

MCSA Administering Microsoft SQL Server 2012/2014 Databases

Microsoft 70-462 Dumps Available Here at:

<https://www.certification-questions.com/microsoft-exam/70-462-dumps.html>

Enrolling now you will get access to 319 questions in a unique set of 70-462 dumps

Question 1

Note: This question is part of a series of questions that use the same set of answers choices. An answer choice may be correct for more than one question in the series.

You administer a Microsoft SQL Server database server that hosts a transactional database and a reporting

database. The transactional database is updated through a web application and is operational throughout the day. The reporting database is only updated from the transactional database.

The recovery model and backup schedule are configured as shown in the following table:

Database	Description
Transactional database	Recovery model: <ul style="list-style-type: none"> • Full Backup schedule: <ul style="list-style-type: none"> • Full database backup: midnight, daily • Differential database backup: on the hour, every two hours starting at 02:00 hours except at 00:00 hours • Log backup: every half hour, except at the times of full and differential backups
Reporting database	Recovery model: <ul style="list-style-type: none"> • Simple Backup schedule: <ul style="list-style-type: none"> • Full database backup: 01:00 hours daily • Differential database backup: 13:00 hours daily Data updates: <ul style="list-style-type: none"> • Changes in data are updated from the transactional database to the reporting database at 00:30 hours and at 12:30 hours • The update takes 15 minutes

At 14:00 hours, you discover that pages 71, 520, and 713 on one of the database files are corrupted on the reporting database.

You also need to ensure that data loss is minimal.

What should you do?

Options:

A. Perform a partial restore.

B. Restore the latest full backup, and restore the latest differential backup. Then, restore each log backup taken before the time of failure from the most recent differential backup.

C. Restore the latest full backup.

D. Restore the latest full backup, and restore the latest differential backup. Then, restore the latest log

backup.

E. Perform a page restore.

F. Restore the latest full backup. Then, restore each differential backup taken before the time of failure from the most recent full backup.

G. Perform a point-in-time restore.

H. Restore the latest full backup. Then, restore the latest differential backup.

Answer: H

Explanation:

Restores a file or filegroup in a multi-filegroup database. Note that under the simple recovery model, the file must belong to a read-only filegroup. After a full file restore, a differential file backup can be restored.

Page restore

Restores individual pages. Page restore is available only under the full and bulk-logged recovery models.

Piecemeal restore

Restores the database in stages, beginning with the primary filegroup and one or more secondary filegroups. A piecemeal restore begins with a RESTORE DATABASE using the PARTIAL option and specifying one or more secondary filegroups to be restored.

Question 2

Note: This question is part of a series of questions that use the same set of answers choices. An answer choice may be correct for more than one question in the series.

You administer a Microsoft SQL Server database server that hosts a transactional database and a reporting

database. The transactional database is updated through a web application and is operational throughout the day. The reporting database is only updated from the transactional database.

The recovery model and backup schedule are configured as shown in the following table:

Database	Description
Transactional database	Recovery model: <ul style="list-style-type: none"> • Full Backup schedule: <ul style="list-style-type: none"> • Full database backup: midnight, daily • Differential database backup: on the hour, every two hours starting at 02:00 hours except at 00:00 hours • Log backup: every half hour, except at the times of full and differential backups
Reporting database	Recovery model: <ul style="list-style-type: none"> • Simple Backup schedule: <ul style="list-style-type: none"> • Full database backup: 01:00 hours daily • Differential database backup: 13:00 hours daily Data updates: <ul style="list-style-type: none"> • Changes in data are updated from the transactional database to the reporting database at 00:30 hours and at 12:30 hours • The update takes 15 minutes

At 16:20 hours, you discover that pages 17, 137, and 205 on one of the database files are corrupted on the transactional database.

You need to ensure that the transactional database is restored.

You also need to ensure that data loss is minimal.

What should you do?

Options:

A. Perform a partial restore.

B. Restore the latest full backup, and restore the latest differential backup. Then, restore each log backup taken before the time of failure from the most recent differential backup.

C. Perform a point-in-time restore.

D. Restore the latest full backup.

E. Restore the latest full backup, and restore the latest differential backup. Then, restore the latest log backup.

F. Perform a page restore.

G. Restore the latest full backup. Then, restore each differential backup taken before the time of failure from the most recent full backup.

