



www.chinatag.com

CHINATAG

70-228

Installing, Configuring and Administering
Microsoft SQL Server 2000 Enterprise Edition

Q&A

DEMO Version

Copyright (c) 2003 Chinatag LLC. All rights reserved.

Important Note Please Read Carefully

For demonstration purpose only, this free version Chinatag study guide contains **10** full length questions selected from our full version products which have more than **200** questions each.

This Study guide has been carefully written and compiled by Chinatag certification experts. It is designed to help you learn the concepts behind the questions rather than be a strict memorization tool. Repeated readings will increase your comprehension.

For promotion purposes, all PDF files are **not** encrypted. Feel free to distribute copies among your friends and let them know Chinatag website. Our IT certification products start at only **\$7.99**.

Study Tips

This product will provide you questions and answers along with detailed explanations carefully compiled and written by our experts. Try to understand the concepts behind the questions instead of cramming the questions. Go through the entire document at least twice so that you make sure that you are not missing anything.

Latest Version

We are constantly reviewing our products. New material is added and old material is revised. Free updates are available for 180 days after the purchase. You should check the products page on the <http://www.chinatag.com> website for an update 3-4 days before the scheduled exam date.

Please tell us what you think of our products. We appreciate both positive and critical comments as your feedback helps us improve future versions. Feedback on specific questions should be send to feedback@chinatag.com.

Thanks for purchasing our products and look forward to supplying you with all your Certification training needs.

Good studying!

Technical and Support Team
Chinatag LLC.

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 battery-backed 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 restart 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 CLEANTABLE 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.

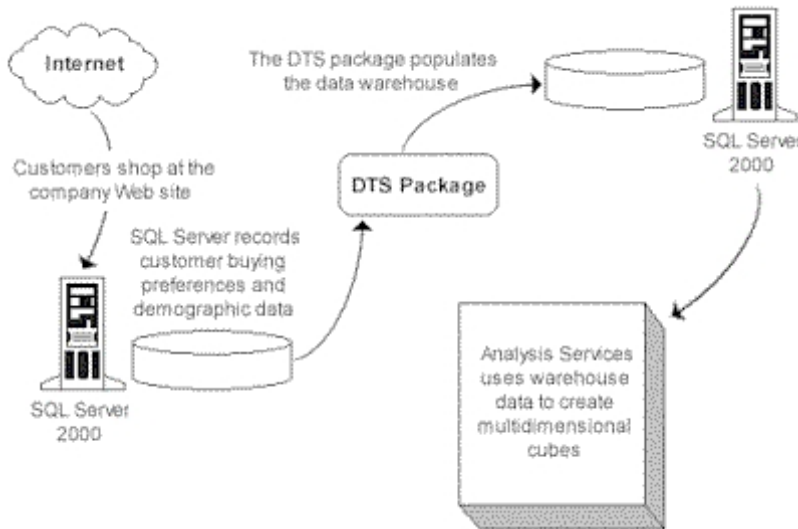
Incorrect Answers:

- 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 CLEANTABLE 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 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 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 bulk-logged 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 back up 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 use. This can track and record and 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: 4

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 resource-intensive to run all the agents at one site or all at the Distributor; and when Subscribers are autonomous, disconnected, and/or mobile.

