

100% Money Back
Guarantee

Vendor: Microsoft

Exam Code: 70-450

Exam Name: PRO:MS SQL Serv 08,Design,Optimize,
and Maintain DB Admin Solu

Version: Demo

Question No : 1

You administer a SQL Server 2008 instance.

You plan to deploy a new database that has the following capacity requirements:

275 GB for the database data file

50 GB for the transaction log file

The storage array has six 100-GB disk drives available for the database. The disks are attached to a redundant array of independent disks (RAID) controller that supports RAID levels 0, 1, 5, and 10.

The write performance of the transaction log needs to be maximized. The database and transaction log files must be protected in the event of a drive failure.

You need to design the storage system.

Which storage configuration should you use?

- A. A single RAID 5 volume
- B. A single RAID 10 volume
- C. A RAID 0 volume and a RAID 5 volume
- D. A RAID 1 volume and a RAID 5 volume

Answer: D

Explanation:

Question No : 2

You administer a SQL Server 2008 instance.

The instance is located on a four-processor, quad-core server. The server frequently experiences CPU pressure.

The instance contains a very large mission-critical database that is used continuously.

You need to ensure that online index rebuilds do not consume all available CPU cycles.

Which configuration option should you use?

- A. affinity mask
- B. affinity I/O mask
- C. max degree of parallelism
- D. optimize for ad hoc workloads

Answer: C

Explanation:

Question No : 3

You administer a SQL Server 2008 infrastructure.

An instance runs on a computer that has eight quad-core processors and 128-GB RAM. Four different applications use the instance.

The instance experiences a low number of CXPACKET waits. The instance also experiences a large number of lazy writer waits.

You need to optimize the performance of the instance.

What should you do?

- A. Configure the Resource Governor.
- B. Configure the Windows System Resource Manager.
- C. Configure software non-uniform memory access (soft-NUMA).
- D. Configure an increase in the maximum degree of parallelism option.

Answer: C

Explanation:

Failover clustering in SQL Server provides high—availability support for an entire SQL Server instance. For example, you can configure a SQL Server instance on one node of a failover cluster to fail over to any other node in the cluster during a hardware failure, operating system failure, or a planned upgrade.

A failover cluster is a combination of one or more nodes (servers) with two or more shared disks, known as a resource group. The combination of a resource group, along with its network name, and an internet protocol (IP) address that makes up the clustered

application or server, is referred to as a failover cluster or a failover cluster. A SQL Server failover cluster appears on the network as if it were a single computer, but has functionality that provides failover from one node to another if the current node becomes unavailable. A failover cluster appears on the network as a normal application or single computer, but it has additional functionality that increases its availability.

Failover clustering has a new architecture and new work flow for all Setup scenarios in SQL Server 2008. The two options for installation are Integrated installation and Advanced Enterprise installation. Integrated installation creates and configures a single-node SQL Server failover cluster instance. Additional nodes are added using add node functionality in Setup. For example, for Integrated installation, you run Setup to create a single-node failover cluster. Then, you run Setup again for each node you want to add to the cluster. Advanced Enterprise installation consists of two steps. The Prepare step prepares all nodes of the failover cluster to be operational. Nodes are defined and prepared during this initial step. After you prepare the nodes, the Complete step is run on the active nodes—the node that owns the shared disk—to complete the failover cluster instance and make it operational.

When to Use Failover Clustering

Use failover clustering to:

- Administer a failover cluster from any node in the clustered SQL Server configuration. For more information, see [Installing a SQL Server 2008 Failover Cluster](#).
- Allow one failover cluster node to fail over to any other node in the failover cluster configuration.

For more information, see [Installing a SQL Server 2008 Failover Cluster](#).

- Configure Analysis Services for failover clustering. For more information, see [How to: install](#)

[Analysis Services on a failover cluster](#).

- Execute full-text queries by using the Microsoft Search engine with failover clustering. For more information, see [Using SQL Server Tools with Failover Clustering](#).

Question No : 4

You administer a SQL Server 2008 instance.

The instance runs on a computer that has the following features:

You plan to implement a transactional database on the instance. The database is expected to have a high volume of INSERT, UPDATE, and DELETE activities. The activities include creation of new tables.

You need to optimize the performance of the database by maximizing disk bandwidth and reducing the contention in the storage allocation structures.

What should you do?

- A. Create multiple data files for the database.
- B. Place database and log files on the same volume.
- C. Configure the affinity mask option appropriately.
- D. Configure the affinity I/O mask option appropriately.

Answer: A

Explanation:

Question No : 5

You administer a SQL Server 2008 instance.

You plan to deploy a new database to the instance. The database will be subject to frequent inserts and updates. The database will have multiple schemas. One of the schemas will contain a large amount of read-only reference data.

You need to design the physical database structure for optimal backup performance.

What should you do?

- A. Create the database by using a single data file and a single log file.
- B. Create the database by using a single data file and multiple log files.
- C. Create the database by using a single log file and multiple filegroups.
- D. Create the database by using a single log file and a filegroup that has multiple data files.

Answer: C

Explanation:

Reasons to take database snapshots include:

- Maintaining historical data for report generation.

Because a database snapshot provides a static view of a database, a snapshot can extend access to data from a particular point in time. For example, you can create a database snapshot at the end of a given time period (such as a financial quarter) for later reporting. You can then run end-of-period reports on the snapshot. If disk space permits, you can also maintain end-of-period snapshots indefinitely, allowing queries against the results from these periods; for example, to investigate organizational performance.

- Using a mirror database that you are maintaining for availability purposes to offload reporting.

Using database snapshots with database mirroring permits you to make the data on the mirror server accessible for reporting. Additionally, running queries on the mirror database can free up resources on the principal. For more information, see Database Mirroring and Database Snapshots.

- Safeguarding data against administrative error.
- Before doing major updates, such as a bulk update or a schema change, create a database snapshot on the database protects data. If you make a mistake, you can use the snapshot to recover by reverting the database to the snapshot. Reverting is potentially much faster for this purpose than restoring from a backup; however, you cannot roll forward afterward.

Question No : 6

You administer a SQL Server 2008 instance.

You deploy a new database named Engineering. The Engineering database manages large documents that will be revised occasionally.

You need to design a table structure that allows fast read access. You also need to minimize storage space requirements.

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

- A.** Use NTFS file system compression on the volume.
- B.** Use row-level compression on the document table.
- C.** Use varbinary(MAX) data type with FILESTREAM storage.

D. Enable row-level compression on all columns that use the vardecimal() data type.

Answer: A,C

Explanation:

Question No : 7

You administer a SQL Server 2008 Reporting Services (SSRS) instance.

You plan to design an authentication method for SSRS reports that use SQL Server data connections.

You need to ensure that all SSRS reports communicate to the SQL Server instance by using the Kerberos authentication protocol.

What should you do?

- A. Configure the default Web site on the instance to use the Integrated Security type.
- B. Configure the application pool used by the default Web site to use a domain account.
- C. Configure the default Web site on the instance to use the Secure Sockets Layer (SSL) protocol.
- D. Register a service principal name (SPN) in the Active Directory directory service for the SSRS instance.

Answer: D

Explanation:

Question No : 8

You administer a SQL Server 2008 instance.

You plan to design the security requirements for a new database application that will be deployed to the instance.

The new database contains a table that is created by using the following code segment:

```
CREATE TABLE EmpBonusPlan
(
  EmployeeID INT NOT NULL IDENTITY(1,1),
  NumOptions INT NOT NULL,
  BaseSalary MONEY NOT NULL,
  BonusPlan VARBINARY(MAX) FILESTREAM
)
```

You need to ensure that the BonusPlan column is protected from unauthorized access by using the most secure method.

What should you do?

