		
InsuranceCarrierID	int	4		
InsuranceID	int	4		
PolicyTypeID	int	4		
Expiration	datetime	8		

The existing PatientID field is an identity field. According to a new government regulation, the patient IDs must include a three-character prefix before the numeric portion of the ID. The patient's home address defines the prefix.

You want to comply the regulation while minimizing changes to the database. Which two actions should you take? (Each correct answer presents part of the solution. Choose Two)

- A. Drop the FOREIGN KEY constraints. In each table, change the data type of the PatientID field to char.
- B. Add an On Update constraints to the Patients table, and link it to the PatientID field. Update the field in the Patients table.
- C. In the Patients table, add a new field named StatePatientID that has a data type of char. Populate the field with the prefix and the existing PatientID.
- D. In each table, update the PatientID field.
Re-create the FOREIGN KEY constraints.
- E. In the Patients table, update the PatientID field.
- F. Create a trigger that populates the StatePatientID field when a new patient is created.

Answer: C, F.

Explanation: The PatientID column currently holds an integer value. The new requirement demands that the patientID column includes the characters preceding the integer value. We want to implement this change with minimal change of the database. There are basically two approaches to this problem:

Method 1: add a new column StatePatientID:

1. Add the StatePatientID with data type char to the patients table. (C)
2. Populate the field with the three letters and the existing number from the PatientID column. (C)
3. Create a trigger that populates the StatepatientID field when a new patient is created. (F)

Method 2: change the PatientID column to char:

4. Drop foreign key constraints on the PatientID column. (A)
5. Change the data type of the PatientID column to char. (A)
6. Update the PatientID column in each table. We add the three letters to the integer value. (D)
7. Recreate the foreign key constraint. (D)

Method 1 only makes changes to one table meeting the requirement that changes to the database should be minimized. It would on the other hand introduce redundant data since the new StatePatientID column is derived from the PatientID column. It is no requirement to keep a normalized database so method 1 is the preferred solution.

Incorrect Answers:

A, D: Dropping all foreign and primary key constraints on the PatientID column, changing the column PatientID to char and then recreating all foreign and primary key constraints on the PatientID column throughout the database would not minimize changes to the database.

B: An ON UPDATE trigger on the PatientID field doesn't make much sense. The PatientID field is an identity and wouldn't be updated.

E: If the PatientID field is changed in the Patients table then it must be changed in all the other tables as well.

Question No: 7

You are the administrator of a SQL Server 2000 computer. The server contains confidential

information about contracts on which your company has placed bids.

Company policy requires that bid information be removed permanently from the database one year after the bid closes. You need to comply with this policy and minimize server overhead.

What should you do?

A. Create a Data Transformation Services (DTS) package to remove bids that have a closing data older than one year.

B. Create a trigger to delete any bids that have a closing data older than one year.

Bind the trigger to the bids table for all INSERT, UPDATE, and DELETE events.

C. Create a stored procedure to delete any bids that have a closing data order than one year.

Use SQL server agent to schedule the stored procedure to run every night.

D. Create a view that contains a WHERE clause to exclude bids that have a closing date older than one year.

Answer: C

Explanation: This scenario requires us to remove old information on a regular basis. First we create a stored procedure which removes all rows where the bid information date is older than the current date minus one year. Then we schedule this procedure to run every day. We would then be sure to never have any rows older than a year in the table.

Note: Stored procedures also improve performance as the stored procedure is cached after it is run. This results in a reduction of physical I/O. We can further reduce Server overhead by scheduling the stored procedure to run during off-peak hours.

Incorrect Answers:

A: A Data Transformation Services (DTS) package provides a set of tools that allows us to extract, transform, and consolidate data stored in a variety of formats and in different locations into single or multiple destinations supported by DTS connectivity. This solution is inappropriate, as the data exists on the server and not in numerous locations and would thus be correctly formatted for use on an instance SQL Server 2000.

B: Triggers are a special type of stored procedure that executes automatically when an UPDATE, INSERT, or DELETE statement is run against a table or view. In this scenario however, the timestamp of the data in the database must be checked and the data deleted one year after the bid has closed. This cannot be accomplished through the use of constraints or triggers.

D: Creating a view that excludes the bids that have a closing date older than one year does not address the company policy that requires that bid information be removed permanently from the database one year after the bid closes. This data will still exist in the base tables.

Question No: 8

You are the administrator of a SQL Server 2000 computer. The server contains a database named Inventory. The database contains a table that is used to store information about equipment scheduling. The EquipmentSchedules table is configured as shown in the exhibit.

sp_help 'EquipmentSchedules'

Name	Owner	Type	Created_datetime
EquipmentSchedules	dbo	user table	2000-06-11 12:57:06.623

Column_name	Type	Computed	Length	PREC	Scale	Nullable	TrimTrailingBlanks	FixedLen
EquipmentID	int	no	4	10	0	no	(n/a)	(n/a)
CustomerID	int	no	4	10	0	no	(n/a)	(n/a)
Location	varchar	no	50			yes	no	no
StartDate	datetime	no	8			no	(n/a)	(n/a)
EndDate	datetime	no	8			no	(n/a)	(n/a)
Rate	money	no	8	19	4	no	(n/a)	(n/a)

Identity	Seed	Incrment	Not For Replication
EquipmentID	1	1	0

RowGUIDCol
No rowguidcol column defined.

Data located on filegroup
PRIMARY

index name	index description	index keys
PK_EquipmentSchedules	clustered, unique, primary ...	EquipmentID

constraint_type	constraint name	delete action	update action	Status check
PRIMARY KEY (clustered)	PK_EquipmentSchedules	(n/a)	(n/a)	(n/a)

Users report that some equipment schedules have an end date that is earlier than the start date. You need to ensure that the start date is always earlier than or equal to the end date.

You also want to minimize physical I/O. You do not want users to change the Transact-SQL statements they use to modify data within the database.

What should you do?

- A. Create a constraint that compares the start date to the end date.
- B. Create a trigger that compares the start date to the end date.
- C. Create a rule that compares the start date to the end date.
- D. Create a stored procedure that tests the start and end dates before inserting the row into the database.

Answer: A.

Explanation: We simply put a CHECK constraint on the table that compares the start date to the end date.

Note: In SQL Server 2000 constraints are used to define ways in which the integrity of a database can be enforced automatically. Constraints define rules regarding the values allowed in columns and are the

standard mechanism for enforcing integrity. Using constraints is preferred to using triggers, rules, and defaults. SQL Server 2000 supports five classes of constraints: NOT NULL specifies that the column does not accept NULL values; CHECK constraints enforce domain integrity by limiting the values that can be placed in a column; UNIQUE constraints enforce the uniqueness of the values in a set of columns; PRIMARY KEY constraints identify the column or set of columns whose values uniquely identify a row in a table; and FOREIGN KEY constraints identify the relationships between tables. A CHECK constraint can be used to specify a Boolean search condition that is applied to all values entered for the column; all values that do not evaluate to TRUE are rejected. Multiple CHECK constraints can also be specified for each column.

Incorrect Answers:

B: Triggers should be used when the features supported by constraints cannot meet the functional needs

of the application. In this scenario CHECK constraints should be used as they meet the functional requirements of the application.

Note: Triggers are special types of stored procedure that fire automatically when the data in a specified table is modified. They are invoked in response to an INSERT, UPDATE, or DELETE statement and can query other tables and can include complex Transact-SQL statements. Triggers can cascade changes through related tables in the database; enforce restrictions that are more complex than those defined with CHECK constraints as triggers can reference columns in other tables; evaluate the state of a table before and after a data modification and take action(s) based on that difference; and multiple triggers of the same type on a table can allow multiple, different actions to take place in response to the same modification statement. Constraints and triggers each have benefits that make them useful in certain situations. The primary benefit of triggers is that they can contain complex processing logic that uses Transact-SQL code. Triggers can support all of the functionality of constraints; however, triggers are not always the best method for a given feature. It is recommended that entity integrity be enforced at the lowest level by indexes that are part of PRIMARY KEY and UNIQUE constraints or are created independently of constraints; that domain integrity be enforced through CHECK constraints, and that referential integrity be enforced through FOREIGN KEY constraints, where these features meet the functional needs of the application.

C: Rules are a backward-compatibility feature of SQL Server 2000 and performs the same functions as CHECK constraints. CHECK constraints are the preferred, standard way to restrict the values in a column as they are more concise than rules.


D: We must provide some mechanism that makes the procedure execute. It would not execute by itself.

Note: A stored procedure is a set of Transact-SQL statements compiled into a single execution plan and can return data as output parameters, which can return either data or a cursor variable; return codes, which are always an integer value; a result set for each SELECT statement contained in the stored procedure or any other stored procedures called by the stored procedure; or as a global cursor that can be referenced outside the stored procedure. Stored procedures assist in achieving a consistent implementation of logic across applications as the SQL statements and logic needed to perform a task can be coded and tested once in a stored procedure. Thereafter each application that needs to perform that task can execute the stored procedure. Coding business logic into a single stored procedure also offers a single point of control for ensuring that business rules are correctly enforced.

Question No: 9

You are the administrator of a SQL Server 2000 computer. The server contains a database named Inventory. The database has a Parts table that has a field named InStock.

When the parts have shipped, a table named PartsShipped is updated. When the parts are received, a table named PartsReceived is updated. The relationship of these tables is shown in the exhibit.



You want the database to update the InStock field automatically. What should you do?

- A. Add triggers to the PartsShipped and the PartsReceived tables that update the InStock field in the Parts table.
- B. Create a user-defined function that calculates current inventory by running aggregate queries on the **PartsShipped** and **PartsReceived** tables.
- C. Use a view that creates an InStock field as part of an aggregate query.
- D. Create stored procedures for modifying the **PartsShipped** and **PartsReceived** tables that also modify the InStock field in the **Parts** table.

Use these procedures exclusively when modifying data in the PartsShipped and **PartsReceived** tables.

Answer: A.

Explanation: The Instock column of the Parts table must be updated when records are changed in the other two tables PartsReceived and PartsShipped. The only way to accomplish this is to use triggers on the PartsReceived and PartsShipped tables.

Note: Triggers are a special class of stored procedure defined to fire automatically when an UPDATE, INSERT, or DELETE statement is issued against a table or view. They are powerful tools that can be used to enforce business rules automatically when data is modified. Triggers can extend the integrity checking logic of SQL Server constraints, defaults, and rules, although constraints and defaults should be used instead whenever they provide all the needed functionality. In this scenario an AFTER UPDATE trigger can be used to update the tables to update the PartsShipped and the PartsReceived tables that update the InStock column in the parts table.

Incorrect Answers:

B: A function cannot, by itself, be used to update the contents of one table when changes are made to another table.

Note: A user-defined function is a subroutine that is made up of one or more Transact-SQL statements and can be used to encapsulate code for reuse. It can be used to run aggregate calculations. However, a simple update of the InStock column, and not a aggregate function, is required when changes are made to the PartsShipped and PartsReceived tables. This can be accomplished through the use of triggers.

C: A view cannot be used to update the contents of one table when changes are made to another table.

Note: A view is a virtual table that allows data to be accessible through it rather than through the base table. A view can be referenced through Transact-SQL statements in the same way a table is. In addition a table can be used to restrict a user to specific rows and/or columns in a table, join columns from multiple tables so that they appear as a single virtual table and can aggregate information instead of supplying details to the user.

D: A function cannot, by itself, be used to update the contents of one table when changes are made to another table.

Note: A stored procedure is a group of Transact-SQL statements compiled into a single execution plan. Stored procedures can return data as output parameters; return codes; a result set for each SELECT statement contained in the stored procedure or any other stored procedures called by the stored procedure; or a global cursor that can be referenced outside the stored procedure. Stored procedures assist in achieving a consistent implementation of logic across applications and can improve performance.

PART III: Managing, Monitoring, and Troubleshooting SQL Server 2000 Databases (38 questions)

SECTION 1: Optimize database performance. Considerations include indexing, locking, and recompiling. (13 questions)

Question No: 1

You are the administrator of a SQL Server 2000 computer. Your company modifies a sales application it uses so that it can access data from a SQL server database rather than a Microsoft Access database. You configure the new SQL server database. The tables are configured as shown in the table schema exhibit.

Users report that queries to this database respond slowly. You use system manager to analyze database activity and receive the results shown in the exhibit.

You need to modify the database to accelerate query response time. What should you do?

- A. Place primary keys and foreign key indexes on the table.
- B. Remove all indexes from the tables.
- C. Replace the table joins with correlated subqueries.
- D. Use server-side cursor to retrieve data.

Answer: A.

Explanation: Indexes on primary and foreign keys would most likely improve performance of queries, especially joins.

Note primary key, foreign key: A PRIMARY KEY is a set of columns from a table that are guaranteed to have unique values for each row of that table and also called a PRIMARY KEY constraint, because it effectively constrains the values can be add to the table, i.e. it prevents the adding of a row to the table whose PRIMARY KEY columns are equal to the corresponding values of some other row in that table. A

FOREIGN KEY is the correspondence between a set of columns in one table and the set of PRIMARY KEY columns in some other table and is also called a FOREIGN KEY constraint because it constrains table rows by ensuring that any row added to the foreign-key table has a corresponding row in the primary-key table, i.e., it requires that any row added to the foreign-key table have values in the foreign-key column that correspond to the respective values of the primary key columns for some row in the primary-key table.

Note index: In databases, indexes improve query response time by allowing for quicker retrieval of data in a table by not having to scanning the entire table. An index in a database is a list of values in a table with the storage locations of rows in the table that contain each value. Indexes can be created on either a single column or a combination of columns in a table.

Incorrect Answers:

B: Removing all indexes from the tables in a database would thus force queries to use full table scans and would hamper rather than enhance query performance.

C: No performance would be gained by replacing joins by correlated subqueries.

Note join: Joins are used to retrieve data from two or more tables based on logical relationships between the tables. They indicate how SQL Server 2000 should use data from one table to select the rows in another table and can specify the column from each table to be used for the join, typically specifying a foreign key from one table and its associated key in the other table; or it can specify a logical operator to be used in comparing values from the columns. Joins can be specified in either the FROM or WHERE clauses. The join conditions combine with the WHERE and HAVING search conditions to control the rows that are selected from the base tables referenced in the FROM clause, while specifying the join conditions in the FROM clause helps separate them from any other search conditions that may be specified in a WHERE clause, and is the recommended method for specifying joins.

Note correlated subquery: A correlated subquery is an extension of a subquery, which is a SELECT query that returns a single value and is nested inside a SELECT, INSERT, UPDATE, or DELETE statement, or inside another subquery and can be used anywhere an expression is allowed. With a correlated subquery, many queries can be evaluated by executing the subquery once and substituting the resulting value or values into the WHERE clause of the outer query.

D: A server side cursor would not improve performance in this scenario. A server-side cursor could improve performance in scenarios where only a few rows should be sent to the user. Instead of sending all rows a server-side cursor would be able to send only the selected rows. Bandwidth would be saved.

Note: SQL Server 2000 returns result sets as default result sets or as server cursors. Default result sets minimizes overhead; provides maximal performance in fetching data; supports only the default forward-only, read-only cursor functionality; returns one row at a time; supports only one active statement at a time on a connection; and supports all Transact-SQL statements, while server cursors support all cursor functionality; can return blocks of rows; supports multiple active statements on a single connection; balances cursor functionality against performance; does not support any Transact-SQL statement that returns more than a single result set.

Question No: 2

You are the administrator of a SQL Server 2000 computer. The server contains a database named MedicalRecords. Users query the database using the following stored procedure:


```
CREATE PROC GetManyVisits
(@PatientID INT,
@MadeTable TINYINT OUTPUT)
```

```

AS
DECLARE @Visits INT
SET @Visits = (SELECT COUNT(*) FROM Cases
WHERE PatientID = @PatientID)
IF @Visits > 5
BEGIN
SELECT P.PatientName, D.PhysicianName,
V.DateVisited, V.VisitDuration,
INTO ##ClientVisits
FROM Patients P
JOIN Physician D ON P.PhysicianID = D.PhysicianID
JOIN Visits V ON P.PatientID = V.PatientID
WHERE P.PatientID = @PatientID
SET @MadeTable = 1
END
ELSE SET @MadeTable = -1

```

Users report that queries to the database respond slowly. You use System Monitor to measure the system and receive the results shown in the exhibit.



You need to modify the server to accelerate query response time. What should you do?

- A. Execute sp_recompile on the stored procedure.
- B. Replace the stored procedure with a user-defined function.
- C. Convert the stored procedure into two procedures, one procedure to decide if the ##ClientVisits table should be populated and one to populate the ##ClientVisits table.
- D. Replace the stored procedure with a temporary stored procedure created by the user.

Answer: A.

Explanation: The sp_recompile stored procedures causes stored procedures and triggers to be recompiled the next time they are run. Queries used by stored procedures and triggers are optimized only when they are compiled. As indexes or other changes that affect statistics are made to the database, compiled stored procedures and triggers may lose efficiency. By recompiling stored procedures and triggers that act on a table, the queries can be reoptimized.

Note: A stored procedure is a group of Transact-SQL statements compiled into a single execution plan that can return data as OUTPUT parameters, which can return either data such as an integer or character value or a cursor variable; as return codes, which are always an integer value; as a result set for each SELECT statement contained in the stored procedure or any other stored procedures called by the stored procedure; and as a global cursor that can be referenced outside the stored procedure. Stored procedures assist in achieving a consistent implementation of logic across applications and can also improve performance as it contains many tasks that would otherwise be implemented as a series of SQL statements.

Incorrect Answers:


B: A user-defined function is similar to a store procedure in that they are subroutines made up of one or more Transact-SQL statements that can be used to encapsulate code for reuse.

C: In this solution the first stored procedure would perform the same function as the IF statement in the existing procedure. This would thus not improve the performance of the stored procedure.


D: SQL Server 2000 also supports temporary stored procedures that are dropped automatically when the user disconnects from the server. Temporary stored procedures are stored in tempdb system database and are useful when connected to earlier versions of SQL Server. Temporary stored procedures can be used when an application builds dynamic Transact-SQL statements that are executed several times. Rather than have the Transact-SQL statements recompiled each time, a temporary stored procedure that is compiled on the first execution can be created, and can then execute the precompiled plan multiple times. However, heavy use of temporary stored procedures will lead to contention on the system tables in tempdb system database.

Question No: 3

You are the administrator of a SQL Server 2000 computer. The server contains a database named Sales. The database is configured as shown in the exhibit.



In the last six months, the database has grown by 4 GB. Users report that query response time has slowed. The options set up on the database are shown in the Database Options exhibit.



You want to accelerate query response time. What should you do?

- A. Update the database statistics.
- B. Add indexes to the foreign key fields.
- C. Truncate the transaction log.
- D. Run the database maintenance plan wizard.
- E. Drop primary keys from all tables.

Answer: B.

Explanation: A FOREIGN KEY is a column or combination of columns used to establish and enforce a link between the data in two tables and should be indexed because changes to PRIMARY KEY constraints are checked with FOREIGN KEY constraints in related tables and FOREIGN KEY columns are often used in join criteria when the data from related tables is combined in queries by matching the column(s) in the FOREIGN KEY constraint of one table with the primary or unique key column(s) in the other table. These indexes allow SQL Server 2000 to find related data in the foreign key table quicker. Although creating this index is not a requirement, a foreign key relationship between two tables indicates that the two tables have been optimized to be combined in a query that uses the keys as its criteria.

Incorrect Answers:

A: SQL Server keeps statistics about the distribution of the key values in each index and uses these statistics to determine which index would be the best index to use for a particular query process. As the data in a column changes, index and column statistics can become out-of-date. This can hinder the query optimizer's ability to make optimal decisions on how to process a query. As a result query

performance will deteriorate. When statistics become out-of-date they can be updated automatically if the auto update statistics check box on the tables property dialog box is checked. The exhibit shows that this check box is checked; therefore it is not out-of-date statistics that is hampering query performance.

C: Old transaction log records that are no longer necessary for recovering or restoring a database should be deleted to make way for new log records. The process of deleting these old log records to reduce the size of the logical log is called truncating the log. However, the active portion of the transaction log can never be truncated. The active portion of the log is the part of the log needed to recover the database at any time, so must have the log images needed to roll back all incomplete transactions. It must always be present in the database in case the server fails because it will be required to recover the database when the server is restarted. This truncation occurs automatically at the completion of any BACKUP LOG statement or every time a checkpoint is processed, provided the database is using the Simple Recovery model. This includes both explicit checkpoints resulting from a CHECKPOINT statement and implicit checkpoints generated by the system.


D: The Database Maintenance Plan Wizard can be used to set up the core maintenance tasks necessary to ensure that a database performs well, is regularly backed up, and is checked for inconsistencies. The Database Maintenance Plan Wizard creates a SQL Server 2000 job that performs these maintenance tasks automatically at scheduled intervals and can be used to reorganize the data on the data and index pages by rebuilding indexes with a new fill factor; compress data files by removing empty database pages; update index statistics to ensure the query optimizer has up-to-date information about the distribution of data values in the tables; perform internal consistency checks of

the data and data pages within the database to ensure that a system or software problem has not damaged data; backup the database and transaction log files; and to set up log shipping.

E: Indexes are used to speed up the retrieval of rows for SELECT statements. Because indexes control how the data rows of the table are ordered and have pointers to locate specific data rows they can also increase the speed of updates and deletes. This is because SQL Server must first find a row before it can update or delete the row. Removing the indexes would thus hamper query performance.

Question No: 4

You are the administrator of SQL server 2000 computer. The server contains a database named MedicalRecords. Users access medical records by using the PatientID field. This field is the clustered primary key for the Patients table. When users try to access medical records, the database responds slowly. You examine the database options as shown in the exhibit.



You want to accelerate query response time and minimize administrative overhead. How should you reconfigure the database?

- A. Create a SQL Server Agent job to execute the UPDATE STATISTICS statement, and schedule the job to run weekly.
- B. Select the Auto Update Statistics check box.
- C. Run the database maintenance plan wizard, and accept the default settings.
- D. Rebuild the primary key as a nonclustered primary key.
- E. Clear the Auto Create Statistics check box.

Answer: B.

Explanation: The **Auto Update Statistics** option is selected by default but in the Exhibit it has been disabled. Apparently the statistics are out-of-date and the queries are running slowly. The **Auto Update Statistics** option should be enabled to increase performance and minimize administrative overhead.

Note: SQL Server keeps statistics about the distribution of the key values in each index and uses these statistics to determine which index would be the best index to use for a particular query process. As the data in a column changes, index and column statistics can become out-of-date. This can hinder the query optimizer's ability to make optimal decisions on how to process a query. As a result query performance will deteriorate. When statistics become out-of-date they can be updated automatically if the **Auto Update Statistics** check box on the table property dialog box is checked.

Incorrect Answers:

A: It would require some administrative effort to set up job that update the statistics. Just selecting the **Auto Update Statistics check** is a better solution.

C: The Database Maintenance Plan Wizard cannot be used to recreate the statistics, which are the main problem at the moment.

Note: The Database Maintenance Plan Wizard can be used to set up the core maintenance tasks necessary to ensure that a database performs well, is regularly backed up, and is checked for inconsistencies. The Database Maintenance Plan Wizard creates a SQL Server 2000 job that performs these maintenance tasks automatically at scheduled intervals and can be used to reorganize the data on the data and index pages by rebuilding indexes with a new fill factor; compress data files by removing empty database pages; update index statistics to ensure the query optimizer has up-to-date information about the distribution of data values in the tables; perform internal consistency checks of the data and data pages within the database to ensure that a system or software problem has not damaged data; backup the database and transaction log files; and to set up log shipping.


D: Looking at the exhibit we see that this is a statistics problem, not a index problem.

Note: A nonclustered index is similar to an index in a textbook. With this type of index, the data is stored in one place and the index in another. These indexes, however, point to the storage location of the data. Nonclustered indexes should be used on columns that contain a large number of distinct values; for queries that do not return large result sets; and for columns that are frequently involved in search conditions of a query that return exact matches. A clustered index, on the other hand, determines the physical order of data in a table. The PRIMARY KEY index automatically creates a clustered index as it is used to identify a row unless a nonclustered index is specified.


E: When set the Auto Create Statistics check box is checked, statistics are automatically created on columns that are used in a predicate. Adding statistics improves query performance because the SQL Server query optimizer can better determine how to evaluate a query. If the statistics are not used, SQL Server automatically deletes them. When this check box is cleared, statistics are not automatically created.

Question No: 5

You are the administrator of SQL Server 2000 computer. The server contains a database named Inventory. The database has a set of tables that track the location of parts. The tables are configured as shown in the Database Schema exhibit.



You examine the indexes in the PartsLocations table. The indexes are configured as shown in the Indexes exhibit.



	index name	index description	index keys
1	IX_PartsLocations_LocationID_FK	nonclustered located on PRIMARY	LocationID
2	IX_PartsLocations_PartID_FK	nonclustered located on PRIMARY	PartID
3	PK_PartsLocations	clustered, unique, pri...	PartID, LocationID

You also examine the contiguity of the PK_partsLocations index and receive the following results:

DBCC SHOWCONTIG scanning 'PartsLocations' table...

Table: 'PartsLocations' (1); index ID: 1, database ID: 6

TABLE level scan performed.

- Pages Scanned.....: 3000
- Extents Scanned.....: 750
- Extent Switches.....: 502
- Avg. Pages per Extent.....: 4.0
- Scan Density [Best Count:Actual Count].....: 33.33% [1:3]
- Logical Scan Fragmentation.....: 33.33%
- Extent Scan Fragmentation.....: 33.33%
- Avg. Bytes Free per Page.....: 1166.0
- Avg. Page Density (full).....: 85.77%

You need to improve the performance of the queries that join this table to the Parts table and the Locations table. What should you do?

- A. Re-create the PK_PartsLocations index as a clustered index.
- B. Re-create the IX_PartsLocations_PartsID as a clustered index.
- C. Re-create the IX_PartsLocations_LocationId as a clustered index.
- D. Re-create the PK_PartsLocations index as a nonclustered index.

Answer: A.

Explanation: The output of the DBCC SHOWCONTIG shows that the clustered index of the table is fragmented. Both Logical Scan Fragmentation and Extent Scan Fragmentation are above 10%. The clustered index PK_PartsLocations should be rebuilt.

Note: In general, you need to do very little in terms of index maintenance. However, indexes can become fragmented. You can use the DBCC SHOWCONTIG command to report on the fragmentation of an index.

You can use either the Extent Scan Fragmentation value or the Logical Scan Fragmentation value to determine the general level of fragmentation in a table. Logical Scan Fragmentation and, to a lesser extent, Extent Scan Fragmentation values give the best indication of a table's fragmentation level. Both these values should be as close to zero as possible (although a value from 0% through 10% may be acceptable).

Incorrect Answers:


B, C: There already exists a clustered index, the PK_PartsLocations index. A table can only contain one clustered index.

D: A clustered index is faster than a non-clustered index. We should rebuild the clustered index, not

create a nonclustered index instead.

Question No: 6

You are the administrator of a SQL Server 2000 computer. The server contains a database named Inventory. In this database, the Parts table has a primary key that is used to identify each part stored in the company's warehouse. Each part has a unique UPC code that your company's accounting department uses to identify it. The database is configured as shown in the exhibit.



You want to maintain the referential integrity between the Parts table and the OrderDetails table.

You want to minimize the amount of physical I/O that is used within the database.

Which two Transact-SQL statements should you execute? (Each correct answer represents part of the solution. Choose two.)

- A. CREATE UNIQUE INDEX IX_UPC On Parts(UPC)
- B. CREATE UNIQUE INDEX IX_UPC On OrderDetails(UPC)
- C. CREATE TRIGGER UPCRI On OrderDetails
FOR INSERT, UPDATE As
If Not Exists (Select UPC From Parts
Where Parts.UPC = inserted.UPC) BEGIN
ROLLBACK TRAN
END
- D. CREATE TRIGGER UPCRI On Parts
FOR INSERT, UPDATE As
If Not Exists (Select UPC From Parts
Where OrderDetails.UPC = inserted.UPC) BEGIN
ROLLBACK TRAN
END
- E. ALTER TABLE dbo.OrderDetails ADD CONSTRAINT
FK_OrderDetails_Parts FOREIGN KEY(UPC)
REFERENCES dbo.Parts(UPC)
- F. ALTER TABLE dbo.Parts ADD CONSTRAINT
FK_Parts_OrderDetails FOREIGN KEY (UPC)
REFERENCES dbo.Parts(UPC)

Answer: A, E.

Explanation:

A: The unique index on the UPC column must be created in the Parts table since every part has an unique UPC code.

E: One Part row can be contained in many OrderDetails, but each OrderDetail row only corresponds to one specific row in the Part table. There is a one-to-many relation between the Part table and the OrderDetails table. Therefore the foreign key constraint must be defined on the OrderDetails table referencing the UPC column of the Parts table.

Incorrect Answers:

B: Each part has a unique UPC code, but each part can be included in many rows in the OrderDetails table since many different orders can contain the same part. We cannot use a unique index on the UPC column in the OrderDetails table.

C: Referential integrity is best implemented with foreign key constraints not by triggers.

D: Referential integrity is best implemented with foreign key constraints not by triggers.

F: The foreign key constraint must be from the OrderDetails table referencing the UPC column in the Parts table.

If we have a foreign key constraint from the Parts table referencing the UPC column in the OrderDetails table we would have that each Part only can be part of one OrderDetails, but that one OrderDetails row could correspond to several rows. But it should be the opposite.

Question No: 7

You are the administrator of a SQL server 2000 computer. The server contains a database named SalesTransactions. The database handles sales transactions for a network of retail outlets. The primary table is shown below:

	Column Name	Data Type	Length	Allow Nulls
?	SalesTransID	int	4	
	CustomerName	varchar	50	✓
	CustomerAddress	varchar	50	✓
	CustomerCity	varchar	50	✓
	CustomerState	char	10	✓
	CustomerZip	char	10	✓
	DateSold	datetime	8	✓
	PaymentMethod	char	10	✓
	Item1SKU	char	10	✓
	Item1Qty	int	4	✓
	Item2SKU	char	10	✓
	Item2Qty	int	4	✓
	Item2Price	money	8	✓
	Item3SKU	char	10	✓
	Item3Qty	int	4	✓
	Item3Price	money	8	✓
	SalesTax	money	8	✓

Users report that when they add new sales, their inserts often time out. You examine the indexes on the table. The indexes are shown below:

Index_name index_description

IX_CustomerName nonclustered located on PRIMARY

IX_CustomerState nonclustered located on PRIMARY

IX_CustomerZip nonclustered located on PRIMARY
IX_PaymentMethod nonclustered located on PRIMARY
PK_SalesTransactions_1 nonclustered, unique, primary key located on PRIMARY

You want to decrease the time required to store purchases made at the retail outlets. What are two possible ways to achieve this goal?

- A. Normalize the data structure of the table.
- B. Drop the nonclustered indexes of the table.
- C. Drop the primary key.
- D. Add nonclustered indexes to the item1SKU, item2SKU, and item3SKU fields in the table.
- E. Re-create the PK_Sales Transactions index as a clustered index.

Answer: A, B.

Explanation:

A: This primary table of the database is highly denormalized. The columns Item1SKU, Item1Qty, Item2SKU, Item2Qty, Item2SKU, Item2Qty should be moved to a separate item table. The denormalization results in a lot of NULL values and performance is decreased.

B: Every time a new row is inserted into a table all the indexes of the table has to be updated. This slows the insertion process. Dropping the non-clustered indexes of the sales table would improve insertion operations of the table. The time required to store purchases would decrease.

Incorrect Answers:

C: Dropping the primary key might improve insertion performance somewhat. Dropping all nonclustered indexes is more beneficial though.

D: Adding more non-clustered indexes would degrade the performance of insertions of new rows into the table.

E: Every table should have one clustered index. A clustered index is the most efficient index since the rows are physically stored in order and no pointers are used. A clustered index is beneficial when the database is queried, for example in join statements and in WHERE clauses. A clustered index would not improve performance for adding new rows into the sales table which is the requirement of this scenario.

Question No: 8

You are the administrator of a SQL Server 2000 database named Articles. A large city newspaper uses this database to store its archived newspaper articles. Journalists query the database as part of their research. You have created full-text indexes so that users can query the articles quickly. A hardware failure forces you to restore the entire database from backup. After you restore the database, the journalists report that they can no longer run queries without receiving errors. You want to ensure that the journalists can query the database. What should you do?

- A. Create a new full-text catalog for the Articles database
- B. Rebuild the full-text catalog for the Articles database
- C. Create a new full-text index on the table that holds the text for the articles
- D. Repopulate the full-text index on the table that holds the text for the articles

Answer: D.

Explanation: Backing up a database does not back up full-text index data in full-text catalogs. The full-text catalog files are not recovered during a Microsoft SQL Server recovery. However, if full-text indexes have been defined for tables, the metadata for the full-text index definitions are stored in the system tables in the

database containing the full-text indexes. After a database backup is restored, the full-text index catalogs can be re-created and repopulated.

Note 1:

Setting up a full-text search is a three-step procedure:

1. Make sure that the full-text indexing feature has been installed with your SQL Server.

The full-text index is installed as a service in the system. It would not have to be reinstalled after restoring a database.

2. Create a catalog in the operating system to store the full-text indexes.

A catalog can contain the indexes from one or more tables, and it is treated as a unit for maintenance purposes. In general, you shouldn't need more than one catalog per database.

The catalog is saved outside the database on separate files

3. Before you can implement a full-text search in a given database, you must ensure that the full-text search indexes are populated regularly.

In this scenario we must recreate the index by repopulating the full-text catalog.

Note 2: A SQL Server 2000 full-text index provides support for sophisticated word searches in character string data. The full-text index stores information about significant words and their location within a given column. This information is used to quickly complete full-text queries that search for rows with particular words or combinations of words. The full-text indexes are contained in full-text catalogs. Full-text catalogs and indexes are not stored in the database to which they belong. The catalogs and indexes are managed separately by the Microsoft Search service. Backing up a database does not back up full-text index data in full-text catalogs. However, if full-text indexes have been defined for tables, the meta data is backed up when a database backup is created.

Reference: INF: How to Move, Copy and Backup SQL 7.0 Full-Text Catalog Folders and Files (Q240867)

Incorrect Answers:


A: The catalog already exists. The catalog is stored outside SQL server and would not be affected by the failure and restoration of the database.

B: Rebuilding the full text catalog will not enable the Journalist to run the full-text queries. The catalogs would have to be repopulated as well.

C: The index should be repopulated since it already exists from the backup. There is no need to recreate it.

Question No: 9

You are the administrator of a SQL Server 2000 computer. The server contains a database named Sales. The database is used to analyze trends. The database is configured as shown in the Database Schema exhibit.



Analysts examine customer buying patterns and frequently use the following query.

```
SELECT Customers.CustomerID,
Invoices.InvoiceID,
Invoices.InvoiceDate,
OrderDetail.Quantity, Products.SKU
FROM Customers
INNER JOIN Invoices ON Customers.CustomerID = Invoices.CustomerID
INNER JOIN OrderDetail ON Invoices.InvoiceID =
OrderDetail.InvoiceID
INNER JOIN Products ON OrderDetail.ProductID =
Products.ProductID
WHERE CustomerID = ?
```

You analyze the indexes in the Invoices table. The indexes are configured as shown in the Indexes Exhibit.

	index name	index description	index keys
1	IX_Customers_FK	nonclustered located on PRIMARY	CustomerID
2	IX_SalesPerson_FK	nonclustered located on PRIMARY	SalesPersonID
3	PK_Invoices	clustered, unique, primary key located on PRIMARY	InvoiceID

You want to optimize the database schema so that it responds quickly to the analysts' query. Which two actions should you take? (Each correct answer represents part of the solution. Choose two.)

- A. Re-create the PK_Invoices index as a nonclustered index
- B. Re-create the IX_Customers_FK index as a clustered index
- C. Re-create the IX_SalesPerson_FK index as a clustered index
- D. Create a composite index on CustomerID fields and SalesPersonID fields
- E. Create a composite index on InvoiceID fields and CustomerID fields

Answer: A, B

Explanation:

Indexing should be used on columns that have a foreign key constraint (or are used in joins) or columns that are used in a WHERE clause. The CustomerID column is used in a JOIN statement in the SQL batch. It is also used in the WHERE clause. We should create an index on this column.

A clustered index is faster than a non-clustered index since the rows are stored directly in physical order with no extra pointers at the leaf level. We should therefore create a clustered index on the CustomerID column.

A: In order to put a clustered index on the CustomerID column we must first recreate the PK_Invoices index as a non-clustered index. There can only be one clustered index in a table.

B: The CustomerID column is used both in a JOIN clause and the WHERE clause in the SQL batch. The best performance would therefore be gained by using a clustered index on that column.

Incorrect answers:


C: The SalesPersonID column isn't used in any joins. There would be no benefit in recreating a index on the SalesPersonID column.

D: The SalesPersonID column isn't used in any joins. There would be no benefit in recreating a index on the SalesPersonID column.

E: To create a composite a composite index on the InvoiceID fields and CustomerID fields the index PK_Invoices index must be dropped first.

Question No: 10

You are the administrator of a SQL Server 2000 computer. The server contains a database named Sales. The database has a group of tables that analysts use to examine sales trends. The database options are set to their default values, and the tables are configured as shown in the exhibit.



Analysts who use the database report that query performance has become slower. You analyze the clustered primary key on the Invoices table and receive the following results.

DBCC SHOWCONTIG scanning 'Invoices' table...

Table: 'Invoices' (21575115); index ID: 1, database ID: 6

TABLE level scan performed.