H. Restore the latest full backup. Then, restore the latest differential backup.

Answer: F

Explanation:

Requirements for Restoring Pages

A page restore is subject to the following requirements:

- The databases must be using the full or bulk-logged recovery model. Some issues exist if you are using the bulk-logged model. For more information, see the following section.
- Pages in read-only filegroups cannot be restored. Trying to make a filegroup read-only will fail if there is a page restore going on at the same time in the filegroup.
- The restore sequence must start with a full, file, or filegroup backup.-

A page restore requires an unbroken chain of log backups up to the current log file, and they must all be applied so that the page is brought up to date with the current log file.

- As in a file-restore sequence, in each restore step, you can add more pages to the roll forward set.
- A database backup and page restore cannot be run at the same time.

Bulk-logged Recovery Model and Page Restore

For a database that uses the bulk-logged recovery model, page restore has the following additional conditions:

- Backing up while filegroup or page data is offline is problematic for bulk-logged data, because the offline data is not recorded in the log. Any offline page can prevent backing up the log. In this cases, consider using
- DBCC REPAIR, because this might cause less data loss than restoring to the most recent backup.
- If a log backup of a bulk-logged database encounters a bad page, it fails unless WITH
- CONTINUE_AFTER_ERROR is specified.
- Page restore generally does not work with bulk-logged recovery.
- A best practice for performing page restore is to set the database to the full recovery model, and try a log backup. If the log backup works, you can continue with the page restore. If the log backup fails, you either have to lose work since the previous log backup or you have to try running DBCC must be run with the REPAIR_ALLOW_DATA_LOSS option.

Question 3

You administer a Microsoft SQL Server database.

You create an availability group named haContosoDbs. Your primary replica is available at Server01 \Contoso01.

You need to configure the availability group to prevent data loss. In the event of a database failure, the designed secondary database must come online automatically.

Which Transact-SQL statement should you use?

Options:

A.

```
ALTER AVAILABILITY GROUP haContosoDbs
    MODIFY REPLICA ON 'Server01'\Contoso01'
    WITH (AVAILABILITY_MODE=ASYNCHRONOUS_COMMIT,
          FAILOVER_MODE=AUTOMATIC)
```

B.

```
ALTER AVAILABILITY GROUP haContosoDbs
    MODIFY REPLICA ON 'Server01'\Contoso01'
    WITH (AVAILABILITY_MODE=ASYNCHRONOUS_COMMIT,
          FAILOVER_MODE=MANUAL)
```

C.

```
ALTER AVAILABILITY GROUP haContosoDbs
    MODIFY REPLICA ON 'Server01'\Contoso01'
    WITH (AVAILABILITY_MODE=SYNCHRONOUS_COMMIT,
          FAILOVER_MODE=AUTOMATIC)
```

D.

```
ALTER AVAILABILITY GROUP haContosoDbs
    MODIFY REPLICA ON 'Server01'\Contoso01'
    WITH (AVAILABILITY_MODE=SYNCHRONOUS_COMMIT,
          FAILOVER_MODE=MANUAL)
```

Answer: C

Explanation:

Explanation:

Failure conditions can trigger an automatic failover only if both the primary and secondary replicas are configured for automatic failover mode (FAILOVER_MODE = AUTOMATIC) and the secondary replica is currently synchronized with the primary replica.

AUTOMATIC is supported only if you also specify AVAILABILITY_MODE = SYNCHRONOUS_COMMIT.

References: <https://docs.microsoft.com/en-us/sql/t-sql/statements/alter-availability-group-transact-sql>

Question 4

You administer all the deployments of Microsoft SQL Server 2012 in your company.

You need to ensure that an OLTP database that includes up-to-the-minute reporting requirements can be off-loaded from the primary database to another server. You also need to be able to add indexes to the secondary database.

Which configuration should you use?

Options:

A. Two servers configured in different data centers

SQL Server Availability Group configured in Synchronous-Commit Availability Mode

One server configured as an Active Secondary

B. Two servers configured in the same data center

SQL Server Availability Group configured in Asynchronous-Commit Availability Mode

One server configured as an Active Secondary

C. Two servers configured in the same data center

A primary server configured to perform log-shipping every 10 minutes

A backup server configured as a warm standby

D. Two servers configured in different data centers

SQL Server Availability Group configured in Asynchronous-Commit Availability Mode

E. Two servers configured on the same subnet

SQL Server Availability Group configured in Synchronous-Commit Availability Mode

F. SQL Server that includes an application database configured to perform transactional replication

G. SQL Server that includes an application database configured to perform snapshot replication

H. Two servers configured in a Windows Failover Cluster in the same data center

SQL Server configured as a clustered instance

Answer: F

Explanation:

Reference:

<http://msdn.microsoft.com/en-us/library/jj542414.aspx>

Question 5

You administer all the deployments of Microsoft SQL Server 2012 in your company.