- A. Use Transparent data encryption.
- B. Use the Trustworthy option for the database.
- C. Use the Advanced Encryption Standard encryption on all columns in the database.
- D. Use the NTFS file system security and limit the access of the database files to the SQL Server 2008 Service account.

Answer: D

Explanation:

Question No : 9

You administer a SQL Server 2008 instance.

The instance hosts a database that is used by a Web-based application. The application processes 15,000 transactions every minute.

A table in the database contains a column that is used only by the application. This column stores sensitive data.

You need to store the sensitive data in the most secure manner possible. You also need to ensure that you minimize the usage of memory space and processor time.

Which encryption type should you use?

-
- A. Symmetric key encryption
 - B. Asymmetric key encryption
 - C. Certificate-based encryption
 - D. Transparent data encryption

Answer: A

Explanation:

---Jeff---

Symmetric key encryption is the most commonly recommended option for SQL Server 2008 encryption.

It provides security without overly taxing the server like asymmetric encryption might.

Certificates have a similar architecture (public-private key).

TDE is not appropriate, as it would encrypt the entire database.

<http://technet.microsoft.com/en-us/library/cc278098%28SQL.100%29.aspx?ppud=4>

<http://dotnetslackers.com/articles/sql/IntroductionToSQLServerEncryptionAndSymmetricKeyEncryptionTutorial.aspx>

Question No : 10

You administer a SQL Server 2008 instance.

You plan to deploy a third-party database application to the instance. The application uses stored procedures that are developed by using SQL CLR integration.

The application must be configured to enable the EXTERNAL_ACCESS code access security setting.

You need to ensure that the application is deployed to the instance without loss of functionality.

What should you do first?

- A. Replace read-write non-static fields in the code with static fields.
- B. Replace read-only static fields in the code with read write static fields.
- C. Use the peverify.exe PEVerify tool to verify whether the code meets the type-safety requirements.
- D. Use the regasm.exe assembly registration tool to register the assembly on the server before deployment.

Answer: C

Explanation:

Question No : 11

You administer a SQL Server 2008 instance.

You plan to design the security requirements for a new database application. The application uses a code segment that contains the following components:

A method that accesses the registry on the SQL Server

A method that accesses the file system on a network file server

A class definition that uses public static fields

The code segment uses SQL CLR integration and is implemented as a single assembly.

You need to ensure that the application is successfully deployed to the instance.

What should you do?

- A. Use the SAFE code access security for the assembly.
- B. Replace all public static fields with public fields.
- C. Replace all public static fields with public static read-only fields. Use the EXTERNAL_ACCESS code access security for the assembly.
- D. Replace all public static fields with public static read-only fields. Register the assembly by using the regasm.exe utility before deployment.

Answer: C

Explanation:

Question No : 12

You administer a SQL Server 2008 instance that will host a new database application.

You plan to design the security requirements for the application. Each application user has a unique login to the SQL Server 2008 server.

The application database contains stored procedures to execute stored procedures in the MSDB database.

The stored procedures in the MSDB database schedule SQLAgent jobs.

You need to ensure that the stored procedures in the MSDB database are executed by using the security context of the application user.

What should you do?

- A. Add each user to the public role in the MSDB database.
- B. Add each user to the db_dtsltduser database role in the MSDB database.
- C. Configure the new database to use the TRUSTWORTHY option, and then add each user to the MSDB database.
- D. Configure the MSDB database to use the TRUSTWORTHY option, and then add each user to the MSDB database.

Answer: C

Explanation:

---Jeff---

If this ever actually comes up, you may want to re-think your application design.

Presumably, the application users are logging in to the app database. Since this database is trusted, and the users exist in MSDB, the users will be able to schedule jobs via the mechanism in the question... which probably is not a good idea.

Question No : 13

You administer a SQL Server 2008 instance. The instance hosts a new database application.

You plan to design the data security strategy for the application. You will use Snapshot replication to replicate the data to another instance of SQL Server 2008.

You need to ensure that all stored data is encrypted by using the least amount of administrative effort.

What should you do?

-
- A. Enable encrypted connections between the instances
 - B. Enable Transparent data encryption for the MSDB database on the two instances
 - C. Enable Transparent data encryption for the Publisher, Distribution, and Subscriber databases
 - D. Enable certificate-based encryption for the Publisher, Distribution, and Subscriber databases

Answer: C

Explanation:

Question No : 14

You administer SQL Server 2008 servers on a single site.

The site contains the following two SQL Server 2008 instances:

Each instance hosts a single application.

You need to recommend a high-availability solution for the site to meet the following business requirements:

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

- A. Replication
- B. Log shipping
- C. Failover clustering
- D. Database snapshot

Answer: A,B

Explanation:

---Jeff---

Snapshots are not a high availability solution no matter how many times they are included in questions about high availability.

Failover Clustering would require the addition of a shared storage array.

That leaves log shipping and replication. Log shipping is the weakest of high availability options in SQL Server 2008.

Peer-to-peer replication is viable.

Question No : 15

You administer SQL Server 2008 instances at three sites.

All sites share a mission-critical database. The business requirements specify that users at each site must be able to perform the following tasks:

You need to provide a high-availability solution that meets the business requirements.

Which solution should you implement?

- A. Failover clustering
- B. Peer-to-Peer replication
- C. Asynchronous database mirroring without a witness server
- D. Log shipping to servers at two of the sites to provide read-only copies of data

Answer: B

Explanation:

---Jeff---

To minimize latency, users need a local copy of the database.

Failover clustering doesn't provide multiple copies.

Database mirroring only supports one destination.

Database mirroring and log shipping leave the secondary or secondaries in a restoring state, so data is not easily accessible.

Peer-to-peer replication is the only viable solution.

Question No : 16

You administer a SQL Server 2008 infrastructure.

An instance hosts a business-critical database that must be continuously available to the users without data loss.

The database includes Filestream data.

You need to implement a high-availability solution for the site.

Which solution should you use?

- A. Failover clustering
- B. Database snapshot
- C. Asynchronous database mirroring
- D. Synchronous database mirroring with a witness server

Answer: A

Explanation:

Failover clustering in SQL Server provides high—availability support for an entire SQL Server instance. For example, you can configure a SQL Server instance on one node of a failover cluster to fail over to any other node in the cluster during a hardware failure, operating system failure, or a planned upgrade.

A failover cluster is a combination of one or more nodes (servers) with two or more shared disks, known as a resource group. The combination of a resource group, along with its network name, and an internet protocol (IP) address that makes up the clustered application or server, is referred to as a failover cluster or a failover cluster instance. A SQL Server failover cluster appears on the network as if it were a single computer, but has functionality that provides failover from one node to another if the current node becomes unavailable. A failover cluster appears on the network as a normal application or single computer, but it has additional functionality that increases its availability.

Failover clustering has a new architecture and new work flow for all Setup scenarios in SQL Server 2008. The two options for installation are Integrated installation and Advanced!Enterprise installation. Integrated installation creates and configures a single-node SQL Server failover cluster instance. Additional nodes are added using add node functionality in Setup. For example, for Integrated installation, you run Setup to create a single-node failover cluster. Then, you run

Setup again for each node you want to add to the cluster. Advanced!Enterprise installation consists of two steps. The Prepare step prepares all nodes of the failover cluster to be operational. Nodes are defined and prepared during this initial step. After you prepare the nodes, the Complete step is run on the active nodes—the node that owns the shared disk—to complete the failover cluster instance and make it operational.

When to Use Failover Clustering

Use failover clustering to:

-
- Administer a failover cluster from any node in the clustered SQL Server configuration. For more information, see [Installing a SQL Server 2008 Failover Cluster](#).
 - Allow one failover cluster node to fail over to any other node in the failover cluster configuration.

For more information, see [Installing a SQL Server 2008 Failover Cluster](#).

- Configure Analysis Services for failover clustering. For more information, see [How to: install](#)

[Analysis Services on a failover cluster](#).

