



Microsoft

70-462 Exam

**Microsoft Administering Microsoft SQL Server 2012 Databases
Exam**

Thank you for Downloading 70-462 exam PDF Demo

You can also try our 70-462 practice exam software

Download Free Demo:

<https://www.dumpsschool.com/70-462-exam-dumps.html>

**Questions & Answers
(Demo Version – Limited Content)**

Version: 19.0

Question: 1

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?

- A. Two servers configured in different data centersSQL Server Availability Group configured in Synchronous-Commit Availability ModeOne server configured as an Active Secondary
- B. Two servers configured in the same data centerSQL Server Availability Group configured in Asynchronous-Commit Availability ModeOne server configured as an Active Secondary
- C. Two servers configured in the same data centerA primary server configured to perform log-shipping every 10 minutesA backup server configured as a warm standby
- D. Two servers configured in different data centersSQL Server Availability Group configured in Asynchronous-Commit Availability Mode
- E. Two servers configured on the same subnetSQL 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 centerSQL Server configured as a clustered instance

Answer: F

Question: 2

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?

- A. SQL Server that includes an application database configured to perform transactional replication
- B. Two servers configured in different data centersSQL Server Availability Group configured in Asynchronous-Commit Availability Mode
- C. Two servers configured in different data centersSQL Server Availability Group configured in Synchronous-Commit Availability ModeOne 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 centerSQL Server Availability Group configured in Asynchronous-Commit Availability ModeOne server configured as an Active Secondary
- F. Two servers configured on the same subnetSQL Server Availability Group configured in Synchronous-Commit Availability Mode
- G. Two servers configured in a Windows Failover Cluster in the same data centerSQL 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

Question: 3

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?

- 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. 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: 4

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?

- 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: 5

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 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

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?

- 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: 6

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 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?

- 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: 7

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 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?

- 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:

Bulk-logged Recovery Model and Page Restore

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

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 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?

- 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: 9

DRAG DROP

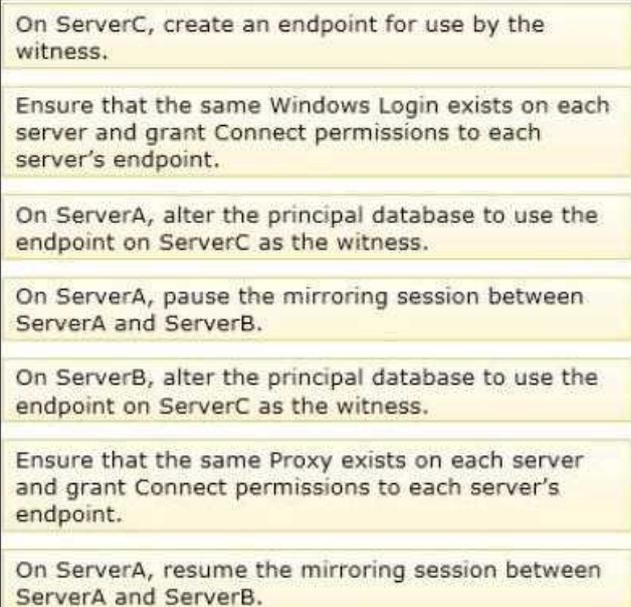
You administer three Microsoft SQL Server 2012 servers named ServerA, ServerB, and ServerC.

ServerA is the acting principal and ServerB is the mirror.