You need to ensure that data changes are sent to a non-SQL Server database server in near real time.

You also need to ensure that data on the primary server is unaffected.

Which configuration should you use?

Options:

- A. SQL Server that includes an application database configured to perform transactional replication
- B. Two servers configured in different data centers
SQL Server Availability Group configured in Asynchronous-Commit Availability Mode
- C. Two servers configured in different data centers
SQL Server Availability Group configured in Synchronous-Commit Availability Mode
One server configured as an Active Secondary
- D. SQL Server that includes an application database configured to perform snapshot replication
- E. Two servers configured in the same data center
SQL Server Availability Group configured in Asynchronous-Commit Availability Mode
One server configured as an Active Secondary
- F. Two servers configured on the same subnet
SQL Server Availability Group configured in Synchronous-Commit Availability Mode
- G. Two servers configured in a Windows Failover Cluster in the same data center
SQL Server configured as a clustered instance
- H. Two servers configured in the same data center
A primary server configured to perform log-shipping every 10 minutes
A backup server configured as a warm standby

Answer: A

Explanation:

Reference:

<http://msdn.microsoft.com/en-us/library/ms151149.aspx>

Question 6

You administer all the deployments of Microsoft SQL Server 2012 in your company.

A database contains a large product catalog that is updated periodically.

You need to be able to send the entire product catalog to all branch offices on a monthly basis.

Which configuration should you use?

Options:

- A. Two servers configured in the same data center
A primary server configured to perform log-shipping every 10 minutes
A backup server configured as a warm standby
- B. SQL Server that includes an application database configured to perform transactional replication

- C. Two servers configured in the same data center
SQL Server Availability Group configured in Asynchronous-Commit Availability Mode
One server configured as an Active Secondary
- D. D Two servers configured in a Windows Failover Cluster in the same data center
SQL Server configured as a clustered instance
- E. SQL Server that includes an application database configured to perform snapshot replication
- F. Two servers configured in different data centers
SQL Server Availability Group configured in Synchronous-Commit Availability Mode
One server configured as an Active Secondary
- G. Two servers configured on the same subnet
SQL Server Availability Group configured in Synchronous-Commit Availability Mode
- H. Two servers configured in different data centers
SQL Server Availability Group configured in Asynchronous-Commit Availability Mode

Answer: E

Question 7

Note: This question is part of a series of questions that use the same set of answer choices. An answer choice may be correct for more than one question in the series.

You administer all the deployments of Microsoft SQL Server in your company.

You need to ensure that an OLTP database that uses a storage area network (SAN) remains available if any of the servers fail.

You also need to minimize the amount of storage used by the database.

Which configuration should you use?

Options:

- A. Two servers configured in different data centers
SQL Server Availability Group configured in Synchronous-Commit Availability Mode
One server configured as an Active Secondary
- B. SQL Server that includes an application database configured to perform transactional replication
- C. Two servers configured in the same data center
SQL Server Availability Group configured in Asynchronous-Commit Availability Mode
One server configured as an Active Secondary
- D. Two servers configured in different data centers
SQL Server Availability Group configured in Asynchronous-Commit Availability Mode
- E. Two servers configured in the same data center

A primary server configured to perform log-shipping every 10 minutes

A backup server configured as a warm standby

F. Two servers configured on the same subnet

SQL Server Availability Group configured in Synchronous-Commit Availability Mode

G. SQL Server that includes an application database configured to perform snapshot replication

H. Two servers configured in a Windows Failover Cluster in the same data center

SQL Server configured as a clustered instance

Answer: H

Question 8

Note: This question is part of a series of questions that use the same set of answers choices. An answer choice may be correct for more than one question in the series.

You administer a Microsoft SQL Server database server that hosts a transactional database and a reporting

database. The transactional database is updated through a web application and is operational throughout the day. The reporting database is only updated from the transactional database.