- Execute full-text queries by using the Microsoft Search engine with failover clustering. For more information, see [Using SQL Server Tools with Failover clustering](#).

Question No : 17

You administer four SQL Server 2008 instances. Each instance hosts a single database application.

You plan to migrate all four instances to a new SQL Server failover cluster.

The four instances are configured as shown in the following table.

Instance Name	Configuration
Instance 1	8 processors that have 16-GB RAM
Instance 2	4 processors that have 8-GB RAM
Instance 3	2 processors that have 16-GB RAM
Instance 4	4 Processors that have 8-GB RAM

The four instances are fully optimized and have no spare CPU cycles or extra memory.

The new cluster will host all four databases on a single virtual cluster IP address.

You need to ensure that the new cluster is configured to handle the workload of all the database applications by using the minimum amount of hardware resources.

Which cluster configuration should you use?

A. Four-node active/active/active/active cluster that has each node containing a minimum of 9 processors and a 25-GB memory

-
- B.** Two-node active/active cluster that has each node containing a minimum of 9 processors and a 25-GB memory
 - C.** Two-node active/passive cluster that has each node containing a minimum of 18 processors and a 50-GB memory
 - D.** Four-node active/active/active/passive cluster that has each node containing a minimum of 18 processors and a 50-GB memory

Answer: C

Explanation:

---Jeff---

Clustering provides protection against server hardware failure by using multiple systems (nodes) that share a disk array.

It does NOT protect against that disk array failing.

Active/Passive means that one node is hosting the SQL Server instance, while the other is sitting around waiting for the first one to fail. It's the preferred solution, but means that both nodes have to be powerful enough to host the entire SQL server instance.

Two-node implementations are the most common and least expensive. More nodes can be added to distribute the load.

Since a requirement mentioned is to minimize hardware usage, a two-node option is preferable. Active/Passive is generally recommended by Microsoft as well. Also, the question states that each instance is currently EXACTLY powerful enough to perform its duties without any overhead. This would make splitting the hardware into two active nodes impossible, as the values for processor and RAM don't add up neatly enough to split the instances 50/50 based on existing requirements.

Required reading: Start at SQL-Server-Performance, move on to MSSQLTips, finish with the white paper (PDF).

<http://www.sql-server-performance.com/2002/clustering-intro/>

<http://www.mssqltips.com/sqlservertip/1554/sql-server-clustering-active-vs-passive/>

<http://www.google.com/url?sa=t&source=web&cd=1&ved=0CBIQFjAA&url=http%3A%2F%2Fdownload.>

[microsoft.com%2Fdownload%2F6%2F9%2FD%2F69D1FEA7-5B42-437A-B3BA-A4AD13E34EF6%](http://microsoft.com%2Fdownload%2F6%2F9%2FD%2F69D1FEA7-5B42-437A-B3BA-A4AD13E34EF6%2FSQLServer2008FailoverCluster.docx&rct=j&q=sql%20server%202008%20cluster%20microsoft%20white%20paper&ei=7T2rTOG8EcWBIAfN8fnEDA&usg=AFQjCNGI1hQlgs4JkiNAfD6zVvVI1-8UuA&cad=rja)

[2FSQLServer2008FailoverCluster.docx&rct=j&q=sql%20server%202008%20cluster%20microsoft%20white%](http://microsoft.com%2Fdownload%2F6%2F9%2FD%2F69D1FEA7-5B42-437A-B3BA-A4AD13E34EF6%2FSQLServer2008FailoverCluster.docx&rct=j&q=sql%20server%202008%20cluster%20microsoft%20white%20paper&ei=7T2rTOG8EcWBIAfN8fnEDA&usg=AFQjCNGI1hQlgs4JkiNAfD6zVvVI1-8UuA&cad=rja)

[20paper&ei=7T2rTOG8EcWBIAfN8fnEDA&usg=AFQjCNGI1hQlgs4JkiNAfD6zVvVI1-8UuA&cad=rja](http://microsoft.com%2Fdownload%2F6%2F9%2FD%2F69D1FEA7-5B42-437A-B3BA-A4AD13E34EF6%2FSQLServer2008FailoverCluster.docx&rct=j&q=sql%20server%202008%20cluster%20microsoft%20white%20paper&ei=7T2rTOG8EcWBIAfN8fnEDA&usg=AFQjCNGI1hQlgs4JkiNAfD6zVvVI1-8UuA&cad=rja)

You administer a SQL Server 2008 instance.

The instance hosts a database configured by using high-safety mirroring operation mode along with a witness server.

The witness server is experiencing memory failure and will be offline for four hours.

You need to reconfigure the mirroring strategy to minimize the risk of database unavailability.

What should you do?

- A. Pause database mirroring.
- B. Use asynchronous operating mode.
- C. Remove database mirroring.
- D. Remove the witness server from the mirroring session.

Answer: D

Explanation:

Question No : 19

You administer a SQL Server 2008 instance.

The instance is one of three servers in a peer-to-peer transactional replication topology. The publication contains a table named Orders. The Orders table contains 200 GB of data across multiple partitions.

A batch process bulk loads 10 GB of data to the Orders table periodically.

You need to design a replication strategy to minimize latency.

What should you do?

- A. Configure the Distributor agent to replicate transactions continually.
- B. Configure the Distributor agent to change the commit batch threshold.
- C. Use the BULK INSERT command in multiple staging tables. Switch partitions into the Orders table.

D. Disable the Distributor agent. Use the BULK INSERT command in the Orders table. Reinitialize the publication.

Answer: C

Explanation:

According to the exam resources, the answer is A). But according to MSDN, partition switching is not supported when the source or destination table is involved in replication. The vagueness of the question doesn't help. If we assume latency refers to the time between data replication, then B) may be the optimal solution. Using the -continuous option on the distributor will ensure that transactions are replicated as quickly as possible. The commit batch threshold refers to replication commands and isn't relevant here.

Note: Answer confirmed on exam as A), regardless of whether this is actually correct or not.

Question No : 20

You administer a SQL Server 2008 infrastructure.

You implement log shipping for several databases on three SQL Server instances.

The logs are shipped to a fourth SQL Server instance. You plan a manual failover.

You need to ensure that the database applications use the secondary server after failover.

You also need to ensure that users can access the most recent data.

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

- A. Redirect client computers to the secondary instance.
- B. Back up all databases on the secondary instance.
- C. Copy all log shipping network shares to the secondary instance.
- D. Back up the tail of the transaction log of primary databases by using the WITH NORECOVERY option.
- E. Apply any unapplied transaction log backups in sequence to each secondary database by using the WITH RECOVERY option on the last log.

Answer: A,D,E

Explanation:

---Jeff---

Official source: <http://msdn.microsoft.com/en-us/library/ms191233.aspx>

While the maturation of clustering and mirroring in SQL Server 2008, log shipping doesn't see the widespread usage it once did. Jose Baretto's blog has a nice concise explanation of the steps in a log shipping failover. They're out of order in the above question, which can make thinking through the process a bit more difficult.

1. Backup of the tail of the log on the primary using NORECOVERY (assuming this is possible).
2. Restore that log backup as the final link to the secondary, and specify RECOVERY.
3. Point any applications to the now-restored secondary server.

Question No : 21

You administer a SQL Server 2008 instance that hosts a large database.

The database experiences high volumes of queries and transactions that are processor-intensive.

You plan to implement an incremental backup strategy for the database. You also plan to use backup compression.

You need to ensure that the backup jobs have a minimal effect on the performance of the server.

What should you do?

- A. Use database snapshots.
- B. Use the Resource Governor.
- C. Reconfigure the Affinity I/O Mask configuration option.
- D. Spread the database backup across multiple backup devices.

Answer: B

Explanation:

Resource governor can be used to limit the resource usage of a given session.

