ValidDumps

ValidDumps

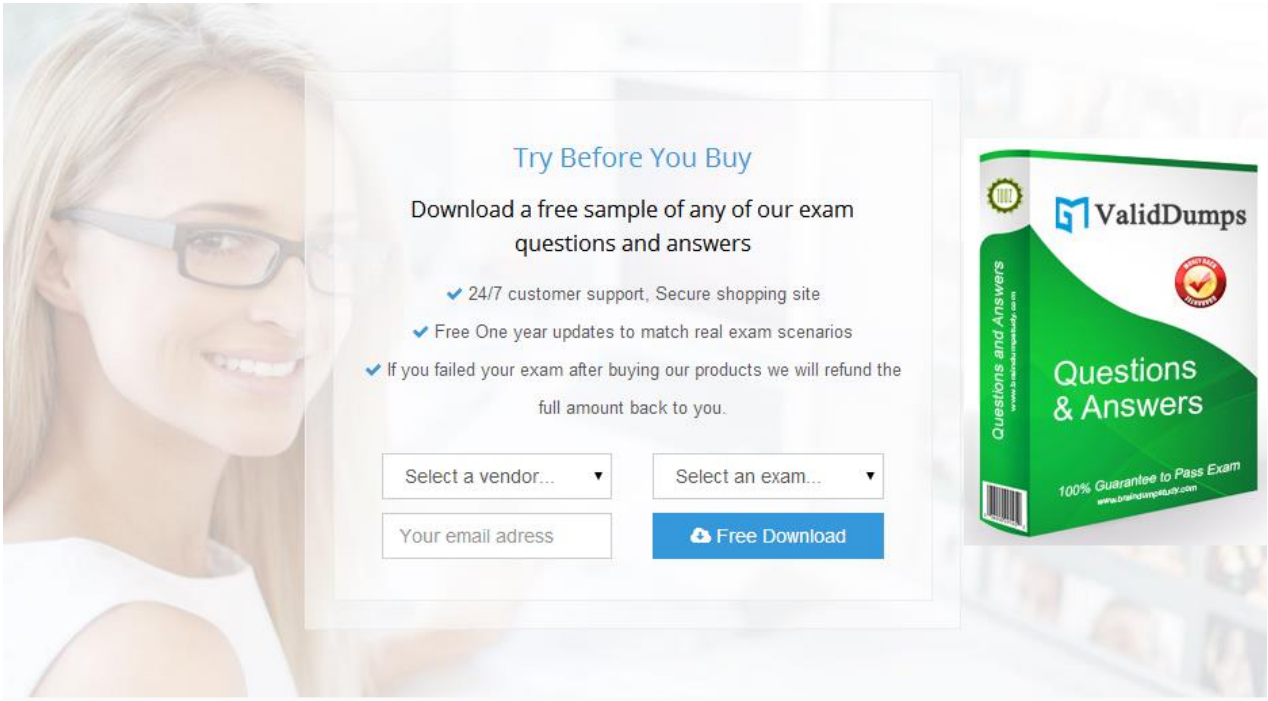

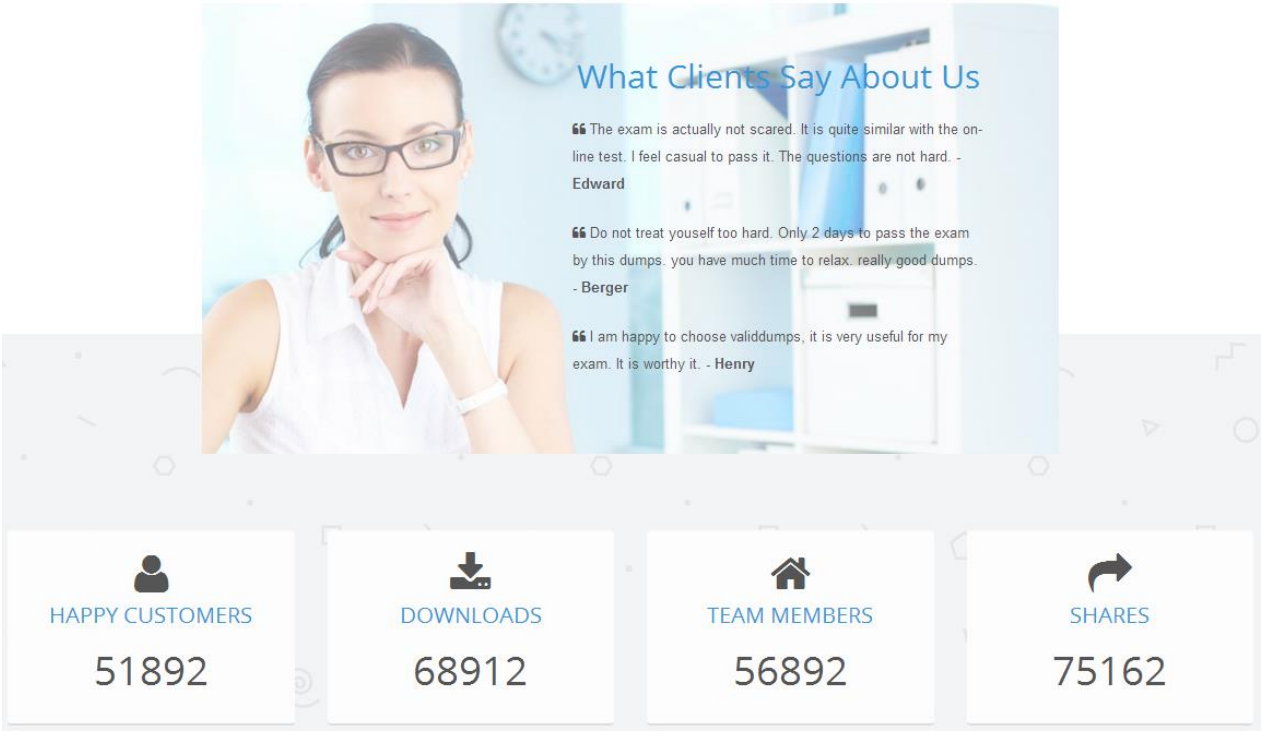

http://www.validdumps.top

The valid and reliable exam dumps are the guarantee for your success

Exam : **70-764**

Title : Administering a SQL Database Infrastructure

Vendor : Microsoft

Version : DEMO

NO.1 Note: This question is part of a series of questions that use the same or similar answer choices. An answer choice may be correct for more than one question in the scries. Each question is independent of the other questions in this scries. Information and details provided in a question apply only to that question.
You are examining information about users, sessions, and processes in an on-premises Microsoft SQL Server
2016 Standard Edition server.
You need to identify waits for resources and return only the following information:
*a list of all databases on the SQL Server instance, along with information about the database files, their paths, and names
*a list of the queries recently executed that use most of memory, disk, and network resources What should you use?
A. SQL Server Data Tools (SSDT)
B. Object Explorer in SQL Server Management Studio (SSMS)
C. Activity Monitor
D. sp_who3 stored procedure
E. SQL Server Data Collector
Answer: B

NO.2 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.

