Microsoft

Exam 70-764

Administering a SQL Database Infrastructure

Verson: Demo

[Total Questions: 10]

Question No : 1

You plan to create a database.

The database will be used by a Microsoft .NET application for a special event that will last for two days.

During the event, data must be highly available.

After the event, the database will be deleted.

You need to recommend a solution to implement the database while minimizing costs. The solution must not affect any existing applications.

What should you recommend?

More than one answer choice may achieve the goal. Select the BEST answer.

A. SQL Server 2016 Enterprise
B. SQL Server 2016 Standard
C. SQL Azure
D. SQL Server 2016 Express with Advanced Services

Answer: B

Explanation:

Programmability (AMO, ADOMD.Net, OLEDB, XML/A, ASSL) supported by Standard and Enterprise editions only.

Reference: Features Supported by the Editions of SQL Server 2016

Question No : 2 DRAG DROP

You plan to deploy SQL Server 2012.

You are designing two stored procedures named USP_1 and USP_2 that have the following requirements:

- Prevent data read by USP_1 from being modified by other active processes.
- Prevent USP_2 from performing dirty reads.

You need to recommend the isolation level for each stored procedure. The solution must maximize concurrency.

Which isolation levels should you recommend?

To answer, drag the appropriate isolation level to the correct stored procedure in the answer area.

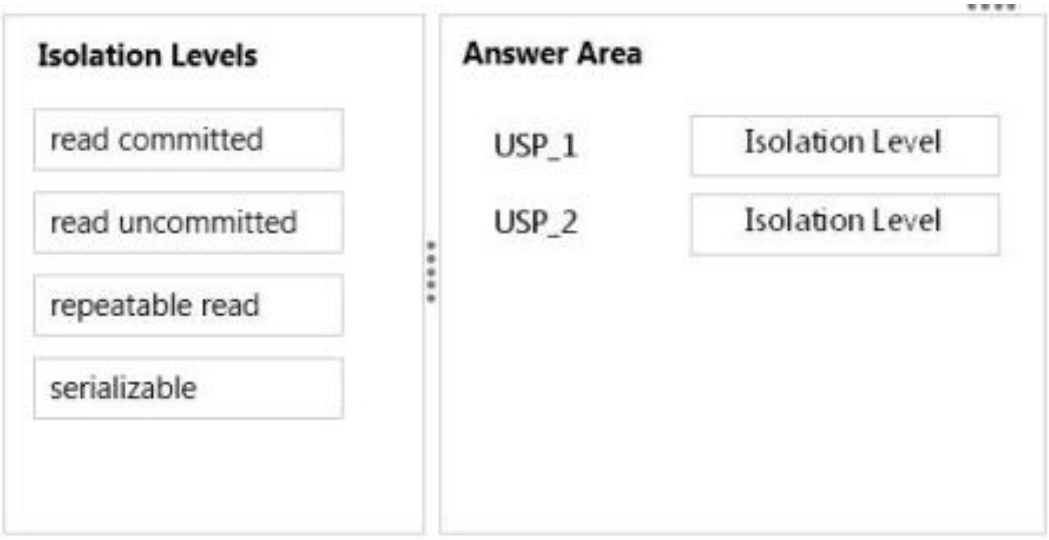

Answer:

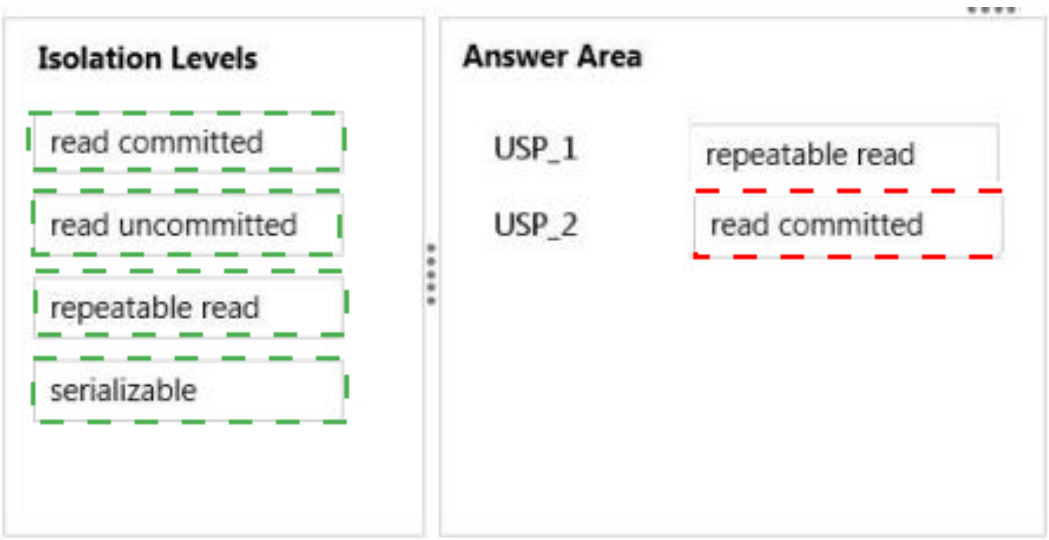

Explanation:

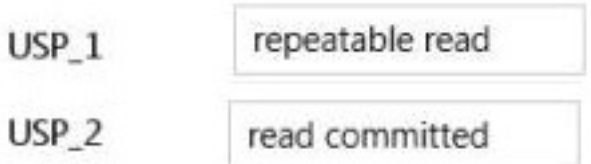

Note:
* REPEATABLE READ
This isolation level includes the guarantees given by SNAPSHOT isolation level. In addition, REPEATABLE READ guarantees that for any row that is read by the transaction, at the time the transaction commits the row has not been changed by any other transaction. Every read operation in the transaction is repeatable up to the end of the transaction.
* Committed Read is SQL Server's default isolation level. It ensures that an operation will never read data another application has changed but not yet committed.

Question No : 3

You deploy a database by using SQL Server 2016. The database contains a table named Table1.

You need to recommend a solution to track all of the deletions executed on Table1. The solution must minimize the amount of custom code required.

What should you recommend?

A. Change data capture
B. Statistics
C. A trigger
D. Master Data Services

Answer: A
Explanation:

Change data capture is designed to capture insert, update, and delete activity applied to SQL Server tables, and to make the details of the changes available in an easily consumed relational format. The change tables used by change data capture contain columns that mirror the column structure of a tracked source table, along with the metadata needed to understand the changes that have occurred.

Reference: About Change Data Capture (SQL Server)

Question No : 4

Note: This question is part of a series of questions that use the same scenario. For your convenience, the scenario is repeated in each question. Each question presents a different goal and answer choices, but the text of the scenario is exactly the same in each question in this series.

You have five servers that run Microsoft Windows 2012 R2. Each server hosts a Microsoft SQL Server instance. The topology for the environment is shown in the following diagram.

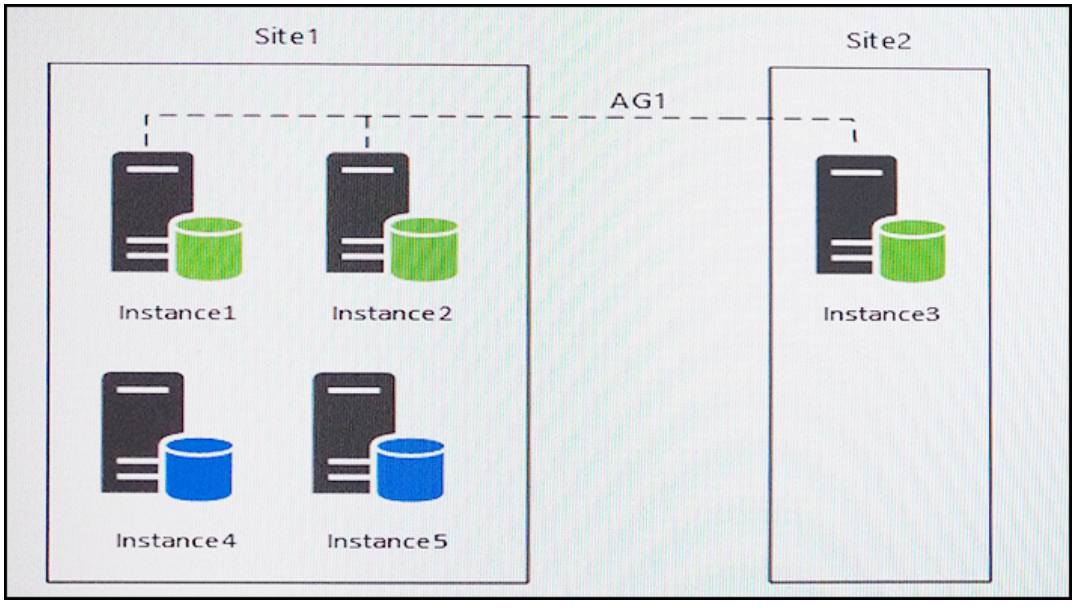

You have an Always On Availability group named AG1. The details for AG1 are shown in the following table.

Instance	Node type
Instance1	Primary
Instance2	Synchronous readable secondary
Instance3	Asynchronous readable secondary

Instance1 experiences heavy read-write traffic. The instance hosts a database named OperationsMain that is four terabytes (TB) in size. The database has multiple data files and filegroups. One of the filegroups is read_only and is half of the total database size.

Instance4 and Instance5 are not part of AG1. Instance4 is engaged in heavy read-write I/O.