<http://www.sql-server-performance.com/2008/resource-governor-in-sql-server-2008/>

Database snapshots are not a backup methodology, as they rely on the existing database.

The affinity mask I/O option is for binding disk operations to certain CPUs and was covered earlier.

Spreading backups across multiple devices is supported in SQL Server 2008. If the resource in demand was disk I/O, this would be a viable solution.

Question No : 22

You administer a SQL Server 2008 instance.

The instance hosts a database that contains sensitive data.

You plan to implement a database backup strategy for the database.

You need to ensure that all backups of the database are encrypted. You also need to prevent unauthorized access of the backups.

What should you do?

- A. Use Windows BitLocker Drive Encryption technology
- B. Use Transparent database encryption
- C. Use the BACKUP statement along with the PASSWORD option
- D. Use the BACKUP statement along with the MEDIAPASSWORD option

Answer: B

Explanation:

Official source: <http://msdn.microsoft.com/en-us/library/bb934049.aspx>

Question No : 23

You administer a SQL Server 2008 instance that hosts a large financial database.

The database has the following backup strategies:

You plan to execute an end-of-year batch process that takes two hours to run.

The batch process will modify about five percent of data within the database.

You need to ensure that if the Finance department does not approve the batch process, the batch operation can be rolled back in minimum possible time.

What should you do prior to starting the batch process?

- A. Perform a differential backup.
- B. Create a database snapshot.
- C. Create a marked transaction. Perform a transaction log backup.
- D. Record the time before the batch operation. Perform a transaction log backup.

Answer: B

Explanation:

Reasons to take database snapshots include:

- Maintaining historical data for report generation.

Because a database snapshot provides a static view of a database, a snapshot can extend access to data from a particular point in time. For example, you can create a database snapshot at the end of a given time period (such as a financial quarter) for later reporting.

You can then run end-of-period reports on the snapshot. If disk space permits, you can also maintain end-of-period snapshots indefinitely, allowing queries against the results from these periods; for example, to investigate organizational performance.

- Using a mirror database that you are maintaining for availability purposes to offload reporting.

Using database snapshots with database mirroring permits you to make the data on the mirror server accessible for reporting. Additionally, running queries on the mirror database can free up resources on the principal. For more information, see Database Mirroring and Database

Snapshots,

- Safeguarding data against administrative error.
- Before doing major updates, such as a bulk update or a schema change, create a database snapshot on the database protects data. If you make a mistake, you can use the snapshot to recover by reverting the database to the snapshot. Reverting is potentially much faster for this purpose than restoring from a backup; however, you cannot roll forward afterward.

You administer a SQL Server 2008 database solution.

All data modifications are performed through stored procedures that use only the INSERT, UPDATE, or DELETE statements.

You are designing a backup strategy.

You need to ensure that the following business requirements are met:

What should you do?

- A. Use hourly database snapshots.
- B. Use the full-recovery model along with differential backups.
- C. Use the simple-recovery model along with differential backups.
- D. Use the full-recovery model along with transaction log backups.

Answer: D

Explanation:

Full-recovery Model

Description

- Requires log backups.
- No work is lost due to a lost or damaged data file.
- Can recover to an arbitrary point in time (for example, prior to application or user error).

Work loss exposure

- Normally none.
- If the tail of the log is damaged, changes since the most recent log backup must be redone.

Recover to point in time? S3.

- Can recover to a specific point in time, assuming that your backups are complete up to that point in time.

Question No : 25

You administer a SQL Server 2008 instance that hosts a database.

The backup strategy used by the database is as shown in the following table.

Backup Type	Frequency	Time of the Backup	Time Taken for the Backup
Full Database	Tuesday, Thursday, Saturday	21:00 hr	60-90 minutes
Differential	Monday, Wednesday, Friday	21:00 hr	30-45 minutes
Transaction Log	Hourly	Hourly	< 5 minutes

You run a batch process on Saturday that lasts from 21:00 hr to 23:00 hr.

You discover that the batch process is invalidated because a user has modified some data at 21:05 hr on the same day.

You need to restore the database to its state at the start of the batch process in the least possible time.

What should you do?

A. Restore the full database backup that was performed on Saturday.

B. Restore the full database backup that was performed on Thursday.

Restore all transaction logs from the time of the full backup on Thursday and stop at 21:00 hr on Saturday.

C. Restore the full database backup that was performed on Saturday.

Restore all transaction logs from the time of the full backup on Saturday and stop at 21:00 hr on Saturday.

D. Restore the full database backup that was performed on Thursday.

Restore the differential backup that was performed on Friday.

Restore all transaction logs from the time of the differential backup on Friday and stop at 21:00 hr on Saturday.

Answer: D

Explanation:

---Jeff---

The backup just before the batch takes 90 minutes to complete.

The full backup will grab any recently completed transactions as it nears completion, so most - if not all - of the batch modifications will be included. It's necessary to go back further than that. Lesson: don't run major changes during your backups.

Question No : 26

You administer a SQL Server 2008 instance that hosts a large database.

The following backup strategy is used for the database:

A power failure on Thursday causes the SQL Server 2008 server to restart at 09:15 hr. Fifteen minutes after the server restarts, the users report that they are unable to execute certain queries that access customer data. You discover that the customer data is unmodified after the power failure.

When you execute the DBCC CHECKDB command on the database, you receive the following error message:

```
Object ID 2814307832, index ID 2, partition ID 83127819437122157, alloc unit ID 82134587923221126 (type In-row data): Page (3421:169) could not be processed. See other errors for details.  
Table error: Object ID 2814307832, index ID 2, partition ID 83127819437122157, alloc unit ID 82134587923221126 (type In-row data), page (3421:169). Test (IS_OFF (BUF_IOERR, pBUF->bstat)) failed. Values are 16928422 and -8.  
CHECKDB found 0 allocation errors and 2 consistency errors in table 'tbl_Customer' (object ID 2814307832).
```

When you execute the `sp_help 'tbl_customer'` stored procedure you receive the following result set:

```
index_name index_description  
-----  
PK          clustered located on PRIMARY  
NCIX        nonclustered located on PRIMARY
```

You need to ensure that the data is available as quickly as possible with minimal effect on users.

What should you do?

- A. Drop and recreate the PK index.
- B. Drop and recreate the NCIX index.
- C. Restore the latest full database backup. Restore all transaction log backups from the latest full database backup.
- D. Restore the latest full database backup. Restore the latest differential backup. Restore all transaction log backups from the latest differential backup.

Answer: B

Explanation:

---Jeff---

The error refers to index 2. Since the clustered index on a table is always index 1, this means there's a problem with the nonclustered index. Dropping and recreating this may fix the seek issue quickly as indicated in the requirements.

In situations such as this, it's always a better idea to restore from backups. There's no way of telling how bad the damage is.

However, to get the database running as quickly as possible, D) is the immediate solution. Any steps taken after that aren't relevant.

Question No : 27

You administer a SQL Server 2008 database solution that is log-shipped for high-availability purposes.

The data files of the database are located on drive D. The transaction log files of the database are located on drive E.

You are designing a recovery test plan to meet the following requirements for the log-shipping solution:

The secondary database is brought online as quickly as possible.

The data loss is minimal.

The data is in a consistent state.

You need to identify the first step for the recovery test plan when drive D fails.

Which step should you choose?

- A. Bring the secondary database online.
- B. Perform the tail-log backup of the primary database.
- C. Execute the DBCC CHECKDB command along with the REPAIR_REBUILD option against the primary database.
- D. Execute the DBCC CHECKDB command along with the REPAIR_ALLOW_DATA_LOSS option against the primary database.

Answer: B

Explanation:

Question No : 28

You administer a SQL Server 2008 instance that hosts a large database.

The database has the following backup strategy:

Full database backups are performed weekly.

Differential backups are performed daily.

Transaction log backups are performed hourly.