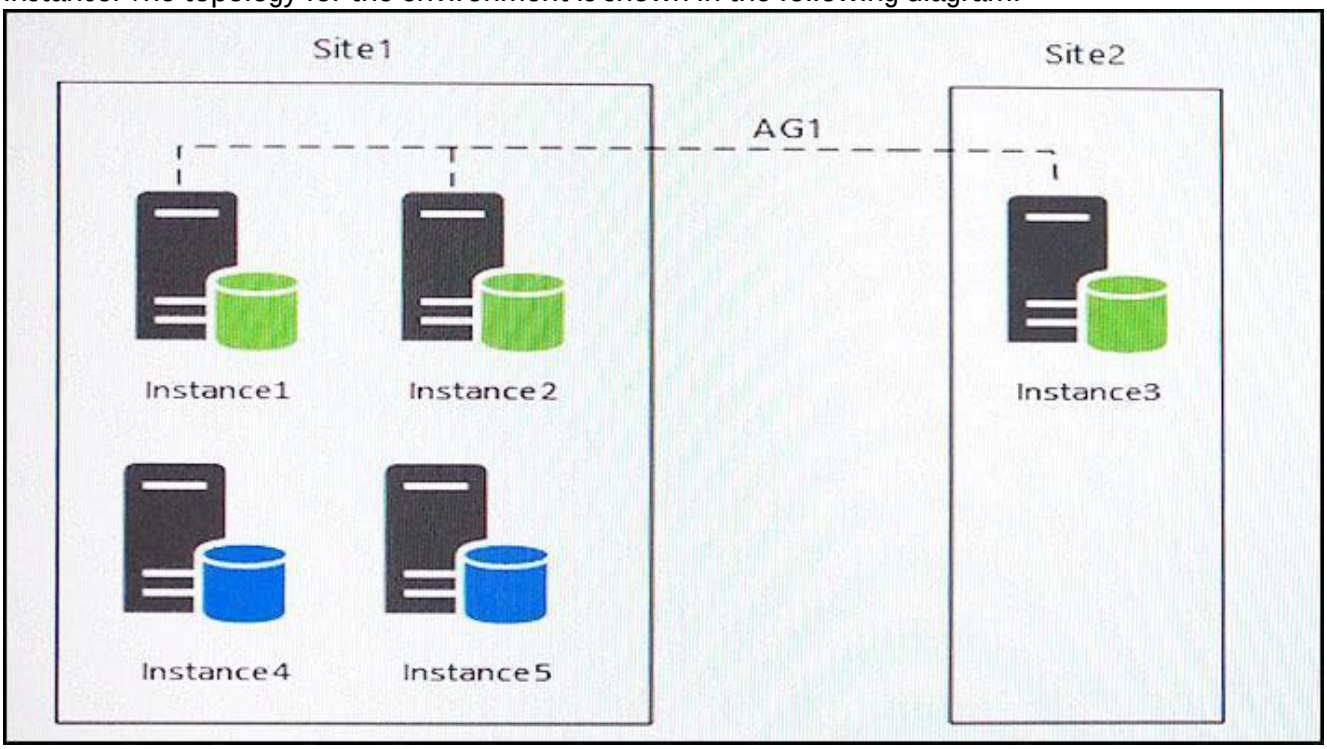

You have an Always On Availability group named AG1. The details for AG1 are shown in the following table.

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.
* Reporting system: This solution accesses data inDB1with a login that is mapped to a database user that is a member of the db_datareader role. The user has EXECUTE permissions on the database. Queries make no changes to the data. The queries must be load balanced over variable read-only replicas.
* Operations system: This solution accesses data inDB1with a login that is mapped to a database user that is a member of the db_datareader and db_datawriter roles. The user has EXECUTE permissions on the database. Queries from the operations system will perform both DDL and DML operations.
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 configure a new replica of AG1 on Instance6.
How should you complete the Transact-SQL statement? To answer, drag the appropriate Transact-SQL statements to the correct locations. Each Transact-SQL segment may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

Transact-SQL segments

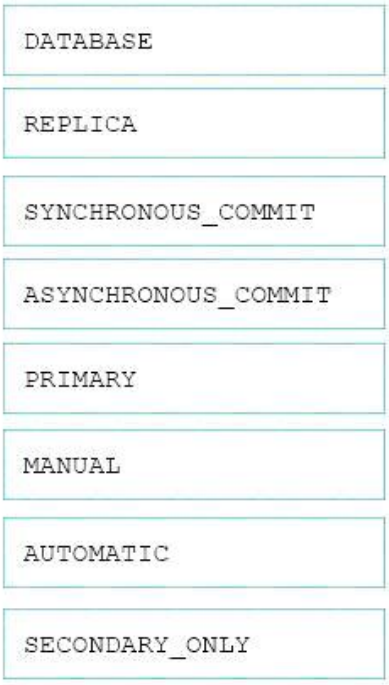

Answer Area

```
ALTER AVAILABILITY GROUP AG_1 MODIFY [ Transact-SQL segment ] ON 'INSTANCE6'

WITH (AVAILABILITY_MODE = [ Transact-SQL segment ] );

ALTER AVAILABILITY GROUP AG_1 MODIFY [ Transact-SQL segment ] ON 'INSTANCE6'

WITH (FAILOVER_MODE = [ Transact-SQL segment ] );
```

Answer:

Transact-SQL segments

DATABASE

REPLICA

SYNCHRONOUS_COMMIT

ASYNCHRONOUS_COMMIT

PRIMARY

MANUAL

AUTOMATIC

SECONDARY_ONLY

Answer Area

```
ALTER AVAILABILITY GROUP AG_1 MODIFY [ REPLICA ] ON 'INSTANCE6'
WITH (AVAILABILITY_MODE = [ SYNCHRONOUS_COMMIT ] );
ALTER AVAILABILITY GROUP AG_1 MODIFY [ REPLICA ] ON 'INSTANCE6'
WITH (FAILOVER_MODE = [ AUTOMATIC ] );
```