Instance5 hosts a database named StagedExternal. A nightly BULK INSERT process loads data into an empty table that has a rowstore clustered index and two nonclustered rowstore indexes.

You must minimize the growth of the StagedExternal database log file during the BULK INSERT operations and perform point-in-time recovery after the BULK INSERT transaction. Changes made must not interrupt the log backup chain.

You plan to add a new instance named Instance6 to a datacenter that is geographically distant from Site1 and Site2. You must minimize latency between the nodes in AG1.

All databases use the full recovery model. All backups are written to the network location \\SQLBackup\. A separate process copies backups to an offsite location. You should minimize both the time required to restore the databases and the space required to store backups. The recovery point objective (RPO) for each instance is shown in the following table.

Instance	Recovery point objective
Instance 1	5 minutes
Instance 2	5 minutes
Instance 3	5 minutes
Instance 4	60 minutes
Instance 5	24 hours

Full backups of OperationsMain take longer than six hours to complete. All SQL Server backups use the keyword COMPRESSION.

You plan to deploy the following solutions to the environment. The solutions will access a database named DB1 that is part of AG1.

The wait statistics monitoring requirements for the instances are described in the following table.

Instance	Description
Instance1	Aggregate wait statistics since the last server restart.
Instance4	Identify the most prominent wait types for all the commands originating from a session, between session connections, or between application pool resets.
Instance5	Identify all the wait types for queries currently running on the server.

You need to reduce the amount of time it takes to backup OperationsMain.

What should you do?

A. Modify the backup script to use the keyword SKIP in the FILE_SNAPSHOT statement.
B. Modify the backup script to use the keyword SKIP in the WITH statement
C. Modify the backup script to use the keyword NO_COMPRESSION in the WITH statement.
D. Modify the full database backups script to stripe the backup across multiple backup files.

Answer: D
Explanation:

One of the filegroup is read_only should be as it only need to be backup up once. Partial backups are useful whenever you want to exclude read-only filegroups. A partial backup resembles a full database backup, but a partial backup does not contain all the filegroups. Instead, for a read-write database, a partial backup contains the data in the primary filegroup, every read-write filegroup, and, optionally, one or more read-only files. A partial backup of a read-only database contains only the primary filegroup.

From scenario: Instance1 experiences heavy read-write traffic. The instance hosts a database named OperationsMainthat is four terabytes (TB) in size. The database has multiple data files and filegroups. One of the filegroups is read_only and is half of the total database size.

References: https://docs.microsoft.com/en-us/sql/relational-databases/backup-restore/partial-backups-sql-server

Question No : 5 DRAG DROP

Note: This question is part of a series of questions that use the same scenario. For your convenience, the scenario is repeated in each question. Each question presents a different goal and answer choices, but the text of the scenario is exactly the same in each question in this series.

You have five servers that run Microsoft Windows 2012 R2. Each server hosts a Microsoft SQL Server instance. The topology for the environment is shown in the following diagram.

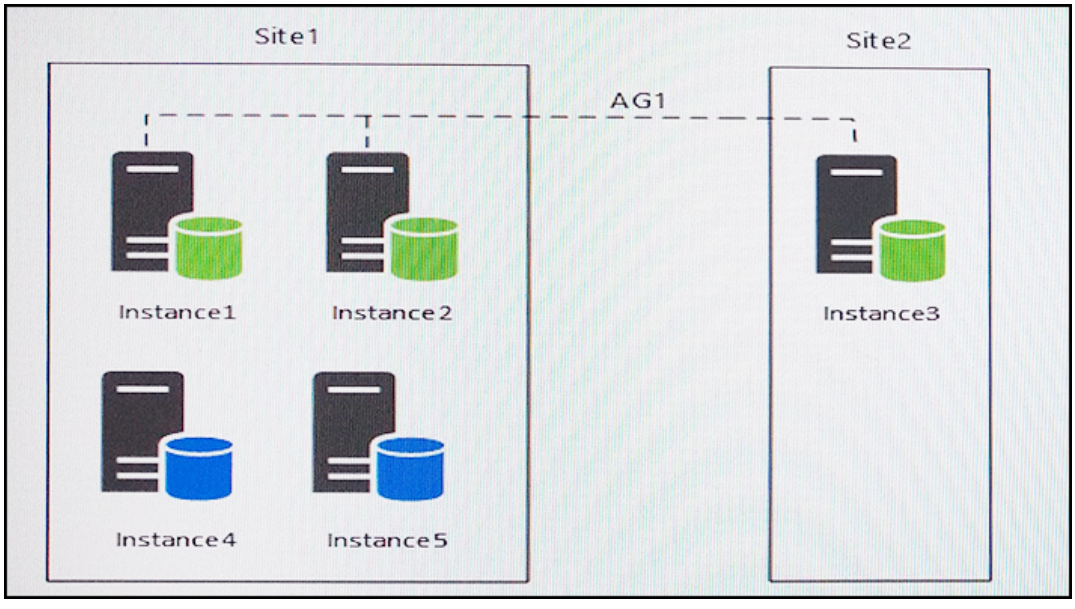

You have an Always On Availability group named AG1. The details for AG1 are shown in the following table.