Pages Scanned.....: 200

- Extents Scanned.....: 50

- Extent Switches.....: 40
- Avg. Pages per Extent.....: 4.0
- Scan Density [Best Count:Actual Count].....: 60.00% [3:5]
- Logical Scan Fragmentation.....: 0.00%
- Extent Scan Fragmentation.....: 40.00%
- Avg. Bytes Free per Page.....: 146.5
- Avg. Page Density (full).....: 98.19%

DBCC execution completed. If DBCC printed error messages, contact your system administrator.

You want to improve performance of queries that join tables to the Invoices table. What are three possible Transact-SQL statements you can execute to achieve this goal? (Each correct answer represents a complete solution. Choose three.)

- A. DBCC UPDATEUSAGE ('Sales','Invoices','PK_Invoices')
- B. CREATE UNIQUE CLUSTERED INDEX PK_Invoices On Invoices(InvoiceID) WITH DROP_EXISTING
- C. DBCC INDEXDEFRAG ('Sales','Invoices','PK_Invoices')
- D. DBCC DBREINDEX ('Sales.dbo.Invoices', 'PK_Invoices')
- E. UPDATE STATISTICS 'Invoices'
- F. DBCC CHECKALLOC (Sales, REPAIR_FAST)

Answer: B, C, D.

Explanation: We want to improve performance on queries that join tables to the Invoices table. One way to keep good performance is keep the indexes defragmented.

The problem could be solved in different ways (note that these are different approaches not one solution consisting of several steps):

- create a unique clustered index PK_invoices on the InvoiceID column while dropping the existing index (B).
- defrag the index with the DBCC INDEXDEFRAG command. The DBCC INDEXDEFRAG command defragments clustered and secondary indexes of the specified table or view (C).
- rebuild the indexes on the tables with the DBCC DBREINDEX command (D).

Note: In general, you need to do very little in terms of index maintenance. However, indexes can become fragmented. You can use the DBCC SHOWCONTIG command to report on the fragmentation of an index. You can use either the Extent Scan Fragmentation value or the Logical Scan Fragmentation value to determine the general level of fragmentation in a table. Logical Scan Fragmentation and, to a lesser extent, Extent Scan Fragmentation values give the best indication of a table's fragmentation level. Both these values should be as close to zero as possible (although a value from 0% through 10% may be acceptable). In our scenario we see that Logical Scan Fragmentation is 0% and acceptable, but the Extent Scan Fragmentation is 40.00% and is a clear that the table is fragmented. . To solve the fragmentation problem the index should be rebuilt. The index shown is the clustered index. Here we have to assume the clustered index is created on the primary key column InvoiceID.

Incorrect Answers:

A: The DBCC UPDATEUSAGE command reports and corrects inaccuracies in the sysindexes table, which may result in incorrect space usage reports by the sp_spaceused system stored procedure. It cannot be used to defragment indexes in a user table.


E: Updating statistics could improve performance in a query. It would give the SQL Server more

accurate information on what optimizations that could be made. In this scenario, however, the DBCC SHOWCONTIG command shows that the index is fragmented and this is the main performance problem.

F: The DBCC CHECKALLOC command checks the consistency of disk space allocation structures for a specified database. It cannot be used to defragment indexes in a user table.

Question No: 11

You are the administrator of a SQL Server 2000 computer. The server contains a database named Sales. The Sales database is configured as shown in the Database Schema exhibit.



The Sales database stores point-of-sale transaction data. The data is later imported into another database for analysis. Recently, users have reported slow response times whenever they enter point-of-sale information.

You examine the indexes on the Invoices table that are configured as shown in the Index Description exhibit.

sp_helpindex 'Invoices'

	index_name	index_description	index_keys
1	IX Customers FK	nonclustered located on PRIMARY	CustomerID
2	IX SalesPerson FK	nonclustered located on PRIMARY	SalesPersonID
3	PK Invoices	clustered, unique, primary key located on PRIMARY	InvoiceID

You want to accelerate database response time when users insert new point-of-sale information. What should you do?

- A. Drop IX_Customers_FK and IX_SalesPerson_FK.
- B. Change IX_Customers_FK to a clustered index.

- C. Change IX_SalesPerson_FK to a clustered index.
- D. Change PK_Invoices to a nonclustered index.

Answer: A.

Explanation: Indexes are used to speed up queries. It is useful to have indexes on columns used in joins and columns that are used in WHERE clauses. In this scenario, however, the Sales database isn't used primarily for queries, it stores point-of-sale transaction data that later is imported into another database for analysis. We want to improve performance when entering data into the Invoices table. Indexes must be updated when new rows are inserted and by dropping unnecessary indexes we could improve the response time when entering new rows.

Incorrect Answers:

- B:** There is already a clustered index. There can only be one clustered index in a table.
- C:** There is already a clustered index. There can only be one clustered index in a table.
- D:** Changing the index of the primary key to a non-clustered index would not improve performance for insertions.

Question No: 12

You are the administrator of a SQL Server 2000 database. Using the default options configures the database. You drop indexes to facilitate a large import of data, and rebuild them after the import is complete.

Users report that response times from a stored procedure have become slower. In the stored procedure, you examine the following code.

```
CREATE PROCEDURE GetCustomer
( @CustomerID Int )
AS
SELECT FirstName, LastName, Address1, Address2, City,
State, Zip, MAX(PurchaseDate) AS LastPurchase
FROM Customers C
JOIN Addresses A ON A.AddressID = C.AddressID
JOIN Purchases P ON P.CustomerID = C.CustomerID
WHERE C.CustomerID = @CustomerID
```

GROUP BY FirstName, LastName, Address1, Address2, City, State, Zip

You want to improve the performance of the procedure by using the least amount of administrative effort. What should you do?

- A. Recompile the stored procedure.
- B. Update the database statistics.
- C. Check the consistency of the database and repair errors.
- D. Create a view and modify the stored procedure to use the view.

Answer: A.

Explanation: By recompiling the stored procedure SQL Server will create a new execution plan for the statements in the procedure. SQL Server will detect that statistics are not up to date and it will automatically recreate the statistics during this process. The execution plan will be optimized for the current distribution of data in the tables.

Note: Stored procedures can be used to improve database query performance. If these SQL statements and conditional logic are written into a stored procedure, they become part of a single execution plan that is cached on the server. When the database is changed by the addition of new indexes or when data in indexed columns are changed, the original query plans and stored procedures used to access the tables should be reoptimized

by recompiling them. This optimization happens automatically the first time a stored procedure is run after SQL Server 2000 is restarted. It also occurs if an underlying table used by the stored procedure changes. But if a new index is added from which the stored procedure might benefit, optimization does not automatically occur until the next time the stored procedure is run after SQL Server is restarted. The stored procedure can be manually recompiled by using the sp_recompile system stored procedure. This forces the stored procedure to be recompiled the next time it is run.

Incorrect Answers:

B: SQL Server detects whether statistics are not up-to-date, and by default it automatically updates them during query optimization. By recompiling the query the statistics will automatically be updated.


Note: SQL Server keeps statistics about the distribution of the key values in each index and uses these statistics to determine which index(es) to use in query processing. Query optimization depends on the accuracy of the distribution steps. If there is significant change in the key values in the index, or if a large amount of data in an indexed column has been added, changed, or removed, or the table has been truncated and then repopulated, UPDATE STATISTICS should be rerun.

C: Consistency of a database is maintained through the use of constraints and triggers. These mechanisms are used to ensure that the database remains consistent by controlling the type of data that can be inserted into a column, through referential integrity checks, and through cascading referential constraints.

D: A view is a virtual table that can be used to restrict users access to the base table, or to join various base tables so that they appear as a single table. However, creating a view that the stored procedure must access would require the creation of indexes on the view and will have to be updated when the base tables are updates. This would require greater administrative effort.

Question No: 13

You are the administrator of a SQL Server 2000 computer. The server contains a database named Sales. The database is configured as shown in the exhibit.



Each evening, you import transactions from your company's retail locations to the sales database. These transactions are imported as batch operations. Analysts use the imported information to decide which products to order for the stores and to predict market trends. The analysts frequently use the following query:

```
SELECT Customers.CustomerName, Invoices.InvoiceID,  
Invoices.InvoiceDate,  
SalesPersons.FirstName+ ' ' + SalesPerson.LastName AS SalesPerson,  
OrderDetail.Quantity,  
Products.Description, Products.RetailPrice,  
OrderDetail.SalesDiscount  
FROM Invoices  
INNER JOIN Customers ON Invoices.Customer ID = Customer.CustomerID  
INNER JOIN OrderDetail ON Invoices.Invoice ID =  
OrderDetail.InvoiceID  
INNER JOIN Products ON OrderDetail.ProductID = Products.ProductID  
INNER JOIN SalesPersons ON Invoices.SalesPersonID =  
SalesPersons.SalesPersonID
```

Recently, the analysts have reported that query response time that has slowed. You want to accelerate query response time, but you do not want to change the way the analysts query the database. What should you do?

- A.** Create indexes on the foreign key fields.
- B.** Create a view that returns the same results as the query
- C.** Create a stored procedure that returns the same results as the query
- D.** Create a summery table that allows a simpler query, and then use triggers to maintain the data in the table.

Answer: D.

Explanation: The query consists of four joins and it is only used for analyzing purposes. New data is imported as a batch job during the nights. This suggests that the database design should be denormalized. We should create a single table from the five tables used in the query. This would allow a simpler query that only uses one table. No joins would be needed. This would improve performance of the query. Updates of the tables would be slower, but that is not important since the table only gets updated once a day in a batch.

Incorrect Answers:


- A:** Creating indexes could improve performance somewhat. But since this database is used for analysis only it would be better to denormalize the database for speed.
- B:** A view would not, by itself, improve the performance of a query.
- C:** A stored procedure could, in a very minor way, improve performance since the execution plans are saved at the server. The stored procedure doesn't have to be recompiled. This also saves some small amount of time. But since this database is used for analysis only it would be better to denormalize the database for speed.

SUBSECTION A: Optimize data storage. (1 question)

Question No: 1

You are the administrator of a SQL server 2000 computer that contains a database named Acct. The

database contains 1.5 GB of data. The server has one 9-GB hard disk that is configured as shown in the Exhibit.



You need to import data into the database without adversely affecting database performance. The data will require an additional 2GB of storage space.

What should you do?

- A. Add another data file on drive E, and then add the file to the PRIMARY filegroup.
- B. Move the transaction log file to drive E, and set the file growth of Acc_Data.mdf by selecting the **Unrestricted file growth** option.
- C. Rebuild all clustered indexes so that they have a fill factor of 100.
- D. Compress drive D.

Answer: A.

Explanation: We must take precautions so that the import process of 2 GB data will run. We add a new file to drive E, which has 3 GB free disk space. By placing the new file in the primary file group we ensure that new data added to the database will be added to new data file. The primary file group is the default file group and new data is automatically added to it.

Incorrect Answers:

B: Moving the transaction log from Drive D to Drive would free up 800MB. This would leave 1.3GB of free disk space on drive D. But the import process will require another 2GB so we will not have enough free disk space.

C: Rebuilding the clustered indexes with a fill factor of 100 makes that the index pages full. This would


degrade performance during the import process due to page splits.

D: Compressing a partition which is used by SQL Server is not a good idea. It would, if at all possible, degrade performance.


SUBSECTION B: Optimize files and filegroups. (4 questions)

Question No: 1

You are the administrator of a SQL Server 2000 computer. The server contains database named Sales. Users report that they cannot add new data to the database. You examine the properties of the database. The database is configured as shown in the Sales Properties exhibit.



You examine drive E. The hard disk is configured as shown in the Local Disk Properties exhibit.



You want the users to be able to add data, and you want to minimize administrative overhead. What should you do?

- A. Increase the maximum file size of Sales_Data to 1,500MB.**
- B. Use the DBCC SHRINKDATABASE statement.**
- C. Set automatic file growth to 10 percent.**
- D. Create another data file named Sales_Data2 in a new SECONDARY filegroup.**

Answer: A.

Explanation: By examining the first exhibit we see that the database has reached its maximum size of 1001 MB. The data file is located on drive E and the second exhibit shows us that there is enough free space on that disk to increase the maximum size of the data file to 1,500 MB.

Incorrect Answers:

- B:** DBCC SHRINKDATABASE could possibly shrink the database and new records could be added for a while, but it would just be a temporary solution.
- C:** The database is already set to automatic file growth, but it has reached the maximum size.
- D:** The primary file has reached its maximum size. A secondary filegroup would help if it was set to the default file group. New data would still be added to the primary filegroup which is at its maximum size.

Question No: 2

You are the administrator of a SQL Server 2000 computer. The server contains a database that is heavily indexed and that company users query extensively. The database has grown and query

**response time has slowed. The database is stored in a single data file.
You want to accelerate query response time. What should you do?**

- A.** On a new hard disk, create a new filegroup.
Drop the existing nonclustered indexes, and then re-create them on the new filegroup.
- B.** On a new hard disk, add a new file to the PRIMARY filegroup.
Drop the existing nonclustered indexes, and then re-create them on the PRIMARY filegroup.
- C.** On the existing hard disk, create a new filegroup.
Drop the existing nonclustered indexes, and then re-create them on the new filegroup.
- D.** On the existing hard disk, add a new file to the primary filegroup.
Drop the existing nonclustered indexes, and then re-create them on the PRIMARY filegroup.

Answer: A.

Explanation: With SQL Server 2000 it is possible to create tables or indexes on specific filegroups. This allows us to control where the database's tables and indexes are physically located as filegroups can be placed on different hard drives. Placing heavily accessed tables in one filegroup and the table's nonclustered indexes in another filegroup on different physical disk can improve database performance because it will allow separate threads to access the tables and indexes. However, a table and its clustered index cannot be separated into different filegroups as the clustered index determines the physical order of the data in the table.

Incorrect Answers:


B: SQL Server 2000 can operate effectively without filegroups. In this case, all files are included in the primary filegroup and SQL Server 2000 can allocate data anywhere in the database. Furthermore, filegroups are not the only method that can be used to distribute I/O across multiple drives. However, the nonclustered index is added to the filegroup rather than the new file in this solution. This is thus not the best solution.

C: This solution does not take advantage of the possible performance gains that can be realized by placing data tables and indexes on separate physical disks.


D: This solution does not take advantage of the possible performance gains that can be realized by placing data tables and indexes on separate physical disks. This would require that a new file group be created.

Question No: 3

You are the administrator of a SQL Server 2000 computer. The server contains a database named MedicalRecords. You have recently added a hard disk to the server and configured the database. The data files are configured as shown in the Database Properties exhibit.



The database is configured as shown in the Database Schema exhibit.



Each primary key is clustered. All foreign keys contain nonclustered indexes. You want to use the space on the new hard disk to make the database respond more quickly to queries. What should you do?

A. Re-create the **Notes** table on the SECONDARY filegroup.

- B. Re-create the nonclustered indexes on the SECONDARY filegroup.
- C. Re-create the **Patients** and **Doctors** tables and their indexes on the SECONDARY filegroup.
- D. Re-create the **Cases** and **Notes** tables and their indexes on the SECONDARY filegroup.

Answer: B.


Explanation: With SQL Server 2000 it is possible to create tables or indexes on specific filegroups. This allows us to control where the database's tables and indexes are physically located as filegroups can be placed on different hard drives. Placing heavily accessed tables in one filegroup and the table's nonclustered indexes in another filegroup on different physical disk can improve database performance because it will allow separate threads to access the tables and indexes. However, a table and its clustered index cannot be separated into different filegroups as the clustered index determines the physical order of the data in the table.

Incorrect Answers:

A, C, D: With SQL Server 2000, placing the tables and their nonclustered indexes in different filegroups that are located on different physical drives provides better performance than placing tables in different filegroups and on different physical drives.

Question No: 4

You are the administrator of a SQL Server 2000 computer. The server contains a database named Acct. You purge the Acct database of old records and perform a full backup. The database now uses 4 GB of space. The database files are configured as shown in the Acct Properties exhibit. The server has two hard disks that are configured as shown in the Server Configuration exhibit.



The 3 GB transaction log file for the Acct database is stored on drive C. You need to make room for a new database that has a 3 GB data file and a 1 GB transaction log file. You want to optimize database performance on both databases. You also want to minimize administrative overhead.

What should you do?

A. Shrink the empty data file on drive **E**.

Place the new data file and the new log file on drive **E**.

B. Shrink the empty data file on drive **E**.

Backup and shrink the log file on drive **C** to 2 **GB**.

Place the new data file on drive **E** and the new log file on drive **C**.

C. Shrink the data file on drive **D** to 1 **GB**, and then shrink the data file on drive **E** to 2 **GB**.

Place the new data file on drive **D** and the new log file on drive **E**.

D. Shrink and delete the data file on drive **C**, and then shrink the database files on drives **D** and **E** so that they are both 2 **GB**.

Place the new data file and the new log file on drive **C**.

Answer: B.

Explanation: We must add a new 3GB database with a 1GB transaction log file. Currently we have database that is using 4GB for data and 3GB for the log file. We should place the log file and the data file on two different physical disks out of performance reasons.

The command DBCC SHRINKFILE with the EMPTYFILE option can be used to shrink a datafile to empty. The data will be migrated to the other files in the filegroup. We use this method to free 4.5 GB on Drive **E**.

Then we backup the log file on **C** and reduce it to 2**GB**. Then we put the database data file on drive (3GB) and put the log file of the new database on drive **C** (1GB).

Incorrect answers:

A: It's better to put the data file and the log file of the new database on two different physical disks out of performance reasons.

C: It's better to put the data file and the log file of the new database on two different physical disks out of performance reasons.

D: It's better to put the data file and the log file of the new database on two different physical disks out of performance reasons.

SECTION 2: Manage database fragmentation. (0 questions)

SECTION 3: Modify the database schema.

SUBSECTION A: Perform disaster recovery operations. (7 questions)

Question No: 1

You are the administrator of a SQL Server 2000 computer. The server contains a database that has the torn page detection database option enabled. Backups of the database are made daily.

The server loses power for one minute. When power is restored, torn pages are detected. You notice in SQL Server Enterprise Manager that the database is marked suspect.

You need to correct the problem. What should you do?

A. Execute the DBCC CHECKDB statement, and then specify the PHYSICAL_ONLY option.

B. Execute the DBCC CHECKDB statement, and then specify the REPAIR_REBUILD option.

C. Execute the **sp_resetstatus** stored procedure.

D. Restore the suspect database from backups.

Answer: D.

Explanation: In SQL Server 2000, the TORN_PAGE_DETECTION option is a database recovery option that allows SQL Server to detect incomplete I/O operations caused by power failures or other system outages. When this option is set to ON, which it is by default, it causes a bit to be reversed for each 512-byte sector in an 8KB database page when the page is written to disk. If a bit is in the wrong state when the page is later read by SQL Server, the page was written incorrectly and a torn page is detected. Using batterybacked disk caches can ensure that data is successfully written to disk or not written at all. If the torn page is detected the database is marked suspect. When this occurs, the database backup should be restored, and any transaction log backups applied, because the database is physically inconsistent.

Incorrect Answers:

A: The DBCC CHECKDB statement checks the allocation and structural integrity of all the objects in the specified database. This statement can specify the PHYSICAL_ONLY option, which limits the checking to the integrity of the physical structure of the page and record headers, and to the consistency between the pages' object ID and index ID and the allocation structures. This check also detects torn pages and common hardware failures that can compromise a user's data. However, the PHYSICAL_ONLY option is not allowed with any of the DBCC CHECKDB statement's repair options.

B: The DBCC CHECKDB statement checks the allocation and structural integrity of all the objects in the specified database. This statement can specify repair options. It can specify the REBUILD_FAST option, which performs minor, nontime-consuming repair actions such as repairing extra keys in nonclustered indexes and can be done quickly and without risk of data loss; and it can specify the REPAIR_REBUILD option, which performs all repairs that can be done by REPAIR_FAST and as well as includes time-consuming repairs such as index rebuilding. These repairs can also be done without risk of data loss.

C: The sp_resetstatus stored procedure is not a recovery option. It turns off the suspect flag on a database by updating the mode and status columns of the named database in sysdatabases. Because this procedure modifies the system tables, the system administrator must enable updates to the system tables before creating this procedure and the SQL Server 2000 must shut down and restate immediately after executing this procedure.

Question No: 2

You are the administrator of a SQL Server 2000 computer. The server contains a database named Sales. You perform full database backups every two days. You also run regular database consistency checks on the server. The most recent check of the Sales database returns the following message. "CHECKDB found 0 allocation errors and 9 consistency errors in the table 'Orders' (object ID 214575782)."

You want to correct the data integrity errors while minimizing the amount of data lost. What should you do?

A. Disconnect users from the **Sales** database.

Enable the **single user database** option.

Execute the DBCC CHECKTABLE statement for the **Orders** table, and specify the REPAIR_REBUILD option.

B. Disconnect users from the **Sales** database.

Enable the **DBO use only database** option.

Execute the DBCC CHECKTABLE statement for the **Orders** table, and specify the REPAIR_REBUILD option.

C. Disconnect users from the **Sales** database.

Execute the RESTORE DATABASE statement for the **Sales** database

- D. Execute the DBCC CLEAN TABLE statement for the **Orders** table.
- E. Execute the **sp_table_validation** stored procedure for the **Orders** table.

Answer: A.

Explanation: We should repair the database with the DBCC CHECKTABLE REPAIR_REBUILD command. We should run this repair statement when the database is configured to single user.

Note: DBCC CHECKTABLE checks the integrity of the data, index, text, ntext, and image pages for the specified table or indexed view. DBCC CHECKTABLE can take a specified repair option to repair the found errors but must be in single-user mode to use a repair option. It can specify the REBUILD_FAST option, which performs minor, non time-consuming repair actions such as repairing extra keys in nonclustered indexes and can be done quickly and without risk of data loss; and it can also specify the REPAIR_REBUILD option, which performs all repairs that can be done by REPAIR_FAST and as well as time-consuming repairs such as index rebuilding. These repairs can also be done without risk of data loss.

Incorrect Answers:

B: The database option **DBO use only**, would only allow the database owner running the database. This might be too restrictive.

C: We are not told when how often consistency checks are performed but assuming that consistency occurs more frequently than the database backups then using RESTORE DATABASE to restore a database from the last full backup would result in the loss of data entered into the database since the last full database backup was performed. This would result in data loss.

D: DBCC CLEAN TABLE is used to reclaim space after a variable length column or a text column is dropped using the ALTER TABLE DROP COLUMN statement.

E: The sp_table_validation stored procedure returns rowcount or checksum information on a table or indexed view, or compares the provided rowcount or checksum information with the specified table or indexed view. This stored procedure is used in replication and checks that the structure of the table being replicated between two tables are identical, i.e., that the tables have the same columns existing in the same order, same data types and lengths, and same NULL/NOT NULL conditions.

Question No: 3

You are the administrator of several SQL Server 2000 computers. You want to retrieve information from an archived inventory database. You have a full tape backup of the database. The backup's header information shows that the backup uses the SQL_Latin1_General_CR437_BIN collation. However, the existing SQL server computers in your office are configured to use the SQL_Latin1_General_CP1_CI_AS collation.

You do not want to join tables in the inventory database with tables in other databases. You need to restore the inventory database to a SQL Server 2000 computer by using the least amount of administrative effort.

What should you do?

- A. Use the rebuildm utility to rebuild the system database on an existing SQL server computer. Configure all the databases on that server to use the SQL_Latin1_General_CR437_BIN collation. Restore the inventory database to the server.
- B. Restore the inventory database to an existing SQL server computer. Accept the SQL_Latin1_General_CR437_BIN collation for that database.
- C. Install a new named instance of SQL Server 2000 on an existing SQL Server 2000 computer.

Configure the named instance to use the SQL_Latin1_General_CR437_BIN collation.

Restore the inventory database to the named instance.

D. Install SQL Server 2000 on a new computer.

Configure the new server to use the SQL_Latin1_General_CR437_BIN collation.

Restore the inventory database to the new server.

Answer: B.

Explanation: Collations control the physical storage of character strings in SQL Server 2000 by specifying the bit patterns that represent each character and the rules by which characters are sorted and compared. Whereas earlier versions of SQL Server supported only one collation for each instance of SQL Server, SQL Server 2000 allows objects that have different collations to be stored in a single database. We can thus restore the database to an existing SQL Server 2000 computer without having to change the collation of that database.

Incorrect Answers:

A: The rebuildm utility rebuilds the master database and can be used to change the collation settings of an instance of SQL Server 2000, or to fix a corrupted master database. Rebuilding the master database involves removing all database objects and data, re-creating the database objects and reloading the data, or reattaching the data and log files using sp_attach_db when the rebuild is complete.

C: Collations control the physical storage of character strings in SQL Server 2000 by specifying the bit patterns that represent each character and the rules by which characters are sorted and compared. Whereas earlier versions of SQL Server supported only one collation for each instance of SQL Server, SQL Server 2000 allows objects that have different collations to be stored in a single database. It is thus not necessary to install a new instance of SQL Server 2000 in order to restore the data. We can instead restore the database to an existing SQL Server 2000 computer without having to change the collation of that database.

D: Collations control the physical storage of character strings in SQL Server 2000 by specifying the bit patterns that represent each character and the rules by which characters are sorted and compared. Whereas earlier versions of SQL Server supported only one collation for each instance of SQL Server, SQL Server 2000 allows objects that have different collations to be stored in a single database. It is thus not necessary to install a new instance of SQL Server 2000 on a new computer in order to restore the data. We can instead restore the database to an existing SQL Server 2000 computer without having to change the collation of that database.

Question No: 4

You are the administrator of SQL Server 2000 computer. The server contains a database named sales that has two database files and one transaction log file. Each data file is located on its own hard disk and exists in its own filegroup.

You perform full database, differential, and transaction log backups on a regular basis. All backups made during a single week are striped across three disk backup devices. A portion of the header information for the current week's backups is shown in the following table:

Backup Name	Backup Type	BackupFinishDate
sales_db_10000625_1	2000-06-25 11:57:04.000	
sales_tl_10000626_1	2000-06-26 11:04:12.000	
sales_tl_10000626_2	2000-06-26 15:04:23.000	
sales_df_20000426_5	2000-06-26 21:15:41.000	

On June 28, 2000, at 1:47 P.M., the hard disk that contains the PRIMARY filegroup fails. You want to recover as much data as possible.

What should you do? (Each correct answer presents part of solution. Choose all that apply)

- A. Backup the transaction log by using the NO_LOG option.
- B. Backup the transaction log by using the NO_TRUNCATE option.
- C. Restore the most recent full database backup.
- D. Restore the most recent differential backup.
- E. Restore all differential backups in sequence.
- F. Restore all transaction log backups in sequence.
- G. Restore all transaction logs since the most recent differential backup.

Answer: B, C, D, G.

Explanation:

We want to recover as much data as possible after the failure. We take the following steps:

1. Backup the transaction log with the NO_TRUNCATE option. (B)

If the failure does not include the log files or if you have mirrored them using RAID 1, you should back up the tail of the log with the NO_TRUNCATE option. This will allow us to recover right up to the point of the failure.

2. Restore the most recent full database backup. (C)
3. Restore the most recent differential backup. (D)

4. Restore all transaction logs since the most recent differential backup. (G)

After the last differential backup is restored, you must restore, in sequence, all the log backups made after the last differential backup was made. This includes the tail of the log backed up after the failure if you were able to make this last backup.

Note:

Like differential backups in Windows 2000, differential database backup records only the data that has changed since the last full database backup. This allows for more frequent backups because differential database backups are smaller and faster than full database backups. This decreases the risk of data loss and can be used to restore the database to the point at which the last differential database backup was completed. To recover to the exact point of failure, transaction log backups must be used. To restore differential database backup we must first restore the most recent full database backup, followed by the last differential database backup and then all transaction log backups created after the last differential database backup in the sequence that they were created.

Incorrect answers:

A: The NO_LOG clause is used to truncate the inactive part of the log without making a backup copy of it. Therefore, the changes recorded in the log since the last backup will not be recoverable.

Note: Old transaction log records that are no longer necessary for recovering or restoring a database should be deleted to make way for new log records. The process of deleting these old log records to reduce the size of the logical log is called truncating the log. However, the active portion of the transaction log can never be truncated. The active portion of the log is the part of the log needed to recover the database at any time, so must have the log images needed to roll back all incomplete transactions. It must always be present in the database in case the server fails because it will be required to recover the database when the server is restarted. This truncation occurs automatically at the completion of any BACKUP LOG statement or every time a checkpoint is processed, provided the database is using the Simple Recovery model. This includes both explicit checkpoints resulting from a CHECKPOINT statement and implicit checkpoints generated by the system.

E: Like differential backups in Windows 2000, differential database backup records only the data that

has changed since the last full database backup. This allows for more frequent backups because differential database backups are smaller and faster than full database backups. This decreases the risk of data loss and can be used to restore the database to the point at which the last differential database backup was completed. To restore differential database backup we must first restore the most recent full database backup, followed by the last differential database backup only and then all transaction log backups created after the last differential database backup in the sequence that they were created. The information in the exhibit indicates that the last transaction log backup was made before the last differential backup. We therefore do not need to restore any transaction log backups.

F: The information in the exhibit indicates that the last transaction log backup was made before the last differential backup. Because we must first restore the most recent full database backup, followed by the last differential database backup and then all transaction log backups created after the last differential database backup in the sequence that they were created in order to restore differential database backup, we do not need to restore any transaction log backups as they were made before the last differential backup.

Question No: 5

You are the database administrator of a SQL Server 2000 computer. The server contains your company's Accounts database. Hundreds of users access the database each day.

Because you have had power interruptions in the past, you want to perfect the physical integrity of the Accounts database. You do not want to slow down server operations.

What should you do?

- A.** Enable the torn page detection database option for each database.
- B.** Disable write caching on all disk controllers.
- C.** Ensure that write caching disk controllers have battery backups.
- D.** Create a database maintenance plan to check database integrity and make repairs each night.

Answer: C.

Explanation: The scenario states that power interruptions have occurred in the past. By buying a battery backup solution for the disk controllers the power interruption problem would be prevented.

Incorrect Answers:

A: Torn page detection is a reactive solution which would slow down performance since restores would have to be made. It would better to use a preventive solution.

Note: Torn page detection: If a torn page is detected, the database must be restored from backup because it will be physically inconsistent.

B: Caching is very import for performance. Disabling write-caching would be a bad idea.

D: Repairing the database every night is not a good solution. Database errors cannot be accepted. It is better to prevent the problem instead of trying to fix it after it has happened.

Question No: 6

You are the database administrator for a retail company. The company owns 270 stores. Every month, each store submits approximately 2,000 sales records, which are loaded into a SQL Server 2000 database at the corporate headquarters.

A Data Transformation Services (DTS) package transforms the sales records, as they are loaded. The package writes the transformed sales records to the Sales table, which has a column for integer primary key values. The IDENTITY property automatically assigns a key value to each transformed sales record.

After loading this month's sales data, you discover that a portion of the data contains errors. You stop

loading data, identify the problem records, and delete those records from the database. You want to reuse the key values that were assigned to the records that you deleted. You want to assign the deleted key values to the next sales records you load. You also want to disrupt users' work as little as possible.

What should you do?

- A. Export all records from the **Sales** table to a temporary table. Truncate the **Sales** table, and then reload the records from the temporary table.
- B. Export all records from the **Sales** table to a text file. Drop the **Sales** table, and then reload the records from the text file.
- C. Use the DBCC CHECKIDENT statement to reseed the Sales table's IDENTITY property.
- D. Set the **Sales** table's IDENTITY_INSERT property to ON. Add new sales records that have the desired key values.

Answer: C.

Explanation: DBCC CHECKIDENT is used to check the current identity value for the specified table. It can be set to correct the identity value by specifying the RESEED option with DBCC CHECKIDENT.

Note: Example:

Assume that we have a table called test with an identity column with seed (1,1) and 1000 rows and that we delete row 901 to 950. We could then issue the command

DBCC CHECKIDENT (test, reseed, 900)

And load the 50 missing rows. These rows would get the values 901, 902, ...,950 in the identity column.

We should then set the identity value to 1000 again:

DBCC CHECKIDENT (test, reseed, 1000)

So that the next new row will have identity 1001.

Reference: BOL, DBCC CHECKIDENT

Incorrect Answers:

A: Exporting the data from the Sales table and then reloading the data will not correct errors contained in the table as these errors will also be exported as they are part of the data already in the table.

B: Exporting the data from the Sales table and then reloading the data will not correct errors contained in the table as these errors will also be exported as they are part of the data already in the table.

D: The SET IDENTITY_INSERT command allows explicit values to be inserted into the identity column of a table. To reuse the old key values from the rows that was deleted would require manual commands and would be awkward. This is not the best solution.

Reference: BOL, SET IDENTITY_INSERT

Question No: 7

You are the administrator of a SQL Server 2000 computer. The server contains a database named CustomerManagement that tracks customer requests for product information.

The CustomerManagement database has two filegroups. The PRIMARY filegroup stores current transactional data, and the SECONDARY filegroup stores historical data.

Your backup strategy includes full database and transactional log backups. A portion of the header information for the current week's backups is shown in the following table.

BackupName BackupType BackupA

custmgmt_full_20000625_1 2000-06-25 21:57:03.000
custmgmt_tlog_20000626_1 2 2000-06-26 10:34:21.000
custmgmt_tlog_20000626_2 2 2000-06-26 14:36:32.000
custmgmt_tlog_20000626_3 2 2000-06-26 18:42:11.000
custmgmt_tlog_20000627_1 2 2000-06-27 10:33:37.000
custmgmt_tlog_20000627_2 2 2000-06-27 14:35:18.000

On June 27, 2000, at 4.18 PM, an application developer accidentally deletes the current customer requests and immediately notifies you of the problem. You want to recover as much data as you can as quickly as possible to minimize server downtime.

What should you do?

A. Back up the transaction log.

Restore the PRIMARY filegroup to a new location.

Restore subsequent transaction logs in sequence.

Use the STOPAT option to restore the final transaction log backup.

Import the deleted data from the new location.

B. Back up the transaction log.

Restore the PRIMARY filegroup to the existing database.

Restore subsequent transaction logs in sequence.

Use the STOPAT option to restore the final transaction log backup.

C. Back up the transaction log.

Restore the entire database to a new location.

Restore subsequent transaction logs in sequence.

Use the STOPAT option to restore the final transaction log backup.

Import the deleted data from the new location.

D. Restore the PRIMARY filegroup to the existing database.

Restore subsequent transaction logs in sequence.

Answer: A.

Explanation:

We want to,

a) restore as much data as possible

b) do it as quickly as possible

c) minimize downtime of the database

The following steps outline a process that achieves these goals:

1. When recovering a database, we must first backup the existing transaction log. This will enable us to restore the database changes that were made up to the point of failure rather than to the point at which the transaction log was last backed up.

2. We would then need to restore the last full database backup. We could either restore it into the existing database or into a new location. By restoring into a new location the downtime of the server would be minimized.

3. A transaction log recovery requires an unbroken chain of transaction log backups from the time of the database backup to the time of the failure. As the transaction log is automatically truncated after it is backed up. Therefore we would need to restore the transaction logs in sequence from the first transaction log backup that was made after the last full database backup until the most recent one. This specifies that the database be restored to the state it was in as of the specified date and time.

4. When restoring the final transaction log back up that was made as part of the recovery process, we

must use STOPAT option. With the STOPAT option we can restore data up to the point of failure.
Note: The Secondary filegroup holds only historical data that is not required for query processing and has not been deleted. It is therefore not necessary to restore the Secondary filegroup. Only the Primary filegroup thus needs be restored. This will reduce the time required to restore the database.
5. Finally we import the deleted data from the new location into the existing database.

Incorrect Answers:

B: Downtime will be minimized by restoring to a new location and later importing the restored data into the database. This allows use of the database during the restoration process.


C: The Secondary filegroup holds only historical data that is not required for query processing and has not been deleted. It is therefore not necessary to restore the Secondary filegroup. Thus, only the Primary filegroup and not the entire database needs be restored. This will reduce the time required to restore the database.

D: When recovering a database, we must first backup the existing transaction log. This will enable us to restore the database changes that were made up to the point of failure rather than to the point at which the transaction log was last backed up. If we do not backup the current transaction log, the changes made to the database since the last transaction log backup cannot be restored. This will result in loss of data.

SUBSECTION B: Perform backups. (4 questions)

Question No: 1

You are the administrator of two SQL Server 2000 computers for an online retailer. The servers receive and utilize customer data as shown in the exhibit.



One server contains a database that records customer data. The second server imports and transforms the data for analysis.

The Data Transformation Services (DTS) package is stored in the Meta Data Services repository on the second server. You want to maximize the amount of lineage data that can be recovered if a data file is damaged or lost.

Which two actions should you take? (Each correct answer represents part of the solution. Choose two.)

- A.** Use the Full Recovery model for the staging database.
- B.** Use the Full Recovery model for the msdb database.
- C.** Back up the transaction log in the staging database by using the NO_TRUNCATE option.
- D.** Back up the transaction log in the msdb database by using the NO_TRUNCATE option.

- E. Back up the multidimensional data cube.
- F. Save the DTS package as a file.

Answer: B, D.

Explanation:

B: The DTS package is saved in the msdb database. The full recovery database model is recommended when backing up the msdb database.

Note 1: Meta Data Services uses msdb as the default repository database. The msdb database is used to store data, including scheduling information and backup and restore history information including backups that were created using custom or third party applications. This information includes who performed the backup, when, and where the backup is stored. This information is used by SQL Server Enterprise Manager to propose a plan for restoring a database and applying any transaction log backups. When the backup and restore history information in msdb is used in recovering user databases, it is recommended that the Full Recovery model be used for msdb.