Explanation

Answer Area

```
ALTER AVAILABILITY GROUP AG_1 MODIFY [ REPLICA ] ON 'INSTANCE6'
WITH (AVAILABILITY_MODE = [ SYNCHRONOUS_COMMIT ] );
ALTER AVAILABILITY GROUP AG_1 MODIFY [ REPLICA ] ON 'INSTANCE6'
WITH (FAILOVER_MODE = [ AUTOMATIC ] );
```

Scenario: 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.
Box 1: REPLICA
MODIFY REPLICA ON modifies any of the replicas ofthe availability group.
Box 2: SYNCHRONOUS_COMMIT
You must minimize latency between the nodes in AG1
AVAILABILITY_MODE = { SYNCHRONOUS_COMMIT | ASYNCHRONOUS_COMMIT }
Specifies whether the primary replica has to wait for the secondary availability group to acknowledge the hardening (writing) of the log records to disk before the primary replica can commit the transaction on a given primary database.
FAILOVER AUTOMATIC (box 4) requires SYNCHRONOUS_COMMIT
Box 3: REPLICA
MODIFY REPLICA ON modifies any of the replicas of the availability group.
Box 4: AUTOMATIC
You must minimize latency between the nodes in AG1
FAILOVER_MODE = { AUTOMATIC | MANUAL }
Specifies the failover mode of the availability replica that you are defining.
FAILOVER_MODE is required in the ADD REPLICA ON clause and optional in the MODIFY REPLICA ON clause.
AUTOMATIC enables automatic failover. 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

NO.3 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 next of the scenario is exactly the same in each question in this series.
Start of repeated scenario
Contoso, Ltd. has Microsoft SQL Server databases that support a custom application. The current SQL Server environment consists of two servers: ContosoSQL1 and ContosoSQL2. These two servers participate in an Always On Availability Group named ContosoAG1 that is configured to use synchronous-commit with automatic failover. The secondary replica is not configured for read-only access.
The application performs both transactional processing and historical data retrieval in a database named ContosoDB. The application includes an inventory management module. The inventory management module and database have experienced performance issues.
Users report that a query named InventoryQuery1 takes a long time to complete. The query is shown as follows:
SELECT ProductNumber, Name, ProductLine
FROM Production.Product
WHERE ProductNumber = N'<specific product>'
The query plan used by SQL Server for this query is shown in the exhibit. (Click the Exhibit tab.)
Various performance issues, including frequent long-term blocking episodes, prevent business users from completing their daily tasks. You suspect the tempdb database resources could be responsible.
You must create Blocking reports for the ContosoDB database to identify issues.
Exhibit.