The recovery model and backup schedule are configured as shown in the following table:

Database	Description
Transactional database	<p>Recovery model:</p> <ul style="list-style-type: none"> • Full <p>Backup schedule:</p> <ul style="list-style-type: none"> • Full database backup: midnight, daily • Differential database backup: on the hour, every two hours starting at 02:00 hours except at 00:00 hours • Log backup: every half hour, except at the times of full and differential backups
Reporting database	<p>Recovery model:</p> <ul style="list-style-type: none"> • Simple <p>Backup schedule:</p> <ul style="list-style-type: none"> • Full database backup: 01:00 hours daily • Differential database backup: 13:00 hours daily <p>Data updates:</p> <ul style="list-style-type: none"> • Changes in data are updated from the transactional database to the reporting database at 00:30 hours and at 12:30 hours • The update takes 15 minutes

The differential backup of the reporting database fails. Then, the reporting database fails at 14:00 hours. You need to ensure that the reporting database is restored. You also need to ensure that data loss is minimal.

What should you do?

Options:

- A. Restore the latest full backup, and restore the latest differential backup. Then, restore the latest log backup.
- B. Perform a point-in-time restore.
- C. Restore the latest full backup.
- D. Restore the latest full backup, and restore the latest differential backup. Then, restore each log backup taken before the time of failure from the most recent differential backup.
- E. Restore the latest full backup. Then, restore the latest differential backup.

F. Restore the latest full backup. Then, restore each differential backup taken before the time of failure from the most recent full backup.

