

Microsoft 70-450

**PRO: Designing, Optimizing and Maintaining a
Database Administrative Solution Using Microsoft SQL
Server 2008
Version: 55.0**

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 IRRSTAR.CE. 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:

- A 64-GB RAM
- Four quad-core processors
- Several independent physical RAID volumes

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.