The recovery plan requires you to perform an unscheduled full backup.

You need to perform a full database backup without interrupting the scheduled backup strategy.

Which backup option should you use?

- A. SKIP
- B. NOINIT
- C. MIRROR TO
- D. COPY_ONLY

Answer: D

Explanation:

Question No : 29

You administer a SQL Server 2008 infrastructure.

You design a corporate backup and recovery strategy that has to be validated.

You need to ensure the successful recovery of any single database from a catastrophic failure without requiring a backup data center in a different location.

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

- A. Store all backup media offsite
- B. Script SQLlogin accounts and credentials
- C. Install all SQL Server instances on a failover cluster
- D. Maintain one list of all Windows logins and passwords
- E. Document the administrative processes and application access requirements

Answer: A,B,E

Explanation:

---Jeff---

There isn't enough information here to make any specific recommendations. It's just a best practices question.

Some dumps suggests other answers, but I don't think a cluster would help in this scenario. It's hard to tell because the type of failure isn't explained. Since a restore operation is needed, it seems like a disk failure.

Question No : 30

You administer a SQL Server 2008 infrastructure.

Your company requires capacity planning information.

You need to design a long-term infrastructure monitoring strategy.

Which two tasks should you include in your design? (Each correct answer presents part of the solution. Choose two.)

- A. Backup all databases every day.
- B. Clear the system log and the application log every hour.
- C. Review system monitor counters on a regular basis.
- D. Baseline the system before you deploy production databases.
- E. Create a maintenance plan that rebuilds indexes every week.

Answer: C,D

Explanation:

---Jeff---

Establishing a performance baseline and keeping an eye on fluctuations is an ideal way to determine current and future resource requirements. Backing up databases and rebuilding indexes are good ideas, but not related to this task.

Question No : 31

You administer a SQL Server 2008 instance.

Customers report server performance degradation because of a newly implemented process.

You use Dynamic Management Views to verify that there are no long running queries.

You need to correlate the operating system performance data with the actual query execution trace by using minimum administrative effort.

What should you do?

- A. Use Data Collector.
- B. Use the SQLdiag.exe utility.
- C. Use SQL Server Profiler and System Monitor.
- D. Use SQL Server Profiler and the tracerpt.exe utility.

Answer: C

Explanation:

<http://support.microsoft.com/kb/298475>

<http://msdn.microsoft.com/en-us/library/ms181091.aspx>

Question No : 32

You administer a SQL Server 2008 instance. The instance is using a 32-bit version on a Windows Server 2008 64-bit server. The awe enabled option is enabled.

The instance will experience a predictable increase in query activity.

You plan to ascertain the appropriate time when the migration of the databases to a 64-bit SQL Server 2008 server on the same hardware is beneficial.

You need to identify a data collector type that provides the appropriate information.

Which collector should you use?

-
- A. SQL Trace collector
 - B. T-SQL Query collector
 - C. Query Activity collector
 - D. Performance Counters collector

Answer: D

Explanation:

This question is referencing SQL 2008's management data warehouse feature, which allows the collection of performance-related SQL statistics.

The T-SQL Query collector allows you to write a custom query. The query is executed at intervals and the results are logged.

The SQL Trace collector leverages SQL Profiler to capture trace data.

The performance counters collector tracks OS and SQL-related counters in the same way PERFMON does.

The query activity collector tracks DMV results. It is generally not used.

As far as the question goes, it seems logical to pick the time when the server is experiencing the least activity.

The performance counters collection makes the most sense, as it's possible to determine activity based on disk I/O, processor, memory, etc.

<http://www.sql-server-performance.com/2008/Management-Data-Warehouse/>

http://www.sql-server-performance.com/articles/per/System_Data_Collection_Reports_Install_p1.aspx

Question No : 33

You administer a SQL Server 2008 infrastructure.

You discover that an instance experiences performance degradation for the following reasons:

You need to design a monitoring solution that can provide data, including detailed deadlock information, to monitor and troubleshoot performance issues.

You want to achieve this goal by using the minimum amount of administrative effort.

What tool should you use?

-
- A. Extended Events
 - B. Resource Governor
 - C. Database Engine Tuning Advisor
 - D. Performance Monitor (SYSMON)

Answer: A

Explanation:

---Jeff---

Is a broad error-handling system that can correlate data between SQL server, the OS, and other applications.

Events from SQL Server can be collected, then passed elsewhere (a "target") for reporting or event response.

Resource governor is for limiting resource consumption.

The DTA (Database Tuning Advisor) makes index recommendations.

PERFMON monitors operating system and sql server counters.

Official source: <http://msdn.microsoft.com/en-us/library/bb630354.aspx>

Extra reading: <http://www.sqlteam.com/article/introduction-to-sql-server-2008-extended-events>

Question No : 34

You administer a SQL Server 2008 instance.

You plan to design a monitoring solution for the instance to monitor object usage statistics.

The solution must identify a list of first 10 objects for each of the following components:

You need to implement the monitoring solution to minimize performance effect by using the least amount of administrative effort.

What should you do?

- A. Use dynamic management views.
- B. Use a System Monitor counter log.
- C. Use a client-side SQL Server Profiler trace.
- D. Use a server-side SQL Server Profiler trace.

Answer: A

Explanation:

Background info:

<http://www.simple-talk.com/sql/performance/which-of-your-stored-procedures-are-using-the-most-resources/>

<http://www.sql-server-performance.com/2008/monitor-stored-procedure-performance/>

Question No : 35

You administer a SQL Server 2008 infrastructure.

The instance contains a database required for the day-to-day business of your company. Users experience slow response time when they execute reports.

You plan to design a performance monitoring strategy that captures and stores the following data:

You need to implement the monitoring process by using the minimum amount of administrative effort.

What should you do?

- A. Use the data collector.
- B. Use the client-side profiler trace.
- C. Use the dynamic management views.
- D. Use the System Monitor counter log trace.

Answer: A

Explanation:

The data collector is a component installed on a SQL Server server, running all the time or on a user-defined schedule, and collecting different sets of data. The data collector then stores the collected data in a relational database known as the management data warehouse.

The data collector is a core component of the data collection platform for SQL Server 2008 and the tools that are provided by SQL Server. The data collector provides one central point for data collection across your database servers and applications.

This collection point can obtain data from a variety of sources and is not limited solely to performance data, unlike SQL Trace. w

The data collector enables you to adjust the scope of data collection to suit your test and production environments. The data collector also uses a data warehouse, a relational database that enables you to manage the data you collect by setting different retention periods for your data. The data collector supports dynamic tuning for data collection and is extensible through its API.

Question No : 36

You administer a SQL Server 2008 infrastructure.

You plan to design a maintenance strategy for a mission-critical database that includes a large table named Orders. The design plan includes index maintenance operations.

You must design the strategy after considering the following facts:

You need to implement the strategy so that the performance of the queries on the table is optimized.

What should you do?

- A. Drop the clustered index of the Orders table.
- B. Rebuild the clustered index of the Orders table offline once a month.
- C. Reorganize the clustered index of the Orders table by decreasing the fill factor.
- D. Exclude the clustered index of the Orders table from scheduled reorganizing or rebuilding operations.

Answer: D

Explanation:

Since the clustered index never has any significant fragmentation, there's no reason to rebuild or reorganize it.

Question No : 37

You administer a SQL Server 2008 infrastructure.

An instance contains a database that includes a large table named OrderDetails. The application queries only execute DML statements on the last three months data. Administrative audits are conducted monthly on data older than three months.

You discover the following performance problems in the database.

The performance of the application queries against the OrderDetail table is poor.

The maintenance tasks against the database, including index defragmentation, take a long time.

You need to resolve the performance problems without affecting the server performance.

What should you do?

A. Create a database snapshot for the OrderDetails table every three months.

Modify the queries to use the current snapshot.

