Vendor: Microsoft

Exam Code: 70-450

Exam Name: PRO: Designing, Optimizing and Maintaining a Database Administrative Solution Using Microsoft SQL Server 2008

Version: Demo

Q & A: 132

QUESTION 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

Correct Answer: D

QUESTION 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

Correct Answer: C

QUESTION 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.

Correct 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 lrrstar.穋e. 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 AdvancedlEnterprise 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. AdvancedlEnterprise 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 nods--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 segilce with failover clustering. For more information, see Using SQL Server Tools with Failover stering.

QUESTION 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.

Correct Answer: A

QUESTION 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.

Correct 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 sewer 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 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.

Correct Answer: AC

QUESTION 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.

Correct Answer: D

QUESTION 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.

Correct Answer: D

QUESTION 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

Correct Answer: A
Explanation:
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/IntroductionToSQLServerEncryptionAndSymmetricKeyEncryp tionTutorial.aspx

QUESTION 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.

Correct Answer: C

QUESTION 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.

Correct Answer: C

QUESTION 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.

Correct Answer: C
Explanation:
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 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.

Correct Answer: C

QUESTION 14
You administer SQL Server 2008 servers on a single site. The site contains the following two SQL Server 2008 instances:

- An Enterprise Edition server on a server that has a redundant array of independent disks (RAID) 10 disk system.
- A Standard Edition server on a server that has a RAID 5 disk system Each instance hosts a single application.

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

- The solution can be implemented on the existing systems.
- The database is available with minimal downtime.
- Data loss is minimal.
- There is minimal effect on the existing system.

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

Correct Answer: AB
Explanation:
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 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:

- Access and modify data on all sites with minimal latency.
- Minimize data loss in the event of a server failure.

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

Correct Answer: B
Explanation:
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 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

Correct 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 AdvancedlEnterprise 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.
AdvancedlEnterprise 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 nods--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 segilce with failover clustering. For more information, see Using SQL Server Tools with Failover stering.

QUESTION 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

Correct Answer: C
Explanation:
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%2Fdownl oad.
microsoft.com%2Fdownload%2F6%2F9%2FD%2F69D1FEA7-5B42-437A-B3BA-A4AD13E34EF6%
2FSQLServer2008FailoverCluster.docx&rct=j&q=sql%20server%202008%20cluster%20microsoft%20 white%
20paper&ei=7T2rTOG8EcWBlAfN8fnEDA&usg=AFQjCNGI1hQIgs4JkiNAfD6zVvVl1-8UuA&cad=rja

QUESTION 18
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.

Correct Answer: D

QUESTION 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.

Correct 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 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.

Correct Answer: ADE
Explanation:
Official source: http://msdn.microsoft.com/en-us/library/ms191233.aspx Explanation: 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 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.

Correct 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 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.

Correct Answer: B
Explanation:
Official source: http://msdn.microsoft.com/en-us/library/bb934049.aspx

QUESTION 23
You administer a SQL Server 2008 instance that hosts a large financial database. The database has the following backup strategies:

- A full database backup is performed once a week.
- A differential backup is performed every day.
- A transaction log backup is performed every hour.

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.

Correct 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 sewer 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. Fleverting is potentially much faster for this purpose than restoring from a backup; however, you cannot roll forward afterward.

QUESTION 24
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:

- The backup strategy supports point-in-time recovery for failure at any time of day.
- The transaction log uses the least amount of disk space.

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.

Correct 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.
- Il 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 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.

Correct Answer: D
Explanation:
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 26
You administer a SQL Server 2008 instance that hosts a large database. The following backup strategy is used for the database:

- A full database backup is performed at 02:00 hr every Monday.
- A differential backup is performed at 23:00 hr every day.
- A transaction log backup is performed on an hourly basis.

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.

Correct Answer: B
Explanation:
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 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.

Correct Answer: B

QUESTION 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

Correct Answer: D

QUESTION 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 SQL login 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.

Correct Answer: ABE
Explanation:
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 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.

Correct Answer: CD
Explanation:
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.

EnsurePass.com Members Features:

1. Verified Answers researched by industry experts.
2. Q&As are downloadable in PDF and VCE format.
3. 98% success Guarantee and **Money Back** Guarantee.
4. Free updates for **180** Days.
5. **Instant Access to download the Items**

View list of All Exam provided:
http://www.ensurepass.com/certfications?index=A

To purchase Lifetime Full Access Membership click here:
http://www.ensurepass.com/user/register

Valid Discount Code for 2015: JREH-G1A8-XHC6

To purchase the HOT Microsoft Exams:

Microsoft			
70-243	70-347	70-466	70-515
70-246	70-410	70-467	70-516
70-247	70-411	70-480	70-519
70-321	70-412	70-483	70-583
70-331	70-413	70-484	70-640
70-332	70-414	70-485	70-649
70-336	70-417	70-486	70-668
70-337	70-461	70-487	70-680
70-341	70-462	70-488	70-687
70-342	70-463	70-489	70-688
70-346	70-464	70-513	70-689