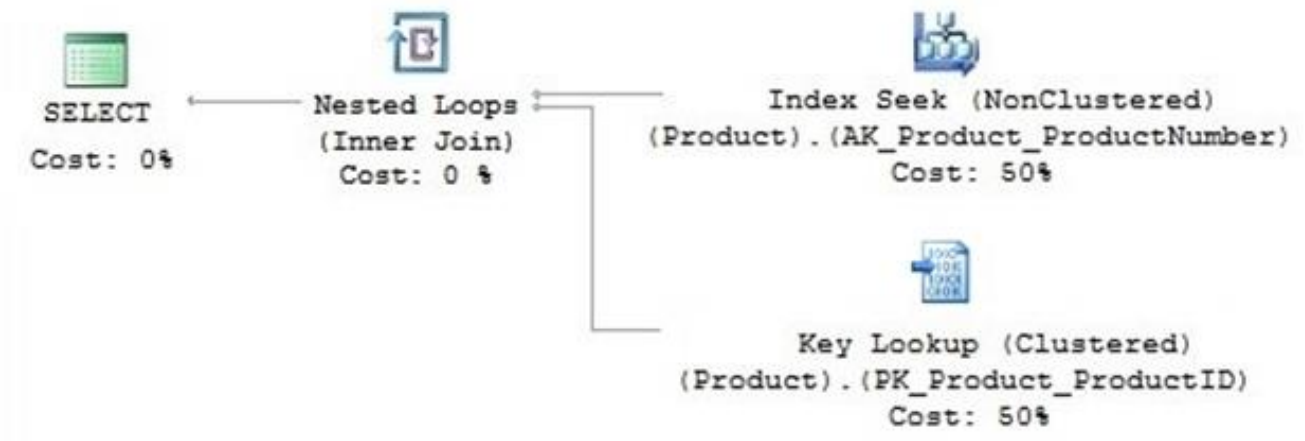

You plan to use Extended Events to review all Transact-SQL statements that are run against the ContosoSQL1 instance. The output from the Extended Events session must contain both start and stop events and must be written to a file. You must configure the Extended Events session to minimize possible data loss and reduce the effect on server performance.
You plan to deploy an additional secondary replica named ContosoSQL3 to ContosoAG1. Read-only traffic must be load-balanced between the two secondary replicas, regardless of which instance is the primary replica.
Contoso plans to add an additional dedicated reporting system that will rely on real-time data from the transactional databases.
The company plans to improve their high availability/disaster recovery (HA/DR) solution. As part of the planned improvements, you will back up all databases from ContosoSQL1 directly to an off-site location.
End of repeated scenario
You need to configure the backup process for ContosoSQL1.
What should you do?

A. Set the recovery model to Simple

B. Perform a backup to a tape device.

C. Perform mirrored backups to a DR datacenter.

D. Create a new backup set.

Answer: C

Explanation
Scenario:
The company plans to improve their high availability/disaster recovery (HA/DR) solution. As part of the planned improvements, you will back up all databases from ContosoSQL1 directly to an off-site location.
One of the features found in the Enterprise Edition of SQL Server is the ability to take mirrored backups.
Basically, taking a mirrored backup means creating additional copies of the backup media (up to three) using a single BACKUP command, eliminating the need to perform the copies with copy or robocopy.
The idea behind is that you can backup to multiple locations and increase the protection level by having additional copies of the backup set. In case one of the copies gets lost or corrupted, you can use the mirrored copy to perform a restore.
Another possible scenario for a mirrored backup is deferred tape migration: you can backup to a local

disk and mirror to a shared folder on a file server. That way you could have a local copy of the backup set and restore it in case of need and let the mirrored copy migrate to tape when the disk backup software processes the file server's disks.
References:
https://www.mssqltips.com/sqlservertip/1779/mirrored-database-backup-feature-in-sql-server-2005-and-sql-serv

NO.4 You need to recommend the actions that are required to partition a table.
In which order should the four actions be performed? To answer, move the actions from the list of actions to the answer area and arrange them in the correct order.

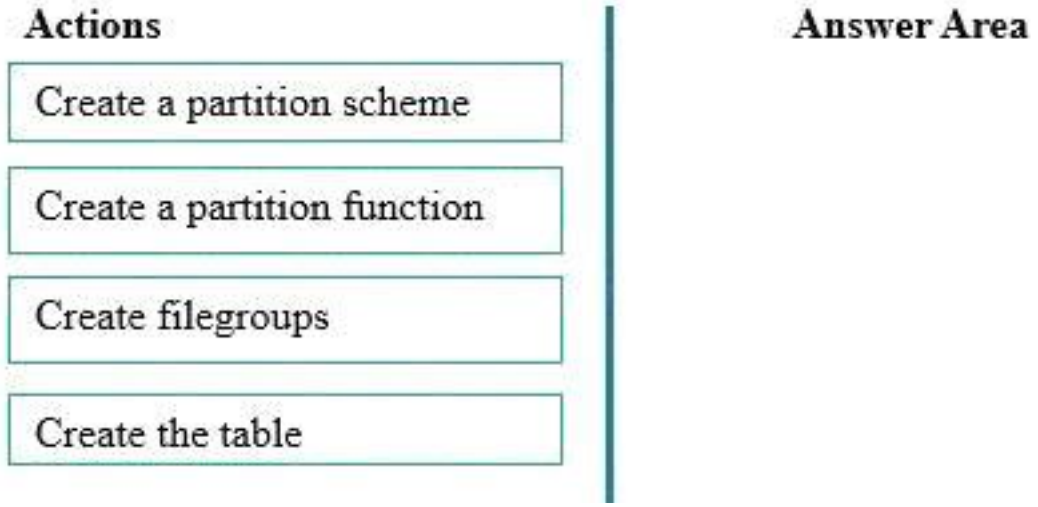

Answer:

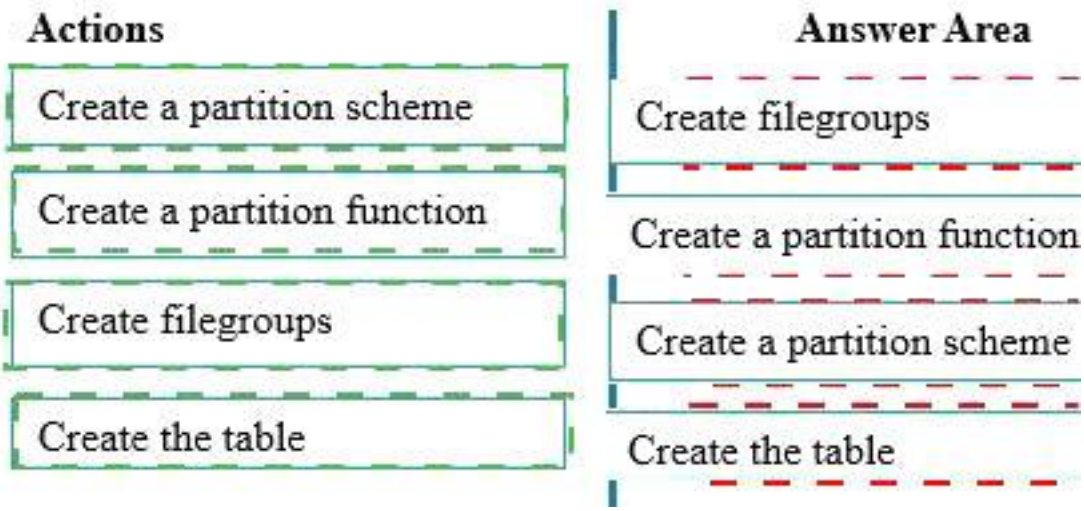

Explanation

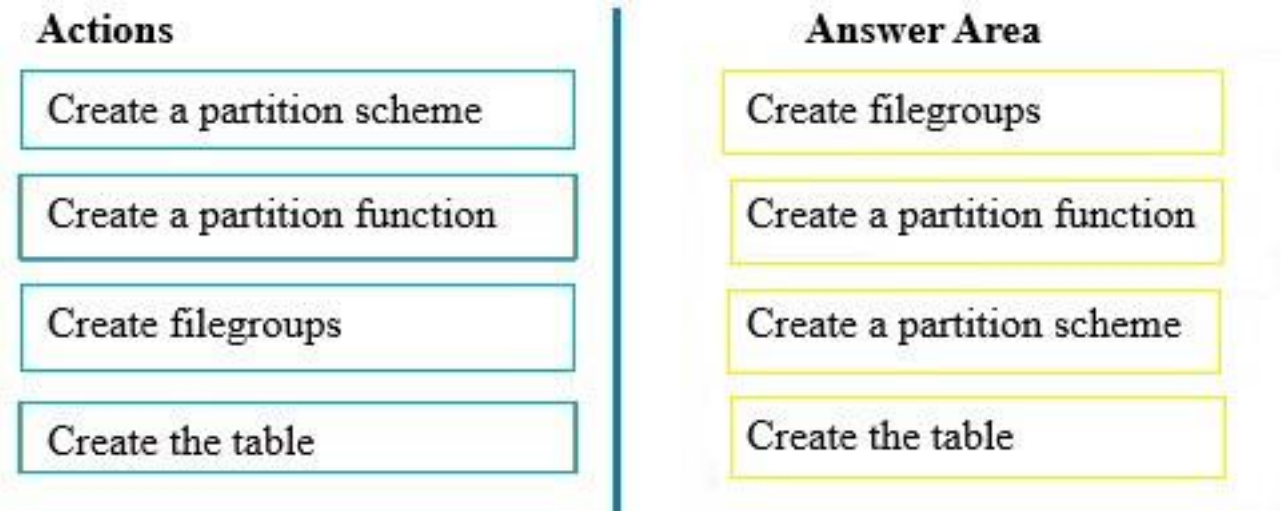

References:
http://technet.microsoft.com/en-us/library/ms188730.aspx

NO.5 You have a SQL Server 2014 environment That contains 20 servers.
The corporate security policy states that all SQL Server 2014 instances must meet specific security standards.
You need to recommend a management strategy for the SQL Server 2014 servers.
What should you include in the recommendation? More than one answer choice may achieve the goal. Select the BEST answer.