B. Create an additional table named OrderDetailsHistory for data older than three months.

Partition the OrderDetails and OrderDetailsHistory tables in two parts by using the OrderDate column.

Create a SQL Server Agent job that runs every month and uses the ALTER

TABLE...SWITCH Transact-SQL statement to move data that is older than three months to the OrderDetailsHistory table.

C. Create an additional table named OrderDetailsHistory for data older than three months.

Create a SQL Server Agent job that runs the following Transact-SQL statement every

month. INSERT INTO OrderDetailsHistory

SELECT * FROM OrderDetails

WHERE DATEDIFF(m,OrderDate,GETDATE())>3

D. Create an additional table named OrderDetailsHistory for data older than three months.

use the following Transact-SQL statement.

CREATE TRIGGER trgMoveData

ON OrderDetails

AFTER INSERT

AS

INSERT INTO OrderDetailsHistory

SELECT * FROM OrderDetails

WHERE DATEDIFF(m,OrderDate,GETDATE())>3

Answer: B

Explanation:

Question No : 38

You administer a SQL Server 2008 infrastructure.

Humongous Insurance has 20 branch offices that store customer data in SQL Server 2008 databases.

Customer data that is stored across multiple database instances has to be security compliant.

You plan to design a strategy for custom policies by using the Policy-Based Management feature. Custom policies are in XML format.

The strategy must meet the following requirements:

You need to implement the strategy by using the least amount of administrative effort.

What should you do?

- A. Use a configuration server.
- B. Use the Distributed File System Replication service.
- C. Distribute the policies by using Group Policy Objects.
- D. Distribute the policies by using the Active Directory directory service.

Answer: A

Explanation:

---Jeff---

Configuration servers are the original name for central management servers, which allow the administration and enforcement of SQL Server 2008 policies for multiple servers to be centralized.

Question No : 39

You administer a SQL Server 2008 infrastructure.

An instance contains a database that includes a table named EmployeeData.

The EmployeeData table has a column named Profile. The column is used to store Microsoft Office Word 2003 documents.

You need to minimize the space occupied by the Profile column. You also need to maintain a reasonable performance to retrieve the content of the column.

What should you do?

- A. Use ROW compression for the EmployeeData table.
- B. Use PAGE LEVEL compression for the EmployeeData table.
- C. Implement the Profile column as a CLR user-defined type.
- D. Store the Profile column by using filestream storage, and use the NTFS file system compression.

Answer: D

Explanation:

Question No : 40

You administer a SQL Server 2008 infrastructure.

You partition the largest table in a database into four equal-sized parts. The first two partitions are rarely queried.

The third partition is frequently queried by several reports. The last partition is heavily queried and frequently modified.

You need to minimize the disk space of the table and ensure that this has minimum effect on the query performance.

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

- A. Enable ROW compression for the last partition
- B. Enable ROW compression for the third partition
- C. Enable ROW compression for the first two partitions
- D. Enable PAGE compression for the last partition

-
- E. Enable PAGE compression for the third partition
 - F. Enable PAGE compression for the first two partitions

Answer: B,F

Explanation:

Question No : 41

You administer a SQL Server 2008 infrastructure.

You plan to design a solution to obtain hardware configurations, such as the number of processors on a computer and the processor type of all SQL Server 2008 computers.

The solution must meet the following requirements:

You need to select a technology that meets the requirements by using the minimum amount of development effort.

What should you do?

- A. Use the Invoke-Sqlcmd cmdlet in SQL Server PowerShell cmdlet.
- B. Define policies based on conditions by using the ExecuteSql function.
- C. Define policies based on conditions by using the ExecuteWQL function.
- D. Use the Windows Management Instrumentation (WMI) provider for the server events.

Answer: C

Explanation:

---Jeff---

ExecuteWQL is a relatively straightforward way to query operating system data from SQL server.

It can then be stored in a database for analysis.

The Invoke-Sqlcmd cmdlet is a powershell cmdlet for executing sql commands. It doesn't apply well to the question.

WMI is a driver extension with scripting language and could theoretically be used to accomplish the goal, but with a much more complex development process.

EXECUTESQL is a SQL command for running a pre-built SQL statement

Question No : 42

You administer a SQL Server 2008 infrastructure.

Your company wants to enforce naming standards for the database objects.

You need to design a solution to enforce naming standards by using the minimum amount of administrative and programming efforts.

What should you do?

- A.** Use event notifications to handle the DDL events.
- B.** Create Policy-Based Management policies, and then distribute the policies by using a configuration server.
- C.** Create DDL triggers for one of the instances, and then use SQL Server Management Objects (SMO) to script the definition of the DDL triggers. Run the script by using Microsoft Windows PowerShell on all instances.
- D.** Create DDL triggers for one of the instances, and then use Microsoft SQL Server Management Studio (SSMS) to script the definition of the DDL triggers. Run the script on all instances by using the registered servers node from SSMS.

Answer: B

Explanation:

Question No : 43

You administer a SQL Server 2008 instance.

You plan to deploy new SQL Server 2008 Integration Services (SSIS) packages that are created by the SQL Server Import and Export Wizard. The packages will connect to external data sources by using stored user names and passwords.

You need to ensure that the user names and passwords are protected from access by unauthorized users.

What should you do?

- A.** Use the SQL Server 2008 server to store the packages, and manage the security in the SQL Server 2008 server.
- B.** Ensure that no data is saved along with the packages, and save the files to the most

secure directory.

C. Ensure that no sensitive data is saved along with the packages, and save the packages in the SQL Server 2008 server.

D. Set a common password to make all information in the package as secure as possible when each package is saved.

Answer: A

Explanation:

Question No : 44

You administer a SQL Server 2008 instance for your company. Your company has a team of database administrators.

A team of application developers create SQL Server 2008 Integration Services (SSIS) packages on the test server in a shared project. One of the packages requires a fixed cache file. On completion of development, the packages will be deployed to the production server.

Only the database administrators can access the production server.

You need to ensure that the application developers can deploy the project successfully to the production server.

What should you do?

A. Use the Import and Export Wizard to save packages.

B. Create a deployment utility for the SSIS project.

C. Create a direct package configuration for each package.

D. Create an indirect package configuration for all packages.

Answer: B

Explanation:

---Jeff---

This is a strange question. The underlying lesson is that deployment utilities make SSIS package deployment easier, especially in situations where limited access may be available.

Direct and indirect package configurations are explained on MSDN.

Official source: <http://msdn.microsoft.com/en-us/library/ms141682.aspx>

Question No : 45

You administer a SQL Server 2008 infrastructure.

You administer two SQL Server 2008 instances named Instance1 and Instance2. Instance1 contains the Sales database, and Instance2 contains the Accounts database.

A procedure in the Sales database starts a transaction. The procedure then updates the Sales.dbo.Order table and the Accounts.dbo.OrderHistory table through a linked server.

You need to ensure that the transaction uses a two-phase commit.

What should you do?

- A. Configure the linked server to use distributed transactions.
- B. Configure a Service Broker to enable the appropriate transaction control.
- C. Ensure that the linked server is appropriately configured for delegation.
- D. Ensure that the linked server is appropriately configured for impersonation.

Answer: A

Explanation:

A distributed transaction can be executed via a linked server using BEGIN DISTRIBUTED TRANSACTION.

This will guarantee that the entire batch will complete or fail. Two-phase commit is a transaction protocol employed by the MSDTC to ensure this is possible.

The service broker is for distributing heavy workloads across multiple database instances.

Delegation and impersonation are for passing credentials from one server/instance to another. Delegation passes windows credentials;

Impersonation passes a SQL Server login.

Question No : 46

You administer a SQL Server 2008 instance named Instance1 at the New York central site.

Your company has a sales team to fulfill purchase orders for customer requests.