Instance	Node type
Instance1	Primary
Instance2	Synchronous readable secondary
Instance3	Asynchronous readable secondary

Instance1 experiences heavy read-write traffic. The instance hosts a database named OperationsMain that is four terabytes (TB) in size. The database has multiple data files and filegroups. One of the filegroups is read_only and is half of the total database size.

Instance4 and Instance5 are not part of AG1. Instance4 is engaged in heavy read-write I/O.

Instance5 hosts a database named StagedExternal. A nightly BULK INSERT process loads data into an empty table that has a rowstore clustered index and two nonclustered rowstore

indexes.

You must minimize the growth of the StagedExternal database log file during the BULK INSERT operations and perform point-in-time recovery after the BULK INSERT transaction. Changes made must not interrupt the log backup chain.

You plan to add a new instance named Instance6 to a datacenter that is geographically distant from Site1 and Site2. You must minimize latency between the nodes in AG1.

All databases use the full recovery model. All backups are written to the network location \\SQLBackup\. A separate process copies backups to an offsite location. You should minimize both the time required to restore the databases and the space required to store backups. The recovery point objective (RPO) for each instance is shown in the following table.

Instance	Recovery point objective
Instance 1	5 minutes
Instance 2	5 minutes
Instance 3	5 minutes
Instance 4	60 minutes
Instance 5	24 hours

Full backups of OperationsMain take longer than six hours to complete. All SQL Server backups use the keyword COMPRESSION.

You plan to deploy the following solutions to the environment. The solutions will access a database named DB1 that is part of AG1.

The wait statistics monitoring requirements for the instances are described in the following table.

Instance	Description
Instance1	Aggregate wait statistics since the last server restart.
Instance4	Identify the most prominent wait types for all the commands originating from a session, between session connections, or between application pool resets.
Instance5	Identify all the wait types for queries currently running on the server.

You need to propose a new process for the StagedExternal database.

Which five actions should you recommended be performed in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

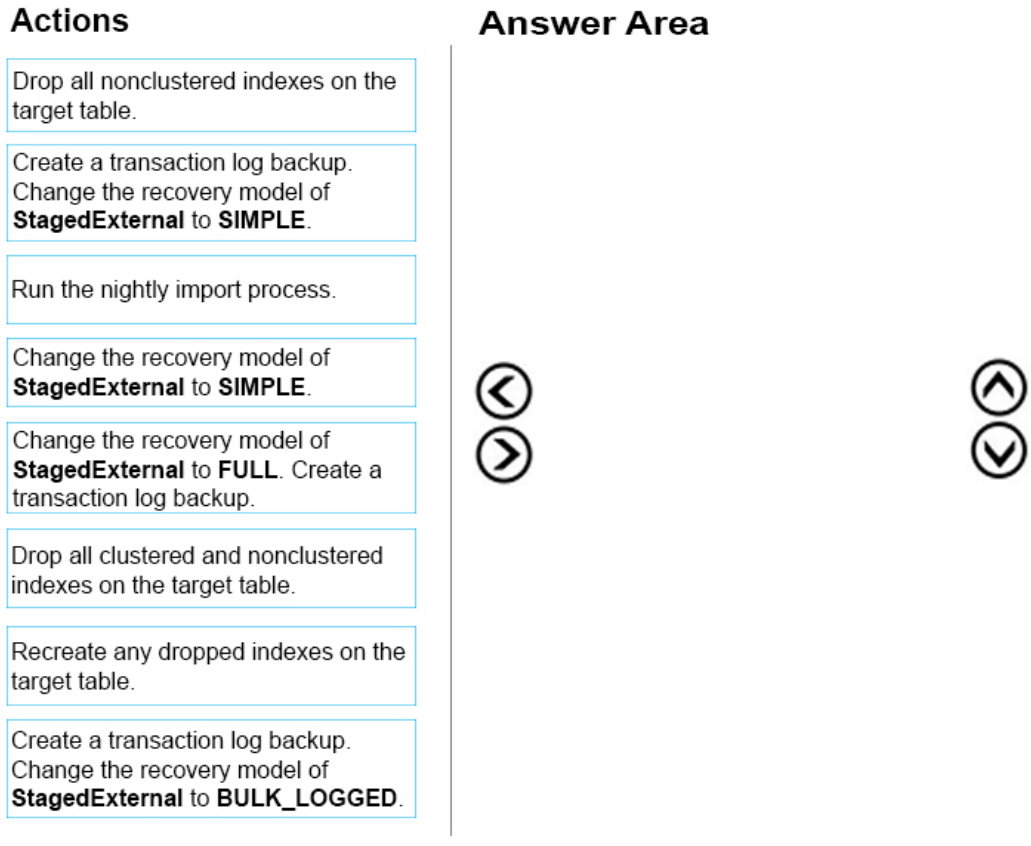

Answer:

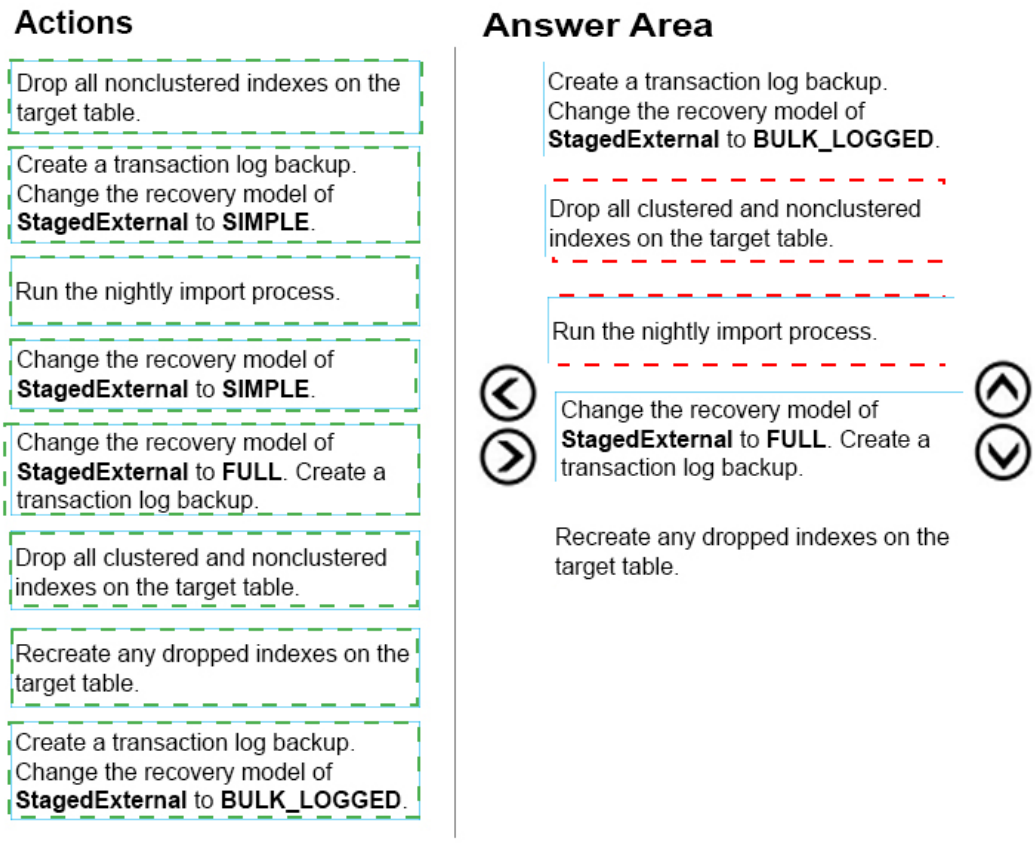

Explanation:

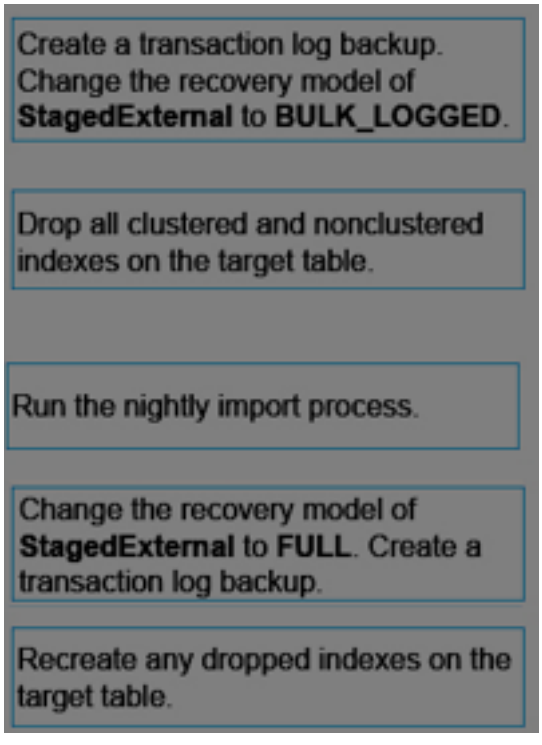

From scenario: Instance5 hosts a database named StagedExternal. A nightly BULK

INSERT process loads data into an empty table that has a rowstore clustered index and two nonclustered rowstore indexes.
You must minimize the growth of the StagedExternaldatabase log file during the BULK INSERT operations and perform point-in-time recovery after the BULK INSERT transaction. Changes made must not interrupt the log backup chain.