A. Maintenance plans

B. Policy-Based Management

C. Common criteria compliance

D. Multi server jobs

Answer: B

Explanation
Policy-Based Management is a system for managing one or more instances of SQL Server. When SQL Server policy administrators use Policy-Based Management, they use SQL Server Management Studio to create policies to manage entities on the server, such as the instance of SQL Server, databases, or other SQL Server objects.

NO.6 You are implementing log shipping. You configure a secondary server in standby mode. You configure log shipping to occur hourly and keep up to seven days of log backups.
Users create reports by querying the secondary server. Transaction logs are not applied to the secondary server.
You need to ensure that the secondary server has current data at the beginning of each business day.
What are two possible ways to achieve the goal? Each answer represents a complete solution.

A. Configure the secondary server as a database mirroring target server.

B. Create a snapshot database from a daily backup to use as the secondary database. Revert the snapshot database back to the previous business day.

C. Disconnect users while transaction logs are applied.

D. Allow restore operations to accumulate. Run a Microsoft SQL Agent job that disconnects all users that are connected to the secondary server.

Answer: A,D

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 section, you will NOT be able to return to it. As a result these questions will not appear in the review screen.
You administer a SQL Server instance at Contoso Ltd.
A new employee named Kim Ralls must be able to read data only from the dbo.order table in the salesDB database.
You create a SQL login named KimRalls. You then create a SQL user account named KimRalls in salesDB.
The user account is associated with the SQL login.
You need to grant KimRalls the necessary permission or permissions. Your solution must follow the principle of least privilege.
Solution: You perform the following actions:
1. Add KimRalls to the db_denydatawriter database role in SalesDB:
2. Add KimRalls to the db owner database role in SalesDB.
Does the solution meet the goal?
A. No
B. Yes
Answer: A
Explanation
Need to set up select permission on dbo.Order table.

NO.8 Overview
You are a database administrator for a company named Litware, Inc.
Litware is a book publishing house. Litware has a main office and a branch office.
You are designing the database infrastructure to support a new web-based application that is being developed.
The web application will be accessed at www.litwareinc.com. Both internal employees and external partners will use the application.
You have an existing desktop application that uses a SQL Server 2008 database named App1_DB.
App1_DB will remain in production.
Requirements
Planned Changes
You plan to deploy a SQL Server 2014 instance that will contain two databases named Database1 and Database2.
All database files will be stored in a highly available SAN.
Database1 will contain two tables named Orders and OrderDetails.
Database1 will also contain a stored procedure named usp_UpdateOrderDetails.
The stored procedure is used to update order information. The stored procedure queries the Orders table twice each time the procedure executes.
The rows returned from the first query must be returned on the second query unchanged along with any rows added to the table between the two read operations.
Database1 will contain several queries that access data in the Database2 tables.
Database2 will contain a table named Inventory.

Inventory will contain over 100 GB of data.
The Inventory table will have two indexes: a clustered index on the primary key and a nonclustered index.
The column that is used as the primary key will use the identity property.
Database2 wilt contains a stored procedure named usp_UpdateInventory. usp_UpdateInventory will manipulate a table that contains a self-join that has an unlimited number of hierarchies. All data in Database2 is recreated each day ad does not change until the next data creation process. Data from Database2 will be accessed periodically by an external application named Application1. The data from Database2 will be sent to a database named Appl_Dbl as soon as changes occur to the data in Database2. Litware plans to use offsite storage for all SQL Server 2014 backups.
Business Requirements
You have the following requirements:
* Costs for new licenses must be minimized.
* Private information that is accessed by Application must be stored in a secure format.
* Development effort must be minimized whenever possible.
* The storage requirements for databases must be minimized.
* System administrators must be able to run real-time reports on disk usage.
* The databases must be available if the SQL Server service fails.
* Database administrators must receive a detailed report that contains allocation errors and data corruption.
* Application developers must be denied direct access to the database tables. Applications must be denied direct access to the tables.
You must encrypt the backup files to meet regulatory compliance requirements.
The encryption strategy must minimize changes to the databases and to the applications.
You need to recommend a solution to allow application users to perform UPDATE operations on the database tables. The solution must meet the business requirements.
What should you recommend?

A. Create a user-defined database role and add users to the role.

B. Create a Policy-Based Management Policy.

C. Create functions that use EXECUTE AS clauses.

D. Create stored procedures that use EXECUTE AS clauses.

Answer: D
Explanation
- EXECUTE AS Clause (Transact-SQL)
In SQL Server you can define the execution context of the following user-defined modules: functions (except inline table-valued functions), procedures, queues, and triggers.