The sales team uses portable computers to update data frequently in a local database. When the portable computers connect to the central site, the local database must be synchronized with a database named Sales.

You plan to create a replication model to replicate the local database to the Sales database.

The replication model must meet the following requirements:

You need to identify the best model to replicate data by using minimum development efforts.

What should you do?

- A.** Use merge replication along with each portable computer that is set up as a subscriber.
- B.** Use snapshot replication along with each portable computer that is set up as a subscriber.
- C.** Use transactional replication along with each portable computer that is set up as a publisher.
- D.** Use SQL Server Integration Services (SSIS) to push data changes and pull updates to the Sales database along with the SSIS packages, on demand.

Answer: A

Explanation:

<http://www.codeproject.com/KB/database/sql2005-replication.aspx>

Question No : 47

You are a professional level SQL Sever 2008 Database Administrator.

It is reported by the customers that the server performance degraded due to a newly implemented process. Dynamic Management Views is utilized to confirm that no long running queries exist.

The operating system performance data should be correlated with the actual query execution trace, and the least administrative effort should be utilized.

Which action will you perform to finish the task?

-
- A. To finish the task, Data Collector should be utilized.
 - B. To finish the task, the SQLdiag.exe utility should be utilized.
 - C. To finish the task, SQL Server Profiler and the tracerpt.exe utility should be utilized.
 - D. To finish the task, SQL Server Profiler and System Monitor should be utilized.

Answer: D

Explanation:

SQL Server profiler displays data about a large number of SQL Server events. Whereas Windows System Monitor graphically displays information about the server internals. You can merge the two sets of information and walk through a scenario viewing both perspectives using SQL Server Profiler.

To set up the dual-perspective experience, you need to simultaneously capture server performance using both Performance Monitor's Counter Logs and SQL Server Profiler. The steps to do this are listed below:

1. Configure System Monitor with the exact counters you want to view later. Be sure to get the scale and everything just right. Set up the Counter Log to the exact same configuration.
2. Configure Profiler with the right set of trace events. They must include the start and end time data columns so that Profiler can integrate the two logs later. Save the profiler and close profiler.
3. Manually start the Counter Log. Open SQL Profiler trace code to start the server—side trace.
4. When the test is complete, stop both the counter Log and the server side trace. You need to stop the SQL Profiler which is a negative point in this.
5. Open profiler and open the saved trace file.
6. Use the Files> Import Performance Data menu command to import the Counter Log. You have the option of selecting only the important counters from the performance monitor. There will be performance issues if you select too many counters.

Source:

http://www.sql-server-performance.com/faq/How_to_Integrate_Performance_Monitor_and_SQL_Profiler_p1.aspx

Question No : 48

You are a professional level SQL Sever 2008 Database Administrator.

There is a database in the instance, and the day-to-day business of your company requires

the database. When reports are executed, slow response time is experienced by Users.

A performance monitoring strategy will be implemented by you so as to have three aspects of data captured and stored:

A performance monitoring strategy will be implemented by you so as to have three aspects of data captured and stored:

You are required to utilize the least amount of administrative effort to implement the monitoring process.

Which action will you perform to finish the task?

- A. To finish the task, the client-side profiler trace should be utilized.
- B. To finish the task, the dynamic management views should be utilized.
- C. To finish the task, the data collector should be utilized.
- D. To finish the task, the System Monitor counter log trace should be utilized.

Answer: C

Explanation:

SQL Server 2008 provides a data collector that you can use to obtain and save data that is gathered from several sources. The data collector enables you to use data collection containers, which enable you to determine the scope and frequency of data collection on a SQL Server server system.

The data collector provides predefined collector types that you can use for data collection. The collector types provide the actual mechanism for collecting data and uploading it to the management data warehouse. For this release of the data collector, the following collector type is provided.

In This Section

Topic	Description
T-SQL Query Collector Type	Describes the structure and behavior of the T-SQL Query collector type.
SQL Trace Collector Type	Describes the structure and behavior of the SQL Trace collector type.
Performance Counters Collector Type	Describes the structure and behavior of the Performance Counters collector type.
Query Activity Collector Type	Describes the structure and behavior of the Query Activity collector type.

Question No : 49

You are a professional level SQL Sever 2008 Database Administrator.

The computer on which the instance run has the following three features:

64-GB RAM, four quad-core processors, and several independent physical raid volumes A transactional database will be implemented on the instance.

In addition, the database should have a high volume of INSERT, UPDATE, and DELETE activities, creation of new tables is contained by the activities. You need to maximize disk bandwidth and decrease the contention in the storage allocation structures so as to have the performance of the database optimized.

Which action will you perform to finish the task?

- A. To finish the task, database and log files should be placed on the same volume.
- B. To finish the task, the affinity mask option should be configured properly.
- C. To finish the task, multiple data files should be create for the database.
- D. To finish the task, the affinity I/O mask option should be configured properly.

Answer: C

Explanation:

If your database is very large and very busy, multiple tiles can be used to increase performance. Here is one example of how you might use multiple files. Let's say you have a single table with 10 million rows that is heavily queried. It the table is in a single file, such as a single database file, then SQL Server would only use one thread to perform a read of the rows in the table. But if the table were divided into three physical files, then SOL Server would use three threads (one per physical file) to read the table, which potentially could be faster. In addition, if each tile were on its own separate physical disk or disk array, the performance gain would even be greater.

Essentially, the more files that a large table is divided into, the greater the potential performance.

Of course there is a point where the additional threads aren't of much use when you max out the server's UO. But up until you do max out the IXO, additional threads (and files) should increase performance.

SOURCE: http://WWW.sql-server-perlormance.comr'tipsffilegroups_p1.aspx

Question No : 50

You are a professional level SQL Sever 2008 Database Administrator.

After a regular test, you find that performance degradation is experienced by an instance for the three reasons:

Excessive CPU usage, Server processes paging and Deadlocks

A monitoring solution should be implemented to provide data, monitor and troubleshoot performance issues and detailed deadlock information should be contained in the provided data.

You should utilize the least amount of administrative effort to finish the task.

Which tool will you utilize to finish the task?

- A. To finish the task, you should utilize Resource Governor.
- B. To finish the task, you should utilize Database Engine Tuning Advisor.
- C. To finish the task, you should utilize Extended Events.
- D. To finish the task, you should utilize Performance Monitor (SYSMON).

Answer: C

Explanation:

Introducing SQL Server Extended Events

SOL Server Extended Events (Extended Events} is a general event-handling system for server systems. The Extended Events infrastructure supports the correlation of data from SOL Server, and under certain conditions, the correlation of data from the operating system and database applications. In the latter case, Extended Events output must be directed to Event Tracing for Windows (ETW) in order to correlate the event data with operating system or application event data. All applications have execution points that are useful both inside and outside an application.

Inside the application, asynchronous processing may be enqueued using information that is gathered during the initial execution of a task. Outside the application, execution points provide monitoring utilities with information about the behavioral and performance characteristics of the monitored application.

Extended Events supports using event data outside a process, This Etta is typically used by:

- Tracing tools, such as SOL Trace and System Monitor.
- Logging tools, such as the Windows event log or the SOL Server error log.
- Users administering a product or developing applications on a product.

Question No : 51

You are a professional level SQL Sever 2008 Database Administrator.

There are 30 branch offices in DoubleSecurity Insurance, and in the branch offices, customer data are stored in SQL Server 2008 databases. Customer data should be security compliant if it is stored through multiple Database instances.

You intend to utilize the Policy-Based Management feature to design a strategy for custom policies.

And the format of custom policies is XML format. The requirements listed below should be satisfied.

The company distributes custom policies to all instances. In addition, the company enforces the policies on all instances. A strategy should be thought out and the minimum amount of administrative effort should be utilized.

Which action should you perform to finish the task?