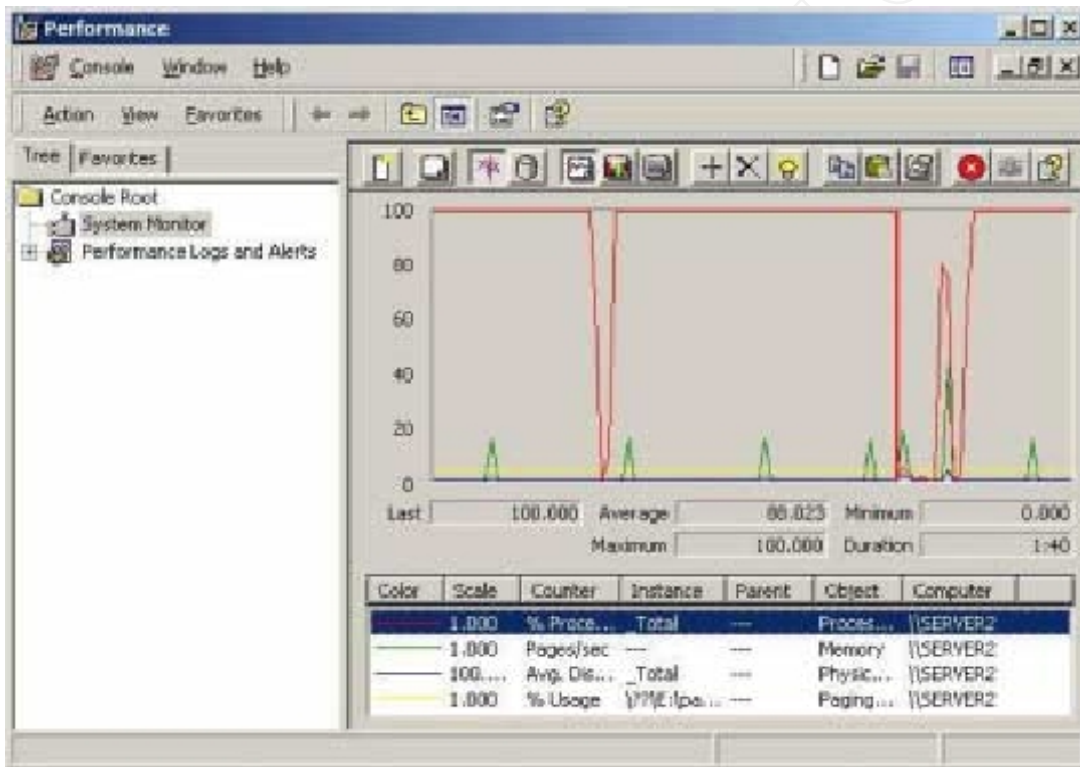
Incorrect Answers:

A: Creating on 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: 5

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: 6

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: 7

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.

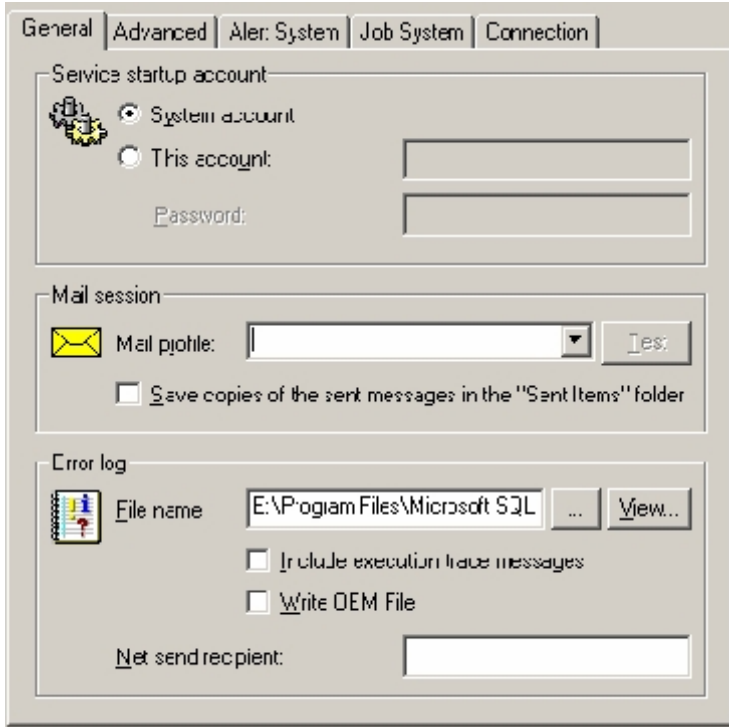
Note correlated subquery: A correlated subquery is an extension of a subquery, which is a SELECT query that returns a single value and is nested inside a SELECT, INSERT, UPDATE, or DELETE statement, or inside another subquery and can be used anywhere an expression is allowed. With a correlated subquery, 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: 8

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 its 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: 9

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.

Question No: 10

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.