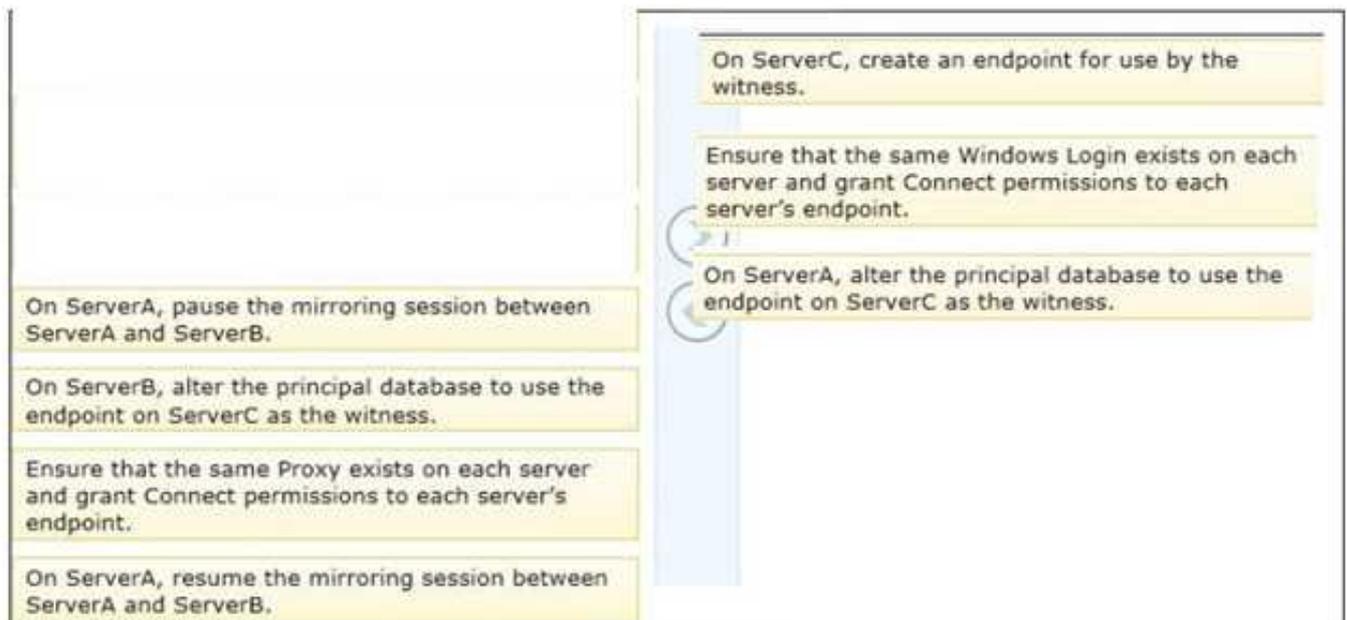
You need to add ServerC as a witness to the existing mirroring session between ServerA and ServerB.

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.)

- 
- 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 ServerB.
 - 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.

Answer:



Question: 10

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?

- A. ALTER AVAILABILITY GROUP haContosoDbsMODITY REPLICA ON 'Server01'\Contoso01'WITH (AVAILABILITY_MODE=ASYNCHRONOUS_COMMIT, FAILOVER_MODE=AUTOMATIC)
- B. ALTER AVAILABILITY GROUP haContosoDbsMODIFY REPLICA ON 'Server01'\Contoso01'WITH (AVAILABILITY_MODE=ASYNCHRONOUS_COMMIT, FAILOVER_MODE=MANUAL)
- C. ALTER AVAILABILITY GROUP haContosoDbsMODITY REPLICA ON 'Server01'\Contoso01'WITH (AVAILABILITY_MODE=SYNCHRONOUS_COMMIT, FAILOVER_MODE=AUTOMATIC)
- D. ALTER AVAILABILITY GROUP haContosoDbsMODIFY REPLICA ON 'Server01'\Contoso01'WITH (AVAILABILITY_MODE=SYNCHRONOUS_COMMIT,FAILOVER_MODE=MANUAL)

Answer: C

Question: 11

DRAG DROP

You administer several Microsoft SQL Server servers. Your company has a number of offices across the world connected by using a wide area network (WAN).

Connections between offices vary significantly in both bandwidth and reliability.

You need to identify the correct replication method for each scenario.

What should you do? (To answer, drag the appropriate replication method or methods to the correct location or locations in the answer are

a. Each replication method may be used once, more than once, or not at all.)

Replication Method	Scenario
Transactional Replication	Multiple databases on the same low-latency subnet must allow applications to write changes locally, and these changes must be replicated to all related databas
Peer-to-Peer Replication	An order summary table is repopulated once a week. This table must be replicated to all databases.
Snapshot Replication	Field offices using unreliable connections keep a local copy of the product catalog and process orders locally. These orders must be periodically replicated to all other
Merge Replication	Information in an order-tracking database must be replicated across a low-latency connection as changes occur to multiple reporting databases.

Answer:

Replication Method	Scenario
	Multiple databases on the same low-latency subnet must allow applications to write changes locally, and these changes must be replicated to all related databas
	An order summary table is repopulated once a week. This table must be replicated to all databases.
	Field offices using unreliable connections keep a local copy of the product catalog and process orders locally. These orders must be periodically replicated to all other
	Information in an order-tracking database must be replicated across a low-latency connection as changes occur to multiple reporting databases.