All databases use the full recovery model.

References:

Question No : 6

You are the administrator for a SQL Server 2016 instance that stores the data for an online transaction processing sales system.

The company takes full backups every week; differential backups on the days with no full backups; and hourly transaction backups. These backups are stored on a backup server in the company's data center.

Every week, the company places the full backup on a tape and sends it to a third-party backup storage system.

The company is worried that a disaster might occur that could destroy their computer center and cause them to lose orders.

You need to determine the best method for providing the smallest amount of data loss and downtime without leasing or purchasing additional physical locations.

What should you do? More than one answer choice may achieve the goal. Select the BEST answer.

A. Set up SQL Server Always On with a SQL Azure database as a replica.
B. Set up SQL Server Always On by using a SQL Server on a Windows Azure Virtual Machine.
C. Put the differential backup on tape and send it to the third-party backup storage system.
D. Use the Microsoft SQL Server Backup to Microsoft Windows Azure Tool to direct all backups to a different geographical location.

Answer: D

Explanation:

SQL Server 2012 was the first version to provide the ability to back up databases to the Cloud, and SQL Server 2016 improves on the process.

Microsoft SQL Server Backup to Windows Azure Tool enables backup to Windows Azure Blob Storage and encrypts and compresses SQL Server backups stored locally or in the cloud.

Reference:

Smart, Secure, Cost-Effective: SQL Server Back Up to Windows Azure - SQL Server Team Blog - Site Home - TechNet Blogs

Question No : 7

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this sections, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

A company has a server that runs Microsoft SQL Server 2016 Web edition. The server has a default instance that hosts a database named DB1.

You need to ensure that you can perform auditing at the database level for DB1.

Solution: You migrate DB1 to a named instance on a server than runs Microsoft SQL Server 2016 Standard edition.

Does the solution meet the goal?

A. Yes
B. No

Answer: B

Explanation:

All editions of SQL Server support server level audits. All editions support database level audits beginning with SQL Server 2016 SP1. Prior to that, database level auditing was limited to Enterprise, Developer, and Evaluation editions.

References: https://docs.microsoft.com/en-us/sql/relational-databases/security/auditing/sql-server-audit-database-engine

Question No : 8 DRAG DROP

You are the database administrator for a Microsoft SQL Server instance. You develop an Extended Events package to look for events related to application performance.

You need to change the event session to include SQL Server errors that are greater than error severity 15.

Which five Transact-SQL segments should you use to develop the solution? To answer, move the appropriate Transact-SQL segments from the list of Transact-SQL segments to the answer area and arrange them in the correct order.

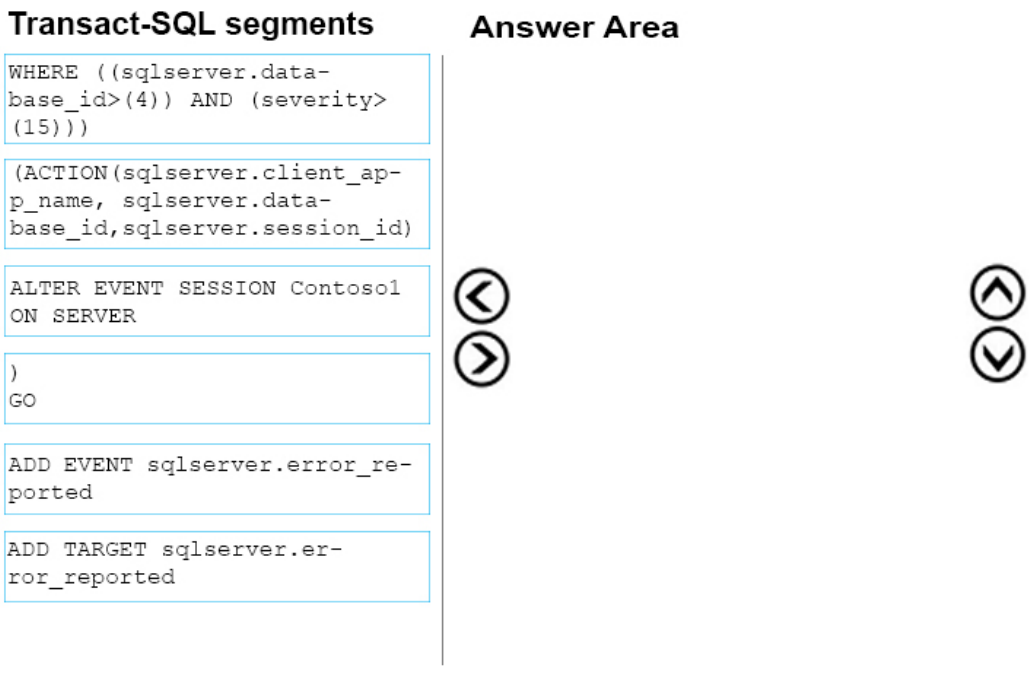

Answer:

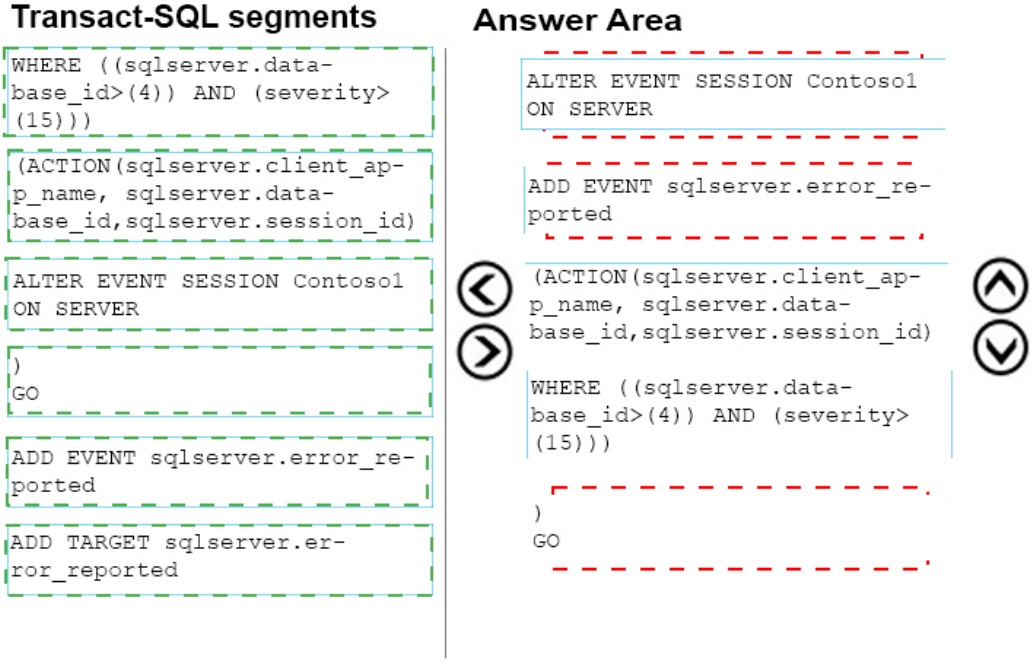

Explanation:

Step 1: ALTER EVENT SESSION Contoso1 ON SERVER

Step 2: ADD EVENT …

Step 3: (ACTION ...

Step 4: WHERE...

Step 5:) GO

Example: To start an Extended Events sessions in order to trap SQL Server errors with severity greater than 10,just run the following script:

CREATE EVENT SESSION [error_trap] ON SERVER
ADD EVENT sqlserver.error_reported
(
ACTION
(package0.collect_system_time,package0.last_error,sqlserver.client_app_name,sqlserver.client_hostname,sqlserver.database_id,sqlserver.database_name,sqlserver.nt_username,sqlserver.plan_handle,sqlserver.query_hash,sqlserver.session_id,sqlserver.sql_text,sqlserver.tsql_frame,sqlserver.tsql_stack,sqlserver.username)
WHERE ([severity]>10)

```
)
ADD TARGET package0.event_file
(
SET filename=N'D:\Program Files\Microsoft SQL
Server\MSSQL11.MSSQLSERVER\MSSQL\XEvents\error_trap.xel'
)
WITH
(
STARTUP_STATE=OFF
)
GO
```

References:

Question No : 9

You have four databases that are accessed by using an Online Transaction Processing (OLTP) application. The databases are stored on a server named SQL1 that has SQL Server 2016 installed.

You plan to deploy an additional server that has SQL Server 2016 installed.

You need to design a high-availability solution for the databases that meets the following requirements:

If SQL1 fails, the databases must be available.

Users must be able to run reports against a secondary copy of the databases.

What should you include in the design?

More than one answer choice may achieve the goal. Select the BEST answer.

A. AlwaysOn availability groups
B. Database mirroring
C. Log shipping
D. Failover Clustering

Answer: A

Explanation:

The AlwaysOn Availability Groups feature is a high-availability and disaster-recovery solution that provides an enterprise-level alternative to database mirroring.
Introduced in SQL Server 2012, AlwaysOn Availability Groups maximizes the availability of a set of user databases for an enterprise. An availability group supports a failover environment for a discrete set of user databases, known as availability databases that fail over together.