Note 2: SQL Server 2000 offers three recovery models: the Simple Recovery model, which allows the database to be recovered to the most recent backup but not to the point of failure or to a specific point in time; the Full Recovery model, which allows the database to be recovered to the point of failure and, if one or more data files is damaged, it can restore all committed transactions while in-process transactions are rolled back; and the Bulk-Logged Recovery model, which allows bulklogged operations. In a Bulk-Logged Recovery model, the data loss exposure for bulk copy operations is greater than in the Full Recovery model. While the bulk copy operations are fully logged under the Full Recovery model, they are minimally logged and cannot be controlled on an operation-by-operation basis under the Bulk-Logged Recovery model. Under the Bulk-Logged Recovery model, a damaged data file can result in having to redo work manually.

D: The DTS package is saved in the msdb database.

Normally when a SQL Server completes a backup of the transaction log, it automatically truncates the inactive portion of the transaction log. This inactive portion contains completed transactions and is no longer used during the recovery process while the active portion of the transaction log contains transactions that are still running and have not yet been completed.

The backup command with the NO_TRUNCATE option allows backing up the log in situations where the database is damaged. This meets the requirement that we should be able to recover as much data as possible if a data file is damaged or lost.

Note 3: When saving a Data Transformation Services (DTS) package, all DTS connections, DTS tasks, DTS transformations, and workflow steps can be saved and the graphical layout of these objects on the DTS Designer design sheet can be preserved. A DTS package can be saved to SQL Server 2000 Meta Data Services. With this save option, the data lineage feature can be used. This can track and record row-level data lineage, which reveals the source of any piece of data and the transformations applied to that data; and column-level data lineage, which provides information about a package version and the database tables and columns the package uses as a source or destination.

Incorrect Answers:

A: The DTS package is saved to the msdb and not the staging database. Therefore the msdb and not the staging database should be backed up to recover DTS packages and transactions.


C: The DTS package is saved to the msdb and not the staging database. Therefore the transaction log in the msdb and not the staging database should be backed up to recover DTS packages and transactions.

E: Cubes are used in online analytic processing (OLAP), which provides fast access to data in a data warehouse. A cube is a set of data that is usually constructed from a subset of a data warehouse and is organized and summarized into a multidimensional structure defined by a set of dimensions and measures.

F: Saving a DTS package to a structured storage file allows you to copy, move, and send a package across the network (such as in a mail message) without storing the package in a database or a repository. However, it would be better to save the DTS package to SQL Server Meta Data Services as this allows you to track package version, meta data, and data lineage (original data source and transformations) information. In this scenario the DTS package has already been saved into the Meta Data Repository. Saving it as a file would not be beneficial.

Question No: 2

You are the administrator of SQL Server 2000 computer. The server contains a database named Medical Records. This database contains clinical records and administrative records. The database consumes about 30GB of disk space. The data files are configured as shown in the exhibit.



You run full database backups each night after business hours. You notice that the backup is not complete by the time the morning shift begins to use the database.

You need to minimize the time needed to restore data in the event of a system, failure. You also want to reconfigure the database to allow the backups to complete during the evening hours.

Which two actions should you take? (Each correct answer represents part of the solution. Choose two)

A. Reorganize the data files into two groups.

Place the system tables in the PRIMARY filegroup and the user-defined tables in the other filegroup.

B. Reorganize the data files into three groups.

Place the system tables and shared objects in the PRIMARY filegroup, the clinical records in a filegroup, and the administrative records in a filegroup.

C. Reorganize the data files into three groups.

Place the system tables in the PRIMARY filegroup, the indexes in a filegroup, and the tables in the other filegroup.

D. Back up the transaction log each night.

Run a filegroup backup on a different filegroup each night.

E. Back up the transaction log each night.

Run a filegroup backup on the database each weekend.

F. Back up the transaction log each night.

Run a differential backup each Sunday, and run a full backup the first day of each month.

Answer: B, D.

Explanation:

B: Split the data into three filegroups: one for system tables and shared object, one for clinical records, and one for administrative records.

The primary filegroup must contain the system tables.

D: By separating the data we can rotate the backup of the filegroups and only backup a single filegroup each night. This will make the backup process quicker and would allow it to complete before business starts in the morning. In order to be able to recover we should also make a log backup every night.

Note: SQL Server 2000 supports backing up or restoring individual files or filegroups within a database, it is possible to reduce backup time by placing data and tables in separate filegroups. Tables, indexes, and text, ntext, and image data can be placed in a specific filegroup. A database can only have one primary filegroup. This filegroup contains the primary data file and any other files not placed in another filegroup. All pages for the system tables are allocated from the primary filegroup. However, a number of user-defined filegroups can be defined for a database. It would thus be possible to reorganize the data files into three groups that can each be backed up separately.

The transaction log holds a record of all modifications that have occurred in the database as well as the transaction that have performed each modification since the transaction log was last backed up. With transaction log backups, it is possible to recover the database to a specific point in time or to the point of failure. They generally use fewer resources than database backups and can thus be backed up more frequently than database backups. This would decrease risk of data loss; however, transaction log backups are used only with the Full and Bulk-Logged Recovery models.

Incorrect Answers:

A: Tables, indexes, and text, ntext, and image data can be placed in specific filegroups. However, the user-defined tables would be much larger than the system tables and might require further partitioning to ensure that the backup of this filegroup is complete.

C: Splitting the indexes and the tables into different filegroups could increase performance for join queries. It would not be a good idea in this scenario however since the separate filegroups are backup up on different days. Therefore a restore operation could give a situation where the indexes and the tables not were up to date. In fact if the indexes of a table span multiple filegroups, all filegroups containing the table and its indexes must be backed up together

E: Running a filegroup backup on the weekend would increase the amount of time between data backups and would increase the time required for recovery as each transaction log backup created since the last filegroup backup would have to be restored in the same sequence in which they were

created.

F: Running a differential backup each Sunday, and run a full backup the first day of each month while backing up the transaction log each night would increase the required time for data recovery. In this solution we would have to back up the active transaction log, restore the most recent database backup, restore the most recent differential backup and then restore each transaction log backup created since the last differential backup in the same sequence in which they were created.

Question No: 3

You are the administrator of a SQL Server 2000 computer. You import a large amount of data into the Contracts database. After the import, you perform a full database backup. You examine the hard disk and discover that the log file has grown from 200MB to 800MB.

You want to recover disk space as quickly as possible without losing backup information. Which two Transact-SQL statements should you execute? (Each correct answer presents part of the solution. Choose two)

- A. BACKUP LOG Contracts WITH no_log
- B. BACKUP LOG Contracts WITH truncate_only
- C. BACKUP LOG Contracts TO disk= 'E:\SQLBackup\Contracts_Log.bkp'
- D. ALTER DATABASE Contracts
(Name= Contract_Log.ldf,
FILENAME= 'E:\Data\Contracts_log.ldf'
SIZE= 200)
- E. DBCC ShrinkDatabase (Contracts, 200)
- F. DBCC ShrinkFile (Contracts_Log, 200)

Answer: C, F.

Explanation: We have the requirement to quickly recover disk space without losing backup information.

C: The log file is backed up to another partition without truncating it. We should not truncate the log file in the backup process. This way we are ensured that we don't lose any information.

F: After the log file is backed up we must truncate it. This is accomplished with the command DBCC SHRINKFILE(*file_name*, *target_size*). We should shrink the logfile to 200MB.

Note 1: The transaction log records all the transactions that have been performed against the database since the transaction log was last backed up. After SQL Server 2000 finishes backing up the transaction log, it automatically truncates the inactive portion of the transaction log. This inactive portion contains completed transactions that are not required for the recovery process, while the active portion of the transaction log contains transactions that are still running and have not yet completed. SQL Server 2000 reuses this truncated, inactive space in the transaction log instead of allowing the transaction log to continue to grow and use more space. The BACKUP LOG statement specifies a backup of the transaction log only. The log is backed up from the last successfully executed BACKUP LOG statement to the current end of the log. Once the log is backed up, the space may be truncated when no longer required by replication or active transactions. However, transaction log truncation does not reduce the physical size of the log file, it reduces the size of the logical log and marks as inactive the virtual logs that do not hold any part of the logical log. The DBCC SHRINKFILE statement is used to reduce the physical size of the log file. The shrinking a log is dependent on first truncating the log, which occurs automatically when the log file is backed up. The DBCC SHRINKFILE can also be used to shrink the file to a size smaller than specified when the file was created.

Note 2: The WITH NO_LOG and the WITH TRUNCATE_ONLY options of the BACKUP LOG statement are used to free up physical disk space by removing the inactive part of the log without making a backup

copy of it and then truncating the log. Specifying a backup device with this option is unnecessary because the log backup is not saved. After backing up the log using the WITH NO_LOG (or the WITH TRUNCATE_ONLY) option, the changes recorded in the log are not recoverable. Therefore, we should execute the BACKUP DATABASE statement immediately after backing up the log with the WITH NO_LOG option to ensure the recoverability of the database. The WITH NO_LOG and WITH TRUNCATE_ONLY options are synonyms. With this solution backup information is lost as there is no option to backup the database.

Incorrect Answers:

A: In order to avoid data loss we must backup the log file and not truncate it.

If you back up the log file to the same partition as the data will not free any space.

B: In order to avoid data loss we must backup the log file and not truncate it.

If you back up the log file to the same partition as the data will not free any space.

D: The ALTER DATABASE statement can be used to alter the attributes of a file. These attributes include the FILENAME, SIZE, FILEGROWTH, and MAXSIZE options, though only one of these attributes can be changed at a time. However, the MODIFY FILE option and the file name must be specified in the ALTER DATABASE statement before the file can be altered. The MODIFY FILE option is not stated in this solution.

E: The DBCC SHRINKDATABASE statement shrinks data files on a per-file basis but shrinks log files as if all the log files existed in one contiguous log pool. The target size for the data and log files cannot be smaller than the minimum size of a file that was specified when the file was originally created, or the last explicit size set with a file size changing operation such as the ALTER DATABASE statement with the MODIFY FILE option or the DBCC SHRINKFILE statement. With this solution we are also attempting to shrink the database to 200 MB and not the transaction log file.

Question No: 4

You are the administrator of a SQL Server 2000 computer. The server contains your company's order processing database. Two-hundred operators take orders by telephone 24 hours a day. Threehundred data entry personnel enter data from orders received by mail.

To ensure that order data will not be lost, your company's disaster recovery policy requires that backups be written to tape. Copies of these tapes must be stored at an off-site company location.

Orders must be entered into the database before they can be filled. If the server fails, you must be able to recover the order data as quickly as possible.

You need to create a backup strategy that meets the company requirements and minimizes server workload. Which two actions should you take? (Each correct answer represents part of the solution. Choose two.)

A. Perform a combination of full database and filegroup backups.

B. Perform a combination of full database and file backups.

C. Perform a combination of full database, differential, and transaction log backups.

D. Back up the data to a local tape drive.

E. Back up the data to a network share, and then use enterprise backup software to write the disk backups to tape.

Answer: C, E.

Explanation:

C: Most OLTP backup schemes include a full database, differential, and transaction log backups

E: In order to recover the data as quickly as possibly we backup the data to network share, and then backup

that share to tape. This minimizes the impact of the backup process on the SQL Server computer.

Note 1: SQL Server 2000 provides three database recovery models. These are the Simple recovery model, the Full Recovery Model and the Bulk-Logged Recovery model. Simple Recovery requires the least administration. In the Simple Recovery model, data is recoverable only to the last full database or differential backup. Transaction log backups are not used, and consequently, minimal transaction log space is required. This Recovery Model is not an appropriate choice for production systems where loss of recent changes is unacceptable. Full Recovery and Bulk-Logged Recovery models provide the greatest protection for data. These models rely on the transaction log to provide full recoverability and to prevent work loss in the broadest range of failure scenarios. The Full Recovery model provides the most flexibility for recovering databases to an earlier point in time. The Bulk-Logged model provides higher performance and lower log space consumption for certain large-scale operations. It does this at the expense of some flexibility of point-in-time recovery. In this scenario we would want to implement the Full Recovery model as this Recovery model provides complete protection against media failure and provides the ability to recover the database to the point of failure or to a specific point in time. To guarantee this degree of recoverability, all operations are fully logged. The backup strategy for full recovery consists of database backups, differential backups, which are optional, and transaction log backups. The differential backups are optional but have less of an

effect on the system while they run, as they are smaller than full database backups. This is because a differential backup creates a copy of all the pages in a database modified after the last database backup. A transaction log backup makes a copy only of the log file and backs up the changes made to the database after the last full or differential database backup. A log file backup cannot be used by itself to restore a database. A log file is used after a database restore to recover the database to the point of the original failure. The full database, differential, and transaction log backups can be written to a hard disk or other disk storage media such as a tape drive. Disk backup devices can be defined on a local disk of a server or on a remote disk on a shared network resource. Tape backup devices are used in the same way as disk devices, with the exception that the tape device must be connected physically to the computer running an instance of SQL Server. SQL Server 2000 thus does not support backing up to remote tape devices.

Note 2, file backup, filegroup backup: SQL Server 2000 supports backing up or restoring individual files or filegroups within a database. This is a complex backup and the restore process that is usually reserved for large databases with high availability requirements. If the time available for backups is not long enough to support backing up the full database, subsets of the database can be backed up at different times. If a disk holding database files or filegroups fails, the site can restore just the lost files or filegroups. This allows for a quicker recovery because only the damaged files or filegroups are restored in the first step, not the entire database.

Incorrect Answers:

A: An OLTP server requires transaction log backups in order to be able to recover up to the point of failure.

B: An OLTP server requires transaction log backups in order to be able to recover up to the point of failure.

D: Backing up to a local tape device would minimize cost but it would compared to a backup to a network it would decrease performance of the server more, and it would take more time to recover.


SUBSECTION C: Recover the system state and restore data. (0 questions)

SECTION 4: Configure, maintain, and troubleshoot log shipping. (1 question)

Question No: 1

You are the administrator of three SQL Server 2000 computers at Five Lakes Publishing. One server, FLPSQL01, stores order data. You want to be able to use the other two servers, FLPSQL02 and

FLPSQL03, to answer queries and run report. The planned network configuration is shown in the exhibit.



You want to use the database Maintenance Plan wizard to configure log shipping from FLPSQL01 to FLPSQL02 and FLPSQL03. You do not want users to add any new data on FLPSQL02 and FLPSQL03, but they must be able to run queries and reports.

Which two actions should you take? (Each correct answer presents part of the solution. Choose two)

- A. Set the database load state for each destination server to No recovery mode.
- B. Set the database load state for each destination server to standby mode.
- C. Enable the allow database to assume primary role option for both destination servers.
- D. Enable the allow database to assume primary role option for one destination server and then disable this option for other destination server.
- E. Disable the allow database to assume primary role option for both destination servers.

Answer: B, E.

Explanation:

B: A standby server can be maintained to take the place of the primary production server, if needed. A standby server is a secondary SQL Server installation that is kept up-to-date with a primary SQL Server installation through the application of transaction log backup files. You can automate the process of maintaining a standby server by creating backup, copy, and restore jobs that are periodically executed by SQL Server Agent on the primary server and on the standby server. This automated process is called log shipping. If you want the destination database to be viewable for

read-only queries, you must select the Standby Mode option, the default is the No Recovery Mode option. The Standby mode specifies that the secondary database be made available for use, but in read-only mode.

E: We don't want users to add any data on FLPSQL02 or FLPSQL03. Therefore we should disable the allow database to assume primary role option for them. If this option were enabled FLPSQL02 and

FLPSQL03 would be able to be used as source servers.

Incorrect Answers:

A: The no recovery mode specifies that the secondary database be made unavailable for use. We need standby mode since we want FLPSQL02 and FLPSQL03 used for queries and reports.

Note: The active portion of the transaction log on the primary server can be backed up using the NORECOVERY backup option. This option leaves the primary database in recovery mode. When the primary server is ready to be brought back online, the transaction logs from the standby server (for the period that users were using the standby server as the primary server) can be applied to the primary server. This avoids the necessity of applying a complete database backup and all applicable transaction logs to restore the primary server. The result is a significant decrease in the time required to bring the primary server back online.

C: The primary role option must be disabled for both destination servers. If not they would be able to change and update data.

D: The primary role option must be disabled for both destination servers. If not they would be able to change and update data.

SECTION 5: Perform integrity checks. Methods include configuring the Database Maintenance Plan Wizard and using the Database Consistency Checker (DBCC). (0 questions)

SECTION 6: Troubleshoot transactions and locking by using SQL Profiler, SQL Server Enterprise Manager, or Transact-SQL. (8 questions)

Question No: 1

You are the administrator of a SQL Server computer. Users report that the database times out when they attempt to modify data. You use the Current Activity window to examine locks held in the database as shown in the following screenshot.



You need to discover why users cannot modify data in the database, but you do not want to disrupt normal database activities. What should you do?

- A.** Use the spid 52 icon in the Current Activity window to discover which SQL statement is being executed
- B.** Use the **sp_who** stored procedure to discover who is logged in as spid 52
- C.** Use SQL Profiler to capture the activity of the user who is logged in as spid 52
- D.** Use System Monitor to log the locks that are granted in the database

Answer: A.

Explanation: The current activity window in SQL Server Enterprise Manager can be used to perform ad hoc monitoring of an instance of SQL Server. This allows for the quick analysis of the volume and general types of activity on the system and can indicate current blocked and blocking transactions, currently connected users on an instance of SQL Server and the last statement executed and locks that are in effect on the database. In the screen shot of current activity window the icon, labeled spid 51, indicates that a database

lock is in place; the second icon, labeled spid 52, can be used to identify the process that is blocking one or more connections; and the third icon, labeled spid 53, can be used to identify the process that is being blocked by another connection

Incorrect Answers:

B: The system stored procedure `sp_lock` can be used to return a result set that contains information about resources that are locked and the `sp_who` stored procedure provides information about current SQL Server 2000 users and processes. The information returned by the `sp_who` procedure can be filtered to return only current non-idle processes. This makes it possible to identify which resources are being locked and which processes are responsible for creating those locks. These stored procedures require more resources and more administrative effort than the Current Activity window. It is therefore not the best answer.


C: The SQL profiler is not the best tool to immediately solve the locking problem. It could be used at a later moment if a long-term solution to the problem is required. SQL profiler would require more system resources and might impact normal database activities.

Note: SQL Profiler is a graphical tool that allows system administrators to monitor events in an instance of SQL Server. Data pertaining to an event can be captured and saved to a file or SQL Server table to be analyzed at a later date. It can be used to step through problem queries to find the cause of the problem; to find and diagnose slow-running queries; to capture the series of SQL statements that led to a problem; and to monitor the performance of SQL Server to tune workloads.

D: SQL Server provides objects and counters that can be used by System Monitor in a Windows 2000 or by Performance Monitor in Windows NT 4.0 to monitor the system performance of computers running an instance of SQL Server. However, these counters are used for statistics and alerts and they cannot, in a direct way, be used to track down the cause of the locking problem.

Question No: 2

You are the administrator of SQL Server 2000 computer. The server contains a database named sales. Pierre reports that he cannot modify data in the database. You use SQL Profiler to capture his activities shown in the exhibit.



EventClass	TestData
TraceStart	
ExistingConnection	-- network protocol: LPC set quoted identifier on set implicit transactions off ...
ExistingConnection	-- network protocol: LPC set quoted identifier on set implicit transactions off ...
SQLBatchCompleted	set transaction isolation level serializable begin transaction select * from ...

set transaction isolation level serializable
begin transaction
select * from Products
update products
set UnitPrice = 10.50
where UnitPrice = 10.00
commit transaction

You want Pierre to be able to modify data. You also want the Transact-SQL statement to execute without compromising data integrity.

What are two possible ways to achieve this goal? (Each correct answer presents a complete solution.)

Choose two)

- A. Change the isolation level of the database connection.
- B. Add indexes to the product table.
- C. Remove the explicit transaction from the Transact-SQL batch.
- D. Create a stored procedure to execute the Transact-SQL batch.
- E. Add a cascading UPDATE trigger to the Products table.

Answer: A, C.

Explanation: This SQL-code is poorly written since the whole batch is one transaction. The explicit transaction should be removed from the code.

The transaction isolation level of serializable is also a bad choice. This is the highest isolation level and would probably lock the whole batch. The isolation level must be lowered.

Incorrect answers:

B: Indexes could improve performance, but this isn't the goals in this scenario. We want to make the SQL batch run and keep data integrity.


D: Putting the code into a stored procedure would keep the problems with the explicit transaction and the high isolation level.

E: The batch is only run on one table. Cascading triggers are used when several tables have to be updated.

Question No: 3

You are the administrator of a SQL Server 2000 computer. The server contains a database named Sales. Your company deploys an upgraded version of a customer relationship management application, which accesses the Sales database.

The Sales database contains a table named Customers, which stores information about your company's customers. Users report that they cannot delete information in the CustomerAddress2 column. The Customers table is configured as shown in the Table Design exhibit.



You use SQL profiler to capture the activity of the application and receive the results shown in the Trace exhibit.

EventClass	TextData	ApplicationName	NTUserName
Audit Login	-- network protocol.	Visual Basic	Admin
SQL:BatchCompleted	Update Customers Set	Visual Basic	Admin
<pre>update Customers set CustomerAddress2= NULL where CustomerID = 1</pre>			

You want to configure the customer's table to support the upgraded application. What should you do?

- A. Alter the table to allow null values in the **CustomerAddress2** column.
- B. Alter the data type of the **CustomerAddress2** column to **sql_variant**.
- C. Create an UPDATE trigger that changes a null value in the **CustomerAddress2** column to an empty string.
- D. Use an INSTEAD OF trigger to change the value of the **CustomerAddress2** column to an empty string whenever the column's value is null.

Answer: A

Explanation: Data cannot be deleted from the CustomerAddress2 column. When data is deleted from a column a NULL value is used instead. But NULL isn't allowed in this column. By simply allowing NULL values the upgraded application would run with no performance impact.

Incorrect Answers:

- B:** Changing the datatype of the column would not do anything about the NULL constraint.
- C:** We must create an INSTEAD OF trigger not an UPDATE trigger, since the UPDATE trigger

wouldn't fire until after the row is updated and then it would be too late since the NULL constraint would not allow the row from being updated.

D: Triggers should only be used when it is necessary. It is better to remove the NULL constraint in this scenario.

Note: By creating an INSTEAD OF trigger that changes NULL values to an empty string whenever the value of the CustomerAddress2 column this problem would be avoided and at the same time the NULL constraint would be kept. However, this would decrease performance and is therefore not the best solution.

Question No: 4

You are the administrator of a SQL Server 2000 computer. The server contains a database named Customers, which is used by a custom client/server application.

The database contains more than 1 million rows. Users report that the database responds slowly when they search for customers by using the application. Users search for customers by scrolling through an alphabetical list of customer names.

You use SQL profiles to capture information about how the application uses the database.

You receive results as shown in the following screen shot.

EventClass	TextData
SQL:BatchCompleted	select * from Customers
	select * from Customers

You need to improve the application's performance. What should you do?

- A.** Create additional nonclustered indexes on the **CustomerInformation** table.
- B.** Rewrite the application so that it uses a user-defined function to return the customer list.
- C.** Rewrite the application so that it uses a query with a WHERE clause to search for customer names.
- D.** Rewrite the application so that it uses a stored procedure to return the customer list.

Answer: C.

Explanation: The WHERE clause of the SELECT statement causes the queries to be selective. Without the WHERE clause all records in the selected table is included in the result set. The WHERE clause tells SQL Server 2000 which data should be returned and thus limits the returned data to only that data that meets the condition set in the WHERE clause. This would improve the efficiency of the application.

Incorrect Answers:

A: Nonclustered indexes are useful for exact match queries because these indexes contain entries describing the exact location in the table of the data values being searched for. However, the application in this scenario does not make use of exact match queries.

B: Since functions only return a single value, not a recordset, they cannot be used in this scenario.

Note: In programming languages, functions are subroutines that are used to encapsulate frequently performed logic. Any code that must perform the logic incorporated in a function can call the function rather than having to repeat all of the function logic. In addition to built-in functions, SQL Server 2000 also allows us to define our own Transact-SQL functions using the CREATE FUNCTION statement. These user-defined functions can take zero or more input parameters, and can return a single value. Some user-defined functions can return a single, scalar data value or decimal value.

D: A stored procedure could be used but would not, by itself, improve any performance gain except that small gain of caching the execution plan.

Note: A stored procedure is a group of Transact-SQL statements that are compiled into a single execution plan. SQL Server 2000 stored procedures return data as output parameters, which can return either data or a cursor variable; return codes, which are always an integer value; a result set for each SELECT statement contained in the stored procedure or any other stored procedures called by the stored procedure; or as a global cursor that can be referenced outside the stored procedure. They assist in achieving a consistent implementation of logic across applications. The SQL statements and logic needed to perform a commonly performed task can be designed, coded, and tested once in a stored procedure. Each application needing to perform that task can then simply execute the stored procedure. Stored procedures can also improve performance.

Question No: 5

You are the administrator of a SQL Server 2000 computer. The server contains a database that stores product data for your company.

You need to execute an existing stored procedure that examines prices for your company's products and can modify them if necessary. You execute the stored procedure after business hours, but it does not complete. You execute the sp_lock stored procedure and receive the following output.

spid dbid ObjId IndId Type Mode

```
61 7 0 0 DB S
64 7 0 0 DB S
72 7 0 0 DB S
72 7 2145623952 1 PAG IS
72 7 2145623952 0 TAB IS
72 7 2145623952 1 KEY S
78 7 0 0 DB S
78 7 2145623952 1 PAG IX
78 7 2145623952 0 TAB IX
78 7 2145623952 1 KEY X
```

You want the stored procedure to complete successfully. What should you do?

- A.** Execute the stored procedure, and specify the WITH RECOMPILE option.
- B.** Execute the DBCC FREEPROCCACHE statement.
- C.** Release the locks that are held by connections 61 and 64.
- D.** Release the locks that are held by connections 72 and 78.

Answer: D.

Explanation: In the exhibit spid 78 has an exclusive a row lock within an index and intent exclusive locks on the entire table as well as on data or index pages. Exclusive and intent exclusive locks prevent access to a resource by concurrent transactions. This will prevent the stored procedure from changing data table. Therefore we need to release the locks held by spid 78. Spid 72 holds share locks on the database, shared row locks within an index, and intent to share locks on the data or index pages and on the entire table.

Note: The sp_lock stored procedure returns information about locks in the form of a table. The returned table includes information used to identify SQL Server 2000 server process ID numbers (spid); the database identification number requesting a lock (dbid); the object identification number of the object requesting a lock (ObjId); the index identification number (IndId); the type of lock (Type); and the lock mode (Mode).

The lock modes used by SQL Server 2000 are shared locks (S), update locks (U), exclusive locks (X), and intent locks, which could be intent shared locks (IS), intent exclusive locks (IX), and shared with intent exclusive locks (SIX). Shared locks allow concurrent transactions to read a resource by using the SELECT statement. While a shared lock exists on the resource, no other transactions can modify the data. Shared locks are released as soon as the data has been read, unless the transaction isolation level is set to repeatable read or higher, or a locking hint is used to retain the shared locks for the duration of the transaction. Update locks prevent a common form of deadlock. A typical update pattern consists of a transaction reading a

record, acquiring a shared lock on the resource, and then modifying the row, which requires lock conversion to an exclusive lock. If two transactions acquire shared-mode locks on a resource and then attempt to update data concurrently, one transaction attempts the lock conversion to an exclusive lock. The shared-mode-to-exclusive

lock conversion must wait because the exclusive lock for one transaction is not compatible with the shared-mode lock of the other transaction; a lock wait occurs. The second transaction attempts to acquire an exclusive lock for its update. Because both transactions are converting to exclusive locks, and they are each waiting for the other transaction to release its shared-mode lock, a deadlock occurs. To avoid this problem, only one transaction is allowed to obtain an update lock on a resource at a time. If a transaction modifies a resource, the update lock is converted to an exclusive lock. Otherwise, the lock is converted to a shared lock. Exclusive locks prevent access to a resource by concurrent transactions. Other transactions are thus prevented from reading or modifying data locked with an exclusive lock. An intent lock indicates that SQL Server 2000 wants to acquire a shared lock or exclusive lock on some of the resources lower down in the hierarchy. Setting an intent lock at the table level prevents another transaction from subsequently acquiring an exclusive lock on the table containing that page. Intent locks improve performance because SQL Server examines intent locks only at the table level to determine if a transaction can safely acquire a lock on that table. This removes the requirement to examine every row or page lock on the table to determine if a transaction can lock the entire table. Intent locks include intent shared, intent exclusive, and shared with intent exclusive. Intent shared locks indicate the intention of a transaction to read some but not all resources lower in the hierarchy by placing shared locks on those individual resources. Intent exclusive locks indicate the intention of a transaction to modify some but not all resources lower in the hierarchy by placing exclusive locks on those individual resources. Shared with intent exclusive locks indicate the intention of the transaction to read all of the resources lower in the hierarchy and modify some but not all of those resources lower by placing intent exclusive locks on those individual resources. Concurrent IS locks at the top-level resource are allowed. There can be only one shared with intent exclusive lock per resource at one time, preventing updates to the resource made by other transactions, although other transactions can read resources lower in the hierarchy by obtaining intent shared locks at the table level. The Type column indicates the type of the resource that is currently locked. RID indicates that a row identifier is used to lock a single row within a table; KEY indicates a row lock within an index, which is used to protect key ranges in serializable transactions; PAG indicates that a data or index page is locked; EXT indicates that contiguous group of eight data pages or index pages are locked; TAB indicates that the entire table, including all data and indexes is locked; and DB indicates that the database is locked.

Incorrect Answers:

A: Recompiling the procedure will not help. This is a locking problem.

Note: As a database changes when indexes are added or data in indexed columns are changed, the original query plans used to access the database tables should be re-optimized by recompiling them. This optimization occurs automatically the first time a stored procedure is run after SQL Server 2000 is restarted. It also occurs if an underlying table used by the stored procedure changes. However, if a new index is added from which the stored procedure might benefit, optimization does not occur

automatically until the next time the stored procedure is run after SQL Server is restarted. A stored procedure can be forced to be recompiled by specifying the WITH RECOMPILE option when stored procedure is executed. This option should be used only if the input parameter is atypical or if the data has changed significantly since the stored procedure was created.

B: DBCC FREEPROCCACHE is used to clear the procedure cache. This would cause an ad-hoc SQL statement or a stored procedure to be recompiled rather than reused from the cache.

C: Spid 61 and spid 64 holds shared locks on the database. These types of locks do not prevent concurrent access. Therefore it is not necessary to release these locks.

Question No: 6

You are the administrator of a SQL Server 2000 computer. The server contains a database that stores product data for your company.

You need to run a report that summarizes data about these products. When you run the report, it does not complete.

Several users are connected to the database. You execute the sp_lock stored procedure and receive the following output.

spid dbid ObjId IndId Type Mode

```
-----  
61 7 0 0 DB S  
64 7 0 0 DB S  
72 7 0 0 DB S  
72 7 2145623952 0 RID S  
72 7 2145623952 0 PAG IS  
72 7 2145623952 0 TAB S  
72 7 2145623952 0 RID S  
78 7 0 0 DB S  
78 7 2145623952 1 RID X  
78 7 2145623952 1 PAG IX  
78 7 2145623952 0 TAB IX  
78 7 2145623952 1 RID X  
78 7 2145623952 1 PAG IX  
78 7 2145623952 1 RID X
```

You want the report to complete successfully. What should you do?

- A.** Execute the sp_releaseapplock stored procedure
- B.** Grant SELECT permission on the Products table to all users.
- C.** Release the locks that are held by connections 72.
- D.** Release the locks that are held by connections 78.

Answer: D.

Explanation: In the exhibit spid 78 has an exclusive a row lock within an index and intent exclusive locks on the entire table as well as on data or index pages. Exclusive and intent exclusive locks prevent access to a resource by concurrent transactions. This will prevent the stored procedure from changing data table. Therefore we need to release the locks held by spid 78. Spid 72 holds share locks on the database, shared row locks within an index, and intent to share locks on the data or index pages and on the entire table.

Note:The sp_lock stored procedure returns information about locks in the form of a table. The returned table includes information used to identify SQL Server 2000 server process ID numbers (spid); the database

identification number requesting a lock (dbid); the object identification number of the object requesting a lock (ObjId); the index identification number (IndId); the type of lock (Type); and the lock mode (Mode). The lock modes used by SQL Server 2000 are shared locks (S), update locks (U), exclusive locks (X), and intent locks, which could be intent shared locks (IS), intent exclusive locks (IX), and shared with intent exclusive locks (SIX). Shared locks allow concurrent transactions to read a resource by using the SELECT statement. While a shared lock exists on the resource, no other transactions can modify the data. Shared locks are released as soon as the data has been read, unless the transaction isolation level is set to repeatable read or higher, or a locking hint is used to retain the shared locks for the duration of the transaction. Update locks prevent a common form of deadlock. A typical update pattern consists of a transaction reading a record, acquiring a shared lock on the resource, and then modifying the row, which requires lock conversion to an exclusive lock. If two transactions acquire shared-mode locks on a resource and then attempt to update data concurrently, one transaction attempts the lock conversion to an exclusive lock. The shared-mode-to-exclusive

lock conversion must wait because the exclusive lock for one transaction is not compatible with the shared-mode lock of the other transaction; a lock wait occurs. The second transaction attempts to acquire an exclusive lock for its update. Because both transactions are converting to exclusive locks, and they are each waiting for the other transaction to release its shared-mode lock, a deadlock occurs. To avoid this problem, only one transaction is allowed to obtain an update lock on a resource at a time. If a transaction modifies a resource, the update lock is converted to an exclusive lock. Otherwise, the lock is converted to a shared lock. Exclusive locks prevent access to a resource by concurrent transactions. Other transactions are thus prevented from reading or modifying data locked with an exclusive lock. An intent lock indicates that SQL Server 2000 wants to acquire a shared lock or exclusive lock on some of the resources lower down in the hierarchy. Setting an intent lock at the table level prevents another transaction from subsequently acquiring an exclusive lock on the table containing that page. Intent locks improve performance because SQL Server examines intent locks only at the table level to determine if a transaction can safely acquire a lock on that table. This removes the requirement to examine every row or page lock on the table to determine if a transaction can lock the entire table. Intent locks include intent shared, intent exclusive, and shared with intent exclusive. Intent shared locks indicate the intention of a transaction to read some but not all resources lower in the hierarchy by placing shared locks on those individual resources. Intent exclusive locks indicate the intention of a transaction to modify some but not all resources lower in the hierarchy by placing exclusive locks on those individual resources. Shared with intent exclusive locks indicate the intention of the transaction to read all of the resources lower in the hierarchy and modify some but not all of those resources lower by placing intent exclusive locks on those individual resources. Concurrent IS locks at the top-level resource are allowed. There can be only one shared with intent exclusive lock per resource at one time, preventing updates to the resource made by other transactions, although other transactions can read resources lower in the hierarchy by obtaining intent shared locks at the table level. The Type column indicates the type of the resource that is currently locked. RID indicates that a row identifier is used to lock a single row within a table; KEY indicates a row lock within an index, which is used to protect key ranges in serializable transactions; PAG indicates that a data or index page is locked; EXT indicates that contiguous group of eight data pages or index pages are locked; TAB indicates that the entire table, including all data and indexes is locked; and DB indicates that the database is locked

Incorrect Answers:

A: Recompiling the procedure will not help. This is a locking problem.

Note: As a database changes when indexes are added or data in indexed columns are changed, the original query plans used to access the database tables should be re-optimized by recompiling them. This optimization occurs automatically the first time a stored procedure is run after SQL Server 2000 is restarted. It also occurs if an underlying table used by the stored procedure changes. However, if a new index is added from which the stored procedure might benefit, optimization does not occur

automatically until the next time the stored procedure is run after SQL Server is restarted. A stored procedure can be forced to be recompiled by specifying the WITH RECOMPILE option when stored procedure is executed. This option should be used only if the input parameter is atypical or if the data has changed significantly since the stored procedure was created.

B: DBCC FREEPROCCACHE is used to clear the procedure cache. This would cause an ad-hoc SQL statement or a stored procedure to be recompiled rather than reused from the cache.

C: Spid 72 holds shared locks on the database. This type of lock does not prevent concurrent access. Therefore it is not necessary to release this lock.

Question No: 7

Exhibit:

```
BULK INSERT testking.dbo.employees
FROM 'tk\africe\fubar512.txt'
WITH DATAFILETYPE = 'char',
FIELD/TERMINATOR = '\t',
ROWTERMINATOR = '\n,'
TABLOCK
```

You work as database administrator at Testig.in. You handle one SQL Server 2000 computer. How should you improve the performance of the query in the exhibit above?

A. Extend the BULK INSERT statement with the BATCHSIZE option set it to equal to 5 percent of the number of rows to be loaded.

B. Extend the BULK INSERT statement with the BATCHSIZE option set it to equal to 10 percent of the number of rows to be loaded.

C. Extend the BULK INSERT statement with the ROWS_PER_BATCH option set it to equal to 5 percent of the number of rows to be loaded.

D. Extend the BULK INSERT statement with the ROWS_PER_BATCH option set it to equal to 10 percent of the number of rows to be loaded.

E. Insert a line before the BULK INSERT statement that drops the indexes for the customers table. Insert a line after the BULK INSERT statement that recreates the indexes for the customers table.

F. The TABLOCK option should be removed.

G. The ROWTERMINATOR line should be removed.

Answer: F.