Question: 12

You administer several Microsoft SQL Server database servers. Merge replication has been configured for an application that is distributed across offices throughout a wide area network (WAN).

Many of the tables involved in replication use the XML and varchar (max) data types. Occasionally, merge replication fails due to timeout errors.

You need to reduce the occurrence of these timeout errors.

What should you do?

- A. Set the Merge agent on the problem subscribers to use the slow link agent profile.
- B. Create a snapshot publication, and reconfigure the problem subscribers to use the snapshot publication.
- C. Change the Merge agent on the problem subscribers to run continuously.
- D. Set the Remote Connection Timeout on the Publisher to 0.

Answer: A

Explanation:

When replication is configured, a set of agent profiles is installed on the Distributor. An agent profile contains a set of parameters that are used each time an agent runs: each agent logs in to the Distributor during its startup process and queries for the parameters in its profile.

For merge subscriptions that use Web synchronization, profiles are downloaded and stored at the Subscriber. If the profile is changed, the profile at the Subscriber is updated the next time the Merge Agent runs. For more information about Web synchronization, see Web Synchronization for Merge Replication.

Replication provides a default profile for each agent and additional predefined profiles for the Log Reader Agent, Distribution Agent, and Merge Agent. In addition to the profiles provided, you can create profiles suited to your application requirements. An agent profile allows you to change key parameters easily for all agents associated with that profile. For example, if you have 20 Snapshot Agents and need to change the query timeout value (the - QueryTimeout parameter), you can update the profile used by the Snapshot Agents and all agents of that type will begin using the new value automatically the next time they run.

You might also have different profiles for different instances of an agent. For example, a Merge Agent that connects to the Publisher and Distributor over a dialup connection could use a set of parameters that are better suited to the slower communications link by using the slow link profile.

Question: 13

DRAG DROP

You administer two Microsoft SQL Server 2012 servers named ServerA and ServerB. You use a database named AdventureWorks.

You need to prepare the AdventureWorks database for database mirroring. ServerB will act as the mirror in a mirroring partnership along with ServerA.

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.)

Back up AdventureWorks on ServerA by using a full backup.	
Back up AdventureWorks on ServerA by using a full backup followed by a transaction log backup by using the NORECOVERY option.	
Back up AdventureWorks on ServerA by using a transaction log backup. Restore the transaction log backup by using the RECOVERY option on ServerB.	
Back up AdventureWorks on ServerA by using a transaction log backup. Restore the transaction log backup by using the NORECOVERY option on ServerB.	
Restore the full database backup of AdventureWorks by using the NORECOVERY option on ServerB as AdventureWorks.	
Restore the full database backup of AdventureWorks by using the RECOVERY option on ServerB as AdventureWorks_Mirror.	

Answer:



Question: 14

You create an availability group named HaContoso that has replicas named Server01/HA, Server02/HA, and Server03/HA.

Currently, Server01/HA is the primary replica.

You need to ensure that the following requirements are met:

Backup operations occur on Server02/HA.

If Server02/HA is unavailable, backup operations occur on Server03/HA.

Backup operations do not occur on Server01/HA.

How should you configure HaContoso?

- A. Set the backup preference of HaContoso to Prefer Secondary. Set the backup priority of Server02/HA to 20. Set the backup priority of Server03/HA to 10.
- B. Set the backup preference of HaContoso to Secondary only. Set the backup priority of Server02/HA to 20. Set the backup priority of Server03/HA to 10.
- C. Set the backup preference of HaContoso to Secondary only. Set the backup priority of Server02/HA to 10. Set the backup priority of Server03/HA to 20.
- D. Set the exclude replica of Server01/HA to true. Set the backup priority of Server02/HA to 10. Set the backup priority of Server03/HA to 20.

Answer: B

Question: 15

You administer a Microsoft SQL Server 2012 instance that has several SQL Server Agent jobs configured.

When SQL Server Agent jobs fail, the error messages returned by the job steps are truncated.

The following error message is an example of the truncated error message:

"Executed as user CONTOSO\ServiceAccount....0.4035.00 for 64-bit Copyright (C) Microsoft Corp 1984-2011. All rights reserved. Started 63513 PM Error 2012-06-23 183536.87 Code 0XC001000E

Source UserImport Description Code 0x00000000 Source Log Import Activity Descript... The package execution f

A. ... The step failed. "

You need to ensure that all the details of the job step failures are retained for SQL Server Agent jobs.