G. Perform a page restore.

H. Perform a partial restore.

Answer: C

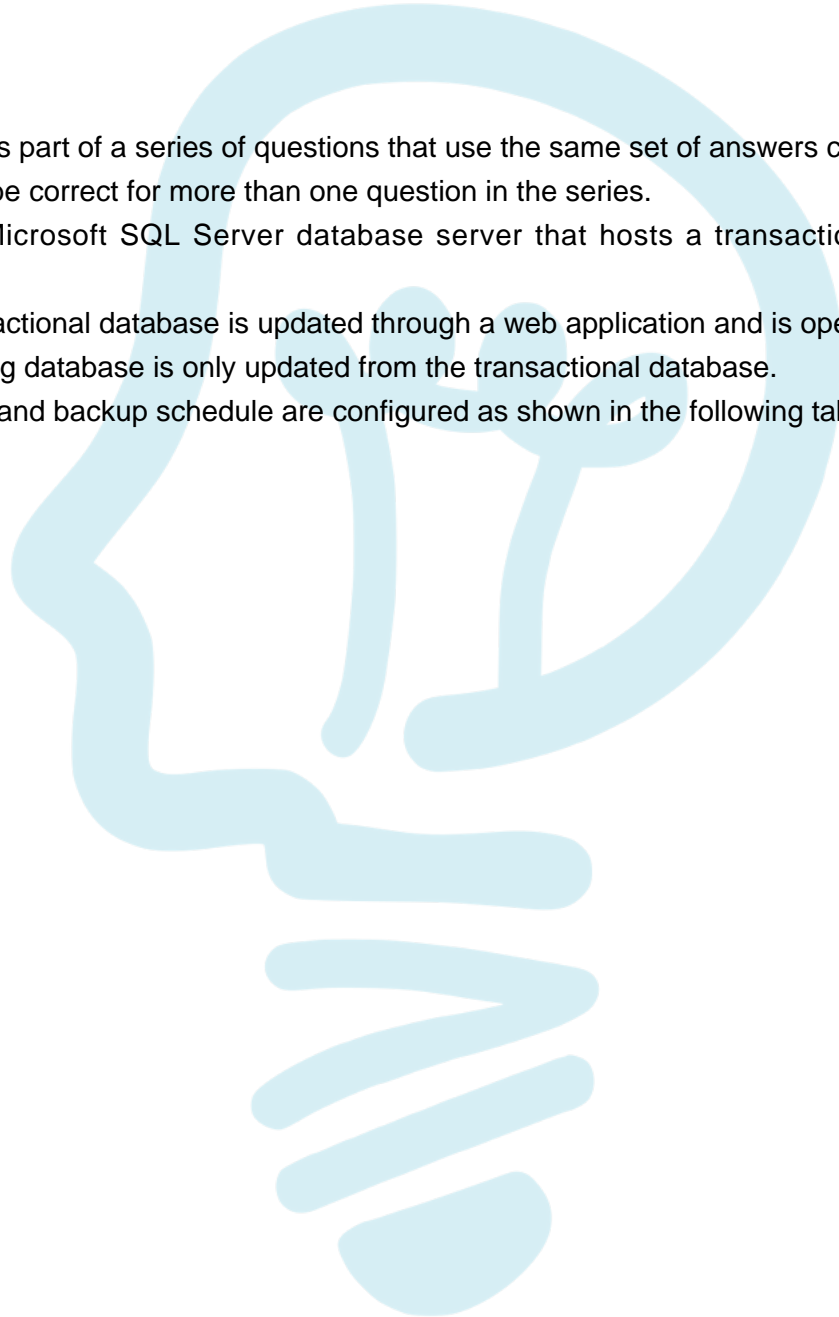
Question 9

Note: This question is part of a series of questions that use the same set of answers choices. An answer choice may be correct for more than one question in the series.

You administer a Microsoft SQL Server database server that hosts a transactional database and a reporting

database. The transactional database is updated through a web application and is operational throughout the day. The reporting database is only updated from the transactional database.

The recovery model and backup schedule are configured as shown in the following table:



Database	Description
Transactional database	Recovery model: <ul style="list-style-type: none"> • Full Backup schedule: <ul style="list-style-type: none"> • Full database backup: midnight, daily • Differential database backup: on the hour, every two hours starting at 02:00 hours except at 00:00 hours • Log backup: every half hour, except at the times of full and differential backups
Reporting database	Recovery model: <ul style="list-style-type: none"> • Simple Backup schedule: <ul style="list-style-type: none"> • Full database backup: 01:00 hours daily • Differential database backup: 13:00 hours daily Data updates: <ul style="list-style-type: none"> • Changes in data are updated from the transactional database to the reporting database at 00:30 hours and at 12:30 hours • The update takes 15 minutes

One of the hard disk drives that stores the reporting database fails at 16:40 hours.

You need to ensure that the reporting database is restored. You also need to ensure that data loss is minimal.

What should you do?

Options:

- A. Restore the latest full backup. Then, restore each differential backup taken before the time of failure from the most recent full backup.
- B. Perform a partial restore.
- C. Restore the latest full backup, and restore the latest differential backup. Then, restore the latest log

backup.

D. Perform a point-in-time restore.

E. Restore the latest full backup.

F. Perform a page restore.

G. Restore the latest full backup, and restore the latest differential backup. Then, restore each log backup taken before the time of failure from the most recent differential backup.

H. Restore the latest full backup. Then, restore the latest differential backup.

Answer: H

Question 10

DRAG DROP

You administer three Microsoft SQL Server 2012 servers named ServerA, ServerB, and ServerC. ServerA is the acting principal and ServerB is the mirror.

You need to add ServerC as a witness to the existing mirroring session between ServerA and ServerB.

You need to achieve this goal without delaying synchronization.

Which three actions should you perform in sequence? (To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.)

Select and Place:

Actions

Answer Area

- On ServerC, create an endpoint for use by the witness
- Ensure that the same Windows Login exists on each server and grant Connect permissions to each server's endpoint
- On ServerA, alter the principal database to use the endpoint on ServerC as the witness
- On ServerA, pause the mirroring session between ServerA and Server B
- On ServerB, alter the principal database to use the endpoint on ServerC as the witness
- Ensure that the same Proxy exists on each server and grant Connect permissions to each server's endpoint
- On ServerA, resume the mirroring session between ServerA and ServerB

Three empty rectangular boxes for the answer area.

Options:

A.

Actions

Answer Area

-
-
-
- On ServerA, pause the mirroring session between ServerA and Server B
- On ServerB, alter the principal database to use the endpoint on ServerC as the witness
- Ensure that the same Proxy exists on each server and grant Connect permissions to each server's endpoint
- On ServerA, resume the mirroring session between ServerA and ServerB

- On ServerC, create an endpoint for use by the witness
- Ensure that the same Windows Login exists on each server and grant Connect permissions to each server's endpoint
- On ServerA, alter the principal database to use the endpoint on ServerC as the witness

Answer: A

Would you like to see more? Don't miss our 70-462 PDF

file at:

<https://www.certification-questions.com/microsoft-pdf/70-462-pdf.html>