Explanation: The TABLOCK hint increases the number of locks during the adding process. This is the reason why response time are slows down during this process. By removing the TABLOCK hint the default more granular row-level lock will be used. This would decrease the scope of the locks which would result in less waiting jobs and performance would improve.

Note: The BULK INSERT statement is used to copy a data file into a database table or view in a format specified by the user. The BULK INSERT statement accepts the TABLOCK hint, which allows the user to specify the locking behavior that the BULK INSERT statement should use. TABLOCK specifies that a bulk update table-level lock is taken for the duration of the bulk copy. If TABLOCK is not specified, the default uses row-level locks.

Incorrect Answers:

A, B, C, D: Effect would not be large

E: Some improvement, but not the best answer.

G: Would make the batch to stop working.

Question No: 8

Exhibit:

```
UPDATE TestKingConsolidated
```

```
SET District = 'Mexico'
```

```
WHERE RtvID = 45
```

You work as database administrator at Testig.in. You handle one SQL Server 2000 computer and one database.

When users run the query in the exhibit above the receive the following error message:

Server: Mrh3612, Level 16, Stats 1, Line 1

View or function 'TestKingConsolidated' is not updateable because it contains aggregates:

How should you enable the query to run?

A. Create a non-normalized TestKingConsolidated table and populate the table with data from the base tables.

B. The appropriate users should be granted UPDATE permissions on each base table.

C. To enable updates on the composite tables create an INSTEAD OF trigger on the view.

D. On each base table add a cascading update trigger.

Answer: C.

Explanation: SQL Server 2000 uses INSTEAD OF triggers and partitioned views to enhance the class of updateable views. INSTEAD OF triggers can be created on a view in order to make a view updateable. The INSTEAD OF trigger is executed instead of the data modification statement on which the trigger is defined. This trigger allows the user to specify the set of actions that need to take place in order to process the data modification statement. Thus, if an INSTEAD OF trigger exists for a view on a given data modification statement, the corresponding view is updateable through that statement. If the view is a partitioned view, the view is updateable, subject to certain restrictions.

Incorrect Answers:

A: Creating a table that combines the composite tables that the view is based on would deny us the ability to use the view to return aggregated column values.

B: The error message is not related to permissions but to the updateable of the table.

D: Cascading referential integrity constraints makes it possible to define the actions SQL Server 2000 takes when a user attempts to delete or update a key to which existing foreign keys point. Cascading referential integrity constraints fire AFTER triggers after all of the cascading referential actions directly caused by the original DELETE or UPDATE are performed.

PART IV: Extracting and Transforming Data with SQL Server 2000 (20 questions)

SECTION 1: Set up IIS virtual directories to support XML. (1 question)

Question No: 1

You are the administrator of two Microsoft Windows 2000 computers. One computer is running Internet Information Services (IIS), and the other is running SQL Server 2000. Company partners need to connect by means of the Internet and query data stored on the SQL Server computer. Some of

these partners have computers that do not use Microsoft operating systems or Web browsers. You need to configure the IIS and SQL Server 2000 computers to allow access to data by means of the IIS virtual directory. IIS uses Basic Authentication to allow access to the virtual directory. Only company partners should have access to the SQL Server computer by means of the Internet. The partners should not be allowed to make ad hoc queries. You want to be able to audit successful and failed logins to the SQL Server computer. You want to allow the partners access to the SQL Server computer while keeping your security scheme as simple as possible. Which two actions should you take? (Each correct answer represents part of the solution. Choose two.)

- A. Configure the IIS virtual directory to allow URL queries.
- B. Configure the IIS virtual directory to allow template queries.
- C. Create a new Windows user account.

Create a corresponding Windows Authenticated login on the SQL Server computer. Configure the IIS virtual directory to always use this login when connecting to the SQL Server computer.

- D. Create a Windows Authenticated login on the SQL Server computer for the IIS Internet Guest Account.

Configure the IIS virtual directory to always use the Internet Guest Account when connecting to the SQL Server computer.

- E. Create a Windows Authenticated login for each company partner on the SQL Server computer. Configure the IIS virtual directory to use Windows Integrated Authentication when connecting to the SQL Server computer.

Answer: B, E.

Explanation:

B: The Allow Template Queries option lets you create XML templates that contain SQL statements, which can be used to return data as XML. No browser would be required, though you could certainly use one. A possible non-browser client could be a Visual Basic application for example.

E: If you force the users to authenticate themselves via Basic authentication, IIS has all information required to make the trusted connection to SQL Server. We set up IIS to use Windows Integrated Authentication when connecting to the SQL Server. Windows Authentication logins must be created on the SQL Server as well.

Reference:

INF: Authentication Methods for Connections to SQL Server in Active Server Pages (Q247931)

HOWTO: IIS and SQL Server on Separate Machines with Trusted Connection (Q176379)

HOWTO: Retrieve XML Data with a Template File from a Visual Basic Client (Q271621)

Incorrect answers:

A: URL queries allow execution of SQL queries directly in the URL. This would require a browser however, but not all partners use browsers.

Note: The Allow URL Queries option lets you type Select statements directly in the browser and return the data as XML.

C: The Internet Guest Account, not a specific Windows account, should be used to connect to the SQL Server from the IIS server.

D: We should not allow access to the SQL Server for the IIS Internet Guest account. That would allow all people connecting to the IIS server access to the SQL Server as well..

SECTION 2: Import and export data. Methods include the Bulk Insert task, the bulk copy program, Data Transformation Services (DTS), and heterogeneous queries. (9 questions)

Question No: 1

You are the administrator of a SQL Server 2000 computer. The server contains a database that stores financial data. You want to use Data Transformation Services packages to import numeric data from other SQL server computers. The precision and scale values of this data are not defined consistently on the other servers.

You want to prevent any loss of data during the import operations. What should you do?

A. Use the ALTER COLUMN clause of the ALTER TABLE statement to change data types in the source tables. Change the data types so that they will use the lowest precision and scale values of the data that will be transferred.

B. Use the ALTER COLUMN clause of the ALTER TABLE statement to change data types in the destination tables. Change the data types to reflect the highest precision and scale values involved in data transfer.

C. Set a flag on each DTS transformation to require an exact match between source and destination columns.

D. Set the maximum error count for each DTS transformation task equal to the number of rows of data you are importing. Use an exception file to store any rows of data that generate errors.

E. Write Microsoft ActiveX script for each DTS transformation. Use the script to recast data types to the destinations precision and scale values.

Answer: B.

Explanation: The ALTER COLUMN clause of the ALTER TABLE statement can be used to change the definition of a column and can be used to alter the precision and scale values of the data types. The destination precision should be changed so no data would be lost. The transformation to a data type with higher precision is automatic and would not require any coding.

Note: The precision of the data type refers to the number of digits in a number while the scale is the number of digits to the right of the decimal point in a number. Thus the number 702.85 has a precision of 5 and a scale of 2. The default maximum precision of numeric and decimal data types is 38.

Incorrect Answers:

A: The ALTER COLUMN clause of the ALTER TABLE statement can be used to change the definition of a column and can be used to alter the precision and scale values of the data types. However, we want to prevent any loss of data therefore we need to ensure that the data being imported correlates with the definition of the destination column. If it does not, any constraint that might exist on the destination column will prevent the data from being inserted.

C: Setting DTS transformation to require an exact match between source and destination columns refers to the data in the columns.

D: When the a maximum error count value for a DTS transformation task specifies the sum of row-level errors detected by the Transform Data task and batch failures. When the Max error count value is exceeded, DTS task execution is terminated. Because rows containing errors detected by the DTS transformation task are discarded before batch submission, these errors do not trigger nor count as batch failures. Errors caught at the destination will fail the batch and add one to the error count regardless of how many rows are in the batch.

E: Recasting data types to the destinations precision could result in loss of data. The destination precision should be changed so no data would be lost.

Note: The capabilities of a Data Transformation Service package can be extended by using Microsoft ActiveX scripts that implement the objects, properties, methods, and collections of the DTS object model. With ActiveX scripts the data can be formatted and transformed as it is copied

from its source to the destination location. An ActiveX script transformation applies to tasks that work on the source data on a row-by-row basis and can be used to replace a two-digit state code in the source data with the legal abbreviation of the state in the destination data; set columns to a default of spaces, validate important columns in the source data and skip records that contain invalid data to prevent them from being copied to the destination, and alter the precision and scale values of data as the data copied from the source to the destination.

Question No: 2

You are the administrator of a SQL Server 2000 computer. Each evening after business hours, you perform database maintenance tasks. You create a Data Transformation Services package to populate the data warehouse. MAPI is not installed on the server.

You want to schedule the DTS package to run each night. You want to attend to other duties while the DTS package is executing. You also want to be notified if the DTS package fails.

What should you do?

- A.** Select the DTS packages log package execution to SQL server check box. Create a SQL server event alert to notify you if the package fails.
- B.** Select the DTS packages write completion status to event log check box. Create a SQL server event alert to notify you if the package fails.
- C.** Configure the DTS package to include an on failure precedence constraint and a send mail task. Configure the send mail task to notify you if the package fails.
- D.** Configure the DTS package to include an on failure precedence constraint and a send mail task. Configure the execute process task to execute a not send command that will notify you if the package fails.

Answer: B.

Explanation: When jobs are completed (successfully or not), the subsystem returns a result status (with optional messages) to SQL Server Agent. SQL Server Agent then records the completion status in both the operating system's application log and the job history table in SQL Server and optionally sends an e-mail report or an SQL Server event alert of the job status to the designated operator.

Incorrect Answers:

A: We are interested to when the package has completed. The completion status, not the log package execution, should be used to trap the completion of a package. The execution status concerns Last run time, Last run status, Last run message, etc.

C, D: The Messaging Application Programming Interface (MAPI) isn't installed. MAPI is needed for SQL Server to send e-mail. No e-mail can be send by SQL Server in this scenario.

Question No: 3

You are the administrator of SQL Server 2000 computer named FABSQL2K01. You create a Data Transformation Services package that contains definition for two transform data tasks.


The tasks gets data from two text files named Customers and Sales. The DTS package is configured as shown in the exhibit.

You do not want the DTS package to add any data to the server unless both transform data tasks complete successfully.

You select the use transactions check box of DTS package and the Join transaction of present check box of both transform data tasks. You then select the fail package on step failure check box of both transform data tasks.

You want to use the DTS package to import data from the text files to the server. You want to ensure

that changes are committed if all imported data is successfully transformed.



What should you do?

- A. Select the commit on successful completion check box of the DTS package.
- B. Select the commit transaction on successful completion of this step check box of the customers transform data task.
- C. Select the commit transaction on successful completion of this step check box of the sales transform data task.
- D. Select the commit transaction on successful completion of this step check box of both transform data tasks.

Answer: A.

Explanation: It is possible to set the Commit on successful package completion check box on the Advanced Tab of the DTS Properties dialog Box. This check box specifies that if each individual SQL statement is a

transaction completes successfully, the transaction will automatically be committed; and if the statement has an error, the statement is rolled back. When this check box is cleared, the DTS connections operate in implicit transaction mode and the first SQL statement begins a transaction that remains in effect until DTS commits it or rolls it back. The next SQL statement executed after any commit or rollback starts a new transaction.

Incorrect Answers:


B: The Commit transaction on successful completion of this step check box treats each step in a DTS package as separate. Thus, selecting the Commit transaction on successful completion of this step check box of the customers transform data task would result in this task being committed regardless of the progress of the sales transformation data task.

C: The Commit transaction on successful completion of this step check box treats each step in a DTS package as separate. Thus, selecting Commit transaction on successful completion of this step check box of the sales transform data task would result in this task being committed regardless of the progress of the customer's transformation data task.

D: The Commit transaction on successful completion of this step check box treats each step in a DTS package as separate. Thus, select the commit transaction on successful completion of this step check box of both transform data tasks would result in the two parts being treated as separate and each part being committed regardless of the progress in the other.

Question No: 4

You are the administrator of a SQL Server 2000 computer. You are creating a database named RetailImport to use as an intermediate data store for a data warehouse. Each night you must import daily sales data into the database from SQL Server 2000 computers in 120 locations. After the data is moved into the data warehouse, the tables are truncated. The database schema is shown in the exhibit.



You want to configure the data import process so that you minimize the time needed to import the sales data and administer the database. What should you do?

- A. Use the simple recovery model and the FOR LOAD option to create the database. Create a Data Transformation Services package that uses the BULK INSERT statement to copy the sales data.
- B. Index the foreign key fields in the child tables.
- C. Create a Data Transformation Services (DTS) package that uses a Data Driven Query task to copy the sales data.
- D. Collect the sales data by using a distributed transaction that inserts the data from all 120 retail locations into the Retailimport database in a single transaction.
- E. On the servers in the retail locations, create stored procedures that submit the updates to the RetailImport database each night as SQLXML updategrams.

Answer: A.

Explanation: We want to optimize for speed and at the same time keep administration down. We accomplish this by:

- Use the simple recovery model. Will avoid logging of the bulk import process. This will increase speed.
- Create a DTS package with BULK INSERT to copy sales data. BULK INSERT is a fast way to import data and a DTS package can be reused and requires little administration.

Note: Recovery Model: A SQL Server 2000 provides three recovery models that can be used to recover database data in case of hardware failure or other eventualities that may compromise data integrity. These

recovery models are: the Simple Recovery model, the Full Recovery model and the Bulk-Logged Recovery model. With the Simple Recovery model, the database can be recovered to the point of the last backup but not to the point of failure or to a specific point in time. To do that, choose either the Full Recovery or Bulk-Logged Recovery model. Furthermore, the FOR LOAD clause of the CREATE DATABASE statement is used for compatibility with earlier versions of SQL Server. The database is created with the dbo use only database option turned on, and the status is set to loading. This is not required in SQL Server version 7.0 because the RESTORE statement can recreate a database as part of the restore operation.

The simple recovery model fits this scenario well since the tables are truncated every night.

Note: Data Warehouse: In a data warehouse, data must be organized for rapid access to information for analysis and reporting, as this is the purpose of the data warehouse. Dimensional modeling is used in the

design of data warehouse databases to organize the data for efficiency of queries that analyze and summarize large volumes of data. The data warehouse schema is usually much simpler than the schema of an OLTP system designed using entity-relation modeling. The verification tables used in OLTP systems that are used to validate data entry transactions are not necessary in the data warehouse database because the data warehouse data has been verified before it was posted to the data warehouse database and the data is not expected to change frequently once it is in the data warehouse. Backup and restore strategies also differ in a data warehouse from those necessary for an OLTP system. Much of the data in a data warehouse is unchanging and does not need repetitive backup. Backup of new data can be accomplished at the time of update, and in some situations it is feasible to do these backups from the data preparation database to minimize performance impact on the data warehouse database. Data to be used in the data warehouse must be extracted from the data sources, cleansed and formatted for consistency, and transformed into the data warehouse schema. The data preparation area is a relational database into which data is extracted from the data sources, transformed into common formats, checked for consistency and referential integrity, and made ready for loading into the data warehouse database. Performing the preparation operations in source databases is not an option because of the diversity of data sources and the processing load that data preparation can impose on online transaction processing systems. Furthermore, attempting to transform data in the data source systems can interfere with online transaction processing (OLTP) performance, and the reconciliation of inconsistencies in data extracted from various sources cannot be accomplished until the data is collected in a common database, at which time data integrity errors can more easily be identified and rectified. The data preparation area is a relational database that serves as a general work area for the data preparation operations. It contains tables that relate source data keys to surrogate keys used in the data warehouse, tables of transformation data, and many temporary tables. It also contains the processes and procedures that extract data from source data systems. The data preparation area should also contain the processes that are used to extract the data from the data sources, the processes that transform and cleanse the data, and the processes that load the data to the data warehouse. These processes may be in the form of SQL queries, stored procedures, DTS packages, or documents of manual instructions. As in the development of any database system, the objective is to automate as much of the process as possible and to manage and maintain the automated tools developed. Storing and maintaining the transformation processes in the data preparation area permits the use of standard database backup and restore mechanisms to preserve them.

Incorrect Answers:

B: Indexes are used for optimizing query response times. In this scenario we are not planning to query the database but must instead import data into the database. The use of indexes is thus not appropriate to this scenario.

C: The Data Driven Query task should be selected for insert operations only if the Transform Data task and the Bulk Insert task do not meet the requirements of the application as these tasks are optimized for insert operations. In this scenario data needs to be inserted into the new table and thus the Data Driven Query task is not the recommended utility to use in this scenario.

Note: A Data Driven Query task allows for flexible Transact-SQL based operations to be performed on the source data. For each row in a source rowset, the Data Driven Query task selects, customizes, and executes one of several SQL statements that are specified via a constant return value set in a Microsoft ActiveX script transformation. Data Driven Query tasks can run a large collection of updates, perform normal file maintenance and can customize SQL Server 2000 transformable subscriptions, which are used to distribute and transform incremental changes during replication.

D: This scenario doesn't really call for a using distributed transaction. We simply want to import data from 120 locations at the highest possible speed. Distributed transaction are used in environments where different databases interact dynamically.

Using a single transaction for the whole copy process from 120 locations seems like a bad idea. It would lock the database unnecessarily.

E: This would not be as fast as BULK INSERT.

Note: Extensible Markup Language (XML) is a hypertext programming language used to describe the contents of a set of data and how the data should be output to a device or displayed in a Web page.

Question No: 5

You are the database administrator for a financial services company. Employees enter data 24 hours a day into a SQL Server 2000 database. These employees report slower response times when new account information is gathered from branch offices and added to the database. You currently use the following BULK INSERT statement to add the account information:

```
BULK INSERT finance.dbo.customers
FROM 'd:\bulk\accts143_10141000.txt'
WITH DATAFILETYPE = 'char',
FIELD/TERMINATOR = '\t',
ROWTERMINATOR = '\n,'
TABLOCK
```

You want to ensure that response times do not slow when new account information is added to the database. What should you do?

A. Drop the indexes for the customers table before the data load, and then re-create the indexes after the data load is complete.

B. Remove the TABLOCK option from the BULK INSERT statement.

C. Add the BATCHSIZE option to the BULK INSERT statement and then set the option equal to 10 percent of the number of rows to be loaded.

D. Add the ROWS_PER_BATCH option to the BULK INSERT statement and then set the option equal to 10 percent of the number of rows to be loaded.

Answer: B.

Explanation: The TABLOCK hint increases the number of locks during the adding process. This is the reason why response time are slows down during this process. By removing the TABLOCK hint the default more granular row-level lock will be used. This would decrease the scope of the locks which would result in less waiting jobs and performance would improve.

Note: The BULK INSERT statement is used to copy a data file into a database table or view in a format specified by the user. The BULK INSERT statement accepts the TABLOCK hint, which allows the user to specify the locking behavior that the BULK INSERT statement should use. TABLOCK specifies that a bulk update table-level lock is taken for the duration of the bulk copy. If TABLOCK is not specified, the default uses row-level locks.

Incorrect Answers:

A: The recommendation to perform the bulk copy operation with the indexes in place depends on the amount of data to be copied into the table compared to the amount of existing data already in the table. The large the amount of data to be loaded into a table, relative to the amount of data already in the table, the more efficient it is to drop all indexes on the table, perform the bulk copy operation, and then re-create the indexes after the data is loaded


C: The BATCHSIZE option of the BULK INSERT statement is used to specify that SQL Server 2000 must treat each batch of rows inserted into the table as a separate transaction. This is useful when errors occur in the bulk copy operation, which would cause the transaction to be rolled back. By specifying the BATCHSIZE option, only the current transaction will be rolled back if an error occurs and the destination table still contains the other batches of rows that were successfully inserted.

D: The ROWS_PER_BATCH option of the BULK INSERT statement SQL Server can be used to indicate the approximate number of rows in a batch. SQL Server 2000 then optimizes the load automatically, according to the batch size value, which may result in better performance. Although

accuracy in the clause is not critical, the larger the batch size is, the better the performance of the bulk copy operation.

Question No: 6

You are the administrator of a SQL Server 2000 computer in your company's personnel department. Employee data is stored in a SQL Server 2000 database. A portion of the database schema is shown in the exhibit.



You want to create a text file that lists these data columns in the following format title, FirstName, LastName, WorkPhone, PositionName, DepartmentName.

You want to create the text file as quickly as possible. You do not expect to re-create this file, and you want to avoid creating new database objects if possible.

What should you do?

A. Use the bcp utility to export data from each table to a separate text file. Use format files to select the appropriate columns. Merge the data from each text file into a single text file.

B. Create a view that joins data from all three tables include only the columns you want to appear in the text file. Use the bcp utility to export data from the view.

C. Create a SELECT query that joins the data from the appropriate columns in the three tables. Add an INTO clause to the query to create a local temporary table. Use the bcp utility to export data from the local temporary table to a text file.

D. Create a SELECT query that joins the data from the appropriate columns in the three tables. Add an INTO clause to the query to create a global temporary table. Use the bcp utility to export data from the global temporary table to a text file.

Answer: D.

Explanation: As the columns that we want in the text file are on different tables in the database, we would have to combine the column into a table or view. Because we will not need to reproduce the file, we can we can combine the column from the various tables into a global temporary table by using the SELECT INTO statement. Then we use the bcp utility to export the data from the global temporary file to a text file.

Note: Temporary tables are not stored in the current database; they are stored in the tempdb system. There are two types of temporary tables: local temporary tables, which are visible only to the connection that created them, and global temporary tables, that are visible to all connections. If global temporary tables are not dropped explicitly before the connection that created them disconnects, they are dropped as soon as all other tasks stop referencing them. No new tasks can reference a global temporary table after the connection

that created it disconnects.

The **bcp** utility copies data between an instance of SQL Server 2000 and a data file in a user-specified format. It can be used to transfer data from a SQL Server table to a data file for use in other programs. Furthermore, data can also be transferred out by specifying a SELECT statement. The bcp utility can use global temporary tables, but not local temporary tables.

Incorrect Answers:

A: Creating a global temporary table by using the SELECT statement and using the bcp utility to export data from that table would require less administrative effort than using bcp to export the required columns from each table into a separate text file and then merging those text files into a single text file.


B: Because we would not need to re-create the data file again, it would be easier to use the SELECT statement to create a global temporary table as these tables are dropped automatically when nothing reference them.

C: The bcp utility cannot use local temporary tables.

Note: Temporary tables are not stored in the current database; they are stored in the tempdb system. There are two types of temporary tables: local temporary tables, which are visible only to the connection that created them, and global temporary tables, that are visible to all connections.

Question No: 7

You are the administrator of a SQL Server 2000 computer. The server contains a database named Sales. The database stored information that employees use to analyze sales data. You configure the database by using the default options. You create the database tables as shown in the exhibit.



You import data as a batch process each evening. Users report that the database responds slowly when they run the following query:

```
SELECT Customers.CustomerName, Invoices.InvoiceID,  
Invoices.InvoiceDate, OrderDetail.Quantity,  
Products.Description, Products.SKU, SalesPersons.FirstName,  
SalesPerson.LastName  
FROM Invoices  
INNER JOIN Customers ON Invoices.Customer ID = Customer.CustomerID  
INNER JOIN OrderDetail ON Invoices.Invoice ID =  
OrderDetail.InvoiceID
```

INNER JOIN Products ON OrderDetail.ProductID = Products.ProductID

INNER JOIN SalesPerson ON Invoices.SalesPersonID =

SalesPerson.SalesPersonID

You want to improve the performance of this query. What should you do?

A. Create indexes on the foreign key fields.

B. Create a view to provide the data that the query will use.

C. Create a stored procedure that executes the query.

D. Create a cascading UPDATE and DELETE triggers to the table.

Answer: A.

Explanation: Indexes contain entries describing the exact location in the table of the data values being searched for in the queries. Nonclustered indexes should be used for columns that contain a large number of distinct values; queries that do not return large result sets; columns frequently involved in search conditions of a query that return exact matches; and decision-support-system applications for which joins and grouping are frequently required. Create multiple nonclustered indexes on columns involved in join and grouping operations, and a clustered index on any foreign key columns.

Incorrect Answers:

B: We could create a view to provide the required data, however, because the view will be based on a number of columns, the view would have to be a partitioned view. Before a partitioned view can be implemented, the tables must first be partitioned horizontally. The original table is replaced with several smaller member tables. Users also have to be granted explicit permissions to use a view.

C: A stored procedure will not improve the performance of the SQL batch significantly. The only gain would be that the execution plan wouldn't need to be recompiled.

Note: A stored procedure is a group of Transact-SQL statements compiled into a single execution plan that can return data as OUTPUT parameters, which can return either data such as an integer or character value or a cursor variable; as return codes, which are always an integer value; as a result set for each SELECT statement contained in the stored procedure or any other stored procedures called by the stored procedure; and as a global cursor that can be referenced outside the stored procedure. Stored procedures assist in achieving a consistent implementation of logic across applications.

D: The query will return a result set that comprises a number of rows joined from the various tables. As the query seeks neither to UPDATE nor DELETE data in any of the tables, placing cascading UPDATE and DELETE triggers on the tables would be inappropriate.

Question No: 8

You are the administrator of a SQL Server 2000 computer. You are creating a data transformation services package. As the first step in this process, you need to load data from text files into a database table. These text files contain data on new stores that join your franchise. The text files list the data columns in the following format StoreID, StoreName, Address, City, State, PostalCode, ManagerID, StoreTypeID, FacilityID.

The destination table is configured as shown in the exhibit.

stores	
	StoreID
	storeName
	address
	city
	state
	postalCode

You want to load the data into the table as quickly as possible. What should you do?

When you use SQL profiler to trace the activity of the application, you receive the results as shown in

- A. Use a Bulk Insert Task to read the data into a temporary table. Use an Execute SQL task to import the appropriate data into the destination table.
- B. Create and edit a format file to select the columns you want to import. Use a Bulk Insert Task, and then specify the format file to import the appropriate data into the destination table.
- C. Use a transform data task to import the data. Use Microsoft ActiveX transformation scripts to write the appropriate data to the appropriate columns in the destination table.
- D. Create and edit a format file to select the columns you want to import. Use a transform data task, and then specify the format file to import the appropriate data into the destination table.

Answer: B.

Explanation: The Bulk Insert Task is the fastest mechanism for copying large amounts of data into a SQL Server 2000 table or view. However, transformations cannot be performed on the data while it is moved from the source file to the table or view, as this would slow down the data copy process. Thus the text file must be formatted before the Bulk Insert Task is used. The format file provides the default information used either to bulk copy the data in the data file into an instance of SQL Server 2000 or to bulk copy data out from the table. The format file also provides a way to bulk copy data selectively from a data file to an instance of SQL Server. This allows for the transfer of data to a table when there is a mismatch between fields in the data file and columns in the table. By using a format file, it is possible to bulk copy data into an instance of SQL Server without having to add or delete unnecessary data, or reorder existing data, in the data file.

Incorrect Answers:

A: The Bulk Insert Task is the fastest mechanism for copying large amounts of data into a SQL Server 2000 table or view. It can make use of a format file that allows for the transfer of data to a table when there is a mismatch between fields in the data file and columns in the table. By using a format file, it is possible to bulk copy data into an instance of SQL Server without having to add or delete unnecessary data, or reorder existing data, in the data file. It thus makes the need to first import the data into a temporary table obsolete.

C: The Transform Data task is used to copy data between a source and destination and to optionally apply column-level transformations to the data. The Transform Data task is the most basic implementation of the data pump engine in Data Transformation Services (DTS) and is optimized for insert-based copying and transforming of column-level data between commercial databases, spreadsheets, and text files. However, the Bulk Insert task provides the best performance for bulk inserting text files into SQL Server 2000, although it cannot be used to transform data.

D: The Bulk Insert Task is the fastest mechanism for copying large amounts of data into a SQL Server 2000 table or view. It can make use of a format file that allows for the transfer of data to a table when there is a mismatch between fields in the data file and columns in the table. By using a format file, it is possible to bulk copy data into an instance of SQL Server without having to add or delete unnecessary data, or reorder existing data, in the data file.

Question No: 9

You are the administrator of a new multiprocessor SQL Server 2000 computer. You are migrating data from other database management systems to this server. You need to load data from two data files into a new table on the server. The table is empty and has no indexes.

You want to complete the data load as quickly as possible. What should you do? (Each correct answer represents part of the solution. Choose all that apply.)

- A. Apply a bulk update lock to the table.
- B. Use separate client connections to load the files into the table simultaneously.
- C. Create a clustered index on the table.
- D. Specify the Full Recovery model for the database.
- E. Specify the Bulk-Logged Recovery model for the database.

Answer: A, B, E.

Explanation:

A: Bulk update locks are taken when bulk copying data into a table and either the TABLOCK hint is specified, or the table lock on bulk load table option is set using sp_tableoption. Bulk update locks allow processes to bulk copy data concurrently into the same table while preventing other processes that are not bulk copying data from accessing the table.

In this scenario the bulk update lock would meet the requirement that data should be loaded as quickly as possible.

B: SQL Server 2000 allows data to be bulk copied into a single table from multiple clients in parallel using the bcp utility or BULK INSERT statement. This can improve the performance of data load operations. To bulk copy data into an instance of SQL Server in parallel, the database must be set to Bulk-Logged Recovery, the TABLOCK hint must be specified, and the table must not have any indexes.

E: By specifying the Bulk-Logged Recovery model, the logging for all SELECT INTO, CREATE INDEX, and bulk loading data operations will be minimal and will therefore require less log space while also improving performance.

Incorrect Answers:

C: When data is bulk copied into a single table from multiple clients in parallel using the bcp utility or BULK INSERT statement to improve the performance of data load operations, the database must be set to Bulk-Logged Recovery, the TABLOCK hint must be specified, and the table must not have any indexes.

D: When data is bulk copied into a single table from multiple clients in parallel using the bcp utility or BULK INSERT statement to improve the performance of data load operations, the database must be set to Bulk-Logged Recovery, the TABLOCK hint must be specified, and the table must not have any indexes.

SECTION 3: Develop and manage DTS packages. (2 questions)

Question No: 1

You are the administrator of a SQL Server 2000 computer. The server contains a Data Transformation Services (DTS) package that queries multiple databases and writes the results to a text file. You run this package by using a Microsoft Windows batch file. The batch file uses the dtstrun utility to execute the DTS package.

You want to ensure that connection properties, such as login names and passwords, cannot be read or modified by users. Which two actions should you take? (Each correct answer represents part of the solution. Choose two.)

- A.** Save the DTS package so that it has an owner password.
- B.** Save the DTS package so that it has a user password.
- C.** Encrypt the DTS package details in the command line of the dtstrun utility.
- D.** Store the DTS package in the Meta Data Services repository.
- E.** Store the DTS package as a Microsoft Visual Basic file.

Answer: A, C.

Explanation: In this scenario we will use a windows batch file with the owner password for maximum security.

A: When running the Dtsrun utility on a package you must either specify a user password or a package password. Specify a password for the package to protect sensitive user name and server password information in the package from unauthorized users with owner password. If the package has an owner password, the data is encrypted with the standard encryption API. This protects protect sensitive user name and server password information in the package from unauthorized user.

C: You can encrypt the command prompt options to be executed by the DTS Run utility, allowing you to create an encrypted dtsrun command for later use. This will ensure that login names, password and other command line options cannot be read or modified by users.

Note: When a package is saved to SQL Server or as a structured storage file, DTS package passwords can be assigned to the package. These passwords are used in addition to the Windows

Authentication or SQL Server Authentication passwords to connect to an instance of SQL Server. There are two types of DTS package passwords that can be set: an owner password and a user password. If an owner password is set, the package user needs the password to edit or run the package and if a user password is set, the package users with access to the user password can run the package.

Reference: BOL dtsrun Utility

Incorrect Answers:

B: There is no need to create a user password. The package will be run with an owner password. The owner password will be protected since the command options are encrypted.

D: The Meta Data Services repository is the storage container for the meta data used by Analysis Services. Meta data is stored in tables in a relational database and is used to define the parameters and properties of Analysis server objects. DTS package information can be saved to Meta Data Services. This saves meta data information about the databases referenced in the package, which can be viewed through the DTS Browser.

E: When a DTS package is saved as a Visual Basic file, these files can be used as templates, or starting points, for user-implemented Visual Basic packages. The generated code sets all properties of all objects referenced in the package to the initial values they will have when package execution begins. This includes those that are set to their default values.

Question No: 2

You are the administrator of a SQL Server 2000 computer. You create a Data Transformation Services (DTS) package. The package is stored as a local package on the server. The package exports data from an online transaction processing (OLTP) database system to an online analytical processing (OLAP) database system that is located on a second SQL Server 2000 computer.

You want the package to execute automatically each evening at 8:00 P.M. What should you do?

A. Use SQL Server Enterprise Manager to create a new job

Create a CmdExec job step that runs the dtsrun utility

Schedule the job to run each night at 8:00 P.M.

B. Use SQL Server Enterprise Manager to create a new job

Create a Transact-SQL job step to query data from the OLTP database

Create a second Transact-SQL job step to launch the DTS package

Schedule the job to run each night at 8:00 P.M.

C. Create a batch file that runs the dtsrun utility

Use the Microsoft Windows Task Scheduler to run the job each night at 8:00 P.M.

Configure the Windows Task Scheduler to use the local system account

D. Move the DTS package to the repository

Configure the SQLServerAgent service on the OLAP database server to update the package each night at 8:00 P.M.

Answer: A.

Explanation: Transfers can be defined and stored in a DTS package, which can then be run from a command line using the dtsrun utility (DTSRUN.EXE). CmdExec job steps are operating system commands or executable programs ending with .bat, .cmd, .com, or .exe.

In this scenario we create a job consisting of a single cmdexec job step which runs the DTS package using the dtstrun utility.

Incorrect answers:

B: Only one job, not two, has to be created at the OLTP database server.

C: If the Task Scheduler is running as the local system account it wouldn't be able to access the remote OLAP database server.

D: The DTS export package should be run at the OLTP database server, not at the OLAP database server. The DTS package is exported from the OLTP server to the OLAP server.

SECTION 4: Manage linked servers.

SUBSECTION A: Manage OLE DB Providers. (2 questions)

Question No: 1

You are the administrator of several SQL Server 2000 computers. A data Transformation Services (DTS) package uses native OLE DB providers to transfer data between the servers. Connection details for the servers are specified in .udl files. The .udl files are frequently updated as connection details change.

You want to distribute the DTS package as a file to developers in your company. You want to make sure connection details are available to developers who receive the DTS package.

Which two tasks should you perform? (Each correct answer presents part of the solution. Choose two)

A. Enable the **Always read** properties from UDL file option in the **Connection Properties** dialog box.

B. Disable the **Always read** properties from UDL file option in the **Connection Properties** dialog box.

C. Delete the .udl files and store connection details in the registry by using system data source names.

D. Delete the .udl files, and store connection details in the registry by using user data source names.

E. Make the .udl files available on a network share.

Answer: A, E.

Explanation:

A: To make the connection information available for the developers we must enable the Always read properties from UDL file option.

E: We provide the developers with access to the .udl file by placing the .udl files in a network share and setting appropriate share and NTFS permissions.

Note: In SQL Server 2000, DTS packages can use Microsoft Data Link (.udl) files to create OLE DB connections and resolve the connections in run time. With this feature we can encapsulate the connection properties from a DTS package into a separate file and can edit the connection string in a .udl file instead of the connection properties in a DTS package in situations where connection information such as the server name, login, or even the OLE DB provider may change. When specifying a data link connection, we must first specify whether to load an existing .udl file or create a data link to save with DTS Designer. If we want to use a .udl file rather than saving the data link with DTS Designer, we can create one either from Windows Explorer or during the data link configuration process in DTS Designer. We must then click Microsoft Data Link in the Data Source list of the Connection Properties dialog box and check the Always read properties from UDL file check box. In this scenario, the .udl file is already specified.

Incorrect Answers:

B: If we clear the Always read properties from UDL file option in the Connection Properties dialog box, we would have to edit the connection properties through the Data Transformation Services (DTS) package. This would increase administrative effort.

C: We cannot delete the .udl files and instead store the connection details in the registry by using system data source names as the connection details change regularly.

D: We cannot delete the .udl files, and instead store connection details in the registry by using user data source names as the connection details change regularly.

Question No: 2

You are a database administrator in the Los Angeles branch office of a specialty foods supplier. A mainframe database at the headquarters contains all company data. Each branch office contains a SQL Server 2000 computer that imports regional data from the mainframe database.

The server in Los Angeles contains a Data Transformation Services (DTS) package that uses OLE DB to connect to the company's mainframe database. The DTS package extracts and transforms data about buyers and products for that region. The DTS package then writes the data to the SQL Server database in Los Angeles.

You need to configure a SQL Server computer for the new branch office in Sydney. You want to copy the Los Angeles package and modify it so that it writes data to the SQL Server database in Sydney. You use the DTS Designer to modify the DTS package so that it imports regional data from the mainframe database to the server in Sydney. The modified DTS package runs successfully on the server in Los Angeles. You save this DTS package to the server in Sydney, but the server in Sydney cannot connect to the mainframe database.

You want to enable the server in Sydney to connect to the mainframe database. What should you do?

- A.** Change the connection properties in the DTS package so that the package uses new login credentials to connect to the mainframe database.
- B.** Modify the workflow in the DTS package so that the server in Sydney is included.
- C.** On the server in Sydney, install an OLE DB provider for the mainframe database.
- D.** On the server in Sydney, delete and then re-create the DTS package.

Answer: C.