What should you do?

- A. Expand agent logging to include information from all events.
- B. Disable the Limit size of job history log feature.
- C. Configure event forwarding.
- D. Configure output files.

Answer: D

Question: 16

DRAG DROP

You administer a Microsoft SQL Server 2012 environment that contains a production SQL Server 2005 instance named SQL2005 and a development SQL Server 2012 instance named SQL2012.

The development team develops a new application that uses the SQL Server 2012 functionality. You are planning to migrate a database from SQL2005 to SQL2012 so that the development team can test their new application.

You need to migrate the database without affecting the production environment.

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.)

Perform a transaction log on SQL2005.	
Perform a full database on SQL2005.	
Perform a VSS backup on the database on SQL2005.	
Restore the VSS backup on SQL2012.	
Restore the database backup and transaction log backup on SQL2012.	
Change the compatibility level of the database to 120 on SQL2012.	
Change the compatibility level of the database to 110 on SQL2012.	

Answer:

Perform a transaction log on SQL2005.		Perform a full database on SQL2005.
Perform a VSS backup on the database on SQL2005.		Restore the database backup and transaction log backup on SQL2012.
Restore the VSS backup on SQL2012.		Change the compatibility level of the database to 110 on SQL2012.
Change the compatibility level of the database to 120 on SQL2012.		

Question: 17

DRAG DROP

You administer three Microsoft SQL Server 2008 R2 instances.

Database mirroring is configured in High-Safety mode with Automatic Failover between the following three servers:

SQL1 is the Principal server.

SQL2 is the mirror server.

SQL3 is the witness server.

You need to upgrade SQL1 and SQL2 to SQL Server 2012. You need to ensure that downtime is minimized during the upgrade.

Which six 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.)

Configure log shipping between SQL1 and SQL2.		
Upgrade SQL1 to SQL Server 2012.		
Upgrade SQL2 to SQL Server 2012.		
Disable log shipping between SQL1 and SQL2.		
Manually failover the database from SQL1 to SQL2.		
Manually failover the database from SQL2 to SQL1.		
Add SQL3 back to the database mirroring solution.		
Remove SQL3 from the database mirroring solution.		

Answer:



Question: 18

You administer a Microsoft SQL Server server that has SQL Server Integration Services (SSIS) installed. You plan to deploy new SSIS packages to the server. The SSIS packages use the Project Deployment Model together with parameters and Integration Services environment variables. You need to configure the SQL Server environment to support these packages. What should you do?

- A. Create SSIS configuration files for the packages.
- B. Create an Integration Services catalog.
- C. Install Data Quality Services.
- D. Install Master Data services.

Answer: B

Question: 19

You administer a Microsoft SQL Server 2012 server. When transaction logs grow, SQL Server must send an email message to the database administrators. You need to configure SQL Server to send the email messages. What should you configure?

- A. SQL Mail
- B. An Extended Events session
- C. Alerts and operators in SQL Server Agent
- D. Policies under Policy-Based Management

Answer: C

Question: 20

You administer two instances of Microsoft SQL Server 2012. You deploy an application that uses a

database on the named instance.

The application is unable to connect to the database on the named instance.

You need to ensure that the application can connect to the named instance.

What should you do?

- A. Use the Data Quality Client to configure the application.
- B. Start the SQL Server Browser Service.
- C. Use the Master Data Services Configuration Manager to configure the application.
- D. Start the SQL Server Integration Services Service.

Answer: B

Question: 21

DRAG DROP

You administer a Microsoft SQL Server database.

You use an OrderDetail table that has the following definition:

```
CREATE TABLE [dbo].[OrderDetail]
  ([SalesOrderID] [int] NOT NULL,
   [SalesOrderDetailID] [int] IDENTITY(1,1) NOT NULL,
   [CarrierTrackingNumber] [nvarchar](25) NULL,
   [OrderQty] [smallint] NOT NULL,
   [ProductID] [int] NOT NULL,
   [SpecialOfferID] [int] NULL,
   [UnitPrice] [money] NOT NULL);
```

You need to create a non-clustered index on the SalesOrderID column in the OrderDetail table to include only rows that contain a value in the SpecialOfferID column.

Which four Transact-SQL statements should you use? (To answer, move the appropriate statements from the list of statements to the answer area and arrange them in the correct order.)