- A. To finish the task, the Distributed File System Replication service should be utilized.
- B. To finish the task, a configuration server should be utilized.
- C. To finish the task, the policies should be distributed by utilizing Group Policy Objects.
- D. To finish the task, the policies should be distributed by utilizing the Active Directory directory service.

Answer: B

Explanation:

Configuration Server or Central Management Server

In SOL Server 2008, you can designate an instance of SQL Server as a Central Management Server. Central Management Servers store a list of instances of SOL Server that is organized into one or more Central Management Server groups. Actions that are taken by using a Central Management Server group will act on all servers in the server group. This includes connecting to servers by using Object Explorer and executing Transact-SOL statements and Policy-Based Management policies on multiple servers at the same time. All Central Management Servers and subordinate servers must be registered by using Windows Authentication. Versions of SOL Server that are earlier than SOL Server 2008 cannot be designated as a Central Management Server.

Question No : 52

You are a professional level SQL Sever 2008 Database Administrator.

All data changes are implemented through stored procedures, and only the INSERT, UPDATE, or DELETE statements are utilized by the procedures. A backup strategy should be implemented.

The business requirements listed below should be satisfied:

Which action should you perform to finish the task?

- A.** To finish the task, hourly database snapshots should be utilized.
- B.** To finish the task, the full-recovery model along with transaction log backups should be utilized.
- C.** To finish the task, the full-recovery model along with differential backups should be utilized.
- D.** To finish the task, the simple-recovery model along with differential backups should be utilized.

Answer: B

Explanation:

Description

- Requires log backups.
- No work is lost due to a lost or damaged data file.
- Can recover to an arbitrary point in time (tor example, prior to application or user error).

Work loss exposure

- Normally none.
- It the tail of the log is damaged, changes since the most recent log backup must be redone.

Recover to point in time?

- Can recover to a specific point in time, assuming that your backups are complete up to that point in time.

Question No : 53

You are a professional level SQL Sever 2008 Database Administrator.

You are experienced in managing databases in an enterprise-level organization, optimizing and sustaining the database life cycle. In the company, your job is to implement solutions on security, troubleshooting, deployment and optimization. A SQL Server 2008 infrastructure is managed by you.

A database is utilized by the instance, and the database is utilized by a Web-based application. 15,000 transactions are processed by the application every minute. A column is contained by a table in the database, and the column is utilized only by the application. Sensitive data is stored in this column. The sensitive data should be stored with the highest security level. In addition, the least amount of memory space and processor time should be utilized.

From the following four encryption types, which one should you utilize?

- A. Asymmetric key encryption should be utilized.
- B. Certificate-based encryption should be utilized.
- C. Symmetric key encryption should be utilized.
- D. Transparent data encryption should be utilized.

Answer: C

Explanation:

At the root of encryption tree is the Windows Data Protection API (DPAPI), which secures the key hierarchy at the machine level and is used to protect the service master key (SMK) for the database server instance. The DIVIK protects the database master key (DMK), which is stored at the user database level and which in turn protects certificates and asymmetric keys. These in turn protect symmetric keys, which protect the data. TDE uses a similar hierarchy down to the certificate. The primary difference is that when you use TDE, the DIVIK and certificate must be stored in the master database rather than in the user database. A new key, used only for TDE and referred to as the database encryption key (DEK), is created and stored in the user database.

This hierarchy enables the server to automatically open keys and decrypt data in both cell-level and database-level encryption. The important distinction is that when cell-level encryption is used, all keys from the DIVIK down can be protected by a password instead of by another key. This breaks the decryption chain and forces the user to input a password to access data. In TDE, the entire chain from DPAPI down to the DEK must be maintained so that the server can automatically provide access to files protected by TDE. In both cell-level encryption and TDE, encryption and decryption through these keys is provided by the Windows Cryptographic API (CAPI). Symmetric keys use the same password to encrypt

and decrypt data, so it is the less space consuming, because one asymmetric key will use one private and one public key.

Question No : 54

You are a professional level SQL Sever 2008 Database.

A new database application is hosted by the instance. The security requirements should be designed for the application.

A unique login to the SQL Server 2008 server is assigned to each application user. Stored procedures are included by the application database to execute stored procedures in the MSDB database. SQLAgent jobs are scheduled by the stored procedures in the MSDB database.

Since you are the technical support, you are required to confirm that the stored procedures in the MSDB database should be executed by utilizing the security context of the application user.

Which action should you perform?

- A. Each user should be added to the public role in the MSDB database.
- B. Each user should be added to the db_dtsltduser database role in the MSDB database.
- C. The MSDB database should be set to utilize the TRUSTWORTHY option, and then each user should be added to the MSDB database.
- D. The new database should be set to utilize the TRUSTWORTHY option, and then each user should be added o the MSDB database.

Answer: D

Explanation:

The TRUSTWORTHY database property is used to indicate whether the instance of SOL Server trusts the database and the contents within it. By default, this setting is OFF, but can be set to ON by using the ALTER DATABASE statement. For example, ALTER DATABASE

AdventureWorks2008R2 SET TRUSTWORTHY ON;

By default msdb has the option TRUSTWORTHY set to True.

Question No : 55

You are a professional level SQL Sever 2008 Database Administrator.

The security requirements should be designed for a new database application. A code segment is utilized by the application, and the code segment includes the following components:

Since you are the technical support, you are required to make sure that the application should be successfully deployed to the instance.

Which action will you perform?

- A. The SAFE code access security should be utilized for the assembly.
- B. All public static fields should be replaced with public fields.
- C. All public static fields should be utilized with public static read-only fields. And then the assembly should be registered by utilizing the regasm.exe utility before deployment.
- D. All public static fields should be replaced with public static read-only fields. And then the EXTERNAL_ACCESS code access security should be utilized for the assembly.

Answer: D

Explanation:

Creates a managed application module that contains class metadata and managed code as an object in an instance of SOL Server. By referencing this module, common language runtime (CLR) functions, stored procedures, triggers, user-defined aggregates, and user-defined types can be created in the database.

`PERMISSION_SET { SAFE | EXTERNAL_ACCESS | UNSAFE }`

Specifies a set of code access permissions that are granted to the assembly when it is accessed by SOL Server. If not specified, SAFE is applied as the default.

We recommend using SAFE. SAFE is the most restrictive permission set. Code executed by an assembly with SAFE permissions cannot access external system resources such as files, the network, environment variables, or the registry.

EXTERNAL_ACCESS enables assemblies to access certain external system resources such as files, networks, environmental variables, and the registry.

Question No : 56

To Read the [Whole Q&As](#), please purchase the [Complete Version](#) from [Our website](#).

Trying our product !

- ★ **100%** Guaranteed Success
- ★ **100%** Money Back Guarantee
- ★ **365 Days** Free Update
- ★ **Instant Download** After Purchase
- ★ **24x7** Customer Support
- ★ Average **99.9%** Success Rate
- ★ More than **69,000** Satisfied Customers Worldwide
- ★ Multi-Platform capabilities - **Windows, Mac, Android, iPhone, iPod, iPad, Kindle**

Need Help

Please provide as much detail as possible so we can best assist you.

To update a previously submitted ticket:



 One Year Free Update <p>Free update is available within One Year after your purchase. After One Year, you will get 50% discounts for updating. And we are proud to boast a 24/7 efficient Customer Support system via Email.</p>	 Money Back Guarantee <p>To ensure that you are spending on quality products, we provide 100% money back guarantee for 30 days from the date of purchase.</p>	 Security & Privacy <p>We respect customer privacy. We use McAfee's security service to provide you with utmost security for your personal information & peace of mind.</p>
---	---	--

[Guarantee & Policy](#) | [Privacy & Policy](#) | [Terms & Conditions](#)

Any charges made through this site will appear as Global Simulators Limited.

All trademarks are the property of their respective owners.

Copyright © 2004-2015, All Rights Reserved.