Explanation: OLE DB is an API that allows COM applications to use data from OLE DB data sources, which includes data stored in different formats. An application uses an OLE DB provider to access an OLE DB data source. An OLE DB provider is a COM component that accepts calls to the OLE DB API and performs the necessary processing required by request against the data source. In this scenario the OLE DB source is the company's mainframe database, which contains all company data. As each branch office contains a SQL Server 2000 computer that imports regional data from the mainframe database, all existing and future branch office SQL Server 2000 servers will require an OLE DB provider to access the company's mainframe database.

Incorrect Answers:

A: The DTS package requires an OLE DB provider to access an OLE DB data source. It needs this provider to connect to the company's mainframe database. Resetting the connection properties in the DTS package so that the package uses new login credentials to connect to the mainframe database will not provide the DTS package with the required access to the mainframe database.

B: A separate DTS package must be created for the Sydney branch office. This DTS package must extract data from the company's mainframe computer that is relevant to the branch office in Sydney. The existing DTS package used to extract data from the company's mainframe for the Los Angeles SQL Server 2000 server can be modified to serve the same function for the Sydney office.

D: Re-creating the DTS package on the SQL Server 2000 server in the branch office in Sydney is unnecessary the that SQL Server 2000 server would still require an OLE DB provider to access the company's mainframe database. It would require less administrative effort to modify the existing DTS package.

SUBSECTION B: Configure security mapping. (0 questions)

SECTION 5: Convert data types. (1 questions)

Question No: 1

You are the administrator of a SQL server 2000 computer. The server contains a database named

Inventory.

Developers at your company upgrade an inventory tracking application. Users report that when they insert new information in the Locations table, the upgrade application returns the following error message 'string or binary data would be truncated. The statement has been terminated.'

When you use SQL profiler to trace the activity of the application, you receive the results as shown in the trace exhibit.

EventClass	TextData	ApplicationName
SQL:BatchCompleted	Set implicit_transactions off SET T...	MS SQLEN
SQL:BatchCompleted	Use Inventory	MS SQLEN
Audit Login	--Network protocol: LPC set quoted	Visual Basic
SQL:BatchCompleted	Insert into locations (Description....	Visual Basic

You examine the design of the locations table as shown in the table design exhibit.

Table Schema

Customers - CustomerID - CompanyName - ContactName - ContactTitle - Address - City - Region - PostalCode - Country - Phone - Fax	Orders - OrderID - CustomerID - EmployeeID - OrderDate - RequiredDate - ShipVia - Freight - ShipName - ShipAddress - ShipCity - ShipRegion	Employees - EmployeeID - LastName - FirstName - Title - TitleOfCourtesy - BirthDate - HireDate - Address - City - Region - PostalCode - Country
Suppliers - SupplierID - CompanyName - ContactName - ContactTitle - Address - City - Region - PostalCode - Country - Phone - Fax - HomePage	Products - ProductID - ProductName - SupplierID - CategoryID - QuantityPerUnit - UnitPrice - UnitsInStock - UnitsOnOrder - ReorderLevel - Discontinued	OrderDetails - OrderID - ProductID - UnitPrice - Quantity - Discount
		Categories - CategoryID - CategoryName - Description - Picture

You need to configure the database to support both versions of the application without affecting server performance. What should you do?

- A. Alter the data type of the description field to varchar(50).
- B. Alter the data type of the Special field to varchar(50).
- C. Alter the data type of the CubicFeet field to float.
- D. In the locations table, create an INSTEAD OF trigger that truncates the description field at 10 characters when the record is updated.

E. In the locations table, create an INSTEAD OF trigger that truncates the Special field at 10 characters when the record is updated.

Answer: A.

Explanation: By examining the first exhibit we see that the Visual Basic application tries to update the Description field. Apparently there is a problem with this field. The description field must be adapted. We increase the field to accommodate for both versions of the application.

Incorrect Answers:

B: The description field, not the special field, is the problem.

C: The error message doesn't indicate a problem in float column, more likely a character column.

D: This would lose data.

E: The description field, not the special field, is the problem.

SECTION 6: Configure, maintain, and troubleshoot replication services. (5 questions)

Question No: 1

You are the administrator of a SQL Server 2000 computer at your company's warehouse. All product orders are shipped from this warehouse. Orders are received at 30 sales offices. Each sales office offers a range of products specific to its region.

Each sales office contains one SQL Server 2000 computer. These servers connect to the warehouse through dial-up connections as needed, typically once a day. Each sales office needs data pertaining only to its region.

You need to replicate inventory data from the server at the warehouse to the servers at the sales offices. You want to minimize the amount of time needed to replicate the data.

Which three actions should you take? (Each correct answer represents part of the solution. Choose three.)

A. Create one publication for each Subscriber.

B. Create one publication for all Subscribers.

C. Enable horizontal filtering.

D. Enable vertical filtering.

E. Use pull subscriptions.

F. Use push subscriptions.

Answer: B, C, E.

Explanation:

B: All subscribers will receive the same type of information, therefore only one publication for all Subscribers is needed.

C: To save bandwidth and connection costs we should only replicate the rows of interest. They require data pertaining only to its region. In a table this data will be located in different rows. Therefore, horizontal filtering is required.

E: The sales office uses the dial-up connections when they need new information from the warehouse. They pull the information from the warehouse.

Note:

The Publisher is a server that makes data available for replication to other servers. Is used to specify which data is to be replicated and can detect which of the data that has been data replicated has changed. It also maintains information about all publications. Usually, any data element that is replicated has a single Publisher, even if it may be updated by several Subscribers or republished by a Subscriber.

Publication **data filtering** has a number of advantages. These includes: minimizing the amount of data sent over the network; reducing the amount of storage space required at the Subscriber; customizing publications and applications based on individual Subscriber requirements; and avoiding or reducing conflicts because

the different data partitions sent to different Subscribers. There are four types of filters that can be applied: horizontal, vertical, dynamic, and join filters. Horizontal and vertical filtering refers to the filtering of row and column respectively. These filters can be used with snapshot, transactional, and merge publications. Horizontal filters, which filter rows, use the WHERE clause of an SQL statement and restrict the rows included in a publication based on specific criteria. Vertical filters, which filter columns, restrict the columns that are included in a publication. Dynamic and join filters extend the capabilities of merge replication. Dynamic filters are row filters that use a function to retrieve a value from the Subscriber and filter data based on that value. The filter is defined once for a publication, but the qualifying result set can be different for each Subscriber and allows the user at a Subscriber to receive only the subset of data customized for their needs. Join filters extend a row filter from one published table to another. A join filter defines a relationship between two tables that will be enforced during the merge process and is similar to specifying a join between two tables.

Push subscriptions simplify and centralize subscription administration, as each Subscriber does not need to be administered individually. Push subscriptions are created at the Publisher, and the replication agents propagate data and updates it to a Subscriber without the Subscriber requesting it. Changes to the replicated data can also be pushed to Subscribers on a scheduled basis.

Push subscriptions should be used when data is typically synchronized on demand or on a frequently recurring schedule; when publications require near real-time movement of data without polling. When the higher processor overhead at a Publisher using a local Distributor does not affect performance; and when easier administration from a centralized Distributor is desired. Pull subscriptions are created at the Subscriber, and the Subscriber requests data and updates made at the Publisher. Pull subscriptions allow the user at the Subscriber to determine when the data changes are synchronized, which can also be on demand or scheduled. Pull subscriptions should be used when the administration of the subscription will take place at the Subscriber; when the publication has a large number of Subscribers; when it would be too resourceintensive

to run all the agents at one site or all at the Distributor; and when Subscribers are autonomous, disconnected, and/or mobile.

Incorrect Answers:

A: Creating one publication per Subscriber is not the best answer. This would increase the processor workload on the Distributor as data changes would need to be tracked to individual publications. It is also a more complex procedure and would require a larger number of Transact-SQL statements to produce. Creating one Publication for all Subscribers and using horizontal filtering would be the better option here.

D: We need horizontal filtering, not vertical filtering, since we want to filter different rows, not different columns, to the different sales offices.

F: Push subscriptions cannot be utilized as the SQL Server 2000 servers connect to the warehouse through dial-up connections as needed. This is usually once a day. Therefore, the subscriber must determine when replication is to be synchronized.

Question No: 2

You are the administrator of a SQL Server 2000 computer. Your company purchased an accounting application from a vendor. The application stores its data in a database named Accounting on the server. The tables in this database contain columns that function as primary keys, but PRIMARY KEY and FOREIGN KEY constraints are not used.

You need to replicate data from this database to another SQL Server computer. This server will use the replicated data to generate reports. Most reports will run each month, but the accounting department needs to have the ability to run reports at any time. Reports should be accurate through the last full working day.

You cannot make any changes to the database, but you need to implement replication. Which two

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

- A. Implement merge replication.
- B. Implement snapshot replication.
- C. Implement transactional replication.
- D. Schedule replication to run continuously.
- E. Schedule replication to run during off-peak hours.

Answer: B, E.

Explanation:

B: Snapshot replication is the best solution since we will not be making changes to the Accounting database but need to replicate data from the database to another SQL Server computer and as the accounting department needs to have the ability to run reports that only need to be accurate through the last full working day.

E: Only one snap-shot replication would be necessary per day, and to make as little impact on the daily use of the server this replication should be scheduled during off-peak hours.

Note: SQL Server 2000 provides three types of replication that can be used in distributed applications. These are snapshot replication, transactional replication and merge replication. Each type provides different capabilities.

Snapshot replication is the process of copying and distributing data and database objects exactly as they were at the time the snapshot was taken. Snapshot replication does not require continuous monitoring of changes because changes made to published data are not propagated to the Subscriber incrementally. Subscribers are updated with a complete data set and not individual transactions. As snapshot replication replicates an entire data set at one time, it may take longer to propagate data modifications to Subscribers. This type of replication is helpful in situations when data is mostly static and does not change often; it is acceptable to have copies of data that are out of date for a period of time; replicating small volumes of data; and when sites are often disconnected and high latency (the amount of time between when data is updated at one site and when it is updated at another) is acceptable.

In **transactional replication**, an initial snapshot of data is propagated to Subscribers, and then when data modifications are made at the Publisher, the individual transactions are captured and propagated to Subscribers. SQL Server 2000 monitors INSERT, UPDATE, and DELETE statements, and changes to stored procedure executions and indexed views. SQL Server 2000 stores the transactions affecting replicated objects and then it propagates those changes to Subscribers continuously or at scheduled intervals. Transaction boundaries are preserved. Transactional replication is used when data modifications must be propagated to Subscribers, often as they occur; transactions must be atomic; subscribers are mostly connected to the Publisher and when the application will not tolerate high latency for Subscribers receiving changes.

Merge replication allows various sites to work autonomously and merge data modifications made at multiple sites into a single, uniform result at a later time. The initial snapshot is applied to Subscribers and then SQL Server 2000 tracks changes to published data at the Publisher and at the Subscribers. The data is synchronized between servers either at a scheduled time or on demand. Updates are made independently at more than one server, so the same data may have been updated by the Publisher or by more than one Subscriber. Therefore, conflicts can occur when data modifications are merged. This type of replication is helpful when multiple Subscribers need to update data at various times and propagate those changes to the Publisher and to other Subscribers; Subscribers need to receive data, make changes offline, and synchronize changes later with the Publisher and other Subscribers; the application latency requirement is either high or low, and when site autonomy is critical.

Incorrect Answers:

A: No data at the second server will not be changed so there is no need to use merge replication.

C: We don't need the accuracy of transactional replication. In this scenario high latency is permissible

as the reports must be accurate only through the last full working day.

Note: Transactional replication is used when data modifications must be propagated to Subscribers, often as they occur; transactions must be atomic; subscribers are mostly connected to the Publisher and when the application will not tolerate high latency for Subscribers receiving changes.

D: As the reports must be accurate through the last full working day, it not necessary to run replication continuously. Running replication continuously will place locks on the destination database, which would hamper the performance of queries and actions run against the destination database.

Question No: 3

You are the administrator of a SQL Server 2000 computer. You want to set up snapshot replication on the server. The server will serve as Publisher and Distributor for a minimum of 40 Subscribers. Currently, you want to publish 3GB of data, but the data is expected to grow over time. Subscribers will receive a new snapshot each month.

You want to minimize the workload on the Publisher/Distributor. Which two actions should you take to configure snapshot replication? (Each correct answer presents part of the solution. Choose two)

- A. Store the snapshot in the default folder on the Publisher/Distributor.
- B. Store the snapshot in an alternative folder on the Publisher/Distributor.
- C. Store the snapshot in a shared folder on a file server.
- D. Create pull subscriptions.
- E. Create push subscriptions.

Answer: C, D.

Explanation:

C: In SQL Server 2000 it is possible to store snapshot files in a location other than or in addition to the default location. Alternate locations can be on another server, on a network share, or on removable media. Saving snapshot files in an alternate location can alleviate disk overhead on the Distributor. However, Subscribers running earlier versions of SQL Server cannot use the alternate snapshot location. When specifying the snapshot location on a network share, it is recommended that that share folder be dedicated for snapshot storage and files. Subscribers that need to access the snapshot files will need Read permission to the snapshot location and appropriate folders. If the snapshot folder is not shared for the appropriate Subscribers, the replication agents cannot access the folder and the replication will fail.

D: With pull subscriptions the Subscriber requests data and updates that are made at the Publisher. Thus allows the user at the Subscriber to determine when the data changes are synchronized. With pull subscriptions the administration of the subscription takes place at the Subscriber. A feature of pull subscriptions is that the Distribution Agent for snapshot and transactional publications and the Merge Agent for merge publications all run at the Subscriber. This can result in a reduction of the amount of processing overhead on the Distributor. However, if needed, the Distribution Agent or Merge Agent can be run at the Distributor to offload agent processing from the Subscriber.

Incorrect Answers:

A: The default location for snapshot replication files are specified in the Specify Snapshot Folder dialog box that appears in the Configure Publishing and Distribution Wizard and in the Create Publication Wizard. The default location is usually on the Distributor.

B: In SQL Server 2000 it is possible to store snapshot files in a location other than or in addition to the default location. Alternate locations can be on another server, on a network share, or on removable media. Saving snapshot files in an alternate location can alleviate disk overhead on the Distributor, particularly when it is located on another server, or a file server.

E: Push subscriptions are created at the Publisher. This places higher processor overhead at a Publisher using a local Distributor as the Distribution Agent for the snapshot publication runs at the

Distributor.

Question No: 4

You are the administrator of a SQL Server 2000 computer at a bank's headquarters. Customers can access their accounts from headquarters or from branch offices. The server at headquarters maintains a master copy of all account data for all customers.

When a customer deposits or withdraws funds at a branch office, the customer's new balance must be reflected on the computers at that office and at headquarters. You want to minimize latency between the headquarters and branch offices.

You want to copy account data between the server at headquarters and the servers at the branch offices. What should you do?

- A.** Implement merge replication. Use the DATETIME (Later Wins) conflict resolver. Schedule replication for off-peak hours.
- B.** Implement transactional replication. Allow immediate-updating subscribers. Schedule replication to run continuously.
- C.** Create a data transformation services package on the server at headquarters. Configure the DTS package to import new account data from each of the servers at the branch offices and write it to all other servers. Schedule the DTS package to run during off-peak hours.
- D.** Add all branch office servers to the headquarters server as linked servers. Create a job on the headquarters server. Configure the job to query each linked server for new account data and write this data to the other servers. Schedule the job to run every three minutes.
- E.** Configure log shipping between the headquarters server and the branch office servers. Configure the headquarters server as the primary server and the branch office servers as secondary servers. Schedule log backups and log shipping to run every 10 minutes.

Answer: B.

Explanation: In this scenario we need the data to be updated centrally as quickly as when it is modified at the branch office. We must therefore implement transactional replication. This will enable immediate updating of Subscribers that must run continuously to ensure low latency, i.e. the amount of time between the change at the Publisher and the update at the Subscriber.

Note: SQL Server 2000 provides the three types of replication that can be used for distributed applications. These are snapshot replication, transactional replication and merge replication. Each type provides different capabilities. Snapshot replication distributes data exactly as it appeared when the snapshot was taken. This type of replication is best used as a method for replicating data that change infrequently or where the most up-to-date values are not a requirement. Merge replication allows users to work and update data autonomously, i.e., while they are disconnected. When servers are reconnected, all sites in the replication topology converge to the same data values. Transactional replication maintains transactional consistency, but Subscriber sites are not as autonomous as they are in merge replication because with this type of replication, an initial snapshot of data is applied at the Subscribers, and then, as data modifications are made at the Publisher, the individual transactions are captured and propagated to Subscribers.

Incorrect Answers:

A: With merge replication users work and update data autonomously, i.e., while they are disconnected. When servers are reconnected, all sites in the replication topology converge to the same data values. This type of replication should be used where high latency is not an issue.

C: In this scenario we need to minimize latency, in other words we need to ensure that the updates occur close to real-time as possible. Scheduling the package to run during off-peak hours will not reduce latency.

Note: A DTS package is a collection of connections, DTS tasks, DTS transformations, and workflow constraints assembled either with a DTS tool or a programming package and saved to SQL Server 2000, SQL Server 2000 Meta Data Services, a structured storage file, or a Microsoft Visual Basic

file. Each package contains the steps that are executed sequentially or in parallel when the package is run. When executed, the package connects to the correct data sources, copies data and database objects, transforms data, and notifies other users or processes of events. Packages can be edited, password protected, scheduled for execution, and retrieved by version

D: In this scenario we need to minimize latency, in other words we need to ensure that the updates occurs close to real-time as possible. Scheduling the job to run every 3 minutes intervals will not minimize latency to the fullest as it will allow a latency of at least 3 minutes.

Note: Linked servers allow us to run distributed queries. Each distributed query can reference multiple linked servers and can perform either update or read operations against each individual linked server. A single distributed query can perform read operations against some linked servers and update operations against other linked servers.

E: In this scenario we need to minimize latency, in other words we need to ensure that the updates occurs close to real-time as possible. Therefore, scheduling the log backups and log shipping to run every 10 minutes intervals will not minimize latency to the fullest as it will allow a latency of at least 10 minutes.

Note: In SQL Server 2000 Enterprise Edition, we can use log shipping to copy transaction logs from one database to another on a constant basis. Continually backing up the transaction logs from a source database and then copying and restoring the logs to a destination database keeps the destination database synchronized with the source database

Question No: 5

You are the administrator of a 12 SQL Server 2000 computers. These servers are located in company offices throughout the world.

You configure a multiserver job on the server in your company's New York City office. You configure this server as a master server, and designate the remaining 11 servers as target servers.

After the job completes successfully, you decide to set up a new SQL Server 2000 computer as the master server. The new server will be named JobMaster and the other 12 servers will be target servers.

You need to make the New York City server a target server instead of a master server. What should you do?

A. On the New York City server, defect all 11 target servers.

On the JobMaster server, enlist the 12 servers as target servers.

B. On the JobMaster server, enlist all other servers as target servers.

C. On the New York City server enlist the JobMaster server as a target server.

Transfer the multiserver jobs from the New York City server to the JobMaster server.

D. On the JobMaster server, enlist the New York City server as a target server.

On the JobMaster server, enlist the other 11 servers as target servers.

On the JobMaster server, defect the New York City server.

Answer: A.

Explanation: When setting up a new master server, the old master server must defect all its target servers. When a newly designated master server replaces an existing master server, the target servers must be defected from the existing master server. The target servers can then be enlisted on the new master server. This can include the existing target server.

Note: In SQL Server 2000 Enterprise Edition, multiserver administration is used to automate the administration of multiple instances of SQL Server 2000. It can be used to manage two or more servers and to schedule information flows between enterprise servers for data warehousing. With multiserver administration, there must be at least one master server and at least one target server. A master server distributes jobs to and receives events from target servers. A master server stores the central copy of job definitions for jobs run on target servers. Target servers connect periodically to their master server to update

their list of jobs to perform. If a new job exists, the target server downloads the job and disconnects from the master server. After the target server completes the job, it reconnects to the master server and reports the status of the job. With SQL Server 2000 Enterprise Edition, master servers can enlist target servers and can defect target servers.

Incorrect Answers:

B: This solution does not address the need to demote the New York City server. It is thus still a master server.

C: In this solution the JobMaster server is enlisted as a target server of the New York City server. This solution also does not address the need to demote the New York City server. It is thus still a master server.

D: In this solution the New York City server is not a target server of the JobMaster server as it has been defected from the JobMaster server.

PART V: Managing and Monitoring SQL Server 2000 Security (38 questions)

SECTION 1: Configure mixed security modes or Windows Authentication.

Considerations include client connectivity, client operating system, and security infrastructure. (9 questions)

Question No: 1

You are the administrator of SQL Server computers. One server is named SQL7, and other is named SQL2000. SQL7 is running SQL Server 7.0 and SQL2000 is running SQL Server 2000.

The net-libraries on SQL2000 are configured as shown in the exhibit.

SQL7 is configured so that it has the named pipes, TCP/IP, NWlink, IPX/SPX and multiprotocol Net Libraries.

SQL2000 and SQL7 exchange confidential company information. You need to ensure that unauthorized users cannot access this information.

Which two actions should you take? (Each correct answer presents part of the solution. Choose two)

- A.** On SQL2000, enable the multiprotocol net library.
- B.** On SQL2000, select the force protocol encryption check box.
- C.** On SQL7, select the force protocol encryption check box.
- D.** On SQL2000, install a secure sockets layer (SSL) encryption certificate.
- E.** On SQL2000 and SQL7, enable multiprotocol encryption.

Answer: A, E.

Explanation:

For compatibility with earlier versions of SQL Server, the Multiprotocol Net-Library continues to support its own encryption. This fits this scenario with a SQL Server 7 and a SQL Server 2000.

Note: Network libraries are used to pass network packets between the clients and the server running SQL Server 2000. The network libraries, implemented as dynamic-link libraries (DLLs), perform the network operations required to communicate using specific interprocess communication (IPC) mechanisms. The Multiprotocol Net-Library uses the Windows NT remote procedure call (RPC) facility and can communicate over most IPC mechanisms supported by Windows NT such as TCP/IP Sockets, NWLink IPX/SPX, and Named Pipes; allow the use of Windows Authentication over all protocols that RPC supports; support encryption for user password authentication as well as data; and can offer performance comparable to native

IPC Net-Libraries for most applications. The Multiprotocol encryption is used only for compatibility purposes with existing systems but is not supported on Windows 98. The Secure Sockets Layer (SSL) encryption that can be enabled using the Server Network Utility is a more comprehensive encryption solution. However, SSL encryption works only with instances of SQL Server 2000 running on a computer

that has been assigned a certificate from a public certification authority. The computer on which the application is running must also have a root CA certificate from the same authority.

Incorrect Answers:

B: You can force protocol encryption using SSL. Here Multiprotocol encryption must be used for compatibility purposes. SSL requires a server certificate and is not supported by SQL Server 7.0.

C: You can force protocol encryption using SSL. Here Multiprotocol encryption must be used for compatibility purposes. SSL requires a server certificate and is not supported by SQL Server 7.0.

D: The Secure Sockets Layer (SSL) encryption that can be enabled using the Server Network Utility is a more comprehensive encryption solution than Multiprotocol encryption. However, SSL encryption works only with instances of SQL Server 2000 running on a computer that has been assigned a certificate from a public certification authority. The Multiprotocol encryption must be used for compatibility purposes.

Question No: 2

You are the administrator of a SQL Server 2000 computer. The server is a member of a Windows NT domain and is configured for Windows authentication.

The server contains a database that stores contact information for public officials in your region.

These officials need to access the database by means of the Internet by using the guest login. However, some users report that they cannot connect to the server by means of the Internet.

You need to allow anyone access to the server by means of the Internet. What should you do?

A. Assign the guest login a blank password.

B. Delete the guest login and create a new login that is mapped to the guest domain user account.

C. Create a database user named Anonymous and assign the user the appropriate database permissions.

D. Configure the server for mixed mode authentication.

Answer: D.

Explanation: Some users are not able to connect to the SQL Server through internet. The SQL Server is configured for Windows authentication. The non Windows clients are not able to use Windows Authentication. In order to let these clients can access to the SQL server we must configure the server for mixed mode authentication. The non-Windows clients would then be able to access the SQL Server through the SQL Server guest login.

Note: You can connect to an instance of SQL Server 2000 over the Internet using SQL Query Analyzer or a client application based on ODBC or DB-Library. To share data over the Internet, the client and server must be connected to the Internet. In addition, the TCP/IP or Multiprotocol Net-Libraries must be used. If the Multiprotocol Net-Library is used, ensure that TCP/IP support is enabled. If the server is registered with Domain Name System (DNS), users can connect to it using its registered name. The guest user account allows a login without a user account to access a database. A login assumes the identity of the guest user when the login has access to an instance of SQL Server 2000 but does not have access to the database through his or her own user account and the database contains a guest user account. Permissions can be applied to the guest user as if it were any other user account. Mixed Mode allows users to connect to an instance of SQL Server using either Windows Authentication or SQL Server Authentication. Users who connect through a Windows NT 4.0 or Windows 2000 user account make use of trusted connections in either Windows Authentication Mode or Mixed Mode. SQL Server Authentication is provided for backward compatibility. When a user connects with a login from a nontrusted connection, SQL Server performs the authentication itself by checking to see if a SQL Server login account has been set up. As we are not able to control where users login in SQL Server 2000 from using the guest account, we cannot set the server for Windows Authentication as users who login from machines running operating systems other than Windows NT 4.0 or Windows 2000 would not be able to use this login.

Incorrect Answers:

A: By default, the guest login account does not use a user account and does not have a password

assigned to it. Thus it is not a password error that is preventing the guest users from accessing the SQL Server 2000 server.

B: By deleting the guest login, users without a user account will not be permitted access to the database. Creating a new account that is mapped to the guest domain login account, users will have access to the Windows 2000 domain if they are able to provide the credentials associated with the login account that is mapped to the guest domain account and if SQL Server 2000 is running in Mixed Mode authentication.

C: By creating an Anonymous user account and providing the account with the required permissions to the database, users will be able to access the database if they are able to provide the credentials for the account and if SQL Server 2000 is running in Mixed Mode authentication.

Question No: 3

You are the administrator of a SQL Server 2000 computer. Your company has 150 client computers that are running UNIX. These client computers requires access to the server to use a database named Techinfo. The server is configured with the TCP/IP and named pipes net-libraries.

You need to configure the server to allow the client computers to use the Techinfo database. You also want to minimize the amount of time it takes to configure the server and allow for the tracking of individual user actions.

Which two actions should you take? (Each correct answer presents part of the solution. Choose two)

- A. Configure Windows authentication.
- B. Configure mixed mode authentication.
- C. Create a SQL server login for each user.
- D. Create a domain user group login for all of the users.
- E. Create single SQL server login for all of the users.
- F. Enable the multiprotocol Net-Library.
- G. Disable the named pipes net-library.

Answer: B, C.

Explanation:

B: Only clients that use authentication through Windows 2000 or Windows NT are able to use Windows authentication. UNIX clients require authentication through the SQL server, not the NT domain. We must use mixed mode authentication to be able to SQL Server authentication.

C: When SQL Server authentication is used a SQL server login must be created for each user.

Note: SQL Server 2000 can operate in Windows Authentication Mode, which allows a user to connect through a Windows NT 4.0 or Windows 2000 user or user group account, or Mixed Mode, which allows users to connect to an instance of SQL Server using either Windows Authentication or SQL Server Authentication. SQL Server Authentication is provided for backward compatibility because applications written for SQL Server version 7.0 or earlier may require the use of SQL Server logins and passwords, and for users who connects with a specified login name and password from a nontrusted connection. When the latter occurs, SQL Server 2000 performs the authentication itself if the SQL Server 2000 server is set to use Mixed Mode authentication. If SQL Server does not have a login for the account that the nontrusted connection user specifies, the authentication fails and the user receives an error message. Windows Authentication can only be used with Windows NT or Windows 2000 computers as the Windows NT operating system makes provision for trust relationships only amongst Windows NT based computes. As the company uses nontrusted UNIX computers to connect to the SQL Server 2000 server, the SQL Server 2000 server must be configured for Mixed Mode Authentication. We must therefore create SQL Server 2000 login for each UNIX user. The sp_addlogin stored procedure is used for this purpose. When using the sp_addlogin stored procedure we must specify the user account and the password that is to be used with the account.

Incorrect Answers:

A: Only users authenticated through Windows 2000 or Windows NT can login to the SQL Server 2000

server using Windows Authentication. This type of authentication cannot be extended to UNIX users. For UNIX users, the SQL Server 2000 server must be configured for Mixed Mode Authentication.

D: In Windows NT or Windows 2000 network environments it is possible to create network domain accounts for trusted relations. Windows NT or Windows 2000 makes provision for trust relations

among Windows client computers. Hence, domain accounts can only be used for Windows clients authenticated through Windows NT 4.0 or Windows 2000 servers.

E: A single SQL Server login account cannot be created for all users as SQL Server 2000 server does not permit simultaneous logins on the same SQL Server login account. We must therefore create a login for each UNIX user.

F: Net-Libraries are used for network communication between clients and a server running SQL Server, using specific interprocess communication (IPC) mechanisms. The Multiprotocol Net-Library uses the Windows NT remote procedure call facility. It communicates over most IPC mechanisms supported by Windows NT; supports encryption for user password authentication as well as data and offers performance comparable to native IPC Net-Libraries for most applications. However, UNIX computer can use TCP/IP Net-Library for network communication.

G: Although the UNIX computers will use TCP/IP Net-Library for network communication, the Named Pipes Net-Library will not affect it as a SQL Server 2000 server can monitor multiple network libraries simultaneously.

Note: Network libraries are used to pass network packets between clients and a server running SQL Server. The network libraries, implemented as dynamic-link libraries (DLLs), perform the network operations required to communicate using specific interprocess communication (IPC) mechanisms. The Named Pipes Net-Library is required on Windows NT and Microsoft Windows 2000 installations of SQL Server. Server-side Named Pipes is not supported on Windows 98. By default, SQL Server listens on the standard pipe for Named Pipes Net-Library connections

Question No: 4

On the server, you create a new database named EmployeeData. The EmployeeData database will store confidential information about company employees.

You need to allow only authenticated users to access the EmployeeData database. Which Transact SQL statement should you execute?

- A. EXEC sp_revokelogin 'guest'
- B. EXEC sp_droplogin 'guest'
- C. EXEC sp_revokedbaccess 'guest'
- D. EXEC sp_addrolereader 'db_denydatareader', 'guest'

Answer: C.

Explanation: The guest user account permits a login without a user account to access a database. The login assumes the identity of the guest user when the login has access to an instance of SQL Server 2000 but does not have access to the database through his or her own user account and the database contains a guest user

account. Permissions can be applied to the guest user as with any other user account. The guest user can be deleted and added to all databases except master and tempdb, where it must always exist. The sp_revokedbaccess stored procedure removes a security account from the current database. This automatically removes the permissions and aliases that depend on the account as well.

Incorrect Answers:

A: The sp_revokelogin applies to the whole server, but we only want to prevent access to a particular database.

Note: The sp_revokelogin stored procedure removes the login entries from SQL Server 2000 for a Windows NT user or group that was created with the sp_grantlogin stored procedure or the

sp_denylogin stored procedure. This prevents Windows 2000 users from connecting to a SQL Server 2000 server through their Windows 2000 user accounts. However, the sp_revokelogin stored procedure is applied only to the specified Windows 2000 user. If the user has access to the SQL Server 2000 server through their membership to a Windows 2000 group, if that group has been granted access to SQL Server 2000 server.

B The sp_droplogin stored procedure removes a SQL Server 2000 login and prevents access to the SQL Server 2000 server via that login name, but we only want to prevent access to a particular database.

D The sp_addrolereader stored procedure specified in this solution does not exist. Furthermore, the argument in this solution is incorrect from the syntax of a stored procedure.

Question No: 5

You are the administrator of a SQL Server 2000 computer. The server is a member of a Microsoft Windows NT domain named CORP. The server is configured for Windows Authentication. Several users can connect to the server by using SQL Query Analyzer. These users belong to a domain user group named DataUsers, and the server contains a login for CORP\DataUsers.

Your company purchases a client/server application that implements its own users security and will access the server by using a login named Csapp1.

You create a domain user account named Csapp1 and add it to the DataUsers domain user group.

However, when you try to run the client/server application, you receive the following error message: "The required SQL Server database cannot be accessed."

You need to ensure that the application runs properly. What should you do?

A. Remove the Csapp1 domain user account from the DataUsers domain user group. Create a login for CORP\Csapp1.

B. Configure the server for Mixed Mode authentication. Create a login named Csapp1, and configure it to access the application's database.

C. Create a local user account on the server, and name it Csapp1. Add a login for BUILTIN\Csapp1.

D. Configure an application role in the application's database, and name the role Csapp1. Set the application role password so that it is the same as the Csapp1 user account password.

Answer: B

Explanation: The new client/server application will implement its own security. It will access SQL Server using a login named Csapp1. SQL Server logins are not supported in Windows Authentication however and currently the SQL server is configured for Windows Authentication. Authentication through logins is only supported in mixed-mode authentication. We must therefore configure the server for Mixed Mode and create a SQL Server login named Csapp1 with access to the database of the application.

Note: Getting access to a SQL Server database is a two-step process. First you must get authenticated by SQL server, and then you must be authorized access to databases.

Incorrect Answers:

A: Logins are only supported in mixed-mode authentication.

C: The login must be created on the SQL server not in Windows as a local user account.

D: The application will try to connect by a login name CSapp1, but the SQL Server is configured for Windows Authentication so this will not work.

Question No: 6

You are the administrator of a new SQL Server 2000 computer. Your company has client computers that are running Microsoft Windows 2000 Professional. All client computers require access to the server.

According to company policy, passwords must be at least eight characters long and change every 30 days.

You need to implement these policies and allow client computers to connect. What should you do?

(Each correct answer presents part of the solution. Choose all that apply.)

- A. Configure Mixed Mode authentication.
- B. Configure Windows Authentication.
- C. Create a SQL Server login for each user who will access the server.
- D. Create a login for the Domain Users group.
- E. Set the Windows 2000 domain policy for password expiration to 30 days.
- F. Set the Windows 2000 domain policy to require a password length of at least 8 characters.
- G. Create a scheduled Job that changes the password and runs every 30 days.

Answer: B, D, E, F.

Explanation:

B: Windows 2000 clients are able to use secure connection using Windows authentication.

D: Windows 2000 user accounts can be put to a group and a login could be created for this group. The group would be authorized to selected databases with appropriate permissions through this login.

E: We must configure a Windows 2000 account policy for password expiration to 30 days.

F: We must configure a Windows 2000 account policy to require a password length of at least 8 characters.

Note: SQL Server 2000 can operate in Windows Authentication Mode, which allows a user to connect through a Windows NT 4.0 or Windows 2000 user or user group account, or Mixed Mode, which allows users to connect to an instance of SQL Server using either Windows Authentication or SQL Server Authentication. SQL Server Authentication is provided for backward compatibility because applications written for SQL Server version 7.0 or earlier may require the use of SQL Server logins and passwords and for users who connects with a specified login name and password from a nontrusted connection. With Windows Authentication, when a user connects through a Windows NT 4.0 or Windows 2000 user account, SQL Server revalidates the account name and password by calling back to Windows NT 4.0 or Windows 2000 for the information. We would use the sp_grantlogin stored procedure to authorize the Windows NT or Windows 2000 domain user group account to be used as a SQL Server login to connect to a SQL Server 2000 server and would have to specify the qualified Windows NT or Windows 2000 domain name with the Windows NT or Windows user group name in the form Domain\User.

Incorrect Answers:

A: There is no need of mixed mode here since there are only Windows 2000 clients.

C: It would require less administrative effort to put the selected Windows user into a group and only create one login for that group.

G: We use Windows authentication, not SQL Server logins, so we must configure the password in the Windows environment, not in SQL Server.

Question No: 7

You are the administrator of a SQL Server 2000 computer named DataSvr. The server is a member of a Microsoft Windows NT domain. DataSvr is currently configured for Windows Authentication and contains logins for several domain user groups. The server is configured to use Named Pipes Net-Library.

Your company purchases five UNIX client computers. You want to allow employees who will be using these computers to access DataSvr.

What should you do?

- A. Configure DataSvr for Mixed Mode authentication.
Add a login for a domain user group that includes the five UNIX computers.
- B. Configure DataSvr for Mixed Mode authentication.
Add a login for each of the five UNIX users.
- C. Install the TCP/IP Net-Library.

Add the five UNIX users to a domain user group that has access to DataSvr.

D. Install the TCP/IP Net-Library.

Add a login for each of the five UNIX users.

Answer: B.

Explanation: We must use Mixed Mode authentication since UNIX computers cannot use Windows authentication. We must add a login for each of the five UNIX users.

Note: SQL Server 2000 can operate in Windows Authentication Mode, which allows a user to connect through a Windows NT 4.0 or Windows 2000 user account, or Mixed Mode, which allows users to connect to an instance of SQL Server using either Windows Authentication or SQL Server Authentication. SQL Server Authentication is provided for backward compatibility because applications written for SQL Server version 7.0 or earlier may require the use of SQL Server logins and passwords, and for users who connects with a specified login name and password from a nontrusted connection. In this event SQL Server performs the authentication itself if the SQL Server 2000 server is using Mixed Mode authentication. If SQL Server does not have a login account set, authentication fails and the user receives an error message.

Incorrect Answers:

A: Windows domain user groups can only be specified for Windows NT and Windows 2000 users. It cannot be specified for UNIX users.

C: Installing the TCP/IP Net-Library would allow the UNIX users to communicate with the SQL Server 2000 server using this Net-Library only if they are able to logon to the SQL Server 2000 server. However, the SQL Server 2000 server is setup for Windows Authentication. This type of authentication makes use only of Windows NT and Windows 2000 domain user accounts for SQL Server 2000 logon purposes. This type of account cannot be extended to UNIX users therefore they will not be able to logon to the SQL Server 2000 server.