WHERE	
FILTER ON	
SpecialOfferID IS NOT NULL;	
ON dbo.OrderDetail (SalesOrderID)	
ON dbo.OrderDetail (SalesOrderID) AS FILTERED_INDEX	
CREATE NONCLUSTERED INDEX FIdx_SpecialOfferID	
CREATE NONCLUSTERED FILTERED INDEX FIdx_SpecialOfferID	

Answer:

FILTER ON	CREATE NONCLUSTERED INDEX FIndx_SpecialOfferID
ON dbo.OrderDetail (SalesOrderID) AS FILTERED_INDEX	ON dbo.OrderDetail (SalesOrderID)
CREATE NONCLUSTERED FILTERED INDEX FIndx_SpecialOfferID	WHERE
	SpecialOfferID IS NOT NULL;

According to these references, this answer looks correct.

References:

<http://msdn.microsoft.com/en-us/library/ms188783.aspx>

<http://msdn.microsoft.com/en-us/library/ms189280.aspx>

Question: 22

HOTSPOT

You administer a Microsoft SQL Server database.

The database contains a table that has the following definition:

```
CREATE TABLE [Sales].[Customer] (
    [CustomerID] int NOT NULL,
    [CustomerName] nvarchar(50) NOT NULL,
    [TerritoryID] int NULL,
    [LastContactDate] datetimeoffset NULL,
    [CustomerType] nchar(1) NOT NULL,
    [Notes] varchar(250) NULL
)
```

You want to export data from the table to a flat file by using the SQL Server Import and Export Wizard.

You need to ensure that the following requirements are met:

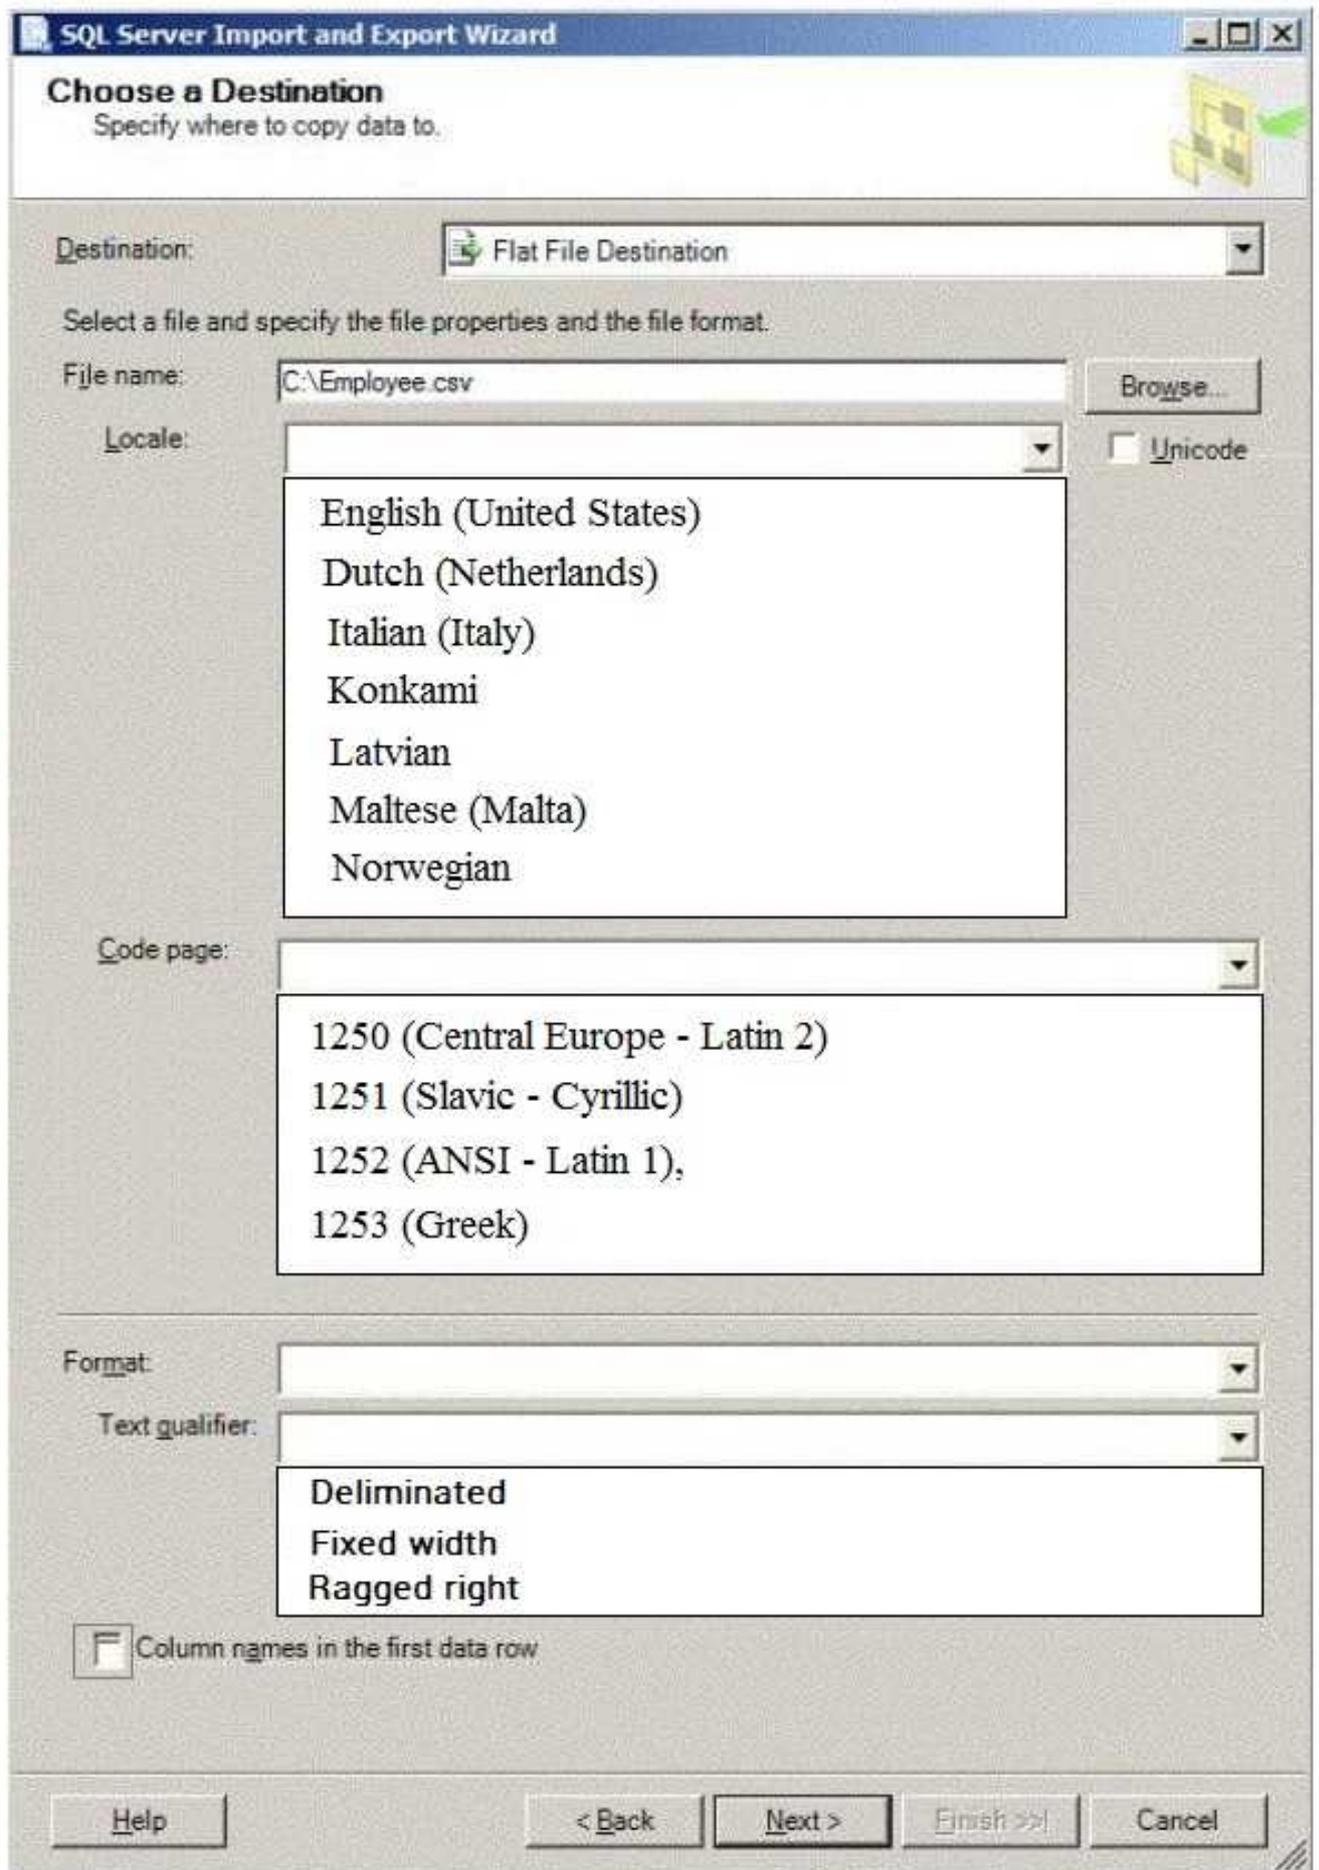
The first row of the file contains the first row of data.