Reference: Failover and Failover Modes (AlwaysOn Availability Groups)

Question No : 10 DRAG DROP

You have a server named SQL1 that has SQL Server 2016 installed. SQL1 hosts a database named Database! Database1 contains a table named Table1. Table1 is partitioned across five filegroups based on the TransactionType field.

The schema of Table1 is configured as shown in the following table.

Column	Data type
ID	Bigint
Account	Bigint
Amount	Decimal
TransactionType	Int
TransactionDate	Date

Table1 contains the indexes shown in the following table.

Index	Type	Column
PK_Table1	Clustered, primary key	ID, TransactionType
IX_Account	Nonclustered	Account
IX_Type	Nonclustered	TransactionType
IX_Date	Nonclustered	TransactionDate
IX_Amount	Nonclustered	Amount

You need to recommend an index strategy to maximize performance for the queries that consume the indexes available to Table1.

Which type of index storage should you recommend?

To answer, drag the appropriate index storage type to the correct index in the answer area.

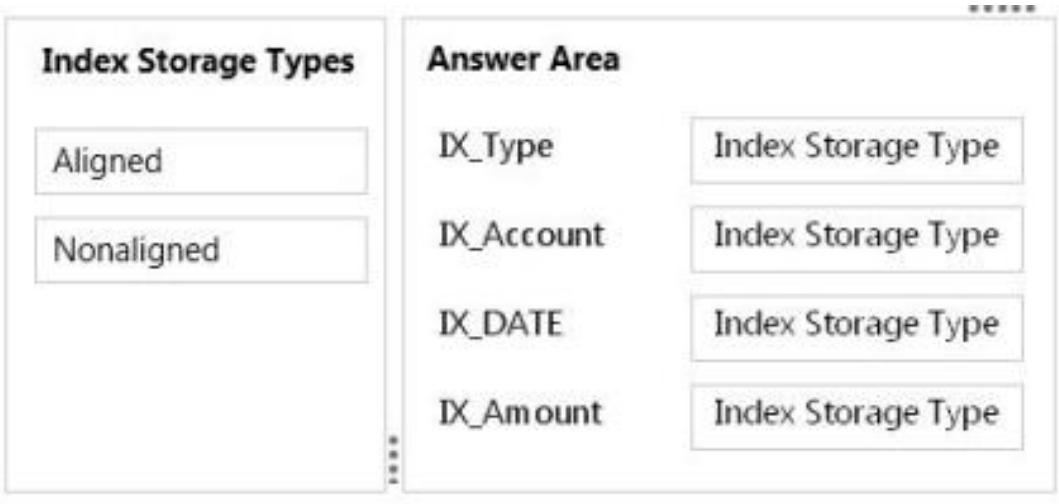

Answer:

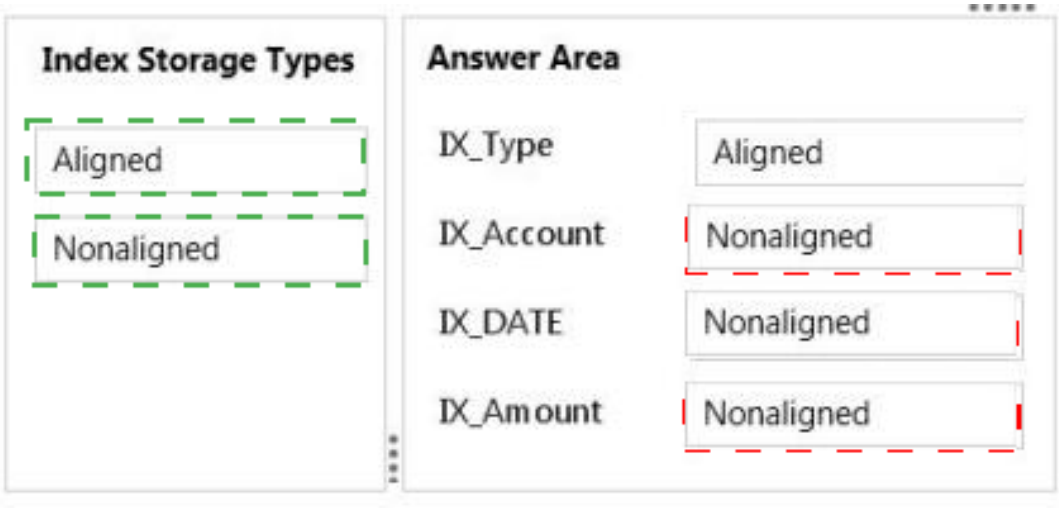

Explanation:

Note:
Index Storage Type
* Designing a partitioned index independently (unaligned) of the base table can be useful in the following cases:
/The base table has not been partitioned.
/The index key is unique and it does not contain the partitioning column of the table.
/ You want the base table to participate in collocated joins with more tables using different join columns.
Partitioned Tables and Indexes