D: Installing the TCP/IP Net-Library would allow the UNIX users to communicate with the SQL Server 2000 server using this Net-Library only if they are able to logon to the SQL Server 2000 server. Although a login for each UNIX user has been created, the SQL Server 2000 server is setup for Windows Authentication. This type of authentication makes use only of Windows NT and Windows

2000 domain user accounts for SQL Server 2000 logon purposes, therefore the UNIX users will not be able to logon to the SQL Server 2000.

Question No: 8

You are the administrator of a SQL Server 2000 computer. Your company has client computers that are running Microsoft Windows 2000 Professional and UNIX. All client computers require access to the server.

Users with domain accounts will access the server by means of the company LAN, and external users will access the server by means of the Internet.

You need to configure the server so that all authorized users can connect to it. What should you do? (Each correct answer presents part of the solution. Choose all that apply.)

A. Configure Windows Authentication.

B. Configure Mixed Mode Authentication.

C. Create a login for the Domain Users group.

D. Create a login for the Guest domain user account.

E. Create a login for each UNIX user, and set a password for each login.

F. Create a login for each external user, and set a blank password for each login.

G. Create a login named Guest, and set a blank password for the login.

Answer: B, C, E.

Explanation:

B: We need to use Mixed Mode Authentication, not Windows Authentication, since there are UNIX clients and they are not able to use Windows authentication.

C: We use Windows authentication for the Domain Users, and simply map each to a login. No password required.

E: We must create a SQL login, which includes user name and password, for every UNIX user.

Note 1: Mixed mode authentication includes both SQL Server authentication and Windows authentication.

Note 2: SQL Server 2000 can operate in Windows Authentication Mode, which allows a user to connect through a Windows NT 4.0 or Windows 2000 user or user group account, or Mixed Mode, which allows users to connect to an instance of SQL Server using either Windows Authentication or SQL Server Authentication. SQL Server Authentication is provided for backward compatibility because applications written for SQL Server version 7.0 or earlier may require the use of SQL Server logins and passwords, and for users who connects with a specified login name and password from a nontrusted connection. When the latter occurs, SQL Server 2000 performs the authentication itself if the SQL Server 2000 server is set to use

Mixed Mode authentication. If SQL Server does not have a login for the account that the nontrusted connection user specifies, the authentication fails and the user receives an error message. As some users are connecting to the SQL Server 2000 server from nontrusted UNIX computers, the SQL Server 2000 server must be configured for Mixed Mode Authentication. Mixed mode Authentication will allow us to create a SQL Server 2000 server login for the Windows NT domain users group, which would connect to the SQL Server 2000 server via Windows Authentication. We would use the `sp_grantlogin` stored procedure to authorize the Windows NT user group account to be used as a SQL Server login to connect to a SQL Server 2000 server using Windows Authentication and would specify the qualified Windows NT domain name with the Windows NT user group name in the form `Domain\User`. However, the Windows NT domain user group cannot include UNIX users as the Windows NT operating system only make provision for the establishment of trust relationships with Windows NT and Windows 2000 users. We must therefore create SQL Server 2000 login for each UNIX user. The `sp_addlogin` stored procedure is used for this purpose. When using the `sp_addlogin` stored procedure we must specify the user account and the password that is to be used with the account.

Incorrect Answers:

A: UNIX clients and they are not able to use Windows authentication. We must use SQL Server logins and they are only available in Mixed mode authentication.

D: We are not required to supply anonymous access through a guest account. This could be a security breach.

Note: The Guest domain user account is a built-in user account in Windows NT computers. People who do not have a domain account use it to access the Windows NT server. A user whose account is disabled can also use the Guest domain account. The Guest domain account does not require a password and rights and permissions for the Guest account can be set like any user account. Creating a SQL Server 2000 logon for the Windows NT domain guest account would permit UNIX users to access the SQL Server 2000 server through the Windows NT guest account. However, other non authorized users may also gain access to the server through the domain guest account as the domain guest account usually does not have a password.

F: We don't want to give external users access to the server by creating user accounts with blank passwords.

G: We are not required to supply anonymous access through a guest account. This could be a security breach.

Question No: 9

You work as database administrator at Testig.in. You handle one SQL Server 2000 computer.

The situation is characterized by:

- the server is a member of a Microsoft Windows NT domain
- Named Pipes Net-Library is configured on the server

• **Ten UNIX client computers are to be used to access the server**
How should you configure the server? (Select all that apply.)

- A. First configure server for Mixed Mode authentication.
- B. First configure server for Windows Server authentication.
- C. Add a login for a domain user group that includes the ten UNIX computers.
- D. Add the ten UNIX users to a domain user group that has access to the Server.
- E. Add a SQL login for each of the then UNIX users.
- F. Install the TCP/IP Net-Library.

Answer: A, E.

Explanation: We must use Mixed Mode authentication since UNIX computers cannot use Windows authentication. We must add a login for each of the ten UNIX users.

Incorrect Answers:

B: We must use Mixed Mode authentication.

C, D: Windows domain user groups can only be specified for Windows NT and Windows 2000 users. It cannot be specified for UNIX users.

F: Not required.

SECTION 2: Create and manage log ons. (9 questions)

Question No: 1

You are the administrator of a SQL Server 2000 computer. The server contains a database named FinanceData that is used by a client/server application. Each employee has a SQL server login that the application uses to access the server.

Your company hires a new employee named Andrew. You create a SQL server login named Andrew. You also create a database user named Andrew in the database and grant the database user full permissions in the database.

Whenever Andrew runs the client/server application, he receives the following error message 'Invalid object name'. Andrew can also use the application. Which transact-SQL statement should you execute?

A. EXEC sp_grantdbaccess 'FinanceData', 'Andrew'

GO

sp_addrolemember 'db_datareader', 'Andrew'

B. GRANT ALL ON FinanceData TO Andrew.

C. EXEC sp_defaultdb 'Andrew', 'FinanceData'

D. EXEC sp_addlogin 'Andrew'

Answer: C.

Explanation: Error 208 produces the message 'Invalid object name' and occurs when an object that does not exist in the current database is referenced. If the object exists in another database, we must use the USE statement to explicitly switch the context to the correct database or we must qualify the object name with the database name. A third option is to use the sp_defaultdb stored procedure. This stored procedure is used to change the default database for a login. When a client connects with SQL Server, the default database defined for its login becomes the current database without an explicit USE statement. The default database can be defined when the login is added with sp_addlogin. When executing sp_addlogin the master database is the default database if a database is not specified. In this scenario a SQL Server 2000 login and a user account has been created for Andrew and he has been given database access and the required permissions.

Incorrect Answers:

A: A security account must be granted access to the database before it can use the database. The sp_grantdbaccess stored procedure is used for this purpose and adds a security account in the

database for a SQL Server 2000 login or Microsoft Windows NT user or group, and enables it to be granted permissions to perform activities in the database. The sp_addrolemember stored procedure is used to add a security account as a member of an existing SQL Server 2000 database role in the database, conferring on the account any permissions applied to the role. While the db_datareader fixed database role allows its members the ability to select all data from any user table in the database. In this scenario a SQL Server 2000 login and a user account has been created for Andrew and he has been given database access and the required permissions.

B: The GRANT statement creates an entry in the security system that allows a user in the current database to work with data in the current database or execute specific Transact-SQL statements while the ALL clause specifies that all applicable permissions be granted. In this scenario a SQL Server 2000 login and a user account has been created for Andrew and he has been given database access and the required permissions.

D: The sp_addlogin stored procedure is used to create a new SQL Server 2000 login that allows a user to connect to an instance of SQL Server using SQL Server Authentication. If the name of a default database is supplied, the user can connect to the specified database without executing the USE statement. In this scenario a SQL Server 2000 login and a user account has been created for Andrew and he has been given database access and the required permissions.

Question No: 2

You are the administrator of a SQL Server 2000 computer. The server is used to store information for your company's sales department.

The permissions on the SalesFigures table are configured as shown in the exhibit.

Users/Database Role...	SELECT	INSERT	UPDATE	DELETE
 Accountants	☒	☒	☒	☒
 Managers	☒	☐	☒	☒
 Salespeople	☒	☐	☐	☐
 guest	☐	☐	☐	☐

Lilly is a user in the sales department. She needs to review the data in the SalesFigures table. She also needs to add new items to the table.

You need to allow Lilly to perform these tasks without giving her additional permissions in the database.

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

- A.** Add Lilly to the Managers database role.
- B.** Add Lilly to the Salespeople database role.
- C.** Grant Lilly INSERT permissions on the table.
- D.** Grant Lilly UPDATE permissions on the table.
- E.** Revoke UPDATE permissions on the table for Lilly.
- F.** Revoke DELETE permissions on the table for Lilly.

Answer: B, C.

Explanation: A user cannot use a database object unless the user is explicitly granted permissions to use the database object. In this scenario Lilly must be able to review and add data to the SalesFigures table. To review the data in the SalesFigures table she will require SELECT permissions and to add data to the table she will require INSERT permissions. The option with the least administrative requirements would be to add Lilly to an existing database role with the required permissions. We could thus add Lilly to the Accounts database role, which has these two permissions. However, this option is not offered as a solution in this scenario. Therefore we need to find another database role that has the required permissions to add Lilly to. We could add Lilly to the Salespeople database role, which has the required SELECT permission but does not have any INSERT permission. We would thus have to grant Lilly INSERT permission on the

SalesFigures table.

Incorrect Answers:

A: Lilly only requires SELECT and INSERT permissions to review and add data on the table and we do not want to grant Lilly more permissions than is required. The Managers database role exceeds the permissions required by Lilly as the Managers database role is granted UPDATE and DELETE

permissions. As a member of this group, Lilly would inherit these permissions. This solution would thus grant Lilly more permissions than is required.

D: Lilly only requires SELECT and INSERT permissions to review and add data on the table and we do not want to grant Lilly more permissions than is required. She does not require UPDATE permission, as she will not be changing data on the table. Therefore, granting Lilly UPDATE permissions on the table would be inappropriate.

E: The REVOKE statement removes a previously granted or denied permission from a user in the current database. As Lilly has not been granted UPDATE permissions it is not necessary to REVOKE this permission.

F: The REVOKE statement removes a previously granted or denied permission from a user in the current database. As Lilly has not been granted DELETE permissions it is not necessary to DELETE this permission.

Question No: 3

You are the administrator of a SQL Server 2000 computer named DataSvr. The server is a member of a Microsoft Windows NT domain named Research.

Several users need access to the database stored on the server. These users all belong to a domain user group named Domain Scientists. They also belong to a local user group named Local Techs on the DataSvr computer.

Your company will be hiring 30 new employees who will be added to the Domain Scientists group by domain administrators. These employees will also need access to DataSvr.

You need to allow all of the appropriate users to access DataSvr. You also want to minimize the amount of time you have spend maintaining security on DataSvr.

Which Transact-SQL statement would you execute?

- A.** EXEC sp_grant login BUILTIN\Local Techs
- B.** EXEC sp_grant login 'Local Techs'
- C.** EXEC sp_grant login [DataSvr\Domain Scientists]
- D.** EXEC sp_grant login [Research\Domain scientists]

Answer: D.

Explanation: The sp_grantlogin stored procedure is used to authorize a Microsoft Windows network account to be used as a SQL Server login for connecting to SQL Server using Windows Authentication. When using this stored procedure the domain name followed by the domain user name or domain group

name must be specified in the form domain\user. In this scenario the domain name is Research and the domain user group that is to be authorized is named Domain Scientists.

Incorrect Answers:

A: The argument of this stored procedure is incorrect. It does not indicate the correct domain name or the correct domain user group name that must be authorized for Windows Authentication. Instead it specifies a domain or local group type and a local user group name.

B: The argument of this stored procedure is incorrect. It does not indicate the both the domain name and the domain user group name. Furthermore, the name specified is neither the domain name nor the domain user group name that must be authorized for Windows Authentication. Instead it indicates a local user group name.

C: Instead of specifying the domain name in the argument of the sp_grantlogin stored procedure, the

SQL Server 2000 computer name is specified. This is incorrect as the domain name followed by the domain user name or domain group name must be specified in the form domain\user must be specified.

Question No: 4

You are the Administrator of a SQL Server 2000 computer. The server is configured for Windows authentication. You add an integrated logon for the domain users group. However, 20 of the users in the domain users group should not have access to the server.

You need to prevent the unauthorized users from accessing the server. You need to accomplish this goal by making the least number of changes to the server.

What should you do?

- A. Add a SQL server login for the unauthorized users, and then deny access to the login.
- B. Add the unauthorized users to a domain user group, add a login for the group and then deny access to the login.
- C. Add a login for each authorized user, and then deny access to each login.
- D. Remove each unauthorized user from the domain users group.

Answer: B.

Explanation: By adding the 20 users that shouldn't have access to the SQL server, the so called unauthorized users, to a domain user group and then deny that group login access the group would still be able to access domain resources but they wouldn't be able to access the SQL server since deny overrides any other access they might be granted.

Incorrect Answers:

A: The server is configured for Windows authentication. A SQL server login wouldn't be used.

C: The authorized users should not be denied access.

D: The unauthorized users should still have access to domain resources. They should only be denied access to the SQL Server.

Question No: 5

You are the administrator of a SQL Server 2000 computer named CorpSQL. The server is a member of a Microsoft Windows NT domain named CORPORATE.

Franz maintains a local user account on CorpSQL and is a member of the local Power Users group. He requests access to the SQL Server databases that are stored on CorpSQL.

You need to allow Franz to access CorpSQL. Which Transact-SQL statement should you execute?

- A. EXEC sp_grantlogin Franz
- B. EXEC sp_grantlogin CORPORATE\Franz
- C. EXEC sp_grantlogin Power Users\Franz
- D. EXEC sp_grantlogin CorpSQL\Franz

Answer: D.

Explanation: The sp_grantlogin stored procedure is used to allow a Windows NT or Windows 2000 user or group account to connect to SQL Server 2000 using Windows Authentication. The Windows NT user or group must be qualified with a Windows NT domain name in the form Domain\User in the argument of the sp_grantlogin stored procedure. In this scenario Franz has a local computer account, not a Windows NT user account, and therefore we must use this local user account in the sp_grantlogin statement. The name of computer is CorpSQL.

Incorrect Answers:

A: The sp_grantlogin stored procedure is used to allow a Windows NT or Windows 2000 user or group account to connect to SQL Server 2000 using Windows Authentication. The Windows NT user or group must be qualified with a Windows NT domain name in the form Domain\User in the argument of the sp_grantlogin stored procedure. In this scenario the domain name is CORPORATE and the user name is Franz. However, this solution does not qualify the domain name.

B: The sp_grantlogin stored procedure is used to allow a Windows NT or Windows 2000 user or group account to connect to SQL Server 2000 using Windows Authentication. The Windows NT user or group must be qualified with a Windows NT domain name in the form Domain\User in the argument of the sp_grantlogin stored procedure. In this scenario Franz has a local computer account, not a

Windows NT user account, and therefore we must use this local user account in the sp_grantlogin statement.

C: The sp_grantlogin stored procedure is used to allow a Windows NT or Windows 2000 user or group account to connect to SQL Server 2000 using Windows Authentication. The Windows NT user or group must be qualified with a Windows NT domain name in the form Domain\User in the argument of the sp_grantlogin stored procedure. In this scenario the domain name is CORPORATE and the user name is Franz. Therefore the correct argument would be CORPORATE/Franz and not Power Users\Franz as Power Users is the local group name and not the domain name.

Question No: 6

You are the administrator of a SQL Server 2000 computer. The server is a member of a Microsoft Windows NT domain named CORPHQ. Several users need access to the server.

These users are members of a domain user group named Receptionists.

The Receptionists group also contains users who should not have access to the server. These users are also members of a domain user group named TempEmployees.

You need to allow only the appropriate users to access the server. Which Transact-SQL statement should you execute?

- A.** EXEC sp_addlogin 'CORPHQ\Receptionists'
EXEC sp_droplogin 'CORPHQ\TempEmployees'
- B.** EXEC sp_addlogin 'CORPHQ\TempEmployees'
EXEC sp_addlogin 'CORPHQ\Receptionists'
REVOKE ALL ON Master FROM CORPHQ\TempEmployees
- C.** EXEC sp_addlogin 'CORPHQ\Receptionists'
EXEC sp_denylogin 'CORPHQ\TempEmployees'
- D.** EXEC sp_addlogin 'CORPHQ\TempEmployees'
EXEC sp_addlogin 'CORPHQ\Receptionists'
GRANT ALL ON Master TO CORPHQ\Receptionists

Answer: C.

Explanation: First we allow all receptions access. Then we deny all Tempemployees access. Deny overrides access so receptionists that are members of the Tempemployees group will not get access.

Note: The sp_addlogin stored procedure is used to create a new SQL Server 2000 login that allows a user to connect to an instance of SQL Server using SQL Server Authentication while the sp_grantlogin stored procedure is used to authorize a Windows NT or Windows 2000 network account to be used as a SQL Server login for connecting to SQL Server using Windows Authentication. A login name and password must

be specified when creating a login using the sp_addlogin stored procedure. The Windows NT or Windows 2000 user or group must be qualified with a domain name in the form Domain\User when the sp_grantlogin stored procedure is used. The sp_denylogin stored procedure is used to prevent a Windows NT or Windows 2000 user or group from connecting to SQL Server 2000. The sp_denylogin stored procedure can only be used only with Windows NT or Windows 2000 accounts in the form Domain\User and cannot be used with SQL Server logins added using the sp_addlogin stored procedure. Although the application in the solution is incorrect, it is the closest to correct answer as it follows the correct logic and uses the correct Domain/User information for Windows Authentication. Windows Authentication would be the preferred form of authentication in this scenario because the server is a member of the Windows NT domain.

Incorrect Answers:

A: This solution would still allow the Receptionists that are member of the Tempemployees group access.

B: The REVOKE ALL statement is used to change deny or grant SELECT, INSERT, UPDATE and DELETE permissions on the specified database. These permissions must already have been assigned. In this scenario, permissions have not been assigned therefore the use of the REVOKE statement is inappropriate. Furthermore, we must prevent the TempEmployees from accessing the SQL Server 2000 server. Revoking permissions does not prevent server access.

D: The GRANT ALL statement assigns SELECT, INSERT, UPDATE and DELETE permissions to specified user account. In this scenario, that user account is the Receptionists group. However, the TempEmployees are also part of the Receptionists group and would thus inherit the assigned permissions.

Question No: 7

You are the administrator of a SQL Server 2000 computer named ProdSql1. The server is a member of a Microsoft Windows NT domain named CORP. ProdSql1 contains a database named Finance. Your company purchases a client/server application that will use a database named Marketing. This database will be stored on ProdSql1.

Maria, Bruno, and Stefan will use the new application. These users belong to a domain user group named Marketers. In the future, several additional users who will also use the new application will be added to the Marketers group.

The logins on ProdSql1 are configured as shown in the exhibit.

Logins 6 Items			
Name	Type	Server Access	
BUILTIN\Administrators	Windows ...	Permit	
sa	Stadard	Permit	
CORP\Maria	Windows ...	Permit	
CORP\Bruno	Windows ...	Permit	
CORP\Stefan	Windows ...	Permit	
Marketers	Standard	Permit	

You need to allow Maria, Bruno, and Stefan access to the Marketing database as quickly as possible. Which batch of Transact-SQL statements should you execute?

A. USE Marketing

GO

EXEC sp_grantdbaccess 'CORP\Maria'

EXEC sp_grantdbaccess 'CORP\Bruno'

EXEC sp_grantdbaccess 'CORP\Stefan'

B. EXEC sp_grantlogin 'CORP\Marketers'

GO

USE Marketing

EXEC sp_grantdbaccess 'CORP\Marketers'

C. USE Marketing

GO

EXEC sp_grantdbaccess 'Marketers'

D. EXEC sp_addlogin 'CORP\Marketers'

GO

USE Marketing

EXEC sp_grantdbaccess 'CORP\Marketers'

Answer: B.

Explanation: In this scenario Maria, Bruno, and Stefan belong to a domain user group named Marketers. The most practical solution, since more users are going to be adding to the Marketers group later, is to grant access to the Marketing group. Therefore we specify the domain name and the domain group account in the sp_grantlogin stored procedure and in the sp_grantdbaccess stored procedure.

Note: Before a user or a user group can access a SQL Server 2000 server, a login must be created for them or their Windows NT or Windows 2000 user or group account must be authorized to connect to the SQL Server 2000 server. The sp_grantlogin stored procedure authorizes a Windows NT or Windows 2000 network user or group account to be used as a SQL Server login to connect to a SQL Server 2000 server using Windows Authentication. The Windows NT or Windows 2000 user or group must be qualified with a Windows NT or Windows 2000 domain name in the form Domain\User. The sp_addlogin stored procedure is used to create and define a login account for SQL Server connections using SQL Server Authentication.

However, a SQL Server 2000 server login does not permit database access on its own. Before a user can make use of a database, that user must also be granted access to the database. The sp_grantdbaccess stored procedure adds a user account in the current database for an existing SQL Server 2000 sever login or Windows NT user or group, and enables it to be granted permissions to perform activities in the database. The database must be specified by using the USE statement before the sp_grantdbaccess stored procedure is used. Because the SQL Server 2000 login must exist before we can use the sp_grantdbaccess stored procedure, we must specify the GO statement. The GO statement specifies that the previous Transact-SQL statements must be executed in a batch before the Transact-SQL statements that follow the GO statement is executed

Incorrect Answers:

A: Adding each user separately works as well, but it would require more administrative effort, especially since more people are going to be added to the Marketing group later.

C: We must specify the full group name which is CORP\Marketers, not just Marketers. We must first grant the Marketers group access.

D: The sp_addlogin statement is used incorrectly. It can only be used to add SQL Server logins, but in this scenario we use Windows authentication.

Note: The sp_addlogin stored procedure is used to create and define a login account for SQL Server connections using SQL Server Authentication. The sp_addlogin stored procedure must specify the user account and the password that is to be used with the account.

Question No: 8

You work as database administrator at Testig.in. You handle one SQL Server 2000 computer named TestKingSrv in a domain named TestKing.

Tess must be granted access to the SQL Server database. She has a local account TestKingSrv and she is a member of the local Power Users group.

What should you do?

A. EXEC sp_grantlogin TESTKINGSRV\Tess

B. EXEC sp_grantlogin TESTKING\Tess

C. EXEC sp_grantlogin Power Users\Tess

D. EXEC sp_grantlogin Tess

Answer: A.

Explanation: The sp_grantlogin stored procedure is used to allow a Windows NT or Windows 2000 user or group account to connect to SQL Server 2000 using Windows Authentication. The Windows NT user or group must be qualified with a Windows NT domain name in the form Domain\User in the argument of the sp_grantlogin stored procedure. In this scenario Franz has a local computer account, not a Windows NT user

account, and therefore we must use this local user account in the sp_grantlogin statement. The name of computer is TestKingSrv.

Incorrect Answers:

B, C, D: This solution does not qualify the domain name.

Question No: 9

You work as database administrator at Testig.in. You handle one SQL Server 2000 computer that is a member of a Windows NT domain TESTKING.

You are required to grant access to the Employees domain user group, however members of the Trainees group should not be granted access. The members of the Trainees group are also member of the Employees group.

Which Transact-SQL statement should you use (Select all that apply.)?

- A.** EXEC sp_addlogin 'TESTKING\Employees'
- B.** EXEC sp_addlogin 'TESTKING\Trainees'
- C.** EXEC sp_droplogin 'TESTKING\Trainees'
- D.** EXEC sp_droplogin 'TESTKING\Employees'
- E.** REVOKE ALL ON Master FROM TESTKING\Trainees
- F.** GRANT ALL ON Master TO TESTKING\Employees
- G.** EXEC sp_denylogin 'TESTKING\Trainees'
- H.** EXEC sp_denylogin 'TESTKING\Employees'

Answer: A, G.

Explanation: First we allow all receptions access. Then we deny all Templemployees access. Deny overrides access so Employeeess that are members of the Trainees group will not get access.

Incorrect Answers:

B: We should not grant access to the Trainees

C, D: If we drop the logins the users would not be able to access the server at all.

E: The REVOKE ALL statement is used to change, deny or grant SELECT, INSERT, UPDATE and DELETE permissions on the specified database. These permissions must already have been assigned. In this scenario, permissions have not been assigned therefore the use of the REVOKE statement is inappropriate. Furthermore, we must prevent the Trainees from accessing the SQL Server 2000 server. Revoking permissions does not prevent server access.

F: The GRANT ALL statement assigns SELECT, INSERT, UPDATE and DELETE permissions to specified user account. In this scenario, that user account is the Employeeess group. However, the Trainees are also part of the Employees group and would thus inherit the assigned permissions.

H: Employees must have access.

SECTION 3: Create and manage database users. (2 questions)

Question No: 1

You are the administrator of a SQL Server 2000 computer. The server will be used to contain data from your company's accounting, engineering, and production departments.

Each department is contained in a Microsoft Windows domain users group. The groups are named Accounting, Engineering, and Production. Members of each department need to read and modify their data. However, users in one department should not have access to data in other departments.

You want to configure the database so that it meets your company's security requirements. You also want to minimize administration time and the consumption of server resources.

Which three actions should you take? (Each correct answer represents part of the solution. Choose three.)

- A.** Create a database for each department.
- B.** Create a database for all departments.
- C.** Create a Windows-Authenticated login for each department.
- D.** Create a named instance of SQL Server 2000 for each department.

E. Map each domain users group to the **processadmin** role.

F. Map each domain users group to the **db_datareader** and **db_datawriter** database roles.

G. Grant each domain users group SELECT, INSERT, UPDATE, and DELETE permissions on the database objects that are used by that department.

Answer: A, C, F.

Explanation:

A: Creating one database for each department would enable us to grant each department access to their own database.

C: This access would be based on Windows Authentication as each department has its own Windows domain user group which could also be used to login to the SQL Server 2000 server.

F: We could then add the departments to the db_datareader and db_datawriter roles. This would give the departments the required permissions to read and modify the data in the databases that pertain to their respective departments only.

Incorrect Answers:

B: Creating one database for all the departments would require more complex administrative duties as the permissions for each department would have to be designated on a per table basis. In other words, we would have to assign the appropriate SELECT, INSERT, UPDATE and DELETE permissions to all the tables that are related to a department. This would mean that the appropriate permissions would have to be specified on each table in the database.

D: Installing an additional instance of SQL Server 2000 for each of the departments would increase the load on the resources and would increase database administration as each instance of SQL Server 2000 would require separate administration.

E: Users are added to the processadmin fixed server role to manage processes running on the SQL Server 2000 server. This is a server wide role and is not database specific. It is therefore inappropriate to this scenario.

G: The appropriate SELECT, INSERT, UPDATE and DELETE permissions would have to be specified for all the tables that are related to a particular department. This would increase administrative duties, as the appropriate permissions would have to be specified on each table in the database.

Question No: 2

You are the administrator of a SQL Server 2000 computer. One of the databases on the server contains a table named SalesData. This table stores information about sales orders that is used to create sales reports and to calculate commissions for sales personnel.

The permissions on the SalesData table are configured as shown in the exhibit.

Users/Database Role...	SELECT	INSERT	UPDATE	DELETE
 Replicators	☒	☐	☐	☐
 SalesManagers	☒	☒	☒	☒
 SalesPeople	☐	☒	☐	☐
 guest	☐	☐	☐	☐

Sales managers monitor the entries made by sales personnel and correct any errors. Because monitoring is time-consuming, the sales department will hire several assistants. The sales managers will make necessary corrections on hard copy, and then the assistants will enter the changes in the table. However, the assistants must not be allowed to delete any data from the table.

You need to configure permissions on the table to support these requirements. You also want to reduce the amount of time needed to manage these permissions.

What should you do?

A. Create a database user account for each assistant.

Add each database user account to the **SalesManagers** database role.

Revoke DELETE permissions for the **SalesManagers** database role.

B. Create a single database user account for all assistants.

Add the database user account to the **SalesManagers** database role.

Revoke DELETE permissions for the **SalesManagers** database role.

C. Create a database user account for each assistant.

Create a database role named **SalesAssistants**. Grant SELECT, UPDATE, and INSERT permissions to the **SalesAssistants** database role.

Add each database user account to the **SalesAssistants** database role.

D. Create a database user account for each assistant.

Create a database role named **SalesAssistants**. Add the role to the **SalesManagers** database role.

Add each database user account to the **SalesAssistants** database role.

Revoke DELETE permissions for the **SalesAssistants** database role.

Answer: C.

Explanation: All permissions and ownership of objects in a database are controlled by the database user account. This user account is specific to a database. A database user account can only permit one access to the database at a time; therefore we need to create a user account for each assistant. To ease administrative duties we could create a database role, assign the appropriate permission to the role and place these users in that role. We would grant that role SELECT, INSERT and UPDATE permissions but not DELETE permissions. As the assistants will only have access to the database through this database role, will not have permissions assigned to their individual accounts and will not be included in any other database role, we do not need to deny them DELETE permissions as these permissions will not be granted to them in another context.

Incorrect Answers:

A: Adding the assistants user accounts to the SalesManagers database role would provide them with the required permissions to SELECT, INSERT and UPDATE information in the database. It would also provide them with permission to DELETE data from the database. We could use REVOKE, which is used to change an already granted or denied permission, to change the DELETE permission. This would prevent the assistants from deleting information in the database. However, the Sales Managers will also not be able to use the DELETE statement as they are part of the SalesManagers database role.

B: The Sales Managers will also not be able to use the DELETE statement as they are part of the SalesManagers database role.

D: In this solution the assistants' database role, SalesAssistants is added to the SalesManagers database role. The members of the SalesAssistants database role will thus inherit the permissions granted to the SalesManagers database role. This would provide them with the required permissions to SELECT, INSERT and UPDATE information in the database. It would also provide them with

permission to DELETE data from the database. Furthermore, as no permissions have been assigned to the SalesAssistants database role, we would not be able to use the REVOKE statement against the SalesAssistants database role as the REVOKE statement is used to change an already granted or denied permission. To prevent the members of the SalesAssistants database role from deleting information in the database, we would have to REVOKE the DELETE permission at the level of the SalesManagers database role as the members of the SalesAssistants database role inherits their permission from the SalesManagers database role. This would prevent the assistants from deleting information in the database but it would also prevent the Sales Managers, who are also members of the SalesManagers database role, from deleting data in the database.

SECTION 4: Create and manage security roles. Roles include application, database, and server. (1 question)

Question No: 1

You are the administrator of a SQL Server 2000 computer that contains multiple databases. One of these databases is named CorpSales.

Your company hires a new employee named Helene. Part of Helene's job will be to create new tables in the CorpSales database. However she should not have access to the other databases on the server. You create a SQL Server login named Helene. Helene reports that she can log on to the server but is unable to access the CorpSales database.

You need to allow Helene access to the CorpSales database. What should you do?

A. Create a new database user in the **CorpSales** database.

Map the database user to Helene's login.

Grant CREATE TABLE permissions to the database user.

B. Create a new database role in the **CorpSales** database.

Add Helene's login as a member of the role.

Grant CREATE TABLE permissions to the database role.

C. Change the server's authentication mode to Windows Authentication.

D. Add Helene's login to the **processadmin** server role on the server.

Answer: A.

Explanation: Before a user can use a database, the user must first be granted access to the database and then must be given the appropriate permissions on that database. The `sp_grantdbaccess` stored procedure is used

to grant a user account access to the specified database. This will map the database user to the appropriate login. Thereafter we must grant the user the appropriate permissions. In this scenario the user, Helene, requires CREATE TABLE permissions.

Incorrect Answers:

B: Database roles are used to group user accounts so as to ease administrative duties. The same set of permissions assigned to a user can be assigned to a database role and vice versa. Therefore it is unnecessary to create a database role for a single user. A database role would not by itself allow access to the database.

C: Changing authentication mode will not give Helene access to the database.

Note: SQL Server 2000 can operate in Windows Authentication Mode, which allows a user to connect through a Windows NT 4.0 or Windows 2000 user account, or Mixed Mode, which allows users to connect to an instance of SQL Server using either Windows Authentication or SQL Server Authentication. SQL Server Authentication is provided for backward compatibility because applications written for SQL Server version 7.0 or earlier may require the use of SQL Server logins and passwords, and for users who connects with a specified login name and password from a nontrusted connection. In this event SQL Server performs the authentication itself if the SQL Server 2000 server is using Mixed Mode authentication. If SQL Server does not have a login account set, authentication fails and the user receives an error message. These authentication modes are used to permit users to login to the SQL Server 2000 server. It does not provide access to the databases on the server.

D: Users are added to the processadmin fixed server role to manage processes running on the SQL Server 2000 server. This is a server wide role and is not database specific. It is therefore inappropriate to this scenario.

SUBSECTION A: Add and remove users from roles. (4 questions)


Question No: 1

You are the administrator of a SQL Server 2000 computer. The server is used to store sales information for your company.

Your company hires a new sales manager named Paulo. You create a user account named Paulo in

your company's Microsoft Windows NT domain, which is named CORPORATE. You create a Windows authenticated login for Paulo on the server.

The permissions on the sales table are configured as shown in the exhibit.



Paulo needs permission to view, add, edit, and remove information in the database tables. Which transact-SQL statement should you execute?

- A. EXEC sp_addrolemember 'SalesManagers', 'Corporate\Paulo'
 - B. GRANT ALL ON sales TO 'Corporate\Paulo'
 - C. EXEC sp_addrolemember 'SalesManagers', 'Paulo'
 - D. EXEC sp_grantdbaccess 'Corporate\Paulo', 'Paulo'
- EXEC sp_addrolemember 'SalesManagers', 'Paulo'

Answer: D.

Explanation: Before a security account can use a database, it must be granted access to the database. The sp_grantdbaccess stored procedure is used to perform this function and adds a security account in the database for a SQL Server 2000 login or Windows NT user or group, and enables it to be granted permissions to perform activities in the database. The sp_addrolemember stored procedure is then used to add a security account as a member of an existing SQL Server 2000 database role in the current database. In this scenario Paulo must be added to the SalesManagers database role as this role has the required permissions that we must grant to Paulo.

Incorrect Answers:

A: The sp_addrolemember stored procedure is then used to add a security account as a member of an existing SQL Server 2000 database role in the current database. In this scenario Paulo must be added to the SalesManagers database role as this role has the required permissions that we must grant to Paulo. However, before a security account can use a database, it must be granted access to the database. The sp_grantdbaccess stored procedure is used to perform this function.

B: The GRANT statement creates an entry in the security system that allows a user in the current

database to work with data in the current database or execute specific Transact-SQL statements while the ALL clause specifies that all applicable permissions be granted. However, before a security account can use a database, it must be granted access to the database. The sp_grantdbaccess stored procedure is used to perform this function.

C: The sp_addrolemember stored procedure is then used to add a security account as a member of an existing SQL Server 2000 database role in the current database. In this scenario Paulo must be added to the SalesManagers database role as this role has the required permissions that we must grant to Paulo. However, before a security account can use a database, it must be granted access to the database. The sp_grantdbaccess stored procedure is used to perform this function.

Question No: 2

You are the administrator of a Microsoft Windows 2000 computer. The computer is also running SQL Server 2000 and contains a database named Sales. The sales department of your company uses a custom application that has just been upgraded.

A user at your company named Pierre reports that when he attempts to run the weekly analysis report, he receives the following error message, “Server.Msg3704, Level 16, State 1 , Line 1 User does not have permission to perform this operation on table ‘Report’.”

You log on to Pierre’s computer, run the program, and receive no error message. You use SQL Profiler to trace database activity while you run the report. You receive the results shown in the exhibit.

EventClass	TextData
SQL:BatchCompleted	Drop Table Report
SQL:BatchCompleted	Create Table Report(ProductName C..

You want to allow Pierre to run the report. You also want to minimize administrative overhead.

What should you do?

- A. Add Pierre’s login to the **db_ddladmin** role in the **Sales** database.
- B. Create an application role for the sales department’s application. Grant the role database owner rights.
- C. Add Pierre to the **db_datawriter** role in the **Sales** database.
- D. Add Pierre to the local administrators group on the server.

Answer: A.

Explanation: A member of the db_ddladmin role can issue any command in the Data Definition Language (DDL). The DDL contains commands that are used to define and manage all the objects in an SQL database such as CREATE, ALTER, and DROP statements. In this scenario Pierre is required to be able to run the DROP and CREATE statements which he would be able to when his login is added to the db_ddladmin role.

Incorrect Answers:

- B:** We don’t want to application to have too much rights. We just want Pierre to be able to run it.
- C:** A member of the db_datawriter rule would be able to modify any data in any user table in the database, but would not be able to drop a table.
- D:** The local administrator group doesn’t not have any particular rights on the SQL Server by default.

Question No: 3

You are the administrator of a SQL Server 2000 computer. Peter is an employee in your company's marketing department. Peter has a SQL Server login named Peter, and he has access to a database named MarketingData.

Peter has been assigned to a research project and needs to view and edit information stored in a database named ResearchData. The database is contained on the same server as the MarketingData database.

You need to configure the appropriate permissions for Peter on the ResearchData database. Which Transact-SQL statement should you execute?