Each record is of the same length.

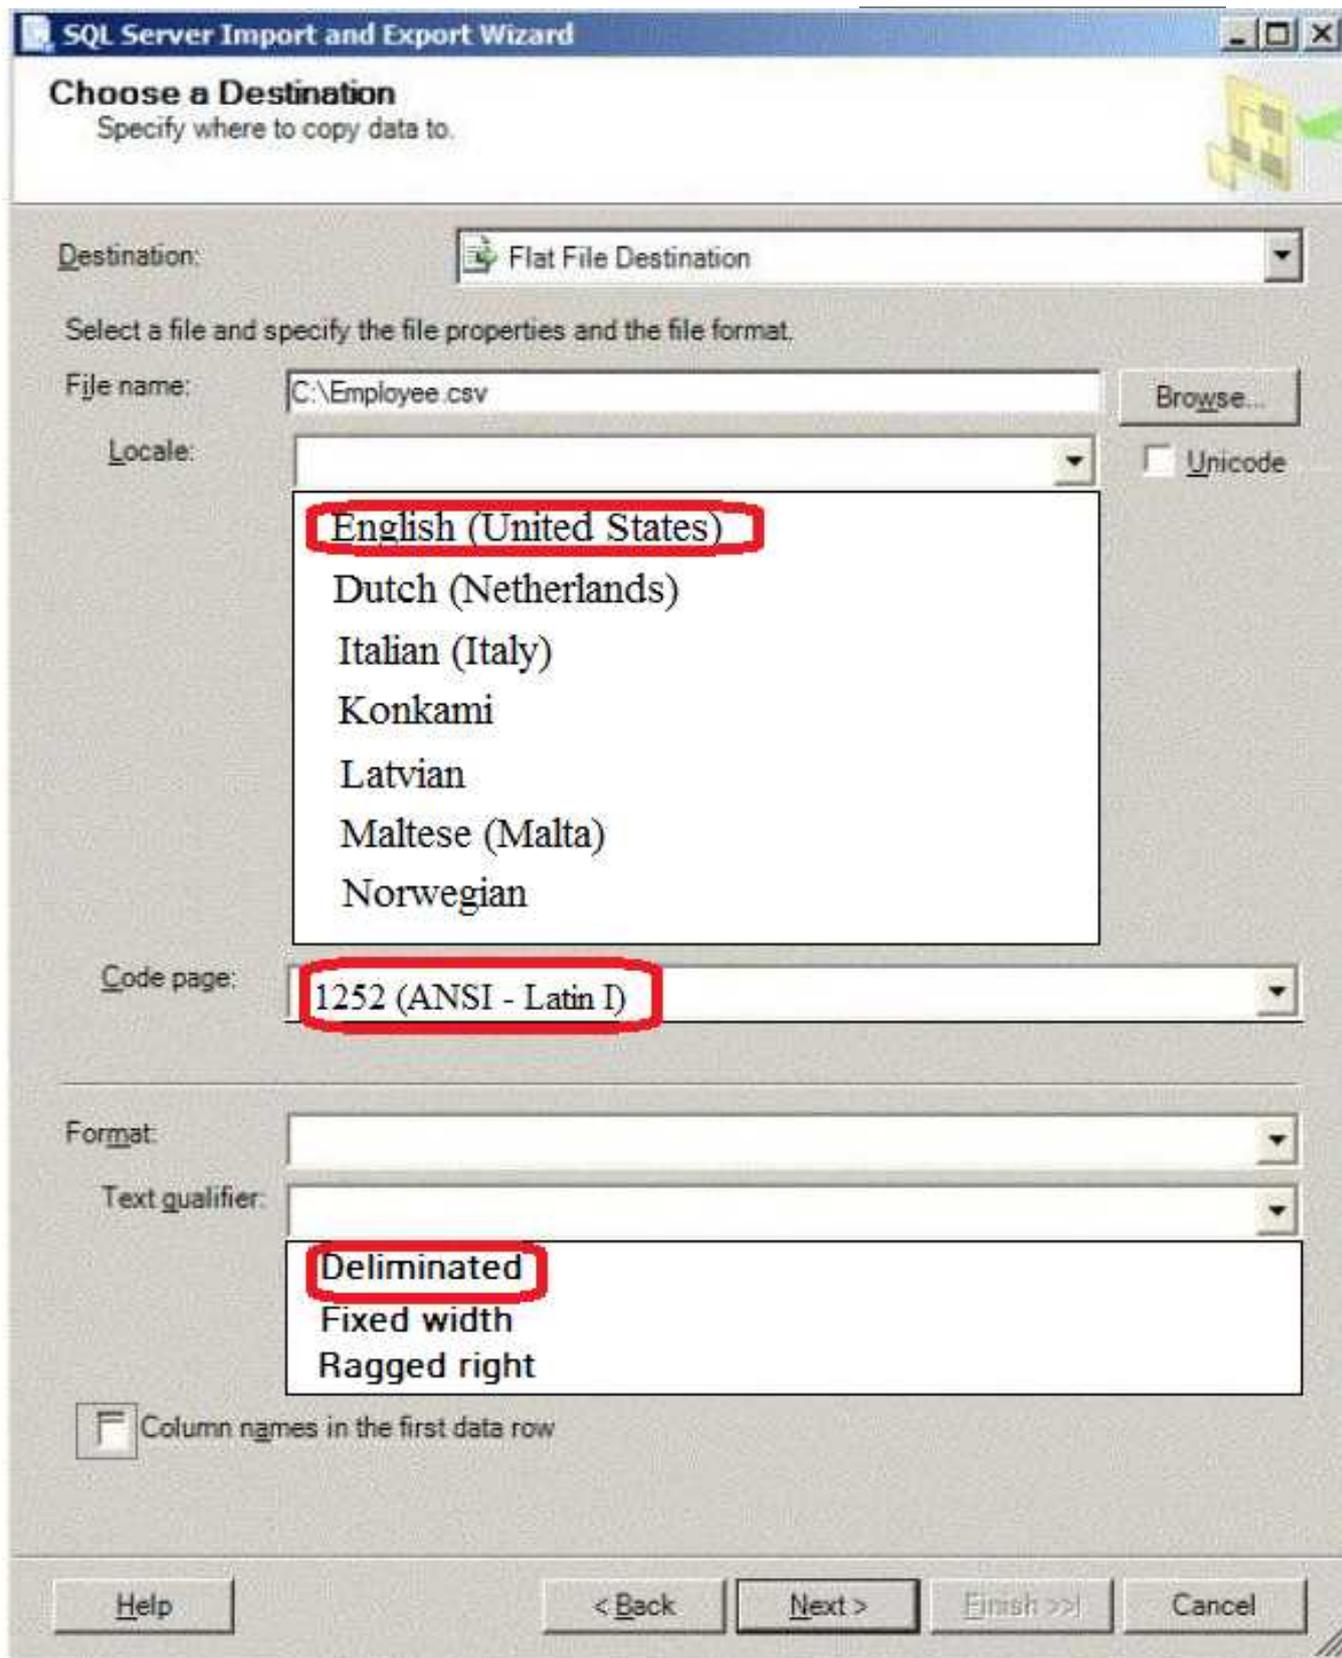
The date follows the U.S. date format.

The file supports international characters.

What should you do? (To answer, configure the appropriate option or options in the dialog box in the answer area.)



Answer:



References:

<http://msdn.microsoft.com/en-us/library/ms178804.aspx>
<http://msdn.microsoft.com/en-us/library/ms187828.aspx>

Thank You for trying 70-462 PDF Demo

To try our 70-462 practice exam software visit link below

<https://www.dumpsschool.com/70-462-exam-dumps.html>

Start Your 70-462 Preparation

Use Coupon “20OFF” for extra 20% discount on the purchase of Practice Test Software. Test your 70-462 preparation with actual exam questions.