A. GRANT ALL ON ResearchData TO 'Peter'
B. GRANT SELECT ON ResearchData TO 'Peter'
GRANT INSERT ON ResearchData TO 'Peter'
GRANT UPDATE ON ResearchData TO 'Peter'

C. EXEC sp_addrolemember 'db_datareader','Peter'
EXEC sp_addrolemember 'db_datawriter','Peter'
D. EXEC sp_grantdbaccess 'Peter',' PeterU'
GO

EXEC sp_addrolemember 'db_datareader',' PeterU'
EXEC sp_addrolemember 'db_datawriter',' PeterU'

Answer: D.

Explanation: Before Peter can use a database he must be granted access to that database. In this scenario Peter has already been granted access to the MarketingData database but not to the ResearchData database. The sp_grantdbaccess stored procedure is used to grant a user account access to a database. As Peter needs to view and edit information in the ResearchData database, we could add him to the db_datareader and db_datawriter fixed database roles. The db_datareader fixed database role allows its members to view all data from all user tables in the database while the db_datawriter fixed database role allows its members to add, change, or delete data from all user tables in the database. The sp_addrolemember is used to add a user account to a database role.

Note: The sp_grantdbaccess stored procedure adds a security account in the current database for a SQL login or Windows user or group, and enables it to be granted permissions to perform activities in the database.

The sp_addrolemember stored procedure adds a security account as a member of an existing database role in the current database.

Incorrect Answers:

A: The GRANT ALL statement specifies that all applicable permissions, namely, SELECT, INSERT, UPDATE and DELETE are being granted to a user account. However, the user must first be granted access to the database before he or she can use the database. The sp_grantdbaccess stored procedure is used to grant a user account access to a database.

B: In addition to the SELECT, INSERT, and UPDATE permissions, a user must first be granted access to the database before he or she can use the database. The sp_grantdbaccess stored procedure is used to grant a user account access to a database. Furthermore, the DELETE permission has not been granted therefore the user cannot fully edit the data, he can only view and change the data. He cannot delete them.


C: As Peter needs to view and edit information in the ResearchData database, we could add him to the db_datareader and db_datawriter fixed database roles by using the sp_addrolemember stored procedure. The db_datareader fixed database role allows its members to view all data from all user tables in the database while the db_datawriter fixed database role allows its members to add, change, or delete data from all user tables in the database. However, these permissions are only applied to the databases to which Peter has been granted access. Before Peter can use a database he must be granted access to that database. In this scenario Peter has already been granted access to the MarketingData database but not to the ResearchData database. The sp_grantdbaccess stored procedure is used to grant a user account access to a database.

Question No: 4

You are the administrator of a SQL Server 2000 computer. The server is used to store information for

your company's finance department.

The permissions on the FinanceInfo database are configured as shown in the exhibit.



Bruno is a member of the Accountants database role. He needs to add a new view to the database. You need to configure the database to allow Bruno to add a new view. Which two actions should you take? (Each correct answer represents part of the solution. Choose two.)

- A. Add Bruno to the **Managers** database role.
- B. Add Bruno to the **Administrators** database role.
- C. Revoke the CREATE TABLE permissions for Bruno.
- D. Revoke the CREATE PROCEDURE permissions for Bruno.

E. Revoke the CREATE VIEW permissions for Bruno.

Answer: A, E.

Explanation: In this scenario Bruno has deny CREATE VIEW permissions on his user account. This prevents Bruno from creating views. However, revoking a deny permission does not automatically provide the user with a corresponding grant permission. It merely makes it possible to change the permission. We therefore need to grant Bruno CREATE VIEW permissions or add him to a database role that has the required permission. In this scenario we can add Bruno to the Managers database role, which would then provide Bruno with the required CREATE VIEW permission. This will not give Bruno CREATE TABLE and CREATE PROCEDURE permissions as Bruno is also a member of the Accountants database role, which has deny CREATE TABLE and CREATE PROCEDURE permissions.

Note: Before a user can use a database, the user must be given the appropriate permissions on that database. The GRANT statement is used to grant users permission to run SELECT, INSERT, UPDATE and DELETE statements against the table. A user can be granted permissions on their user account or on the basis of their membership in a database role. When a user has conflicting permissions as a result of role membership, the most restrictive permission is applied. If the user has database permissions assigned to his or her user

account that must be altered, the REVOKE statement must be used.

Incorrect Answers:

B: Adding Bruno to the Administrators database role would not provide Bruno with the required CREATE VIEW permissions as the Administrators database role only has CREATE TABLE permissions.

C: The REVOKE statement is used to change denied or granted permissions on the specified database. These permissions must already have been assigned. In this scenario no CREATE TABLE permissions have been assigned to Bruno's user account therefore we cannot REVOKE the CREATE TABLE permission, as it does not exist. Furthermore, should we add Bruno to the Managers database role, which has CREATE TABLE, CREATE VIEW and CREATE PROCEDURE permissions, we would not be giving Bruno CREATE TABLE and CREATE PROCEDURE permissions as Bruno is also a member of the Accountants database role, which already has deny CREATE TABLE and CREATE PROCEDURE permissions.

D: The REVOKE statement is used to change denied or granted permissions on the specified database. These permissions must already have been assigned. In this scenario no CREATE PROCEDURE permissions have been assigned to Bruno's user account therefore we cannot REVOKE the CREATE PROCEDURE permission, as it does not exist. Furthermore, should we add Bruno to the Managers database role, which has CREATE TABLE, CREATE VIEW and CREATE PROCEDURE permissions, we would not be giving Bruno CREATE TABLE and CREATE PROCEDURE permissions as Bruno is also a member of the Accountants database role, which already has deny CREATE TABLE and CREATE PROCEDURE permissions.

SUBSECTION B: Create roles to manage database security. (1 question)

Question No: 1

You are the administrator of a SQL Server 2000 computer. The server is used to store sales and product information for your company. A client/server application accesses the server by activating an application role named SalesAppRole.

Several users report that they are unable to use the client/server application. You test the application and discover that it cannot log on to the server.

You examine your security audit logs, which contain the information shown in the exhibit.

You need to enable the application to log on to the server. What should you do?

- A.** Configure the database permissions to deny Maria access to the database.
- B.** Reset the password on the **SalesAppRole** application role.
- C.** Reset the permissions on the **SalesAppRole** application role.
- D.** Add the **SalesAppRole** application role to the **db_owner** database role.
- E.** Remove Bruno's database user account.
- F.** Enable encryption on the password for the **SalesAppRole** application role.

Answer: B.

Explanation: The information in the exhibit indicates that the sp_approlepassword stored procedure was used by Bruno. This stored procedure is used to change an application role password. Only members of the sysadmin fixed server role, and the db_securityadmin and db_owner fixed database roles can execute the sp_approlepassword stored procedure.

The sp_setapprole stored procedure is used to activate the permissions associated with an application role in the current database. After an application role is activated with the sp_setapprole stored procedure, the role cannot be deactivated until the user disconnects from SQL Server.

The application role must be configured with the correct password.

Incorrect Answers:

A: Reconfiguring the Maria's database permissions in this solution would prevent her from accessing, and hence using, the database. It would however not address the problem posed in this scenario, which is the failure of the SalesAppRole to login to the server.

C: The problem posed in this scenario is not a permissions problem. Permissions pertain to using database objects. The problem in this scenario is a server login problem. Changing the permissions assigned to the SalesAppRole would thus be inappropriate to this scenario.

D: Application rules are not added to database rules. Application rules are configured with a password.

E: Removing Bruno's database user account would prevent Bruno from logging on to the server. This however does not address the problem posed in this scenario, which is the failure of the SalesAppRole to login to the server.

F: To protect the application role password, it should be encrypted using the ODBC Encrypt function and specify odbc as the encryption method. It would not, however, enable the application to log on to the server.

SECTION 5: Enforce and manage security by using stored procedures, triggers, views, and user-defined functions. (6 questions)

Question No: 1

You are the administrator of a SQL Server 2000 computer named corpsql. The server is a member of a Microsoft Windows NT domain named CORPDOMAIN. Corpsql is configured for Windows authentication and contains three databases named sales, finance, and research.

Sophie is an administrative assistant in your company. She uses a suite of client/server applications to access all three databases on corpsql.

Sophie transfers to different department within your company and requires access only to the research database. Her former manager requests that you remove Sophia's access to the sales and finance databases.

You need to configure the appropriate permissions for Sophie. Which batch of transact-SQL statements should you execute?

A. USE finance

GO

EXEC sp_revokedbaccess Sophie

USE sales

GO

EXEC sp_revokedbaccess Sophie

B. USE finance

GO

EXEC sp_revokedbaccess Sophie 'db_denydatareader', 'Sophie'

USE sales

GO

EXEC sp_revokedbaccess Sophie 'db_denydatareader', 'Sophie'

C. USE Master

GO

EXEC sp_droplogin 'Corpdomain\Sophie'

GO

EXEC sp_droplogin 'Corpdomain\Sophie'

D. USE Finance

GO

EXEC sp_denylogin 'Sophie'

USE sales

GO

EXEC sp_denylogin 'Sophie'

Answer: A.

Explanation: The sp_revokedbaccess stored procedure is used to remove a database user account from the current database. When the database user account is removed, the permissions and aliases that are dependant on the account are automatically removed.

The solution first makes the finance database current with the **USE** command. Then it revokes Sophie's access to the finance database with the **sp_revokedbaccess 'Sophie'** command. Then her access to the sales database is revoked in the same manner.

Note: Syntax of sp_revokedbaccess

sp_revokedbaccess [@name_in_db =] 'name'

Incorrect Answers:

B: The line with the sp:revokedbaccess statement has an incorrect syntax.

C: Dropping Sophie's login would prevent her from accessing the SQL Server. She must still be able to access the research database. The sp_droplogin command is too restrictive.

Note: The sp_droplogin stored procedure is used to remove a SQL Server login, thereby preventing access to SQL Server using that login name. It is not used to prevent a Windows NT or Windows 2000 user or user group from connecting to the SQL Server 2000 server via Windows Authentication.

D: The sp_denylogin 'Sophie' command would prevent her from accessing the SQL Server. She must still be able to access the research database. The sp_denylogin command is too restrictive.

Note: The sp_denylogin stored procedure is used to prevent a Windows NT or Windows 2000 user or user group from connecting to the SQL Server 2000 server via Windows Authentication. This would prevent Sophie from login on to the SQL Server 2000 server and would thus prevent her from accessing the Research database as well.

Question No: 2

You are the administrator of a SQL Server 2000 computer. The server contains a database named EmployeeData. Your company's human resources department uses this database, which contains several tables as shown in the following table.

Table name Used for

Employees Employee names, addresses, departments, and base salaries

Bonuses Bonuses that have been paid to employees

Awards Awards that have been presented to employees

The human resources department wants to make employee names, addresses, and awards available to anyone who has permission to access the server. However, the human resources department wants to allow only users in the human resources department to access salary and bonus information.

You need to allow company users to access only the appropriate employee information. What should you do?

A. Create a trigger on the **Salary** column of the **Employees** table.

B. Create a stored procedure that retrieves all of the data from the **Employees** and **Awards** tables, inserts the data into a temporary table, and then grants the current user SELECT permissions on the temporary table.

Grant the **Guest** database user EXECUTE permissions on the stored procedure.

C. Move the sensitive information out of the **Employees** table into a new table.

Grant the **Guest** database user SELECT permission on the **Employees** table.

D. Create a view that contains the name, address, and award information.

Grant the Guest database user SELECT permissions on the view.

Answer: D.

Explanation: A view is a virtual table that allows access to the data in the base table. A user can reference the view through Transact-SQL statements in the same way that a table is referenced. A view can, however, be used to restrict a user to specific rows and/or columns in a table while the rest of the table or database is invisible and inaccessible. Permission to access the data in a view must be explicitly assigned, regardless of the set of permissions in force on the underlying table. Thus, granting SELECT permissions on the view allows the user to see nonsensitive columns through the view without having any permissions on the table itself.

Incorrect Answers:

A: Triggers do not fire in response to SELECT statements. They cannot be used to select specific columns from a result set.


Note: Triggers are special types of stored procedure that fire automatically when the data in a specified table is modified. They fire in response to an INSERT, UPDATE, or DELETE statement and can query other tables and can include complex Transact-SQL statements. Triggers can cascade changes through related tables in the database; enforce restrictions that are more complex than those defined with CHECK constraints as triggers can reference columns in other tables; evaluate the state of a table before and after a data modification and take action(s) based on that difference; and multiple triggers of the same type on a table can allow multiple, different actions to take place in response to the same modification statement.

B: The temporary table would contain all columns for the two tables. We want to restrict the result set to specific columns.

C: The creation of a separate table to hold the sensitive is a possible solution to this scenario. Providing that the human resources department has permissions assigned at the database level, additional permissions for the table would not have to be assigned for this table. However, granting explicit permissions to a view is the better solution as it requires less administrative effort.

Question No: 3

You are the administrator of a SQL Server 2000 computer at a doctor's office. The server contains a database named Medical. The database stores patient and treatment information and is configured as shown in the exhibit.



The office accountant needs access to the total number of treatments that have been performed in the

office. The accountant should not have access to any other patient data and should not be able to change any information in the database.

The accountant has a SQL Server Authenticated login and database user account named AccountUser.

You need to make the appropriate information available to the accountant, while protecting the other information in the database. Which batch of Transact-SQL statements should you execute?

A. CREATE VIEW TotalTreatments AS
SELECT COUNT(PatientID) AS Total FROM dbo.Patients
GO

GRANT ALL ON TotalTreatments TO AccountUser

B. CREATE VIEW TotalTreatments AS
SELECT COUNT(*) AS Total FROM dbo.Patients
GO

GRANT ALL ON TotalTreatments TO AccountUser

C. CREATE PROCEDURE TotalTreatments AS
SELECT COUNT(*) AS Total FROM dbo.Patients
GO

GRANT ALL ON TotalTreatments TO AccountUser

D. CREATE PROCEDURE TotalTreatments AS
SELECT COUNT(PatientID) AS Total FROM dbo.Patients
GO

GRANT ALL ON TotalTreatments TO AccountUser

Answer: C.

Explanation: Using a procedure is a good way of protecting from direct access to the user tables. The users would only have to be granted access to the procedure not to the underlying table.

We must use COUNT(*) since COUNT(*) returns the number of rows that match the search conditions specified in the query without eliminating duplicates. It counts each row separately, including rows that contain null values.

COUNT(PatientID) would count distinct rows and would not include rows with NULL values in the PatientID column.

Incorrect Answers:

A: If we should use a view we must grant the users permission on the underlying table Patients. If not there would be a broken ownership ship and the users wouldn't be able to use the view.

COUNT(PatientID) would count distinct rows and would not include rows with NULL values in PatientID.

B: Protecting the underlying table through a view is a possible solution. In general procedures are preferred though since a view in most cases provide the user with more possibilities to access the underlying table.

D: COUNT(PatientID) would count distinct rows and would not include rows with NULL values in PatientID.

Question No: 4

You are the administrator of a SQL Server 2000 computer. The server belongs to a Microsoft Windows NT domain named DOMAIN1. The server contains a database named SalesData. Julia is a sales representative who uses the database. You create a login named DOMAIN1\Julia and manually configure the permissions necessary for her to perform her job.

Julia is temporarily reassigned to the marketing department. During her assignment in the marketing department Julia must not be able to access the server. However, you must restore her access when she returns to the sales department.

You need to configure the appropriate database permissions for Julia. Which Transact-SQL statement should you execute?

A. REVOKE ALL TO SalesData FROM 'Julia'

B. EXEC sp_revokedbaccess 'Julia'

C. EXEC sp_denylogin 'Julia'

D. EXEC sp_droplogin 'Julia'

Answer: C.

Explanation: The sp_denylogin stored procedure is used to temporarily suspend the SQL Server 2000 login access for a Windows NT or Windows 2000 user or group account without removing the permissions and aliases that are dependant on the user account. This prevents a user or group from connecting to an instance of SQL Server while keeping Julia's permissions to the database intact. She will however not be able to use these permissions as she will not have access to the database due to her login to the server being suspended.

Note: The sp_denylogin stored procedure can only be used with Windows NT or Windows 2000 accounts in the form Domain\User and cannot be used with SQL Server logins added with the sp_addlogin stored procedure. The sp_grantlogin stored procedure can be used to reverse the effects of the sp_denylogin stored procedure and restores the user accounts login permission to the instance of SQL Server 2000.

Incorrect Answers:

A: The REVOKE ALL statement specifies that all SELECT, INSERT, UPDATE and DELETE permissions that have already been assigned to a user account must be changed. This could prevent a user from working with the database but does not prevent them from logging on to the SQL Server 2000 server.

B: The sp_revokedbaccess stored procedure is database used to remove a user account from the database. This will automatically removed remove the permissions and aliases that depend on the account. However, the sp_revokedbaccess stored procedure is specific to a database. The user will thus not be able to use a database but they will still be able to login to the SQL Server 2000 server.

D: The sp_droplogin stored procedure removes a SQL Server 2000 login and prevents access to the SQL Server 2000 server using that login name. A login mapped to an existing user in any database cannot be removed unless the database user account has already been removed by using the sp_dropuser stored procedure.

Question No: 5

You are the administrator of a SQL Server 2000 computer. One of the databases on the server contains a table named CustomerInfo. This table is used to store information about your company's customers.

You want some users to be allowed to modify only the table's Address column. Other users should be allowed to modify only the table's CreditCard column. Neither group of users has permission on the CustomerInfo table.

You need to make sure that only authorized users make changes to the Address and CreditCard columns. What should you do?

A. Create a stored procedure and require users to enter information by means of the stored procedure.

B. Create a user-defined function and require users to enter information by means of the function.

C. Create a rule and apply it to the Address and CreditCard columns.

D. Create a constraint and apply it to the Address and CreditCard columns.

Answer: A.

Explanation: We use a stored procedure to customize security. Only the stored procedure is allowed to change the columns, and only authorized users will be granted access to the stored procedure.

Note: A stored procedure is a set of Transact-SQL statements compiled into a single execution plan and can return data as output parameters, which can return either data or a cursor variable; return codes, which are

always an integer value; a result set for each SELECT statement contained in the stored procedure or any other stored procedures called by the stored procedure; or as a global cursor that can be referenced outside the stored procedure. Stored procedures assist in achieving a consistent implementation of logic across applications as the SQL statements and logic needed to perform a task can be coded and tested once in a stored procedure. Thereafter each application or user that needs to perform that task can execute the stored procedure. Coding business logic into a single stored procedure also offers a single point of control for ensuring that business rules are correctly enforced. Stored procedures can also be used to customize security permissions in much the same way as views. Stored procedure permissions can be used to prevent users from modifying data in certain columns and then grant users permissions to execute the procedure.

Incorrect Answers:

B: Functions are not able to return rowsets so they cannot be used here.

Note: In programming languages functions are subroutines used to encapsulate frequently performed logic. Any code that must perform the logic incorporated in a function can call the function rather than having to repeat all of the function logic. SQL Server 2000 supports built-in functions and userdefined functions. Built-in functions operate as defined in the Transact-SQL Reference and cannot be modified. The functions can be referenced only in Transact-SQL statements using the syntax defined in the Transact-SQL Reference. User-defined functions, on the other hand, allows the user to define their own Transact-SQL functions using the CREATE FUNCTION statement. User-defined functions take zero or more input parameters, and return a single value. Some user-defined functions return a single, scalar data value, such as an int, char, or decimal value.

C: Rules cannot be used to configure security.


Note: Rules are a backward-compatibility feature of SQL Server 2000 and performs the same functions as CHECK constraints. CHECK constraints enforce domain integrity by limiting the values that can be placed in a column and are the preferred, standard way to restrict the values in a column, as they are more concise than rules.

D: A constraint cannot be used to configure security.

Note: In SQL Server 2000 constraints are used to define ways in which the integrity of a database can be enforced automatically. Constraints define rules regarding the values allowed in columns and are the standard mechanism for enforcing integrity. Using constraints is preferred to using triggers, rules, and defaults. SQL Server 2000 supports five classes of constraints: NOT NULL specifies that the column does not accept NULL values; CHECK constraints enforce domain integrity by limiting the values that can be placed in a column; UNIQUE constraints enforce the uniqueness of the values in a set of columns; PRIMARY KEY constraints identify the column or set of columns whose values uniquely identify a row in a table; and FOREIGN KEY constraints identify the relationships between tables. These constraints cannot be used to verify the authority of the user inserting the data. It can only control the nature of the data being entered into a column.

Question No: 6

Database design exhibit



You work as database administrator at Testig.in. You handle one SQL Server 2000 computer.

You are required to write a T-SQL batch with the following specifications:

- the total number of treatments should be displayed
- the users should not have access to the tables
- the users has an account Nurses and SQL Server Authentication is used

Which code should you use? (Select three. Select all statements that should be included in the batch)

- A. CREATE VIEW TotalTreatments AS
- B. CREATE PROCEDURE TotalTreatments AS
- C. SELECT COUNT(PatientID) AS Total FROM dbo.Patients
- D. SELECT COUNT(*) AS Total FROM dbo.Patients
- E. GO

GRANT ALL ON TotalTreatments TO Nurses

Answer: B, D, E.

Explanation: Using a procedure is a good way of protecting from direct access to the user tables. The users would only have to be granted access to the procedure not to the underlying table.

We must use COUNT(*) since COUNT(*) returns the number of rows that match the search conditions specified in the query without eliminating duplicates. It counts each row separately, including rows that contain null values.

COUNT(PatientID) would count distinct rows and would not include rows with NULL values in the PatientID column.

Incorrect Answers:

A: If we should use a view we must grant the users permission on the underlying table Patients. If not there would a broken ownership ship and the users wouldn't be able to use the view.

COUNT(PatientID) would count distinct rows and would not include rows with NULL values in PatientID.

C: COUNT(PatientID) would count distinct rows and would not include rows with NULL values in PatientID.

SECTION 6: Set permissions in a database. Considerations include object permissions, object ownership, and statement permissions. (3 questions)

Question No: 1

You are the administrator of a SQL Server 2000 computer. The server is used to store information for your company's accounting department.

The permissions on the accounting table are configured as shown in the exhibit.

Users/Database Role...	SELECT	INSERT	UPDATE	DELETE
Accountants	☒	☒	☒	☒
Bruno	☒	☐	☐	☐
Managers	☒	☐	☒	☒
SalesPeople	☒	☐	☐	☐
guest	☐	☐	☐	☐

Bruno is a member of the Managers and Salespeople database roles. While working in the table, he reports that he can edit and remove information. However, he cannot view any of the information in the table.

You need to allow Bruno to view the information in the Accounting table. Which Transact-SQL statement should you execute?

- A. GRANT PRIVILEGES ON Accounting TO Bruno
- B. GRANT SELECT ON Accounting TO Bruno WITH GRANT OPTION
- C. REVOKE GRANT OPTION FOR SELECT ON Accounting TO Bruno
- D. REVOKE SELECT ON Accounting FROM Bruno

Answer: D.

Explanation: In this scenario Bruno has deny SELECT permissions on his user account. This prevents Bruno from viewing the information in the Accounting table. We must thus REVOKE the SELECT permission on Bruno's user account.

Note: Before a user can use a table, the user must be given the appropriate permissions on the table. The GRANT statement is used to grant users permission to run SELECT, INSERT, UPDATE and DELETE statements against the table. However, if the user has permissions on the table that must be changed, the REVOKE statement must be used. A user can be granted permissions on their user account or on the basis of their membership in a database role. When a user has conflicting permissions as a result of role membership, the most restrictive permission is applied.

Incorrect Answers:

A: The PRIVILEGES keyword is an optional keyword that is used for compliance with SQL-92. PRIVILEGES is not a statement. Therefore this GRANT statement incorrect as it does not have an argument. Furthermore, we need to change an already assigned deny SELECT permission. This is accomplished through the use of a REVOKE statement.

B: An already assigned permission can only be changed through the use of the REVOKE statement. In this scenario Bruno has deny SELECT permissions on his user account. This prevents Bruno from viewing the information in the Accounting table. We must thus REVOKE the SELECT permission on Bruno's user account. The WITH GRANT OPTION at the end of the statement specifies that the account is given the ability to grant the specified object permission to other accounts.

C: The REVOKE GRANT OPTION FOR SELECT ON Accounting TO Bruno is not the correct syntax for the REVOKE statement. The correct REVOKE statement specifies whether the SELECT, UPDATE, INSERT or DELETE statement or ALL permission must be revoked immediately after the REVOKE statement, specifies ON which table the permission should be revoked, and FROM which account.

Question No: 2

You are the administrator of a SQL Server 2000 computer. The server contains seven databases that provide data to several company client/server applications. A different user maintains each database.

You need to configure server and database permissions so that each user who maintains a database has full permissions on that database. What should you do?

A. Create a domain user group for all users.

Add a login for the domain user group.

Add the login to the sysadmin server role.

B. Create a domain user group for all users.

Add a login for the domain user group.

Map the login to a database user in each database.

Add all database users to the **db_owner** database role in each database.

C. Create a login for each user's domain user account.

Map each login to a database user in the appropriate databases.

Add the database users to the **db_owner** database role in each database.

D. Create a login for each user's domain user account.

Add the logins to the **dbcreator** server role.

Answer: C.

Explanation: We must create a specific login for each appropriate domain user account. Then we must map these login to the specific database that particular user should maintain. By adding the login to the db_owner database role the user would have full permissions to that particular database.

Incorrect Answers:

A: A member of the sysadmin role can perform any activity in SQL Server. We need a database role not a server-wide role.

We cannot create one login which applies to all domain users. We must create one role for each user since they only should have permissions to their specific database.

B: We cannot create one login which applies to all domain users. We must create one role for each user since they only should have permissions to their specific database.

D: A member of the dbcreator fixed server role can create, alter, and drop databases. This is a serverwide role, but we only want to assign permissions to a specific database with a database role.


Question No: 3

You are the administrator of a SQL Server 2000 computer. The server is used to store information for your company's finance department.

The permissions on the FinanceData table are configured as shown in the Database Permissions exhibit.

Users/Database Role...	SELECT	INSERT	UPDATE	DELETE
 Accountants	☒	☒	☒	☒
 Administrators	☒	☒	☒	☒
 Reporters	☒	☐	☐	☐
 SalesManagers	☒	☐	☐	☒
 guest	☐	☐	☐	☐

Stefan's database user account is configured as shown in the Login Properties exhibit.



Stefan reports that he is unable to delete information from the FinanceData table. You need to allow him to delete information from the table.

Which Transact-SQL statement should you execute?

- A. EXEC sp_droprolemember 'Accountants', 'Stefan'
- B. EXEC sp_droprolemember 'Administrators', 'Stefan'

C. EXEC sp_addrolemember 'db_datawriter', 'Stefan'

D. EXEC sp_addrolemember 'db_owner', 'Stefan'

Answer: A.

Explanation: In this scenario Stefan has deny DELETE permissions as a member of the Accountants database role. This prevents Stefan from deleting information in the database. We cannot REVOKE the DELETE permission on Accountants database role as this would change the permissions for all member of that role. Instead we should remove Stefan from the Accountants database role. Once he is removed from the Accountants database role, the permissions applied to that role will no longer be applied to Stefan. The sp_droprolemember stored procedure is used to remove as user account from a database role.

Note: Before a user can use a database, the user must be given the appropriate permissions on that database. The GRANT statement is used to grant users permission to run SELECT, INSERT, UPDATE and DELETE statements against the table. If the user has database permissions assigned to his or her user account that must be altered, the REVOKE statement must be used. A user can be granted permissions on their user account or on the basis of their membership in a database role. When a user has conflicting permissions as a result of role membership, the most restrictive permission is applied.

Incorrect Answers:

B: Stefan has deny DELETE permissions as a member of the Accountants database role. This prevents Stefan from deleting information in the database. The permissions he inherits from the Administrators database role are not preventing him from deleting information in the database. We should thus remove Stefan from the Accountants database role and not the Administrators database role.

C: When a user has conflicting permissions as a result of role membership, the most restrictive permission is applied. In this scenario Stefan has deny DELETE permissions as a member of the Accountants database role. This prevents Stefan from deleting information in the database. Adding Stefan to another fixed database role will not provide him with the ability to delete information in the database because the conflicting permissions remain in place and the deny DELETE permissions inherited from the Accountants database role remain. We should remove Stefan from the Accountants database role.

D: When a user has conflicting permissions as a result of role membership, the most restrictive

permission is applied. In this scenario Stefan has deny DELETE permissions as a member of the Accountants database role. This prevents Stefan from deleting information in the database. Adding Stefan to another fixed database role will not provide him with the ability to delete information in the database because the conflicting permissions remain in place and the deny DELETE permissions inherited from the Accountants database role remain. We should remove Stefan from the Accountants database role.

SECTION 7: Manage security auditing. Methods include SQL Profiler and C2 auditing. (3 questions)

Question No: 1

You are the administrator of a SQL Server 2000 computer. The server is used to store confidential company information. Company policy requires that every action and change of permission on the server be logged. Policy also requires that the server can run only when logging is enabled.

You need to configure the server to comply with this policy. What should you do?

A. Use SQL Profiler to capture security events and audit events.

Make sure that file rollover is enabled.

B. On the **Security** tab of the **Server Properties** dialog box, set the **Audit Level** to **All**.

C. Configure the server to use Windows Authentication.

Make sure the Windows security log does not overwrite events.

D. Set the c2 audit mode option to 1.

Restart the MSSQLServer service.

Answer: D.

Explanation: In SQL Server 2000, the c2 audit mode option is used to review both successful and unsuccessful attempts to access statements and objects. This information can be used to document system activity and to check for security policy violations. C2 auditing tracks C2 audit events and records them to a file. To enable C2 auditing, the c2 audit mode option must be set to 1. This setting establishes the C2 audit trace and turns on the option to fail the server should the server be unable to write to the audit file for any reason. After the option has been set to 1, the server must be restarted to begin C2 audit tracing.

Incorrect Answers:

A: SQL Profiler is used to stepping through problem queries to find the cause of the problem; find and diagnosing slow-running queries; capture the series of SQL statements that lead to a problem; and to monitoring the performance of SQL Server to tune workloads. SQL Profiler can also supports auditing the actions performed on instances of SQL Server and meets C2 security certification requirements. However, SQL Profiler cannot prevent the server from running should audit logging cease for any reason.

B: On the Security tab of the SQL Server Properties dialog box we can specify the audit level to ALL. This will result in the auditing of both successful and failed login attempts. Log records for these events appear in the Windows application log, the SQL Server error log, or both, depending on how logging is configured for the SQL Server 2000 server. Once this setting has been specified, the SQL Server 2000 server must be restarted for the auditing to take effect. This however audits logon attempts and not every action and change of permission on the server.

C: The Windows security log contains entries for valid and invalid logon attempts as well as events related to resource use, such as creating, opening, or deleting files or other objects. These events relate to the Windows NT 4.0 or Windows 2000 server and not the SQL Server 2000 server.

Question No: 2

You are the administrator of a SQL Server 2000 computer. The server is used to store confidential information about company employees.

You suspect that at least one user is accessing the employee information database without company approval. You believe the access is occurring after business hours.

You need to discover whether users are accessing the database without approval and which users are accessing the data. You also want to minimize any additional overhead on the server during business hours.

What should you do?

A. Use SQL Profiler to capture security events and audit events.

B. On the **Security** tab of the **Server Properties** dialog box, set the **Audit level** option to **All**.

C. Configure the server to use Windows Authentication.

Make sure the Windows security log does not overwrite events.

D. Set the **c2 audit mode** option to **1**.

Restart the MSSQLServer service.

Answer: A.

Explanation: We configure SQL Profiler to audit security events and audit events. Security events include database access and data access. We will be able to view which users access the database.

Note: SQL Profiler is used to stepping through problem queries to find the cause of the problem; find and diagnosing slow-running queries; capture the series of SQL statements that lead to a problem; and to monitoring the performance of SQL Server to tune workloads. SQL Profiler can also supports auditing the actions performed on instances of SQL Server and meets C2 security certification requirements. SQL Profiler provides the user interface for auditing events. There are several categories of events that can be audited using SQL Profiler, these include end user activity; DBA activity; security events; utility event; server events and audit events.

Incorrect Answers:

B: On the Security tab of the SQL Server Properties dialog box we can specify the audit level to ALL. This will result in the auditing of both successful and failed login attempts. Log records for these events appear in the Windows application log, the SQL Server error log, or both, depending on how logging is configured for the SQL Server 2000 server. Once this setting has been specified, the SQL

Server 2000 server must be restarted for the auditing to take effect. This however audits logon attempts and not every action and change of permission on the server.

C: The Windows security log contains entries for valid and invalid logon attempts as well as events related to resource use, such as creating, opening, or deleting files or other objects. These events relate to the Windows NT 4.0 or Windows 2000 server and not the SQL Server 2000 server.

D: In SQL Server 2000, the c2 audit mode option is used to review both successful and unsuccessful attempts to access statements and objects. This information can be used to document system activity and to check for security policy violations. c2 auditing tracks c2 audit events and records them to a file. To enable c2 auditing, the c2 audit mode option must be set to 1. This setting establishes the c2 audit trace and turns on the option to fail the server should the server be unable to write to the audit file for any reason. After the option has been set to 1, the server must be restarted to begin c2 audit tracing. However, c2 auditing impacts on server performance and is therefore not the best solution.

Question No: 3

You work as database administrator at Testig.in. You handle one SQL Server 2000 computer.

You have been required to ensure that the Server all every change of permission and action on the server is logged and that the SQL Server will only run if the logging is enabled.

What action should you take? (Select all that apply.)

A. Set the c2 audit mode option to 1.

B. Set the c2 audit mode option to 2.

C. Configure the server to use Windows Authentication.

D. Configure the server to use Mixed Mode Authentication.

- E. Use SQL Profiler to capture security events and audit events.
- F. Restart the MSSQLServer service.
- G. Restart the Server.
- H. Make sure that file rollover is enabled.
- I. On the **Security** tab of the **Server Properties** dialog box, set the **Audit Level** to **All**.
- J. Make sure the Windows security log does not overwrite events.

Answer: A, F.

Explanation: In SQL Server 2000, the c2 audit mode option is used to review both successful and unsuccessful attempts to access statements and objects. This information can be used to document system

activity and to check for security policy violations. C2 auditing tracks C2 audit events and records them to a file. To enable C2 auditing, the c2 audit mode option must be set to 1. This setting establishes the C2 audit trace and turns on the option to fail the server should the server be unable to write to the audit file for any reason. After the option has been set to 1, the service must be restarted to begin C2 audit tracing.

Incorrect Answers:

B: It should be set to 1.

C, D: Authentication mode is not an issue here.

E, H: SQL Profiler is used to stepping through problem queries to find the cause of the problem; find and diagnosing slow-running queries; capture the series of SQL statements that lead to a problem; and to monitoring the performance of SQL Server to tune workloads. SQL Profiler can also supports auditing the actions performed on instances of SQL Server and meets C2 security certification requirements. However, SQL Profiler cannot prevent the server from running should audit logging cease for any reason.

G: Only the service has to be restarted.

I: On the Security tab of the SQL Server Properties dialog box we can specify the audit level to ALL. This will result in the auditing of both successful and failed login attempts. Log records for these events appear in the Windows application log, the SQL Server error log, or both, depending on how logging is configured for the SQL Server 2000 server. Once this setting has been specified, the SQL Server 2000 server must be restarted for the auditing to take effect. This however audits logon attempts and not every action and change of permission on the server.

J: The Windows security log contains entries for valid and invalid logon attempts as well as events related to resource use, such as creating, opening, or deleting files or other objects. These events relate to the Windows NT 4.0 or Windows 2000 server and not the SQL Server 2000 server.

PART VI: Managing, Monitoring, and Troubleshooting SQL Server 2000 (22 questions)

SECTION 1: Create, manage, and troubleshoot SQL Server Agent jobs. (6 questions)

Question No: 1

You are the administrator of SQL Server 2000 computer. You create a job that performs several maintenance tasks on the server's databases. You want the job to run whenever the server's processor utilization falls below 5 percent.

You create a new schedule for the job and specify the start whenever the CPU(s) become idle option. After several days, you notice that the job has never executed although the server's processor utilization has fallen below 5 percent several times.

What should you do?

- A. Modify SQL Server Agent properties and specify a smaller idle time.
- B. Modify SQL server agent properties and specify a larger idle time.
- C. Write a stored procedure that executes the job whenever the @@IDLE system variable is less than 5.

D. Write a stored procedure that executes the job whenever the @@IDLE system variable is greater than 1.

Answer: A.

Explanation: In order to make it more likely for the job to start we should specify a smaller idle time for SQL ServerAgent.

Note: Administrative jobs can be scheduled to run automatically when SQL Server Agent starts; when CPU utilization of the computer is at a defined level you have defined as idle; at one time on a specific date and time; on a recurring schedule or in response to an alert. To maximize CPU resources, a CPU idle condition can be defined to determine the most advantageous time to execute jobs. The CPU idle condition is defined as a percentage below which the average CPU usage must remain for a specified time. When the CPU usage level drops below the defined level and remain remains below that level for the specified time, SQL Server Agent starts all jobs that have a CPU idle time schedule. If the CPU usage increases to above the level before the specified time has been exceeded, the monitor is reset. Thus, by specifying a shorter idle time, the length of time that the CPU usage must be below the defined level is reduced. Thus it becomes more likely that the conditions required for the scheduled task to be met increases.

Incorrect Answers:

B: Specifying a larger idle time would require that the CPU usage level remain below the defined level of a longer period of time. Thus it is less likely that the conditions required for the scheduled task to be met will be reached.


C: We cannot create a stored procedure that executes when some condition is met. We schedule jobs not procedures.

Note: SQL Server 2000 keeps track of how much work the server has done since it was last started through a series of functions. The sp_monitor stored procedures is used to display the current values returned by these functions and shows how much they have changed since the last time the procedure was run. @@IDLE is one of these functions and keeps track of the time that the Server has been idle for since the time the server was last started and measures the time in milliseconds.

D: We cannot create a stored procedure that executes when some condition is met. We schedule jobs not procedures.

Question No: 2

You are the administrator of a SQL Server 2000 computer. The SQLServerAgent service on the server is configured as shown in the exhibit.



You configure two SQLServerAgent jobs to perform maintenance tasks. Each job copies files from the server to a file share on a corporate file server. Both jobs fail when you attempt to execute them.

You need both jobs to execute successfully. What should you do?

- A. Ensure that the file copy procedures are in an ActiveX Scripting job step.
- B. Make a member of the **serveradmin** server role the owner of both jobs.
- C. Configure the SQLServerAgent service to use a domain account.
- D. Grant the necessary permissions in the **msdb** database to the job owner's SQL Server login.

Answer: C.

Explanation: By examining the exhibit we see that Service Startup Account, the SQLServerAgent, is configured to run with the system account. This would not give the Jobs network access. The jobs need to access a file share on a corporate file server. We must configure the SQLServerAgent service to be run with a domain account.

Note: The Services Accounts screen is used to assign a logon account to each of the two SQL Server 2000 services, SQL Server and SQL Server Agent. Either the local system or the domain user account can be used. The Local System account does not require a password and does not have network access rights. The Local System account prevents the instance of SQL Server 2000 from interacting with other network servers in. The local system account does allow network access. A domain user account is required to interact with network servers.

Incorrect Answers:

A: ActiveX Script are used to add functionality to Data Transformation Services (DTS) tasks and are script codes used to perform functions that are not available in the other tasks in DTS Designer. The tasks in this scenario are SQLServerAgent tasks and not DTS tasks therefore ActiveX scripts are inappropriate.

B: The sysadmin fixed server role is a server-wide role that allows it members to set server-wide administrative and configuration options. However, the scenario requires access to a file server, which is not part of the SQL Server 2000 server and would thus not fall under the scope of the sysadmin fixed server role.

D: The msdb database is used by SQL Server Agent for scheduling alerts and jobs, and recording

operators.

Question No: 3

You are the administrator of a SQL Server 2000 computer. You create a job to perform several database maintenance tasks. This job must run whenever the server is restarted.

You use SQL Server Enterprise Manager to create a new schedule entry for the job, and specify that the job should start automatically.

You test the job by restarting the server. You check the job's history and discover that the job did not run.

You need the job to run automatically when the server restarts. What should you do?

- A.** Create a stored procedure that uses **sp_start_job** to start the job whenever the MSSQLService service starts.
- B.** Add another schedule entry to the job, specifically that the job start hourly, and then disable the job
- C.** Change the job's category to Database Maintenance
- D.** Configure the SQLServerAgent service to start automatically

Answer: D.

Explanation: Administrative job scheduling is one way of automating administrative tasks. Local jobs or multiserver jobs can be defined to run automatically when SQL Server Agent starts; when CPU utilization of the computer is at a defined level you have defined as idle; at one time on a specific date and time; on a recurring schedule or in response to an alert. SQL Server Agent must be configured to start automatically when SQL Server 2000 starts for a scheduled job to run whenever the server is restarted.

Incorrect Answers:

A: SQL Server Service Manager is used to start, stop, and pause the SQL Server 2000 components on the server while the **sp_start_job** stored procedure instructs SQL Server Agent to execute a job immediately.

B: When a scheduled job is disabled, it is not run until it is re-enabled and a schedule for it to run is defined.

C: The job category under which the job has been created does not affect the scheduling of the job.

Question No: 4

You are the administrator of a SQL Server 2000 computer named CorpSQL. You configure several SQLServerAgent jobs to perform automated maintenance tasks on CorpSQL. The jobs execute properly for several weeks, and then stop executing. You discover that the SQLServerAgent service has stopped.

You attempt to start the SQLServerAgent service. The service does not start. In the Microsoft Windows application event log you discover the following error message: "SQLServerAgent could not be started (reason: Unable to connect to server 'CORPSQL', SQLServerAgent cannot start)". You verify that the MSSQLServer service is running.

You need to start the SQLServerAgent service and execute the maintenance jobs. What should you do?

- A.** Restore the **msdb** database from a backup
Restart the MSSQLServer service, and start the SQLServerAgent service
- B.** Restart the server
Start the MSSQLServer service, and start the SQLServerAgent service
- C.** Configure the SQLServerAgent service to start by using the local system account
- D.** Configure the SQLServerAgent service to start by using the same user account as the MSSQLServer service

Answer: D.

Explanation: The error message in the Windows application event log indicates that SQL Server Agent

could not be started because it was unable to connect to server 'CORPSQL'. As the MSSQLServer service is running, this failure to connect to 'CORPSQL' could be because the SQL Server Agent is attempting to log on to the SQL Server 2000 server from a user account that does not have the required login permissions. If the service does not have the appropriate permissions, certain functionality cannot be accomplished. The MSSQLServer service must run under an account that has all these permissions thus by configuring SQLServerAgent to log on with the same user account we can ensure that the SQLServerAgent will be started.

Note: 1 If services are started by someone who is not a member of the Windows NT local administrators group, the service account must have full control of the main SQL Server directory; full control of the SQL Server database files; the Log on as a service right; full control of registry keys at and below HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\MSSQLServer; selection of the Password Never Expires box; full control of registry keys at and below HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\MSSQLServer; and full control of registry keys at and below HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Perflib.

Incorrect Answers:

A: The msdb database should only be restored from a backup when the msdb database has been damaged. In order to ensure that the SQLServerAgent services do run, we must configure the SQLServerAgent to use the same user account as the MSSQLServer.

B: In order to ensure that the SQLServerAgent services do run, we must configure the SQLServerAgent to use the same user account as the MSSQLServer.

C: If services are started by someone who is not a member of the Windows NT local administrators group, the service account must have full control of the main SQL Server directory; full control of the SQL Server database files; the Log on as a service right; full control of registry keys at and below HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\MSSQLServer; selection of the Password Never Expires box; full control of registry keys at and below HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\MSSQLServer; and full control

of registry keys at and below HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Perflib.

If the service does not have the appropriate permissions, certain functionality cannot be accomplished. The MSSQLServer service must run under an account that has all these permissions thus by configuring SQLServerAgent to log on with the same user account we can ensure that the SQLServerAgent will be started.

Question No: 5

You are the administrator of a SQL Server 2000 computer. You configure SQL server to perform auditing. Audit logs are saved in the C:\MSSQL\Audit folder.

The server functions normally for several days, but then the MSSQL Server Service shuts down without warning. You cannot restart the service.

You need to restart the MSSQL Server service. What should you do?

A. Set the SQLServerAgent service's startup method to **Disabled**.

Start the MSSQLServer service.

B. Start the MSSQLServer service manually by using the **-s** startup option.

C. Start the MSSQLServer service manually by using the **-m** startup option.

D. Delete the audit logs from the C:\MSSQL\Audit folder.

Start the MSSQLServer service.

Answer: D.

Explanation: SQL Server functioned for several days but now it is impossible to start it.

The most likely cause is the size of the auditing log files. The maximum file size for the audit log is fixed at

200 megabytes (MB). When the audit log file reaches 200 MB, a new file will be created and the old file handle will be closed. If the directory fills up, then the instance of Microsoft SQL Server is stopped. The system administrator needs to either free up disk space for the audit log before restarting the instance of SQL Server or restart the instance of SQL Server

Note 1: SQL Server Agent monitors itself and the SQL Server service.

SQL Server Agent can restart the local instance of SQL Server if it has terminated for reasons other than a typical shutdown.

Note 2: Start the MSSQLServer service manually by using the -f startup option would start the service with minimal configuration. It could be used when some configuration errors have been made and the MSSQLServer service isn't able to start.

Incorrect Answers:

A: SQL Server Agent monitors itself and the SQL Server service. Disabling the SQLServerAgent service will not make the MSSQLSERVER service to run.

B: The -s switch allows you to start a named instance of SQL Server 2000. Since the MSSQLServer service isn't able to start it would not start with -s switch either.

C: The audit directory is full. SQL Server would not start with -m option.

Note: Under certain circumstances, you may need to start Microsoft SQL Server in single-user mode. This is the startup option -m. For example, you may want to change server configuration options or recover a damaged master or other system database; both require starting SQL Server in single-user mode.


The following command should be issued to start the MSSQLServer service:

```
net start mssqlserver -m
```

Reference: SQL Server Administrator's companion. Chapter 1 - Starting, Pausing, and Stopping SQL Server

Question No: 6

You are the administrator of a SQL Server 2000 computer. You configure the SQLServerAgent service as shown below:



You configure several SQL Server Agent jobs to perform maintenance tasks. These jobs delete old database records and copy files from the server to file servers on your network.

You discover that none of the jobs will execute. You use SQL Server Enterprise Manager to check the jobs and discover that the SQLServerAgent service will not start. You need to allow the SQL Server Agent service to start and execute the jobs. What should you do?

- A. Configure the SQL Server Agent service to start by using the local system account.
- B. Grant the SQLAgent domain user account **Log on as a service** rights.
- C. Add the SQLAgent domain user account as a Microsoft Windows login account on the server.
- D. Configure permissions on the **master** database to allow full access to the SQLAgent user account.

Answer: B.

Explanation: On Windows NT and Windows 2000 operating systems, SQL Server 2000 and SQL Server Agent are started and run as Windows services. Windows services require 'Log on as a service' rights, which allow a security principal to log on as a service, as a way of establishing a security context. The Local System account on the Windows computer always retains the right to log on as a service. Any service that runs under a separate account must be granted this right.

Incorrect Answers:

A: The local system account does not require a password, does not have network access rights in Windows NT 4.0 and Windows 2000. This restricts SQL Server installation and services from interacting with other network servers.

C: The SQLAgent requires access to the network file server as files must be copied from the SQL Server 2000 server to the file server. Configuring the SQLAgents domain account as a login account on the SQL Server 2000 server does not accomplish this.

D: The SQLAgent requires access to the network file server as files must be copied from the SQL Server 2000 server to the file server. Configuring permissions in the master database does not accomplish this.

SECTION 2: Configure alerts and operators by using SQL Server Agent, System Monitor and Transact-SQL. (4 questions)

Question No: 1

You are the administrator of eight SQL Server 2000 computers. You configure alerts on each server so that various problem conditions will be reported if they occur.

You create 20 operators on one of the servers. You configure these operators by using the e-mail and pager contact information for the employees in your department. You configure the alerts on the server to send e-mail messages and pager messages to the appropriate operators.

You need to configure the same 20 operators on the other seven servers. You want to do this with a minimum amount of administrative time.

What should you do?

A. Detach the **msdb** database from the first server.

Copy the database to the other seven servers.

Attach the database on all eight servers.

B. Use SQL Server Enterprise Manager on a single client computer to create the operators on each server.

C. Use SQL Server Enterprise Manager to script all of the operators.

Use SQL Query Analyzer to run the script on the other seven servers.

D. User snapshot replication to replicate the **msdb** database from the first server to the other seven servers.

Answer: C.

Explanation: Transact-SQL scripts can be used to create defined operators or alerts. Scripting all the predefined operators and alerts and copying them to the servers if the same group of operators is responsible for responding to the same alerts on other servers can save time.

Incorrect Answers:


A: Detaching and attaching databases is useful for moving a database from one computer to another without having to re-create the database and then restore the database backup manually, or to a different physical disk. To move a database to another server or disk detach the database, move the database file(s) to the other server or disk, and attach the database specifying the new location of the moved file(s). When a database is attached, the name and physical location of the primary data file must be specified as it contains the information needed to find the other files comprising the database. Any files that have changed location must be specified in addition to the primary file. The msdb database is a system database that SQL Server Agent uses for scheduling alerts and jobs, and recording operators. Although this could work if the scheduled tasks are the same on all eight servers, it is not the best solution.

B: Using SQL Server Enterprise Manager to create the operators on each server is not the best solution for this scenario, as it would require more administrative time than scripting all of the operators and running the script on the other servers.

D: Replication is used to copy, distribute, and modify data across an enterprise. In snapshot replication the exact data as it was at a specific moment when the snapshot was taken is replicated. Snapshot replication does not monitor the data for updates and is best used as a method for replicating data that changes infrequently or where the most up-to-date data is not a requirement. When synchronization occurs, the entire snapshot is generated and sent to Subscribers.

Question No: 2

You are the administrator of Microsoft Windows 2000 Advanced Server computer. The server is also running SQL Server 2000 and contains a database named Sales. The server is configured as shown in the exhibit.



The sales database is configured by using the default options and increases by 50MB per day. You want an operator to be notified automatically before the database runs out of disk space. You also want to minimize the amount of administrative overhead necessary for this database.

What should you do?

A. Configure systems monitor to log an alert if the hard disk database contains less than 600MB of disk space. Configure SQL server agent to send an e-mail message notifying the operator of the system monitor alert.

B. Create a SQL server agent job that uses the **sp_spaceused** stored procedure each morning. Place the results of the stored procedure in the e-mail message and use SQL mail to send the results to the operator.

C. Set the Maxsize property of sales_data.mdf to 7.5GB and the maxsize property of sales_log.ldf to 1.5GB. Configure SQL server agent to send an e-mail message to an operator when the data file and the transaction log have reached these limits.

D. Use Microsoft Windows explorer to verify the available space on the hard disk each morning. Execute the **sp_spaceused** stored procedure each morning.

Answer: A.

Explanation: The solution with least administrative effort is to set up a system monitor alert to log when disk space is below a certain threshold, here 600MB. Then configure the SQL Server agent to trap this event and send an e-mail notifying the operator when this occurs.

Note: Using SQL Server Agent, you can define alerts and configure responses to alerts.

SQL Server Agent monitors the Windows application log and compares each SQL Server event logged with the alerts that have been defined. If a match is found, an event alert fires. SQL Server Agent can also monitor specific SQL Server performance object counters and fire a performance condition alert when the value for the counter is less than, equal to, or greater than a defined threshold. The Create Alert Wizard can be used to configure the SQL Server agent for alerts.

Incorrect Answers:

B: Getting the size of the database in an e-mail every morning would require some effort of the operators – they would have to read the e-mail. It would be better just to get an e-mail when disk space is starting to run out.

Note: The **sp_spaceused** command displays the number of rows, disk space reserved, and disk space used by a table in the current database, or displays the disk space reserved and used by the entire database.

C: We would like to have an alert before, not when, the database runs out of space. Setting the Maxsize property would not be a good idea when the database increases in size every day.

D: Manually checking the available disk space each morning would require effort, time and cost. It's not solution with minimized administrative effort.

Question No: 3

You are the administrator of a SQL Server 2000 computer. You configure a set of alerts on the server to notify you whenever certain operations fail or encounter errors. The notifications are sent to your Microsoft Windows 2000 Professional computer by using the net send command.

You successfully test the alerts. After several days, however, you stop receiving notifications when operations fail or encounter an error. You verify that the SQL Server Agent service is started.

You need to make sure that you continue to receive alert notifications. What should you do?

- A.** Stop and restart the SQL Server Agent service.
- B.** Clear the application log on the SQL server computer.
- C.** Clear the system log on your Windows 2000 Professional computer.
- D.** Install the SQL Server 2000 administrative tools on your Windows 2000 Professional computer.

Answer: B.

Explanation: The reason for not receiving alerts on the clients could be:

- the SQL Server Agent service isn't running – it is already running here in this scenario.
- the Windows 2000/NT messenger service isn't running. This is not an alternative.
- the event log of the SQL Server computer is full.

Incorrect Answers:

A: You have verified that the SQL Server Agent service is started.

C: The events of the SQL Server are taken place on the server. They are logged in the Event log of the SQL Server computer not the Windows 2000 Professional client computer.

D: There would be no need to install any SQL Server 2000 administrative tools to receive alert notifications.

Question No: 4

You are the administrator of 20 SQL Server 2000 computers. The servers are located in 10 company offices around the world. You work in the New York City office.

You need to configure several dozen alert responses on each server. In the future, you will need to modify these alert responses periodically. You want to reduce the amount of time you spend managing alert responses. You also want to manage the alert responses from the New York City office.

What should you do?

- A.** Use SQL Enterprise manager to manage the alert responses on all company servers.
- B.** Configure alert forwarding on the remote servers.
- C.** Write Transact-SQL scripts that can be used to set up alert responses on the remote servers.
- D.** Write stored procedures that can be used to set up alert responses on the remote servers.

Answer: B.

Explanation: You can forward all SQL Server event messages that meet or exceed a specific error severity level to one instance of SQL Server. You can use event-forwarding to enable centralized alert management for a group of servers.

Reference: BOL: Managing Events

Incorrect Answers:

A: It would be better to administer the alerts on the New York Server only.

C: Event forwarding can be configured without writing any T-SQL code.

D: Event forwarding can be configured without writing any T-SQL code or stored procedures.


SECTION 3: Optimize hardware resource usage. Resources include CPU, disk I/O, and memory.

SUBSECTION A: Monitor hardware resource usage by using System Monitor. (2 questions)

Question No: 1

You are the administrator of a SQL Server 2000 computer that contains a database. Users report that

queries to this database respond slowly. You use system monitor to examine the subsystems on your server and receive the results shown in the exhibit.



You need to modify the server to accelerate query response time. What should you do?

- A. Increase the amount of RAM.
- B. Upgrade to a faster disk subsystem.
- C. Add a faster network adapter.
- D. Add an additional processor.

Answer: D.

Explanation: In Windows 2000 System Monitor, the % Processor Time counter displays the percentage of time that the processor executes a non-Idle thread and is a primary indicator of processor activity. It calculates processor usage by monitoring the time the service was inactive at sample intervals, and then subtracting that value from 100%. A % Processor Time count that is continually above 80% indicates that the CPU is insufficient to cope with the processor load and a CPU upgrade or an additional CPU is required.

Reference: Windows 2000 Server Resource Kit: Performance Monitoring

Incorrect Answers:

A: An average Pages/Sec with 20 or above would indicate that the system would require more memory. By examining the exhibit we that this counter, the green one, only goes over 20 once.

B: A value below 0.3 of the Avg. Disk sec/Transfer counter indicates normal behavior. This seems to be the case in the exhibit.

This counter gives the average disk transfer time.

C: A faulty network adapter could cause the processor to be very busy. This is not the most likely problem though.

Question No: 2

You are the administrator of a Microsoft Windows NT 4.0 computer that is running SQL Server 2000. The computer contains a database named Sales.

The data file and transaction log for the Sales database are located on a 9.1 GB hard disk. Both files are configured for automatic file growth. The data file occupies 7 GB of disk space.

You need to calculate when you will need to add additional disk space. What should you do?

- A.** Configure System Monitor to log disk space utilization
- B.** Configure a database maintenance job to run the DBCC CHECKALLOC statement each night and then to log the job history
- C.** Configure a SQL Server Agent job to execute the **sp_spaceused** stored procedure each day and then to store the results of the procedure to a table in the database
- D.** Configure a SQL Server Agent job to execute the **chkdsk** utility each night at midnight and then to send the results of the utility to the default operator in an e-mail message

Answer: A.

Explanation: The Windows NT 4.0 System Monitor (or Windows 2000 Performance monitor) can be configured to raise an alert when the free disk space drops below a predefined threshold.

Note: System Monitor in Windows 2000 or Performance Monitor in Windows NT 4.0 is used to create charts, alerts, logs, and reports to monitor an instance of SQL Server 2000. Logs can be used to record information on the current activity of selected objects and computers for later viewing and analysis. Data can be collected from multiple systems into a single log file accumulate information on the performance of selected objects on various computers for future analysis. Logs can also be used to track counters over a long period of time, thereby allowing the examination of information more thoroughly and to document system performance.

Incorrect Answers:

B: DBCC CHECKALLOC is used to check the consistency of disk space allocation structures for a specified database. It checks allocation and page usage in a database, including that of indexed views.

C: The sp_spaceused stored procedure displays the number of rows, disk space reserved, and disk space used by a table in the current database, or displays the disk space reserved and used by the entire database.

D: The chkdsk command is a Windows 2000 error detection utility and is used to display a status report for a hard disk. It examines disk space and use for the NTFS and FAT file systems and can also lists errors on the hard disk. If the /f switch is used, chkdsk sends an alert when it encounters an error and corrects the error.

SUBSECTION B: Resolve system bottlenecks by using System Monitor. (2 questions)

Question No: 1

You are the administrator of Microsoft Windows NT server 4.0 computers. The server is the PDC in your domain and also runs SQL Server 2000. The server has four processors and 1GB of RAM. Your network consists of 600 client computers that are running Microsoft Windows millennium edition. Users report that when they log on in the morning their computers respond slowly and occasionally return error messages.

You use system monitor to monitor processor performance and RAM utilization. You discover that the %Total Time and %User Time counters average 90% and the %privileged Time counter averages 20 percent. The available Mbytes counter averages 600MB.

You need to improve server performance for user logons. What should you do?

- A.** Increase the physical RAM on the computer to 2GB.
- B.** Increase the virtual RAM on the computer to 4GB.
- C.** Configure SQL server to use only processors 1, 2, and 3.
- D.** Configure SQL server to use only processors 0, 1 and 2.

Answer: C

Explanation: The system monitor clearly indicates that the processors are overloaded. We need to decide which processors we should use for SQL Server.

Processor 0 is the default CPU for the I/O subsystem. Network Interface Cards (NIC) are assigned to the remaining CPUs, starting from the highest-numbered CPU. The NIC would get processor 3. SQL Server would be using all four processors by default.

The Windows NT/2000/XP operating system use processor 0. In order to avoid the logon problems we should not let SQL Server to use this processor. Instead SQL Server should be configured to use processor 1, 2, and 3.

Note 1: The Affinity mask is used to exclude a processor on a multiprocessor computer from processing SQL Server 2000 threads. Default is equal distribution of SQL Server 2000 processes across all processors.

Note 2: Monitoring the Processor and System object counters provides information about the CPU utilization and helps in determining whether or not a bottleneck exists. The % Total Processor Time is used on multiple CPU Servers to gauge the average activity of the processors and shows the percentage of elapsed time that a processor is busy executing a nonidle thread. Values around 100 percent on a Server computer that processes many client requests indicate that processes are queuing up, waiting for processor time, and causing a bottleneck. Such a sustained high level of processor usage is unacceptable for a server.

Incorrect Answers:


A: The Available Bytes counter indicates the amount of memory that is currently available for use by processes. Low Available Mbytes counts of about 4MB or less indicates an overall shortage of memory or that a program is not releasing memory. In this scenario the Available Mbytes count is at 600MB thus there is no shortage of physical memory.

B: The Available Bytes counter indicates the amount of memory that is currently available for use by processes. Low Available Mbytes counts of about 4MB or less indicates an overall shortage of memory or that a program is not releasing memory. Virtual memory is temporary memory used by a computer to run programs that need more memory than the computer physically has. In this event the paging files on the computer's hard drive are used. In this scenario the Available Mbytes count is at 600MB thus there is no shortage of physical memory and hence there no need to increase the Virtual memory.

D: You could argue that the problem at login time could be solved by configure SQL Server to use processor 0, 1, and 2. This would let the NIC to use processor 3. The operating system is using processor 0 by default and it is more important to decrease the load on this processor instead.

Question No: 2

You are the administrator of a SQL Server 2000 computer that contains a database. Users report that the database responds slowly when they insert data. You use system monitor to examine the server's performance and receive the results shown in the exhibit.



You need to modify the server to improve its performance. What should you do?

- A. Install a larger hard disk.
- B. Increase the amount of RAM.
- C. Increase the size of the virtual memory paging file.
- D. Add an additional processor.

Answer: B.

Explanation: If the Memory Pages/Sec counter value is consistently higher than 5 then the system needs more memory.

Note: Pages/sec is the rate at which pages are read from or written to disk to resolve hard page faults. This counter is a primary indicator of the kinds of faults that cause system-wide delays.

Reference: Optimizing Windows NT for Performance (Q146005)

Incorrect Answers:

A: The system needs more memory not a larger hard disk

A larger hard disk would not improve disk performance.

C: Increasing the size of the page file would not decrease the hard page faults. Hard page faults occur when the system uses the page file instead of the RAM.

D: The system needs more memory not an additional processor.

If Processor %Processor Time value is consistently high (> 80%) suspect the processor.

SECTION 4: Optimize and troubleshoot SQL Server system activity. Activities include cache hits, connections, locks, memory allocation, recompilation, and transactional throughput.

SUBSECTION A: Monitor SQL Server system activity by using traces. (5 questions)

Question No: 1

You are the administrator of a SQL Server 2000 computer. The server contains a database named

Sales. A company e-commerce application uses the Sales database to store sales transactions. The company Web site has grown in popularity, and database utilization has increased. You need to collect data about the utilization of server resources so that you can provide capacity planning. You want to automate the collection process so that information is gathered as quickly as possible.

What should you do?

- A. Configure System Monitor to collect data and store it in a SQL Server table
- B. Create a SQL Server Agent job that executes the **sp_statistics** stored procedure daily and places the results in a text file
- C. Use SQL Profiler to trace server activity and store the results in SQL Server tables
- D. Configure SQL Server alerts to store information in the Windows application event log

Answer: C.

Explanation: SQL Profiler is a graphical tool that allows a system administrator to monitor events in an instance of SQL Server 2000. It can be used to capture and save data about each event to a file or to a SQL Server 2000 table for later analysis. It is possible to filter events so that only events of interest are monitored. This would reduce overhead as monitoring too many events adds overhead to the server and the monitoring process and can cause the trace file or trace table to grow very large, especially when the

monitoring process takes place over a long period of time. SQL Profiler can be used to monitor the performance of an instance of SQL Server; to debug Transact-SQL statements and stored procedures; and to identify slow-executing queries.

Incorrect Answers:

A: SQL Server provides objects and counters that can be used by System Monitor in a Windows 2000 or by Performance Monitor in Windows NT 4.0 to monitor system activity and system performance on computers running an instance of SQL Server. This data cannot be stored in a SQL Server table.

B: SQL Server Agent runs on the server running instances of SQL Server 2000 or earlier versions of SQL Server and is responsible for running SQL Server tasks scheduled to occur at specific times or intervals; detecting specific conditions for which administrators have defined an action, such as alerting someone through pages or e-mail, or a task that will address the conditions; running replication tasks defined by administrators. The **sp_statistics** stored procedure returns a list of all indexes and statistics on a specified table or indexed view.

D: The Windows application logs provide an overall picture of events that occur on the Windows NT 4.0 and Windows 2000 systems, as well as events in SQL Server and SQL Server Agent. Event Viewer can be used to view the Windows application log and to filter the information.

Question No: 2

You are the administrator of a SQL Server 2000 computer. The server contains two databases. A client/server application accesses one of the databases. A different client/server application accesses the other database. Each application uses several stored procedures to query and update its database. Users report that both applications respond slowly.

You use SQL profiler to monitor activity on the server. You receive results of the trace as shown in the following screen shot:

EventClass	TextData
SQL:BatchCompleted	IF @@TRANCOUNT>0 COMMIT TRAN
SQL:BatchCompleted	Update Customer Set CustomerZip=.....
SQL:BatchCompleted	Update Products Set SKU='23232323.....
SQL:BatchCompleted	IF @@TRANCOUNT>0 COMMIT TRAN
SQL:BatchCompleted	Update Customer Set CustomerZip=.....

You need to improve the performance of the applications. What should you do?

- A. Modify the applications so that they use views to query data.
- B. Modify the stored procedures so that they select data into temporary tables.
- C. Re-create the stored procedures so that they do not contain the WITH RECOMPILE option.

D. Disable the Auto update statistics option in both databases.

Answer: B.

Explanation: We see that the same data is queried in two separate statements. This could result in locking and degraded performance. By using temporary tables, when appropriate like in this scenario, the locking could be decreased and the same data could be re-used. Both these benefits would improve performance.

Note: The @@TRANCOUNT function returns the number of active transactions for the current connection.

Incorrect Answers:

A: In general views would not improve performance compared to stored procedures.

C: A stored procedure can be created using WITH RECOMPILE to ensure that its execution plan will be recompiled for every call and will never be reused. Creating a stored procedure WITH RECOMPILE can be useful if procedures take parameters and the values of the parameters differ widely, resulting in a need for different execution plans to be formulated.

Recreating the stored procedure with the WITH RECOMPILE option could improve performance somewhat, but it would be more beneficial to decrease the number of locks.

D: Auto update statistics would benefit performance in databases where the data changes over time.

Disabling auto update statistics would not improve performance.

Question No: 3

You are the administrator of two SQL Server 2000 computers. Each server contains data for one of your company's two regional divisions. Company employees currently use a single data entry application on client computers.

After the new application is installed, query response times slow. You want to monitor server traffic during business hours to help diagnose this problem.

What should you do?

A. Run SQL Profiler, and create one trace to monitor both servers. Configure the trace to include data from the new data entry application only.

B. Run two instances of SQL Profiler, and create a trace to monitor each server. Configure the trace to include data from both data entry applications.

C. Run the sqldiag utility from a command prompt.

Write diagnostic data to a text file.

D. Execute the sp_monitor stored procedure.

Save the output to a text file.

Answer: B.

Explanation: We want to monitor the traffic between the two servers and between the servers and the clients. SQL Profiler can give us detailed information of the performance of SQL Server events, including server traffic between.

Note: SQL Profiler is a graphical tool that allows system administrators to monitor events in an instance of SQL Server 2000. The events monitored by SQL Profiler can be filtered to include only those events that are of interest and can be captured and saved to a file or SQL Server table for analysis at later date. SQL Profiler can be used to monitor the performance of an instance of SQL Server; debug Transact-SQL statements and stored procedures; identify slow-executing queries; and to test SQL statements and stored procedures in the development phase of a project by single-stepping through statements to confirm that the code works as expected. Once captured and saved to file, the captured events can be replayed on a test system. This will allow users to continue using the production system without interference.

Incorrect Answers:

A: SQL Profiler should be configured to trace the old data entry application as well.

C: The sqldiag utility gathers and stores diagnostic information and the contents of the query history trace. The output file includes error logs and output from the sp_configure stored procedure as well as other version information.

D: If sp_monitor should be used then it would have to be used at both SQL Servers, not only at one of them.

Note: SQL Server uses a series of functions to keep track of how much work it has done. The current values returned by these functions are displayed by executing the sp_monitor stored procedure. The values indicate shows how much work was done since the last time the stored procedure was run.

sp_monitor monitors the Volume of work performed during periods of time. Sp_monitor reports statistics, including CPU use, I/O use, and idle time since last execution of sp_monitor.

More specifically, some of the output generated by sp_monitor:

packets_received Number of input packets read by SQL Server.

packets_sent Number of output packets written by SQL Server.

packet_errors Number of errors encountered by SQL Server while reading and writing packets.

total_read Number of reads by SQL Server.

total_write Number of writes by SQL Server.

total_errors Number of errors encountered by SQL Server while reading and writing.

connections Number of logins or attempted logins to SQL Server.

Question No: 4

You are the administrator of a SQL Server 2000 computer. The server contains a database named Orders, which is used to record customer orders for the products your company sells.

Your company's order volume exceeds 1 million orders per day. Each order uses approximately 100 KB of space in the database. Users report that the database responds slowly when they enter new orders.

You use SQL Profiler to monitor the activity on the database and receive the data shown in the exhibit.

You need to modify the database to improve performance. What should you do?

A. Double the size of the data file

B. Configure the database to automatically grow by 10 percent

C. Separate the database into two physical files

D. Increase the size of the transaction log file

Answer: B.

Explanation: The database is growing too frequently. This decreases performance. We must configure it to grow faster and less frequently. This can be done by configuring it to grow automatically by 10 percent.

Incorrect Answers:

A: Doubling the size of the data file would virtually eliminate the file growth of the database, but it too drastic since it would require lot of expensive disk space that wouldn't be used.

C: Splitting the database into two files would still make it grow as frequently as before.

D: The data file not the transaction log file is growing.

Question No: 5

You are the administrator of a SQL Server 2000 computer. The server contains your company's investment-tracking database. Each day, 100 operators make approximately 5,000 changes to customer investments. In addition, daily and monthly reports are created from the investment data. Developers at your company need to optimize the database application. They need a sample of database query activity to discover if they can speed up the transactions. The developers also want to replay the sample on another SQL Server computer.

You need to capture the sample, but you want to minimize any increase to the workload of the server. What should you do?

A. Run SQL Profiler on a client computer.

Configure SQL Profiler to monitor database activity, and log data to a .trc file.

B. Run SQL Profiler on the server.

Configure SQL Profiler to monitor database activity, and log data to a database table.

C. Run System Monitor on a client computer.

Configure System Monitor to monitor database activity, and log data to a .blg file.

D. Start SQL Server from a command prompt.

Specify trace flag 1204 to enable verbose logging.

Answer: A.

Explanation: SQL Profiler can either save the trace to .trc trace file or to an SQL Server table. Saving it to a table would increase the workload of server and in this scenario this should be avoided. Therefore the trace should be saved to a .trc file.

Incorrect Answers:

B: Capturing trace data to a table is slower than capturing to a file. It would increase the workload of the SQL Server.

C: System monitor cannot monitor specific database activity. SQL Profiler is better suited for that purpose.


D: Running SQL Server with trace flag 1204 would return the type of locks participating in the deadlock and the current command affected. SQL Profiles would give more useful information.

SUBSECTION B: Monitor SQL Server system activity by using System Monitor. (3 questions)

Question No: 1

You are the administrator of a SQL Server 2000 computer. The server contains a database named MedicalRecords. The data processing department in your company adds new business rules to the database's stored procedures. After the rules are added, users report that the medical records application receives time-out errors.

You use System Monitor to analyze the server and receive the results shown in the exhibit.



You need to tune the stored procedures to improve performance. What should you do?

- A. Add indexes to the fields that are used in the joins
- B. Remove the WITH RECOMPILE option from the stored procedure
- C. Remove the stored procedures and alter the application so that it uses ad hoc queries to obtain data
- D. Change the transaction isolation level to Repeatable Read

Answer: D.

Explanation: By examining the exhibit we see that there is quite a few Lock Requests/Sec. By changing the isolation transaction level to the highest possible, repeatable read, we will reduce the number of locks by locking at table level instead of at a more granular row level. This is called lock escalation and it improves performance by reducing the number of locks.

Note 1: Lock escalation is the process of converting many fine-grain locks into fewer coarse-grain locks, reducing system overhead. SQL Server automatically escalates row locks and page locks into table locks when a transaction exceeds its escalation threshold.

Note 2: SQL Server has multigranular locking that allows different types of resources to be locked by a transaction. To minimize the cost of locking, SQL Server automatically locks resources at a level appropriate to the task. Locking at a smaller granularity, such as rows, increases concurrency but has a higher overhead because more locks must be held if many rows are locked. Locking at a larger granularity, such as tables, are expensive in terms of concurrency because locking an entire table restricts access to any part of the table by other transactions, but has a lower overhead because fewer locks are being maintained.

Reference:

BOL, Lock Escalation

BOL, Understanding Locking in SQL Server

Incorrect answers:

A: By examining the exhibit we see that only Index Searches are used. No table scans are used at all (Full Scans/sec). There is no need for additional indexes.

B: There is no recompiling problem. By examining the exhibit we see that SQL Re-Compilations/Sec is 0. Removing the WITH RECOMPILE option from the stored procedure would have no effect.


C: Stored procedures are preferred to ad hoc queries. Ad hoc queries wouldn't increase performance, at

the contrary stored procedures would probably be more efficient.

Question No: 2

You are the administrator of a SQL Server 2000 computer. The server contains a database named Sales. Users report that the database times out when they attempt to perform queries.

You use System Monitor to analyze the server and receive the results shown in the exhibit.



You need to discover why database performance is slow. What should you do?

- A. Use the current Activity window to discover who is holding locks in the database
- B. Use SAL Profiler to discover who is scanning tables in the database
- C. Measure disk activity by running the **diskperf /y** command, restarting, then using the **% Disk Time** counter in System Monitor
- D. Use the **sp_monitor** stored procedure to measure the activity of the server

Answer: A.


Explanation: By examining the exhibit we see that there is an excessive amount of Lock Requests/Sec. The current Activity windows in Enterprise Manager is a good place to start troubleshooting locks.

Incorrect Answers:

- B:** This is a locking problem not a table scanning problem.
- C:** This is a locking problem not a disk performance problem.
- D:** The Sp_monitor command displays general statistics about SQL Server. It doesn't provide any information on locks and cannot be used to troubleshoot locks.

Question No: 3

System Monitor Exhibit



You work as database administrator at Testig.in. You handle one SQL Server 2000 computer. You notice that queries times out. You troubleshoot the server through system monitor (see exhibit). Which action should you take? (Select two.)

- A. Use SAL Profiler
- B. Use the current Activity window in the Enterprise manager...
- C. Run the **diskperf /y** command, restart the server, then...
- D. Run the **sp_monitor** stored procedure...
- E. ...to discover who is scanning tables in the database.
- F. ...and use the **% Disk Time** counter in System Monitor.
- G. ...to detect who is holding the locks.
- H. ...to measure the activity of the server

Answer: B, G.

Explanation: By examining the exhibit we see that there is an excessive amount of Lock Requests/Sec. The current Activity window in Enterprise Manager is a good place to start troubleshooting locks.

Incorrect Answers:

- A:** Sal Profiler is not the proper tool here.
- C, F, H:** This is a locking problem not a performance problem.
- D:** The Sp_monitor command displays general statistics about SQL Server. It doesn't provide any information on locks and cannot be used to troubleshoot locks.
- E:** This is a locking problem not a table scanning problem.