NO.9 You are a database administrator for a Microsoft SQL Server 2016 environment.
You want to deploy a new application that will scale out the workload to at least five different SQL Server instances.
You need to ensure that for each copy of the database, users are able to read and write data that will then be synchronized between all of the database instances.
Which feature should you use?

A. Log Shipping

B. Availability Groups

C. Peer-to-Peer Replication
D. Database Mirroring
Answer: C

NO.10 You have a database named DB1 that is configured to use the full recovery model. You have a full daily backup job that runs at 02:00. The job backs up data from DB1 to the file B:\DB1.bak.
You need to restore the DB1 database to the point in time of May 25, 2016 at 02:23 and ensure that the database is functional and starts to accept connections.
Which Transact-SQL statement should you run?

A.
```
BACKUP LOG [DB1] TO DISK = N'B:\DB1Log.bak' WITH RECOVERY
RESTORE DATABASE [DB1] FROM DISK = N'B:\DB1.bak' WITH NORECOVERY
RESTORE LOG [DB1] FROM DISK = N'B:\DB1Log.bak' WITH STOPAT = N'2016-05-25T02:23:00'
```

B.
```
BACKUP LOG [DB1] TO DISK = N'B:\DB1Log.bak' WITH NORECOVERY
RESTORE DATABASE [DB1] FROM DISK = N'B:\DB1.bak' WITH NORECOVERY
RESTORE LOG [DB1] FROM DISK = N'B:\DB1Log.bak' WITH STOPAT = N'2016-05-25T02:23:00'
```

C.
```
BACKUP LOG [DB1] TO DISK = N'B:\DB1Log.bak' WITH NORECOVERY
RESTORE DATABASE [DB1] FROM DISK = N'B:\DB1.bak' WITH NORECOVERY
RESTORE LOG [DB1] FROM DISK = N'B:\DB1Log.bak' WITH STOPAT = N'2016-05-25T02:23:00', NORECOVERY
```

D.
```
RESTORE DATABASE [DB1] FROM DISK = N'B:\DB1.bak' WITH STOPAT = N'2016-05-25T02:23:00', RECOVERY
```

A. Option B
B. Option D
C. Option C
D. Option A
Answer: A

NO.11 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 are a database administrator for a company that has an on-premises Microsoft SQL Server environment and Microsoft Azure SQL Database instances. The environment hosts several customer databases, and each customer uses a dedicated instance. The environments that you manage are shown in the following table.

Customer	Cloud Type	Description
AdventureWorks Cycles	Private	The environment includes a database named **Adventureworks** that contains a single schema named ADVSchema. You must implement auditing for all objects in the ADVSchema schema. You must also implement auditing to record access to data that is considered sensitive by the company.
Tailspin Toys	Private	Tailspin Toys has a custom application that accesses a hosted database named **TSpinDB**. The application will monitor **TSpinDB** and capture information over time about which database objects are accessed and how frequently they are accessed.
Contoso, Ltd.	Private	The environment has a database named **ConDB** that was recently upgraded to Microsoft SQL Server 2016. Contoso reports that **ConDB** is slow to return results when the server is busy. You must modify the startup parameters to **ConDB** to optimize performance.
Wingtip Toys	Private	Wingtip Toys has a database named **WingDB**. All tables in the database have indexes. Users report system response time is slow during peak activity periods. You observe that the performance issues are related to locking. Wingtip Toys receives data updates from suppliers each week. You must implement a process for importing the data into **WingDB**. You must use minimal logging and minimized data loss during import process.
Wide World Importers	Public	The environment includes a database named **WDWDB**. Neither auditing nor statistics are configured for **WDWDB**. You must log any deletion of views and all database record update operations.

You need to configure auditing for WDWDB.
In the table below, identify the event type that you must audit for each activity.

Answer Area

Event type	View deletions	Update operations
Data changes	○	○
Schema changes	○	○
SQL batch	○	○
Data access	○	○

Answer:

Answer Area

Event type	View deletions	Update operations
Data changes	○	○
Schema changes	○	○
SQL batch	○	○
Data access	○	○

Explanation

Answer Area

Event type	View deletions	Update operations
Data changes	○	○
Schema changes	○	○
SQL batch	○	○
Data access	○	○

NO.12 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.
Start of repeated scenario.
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.

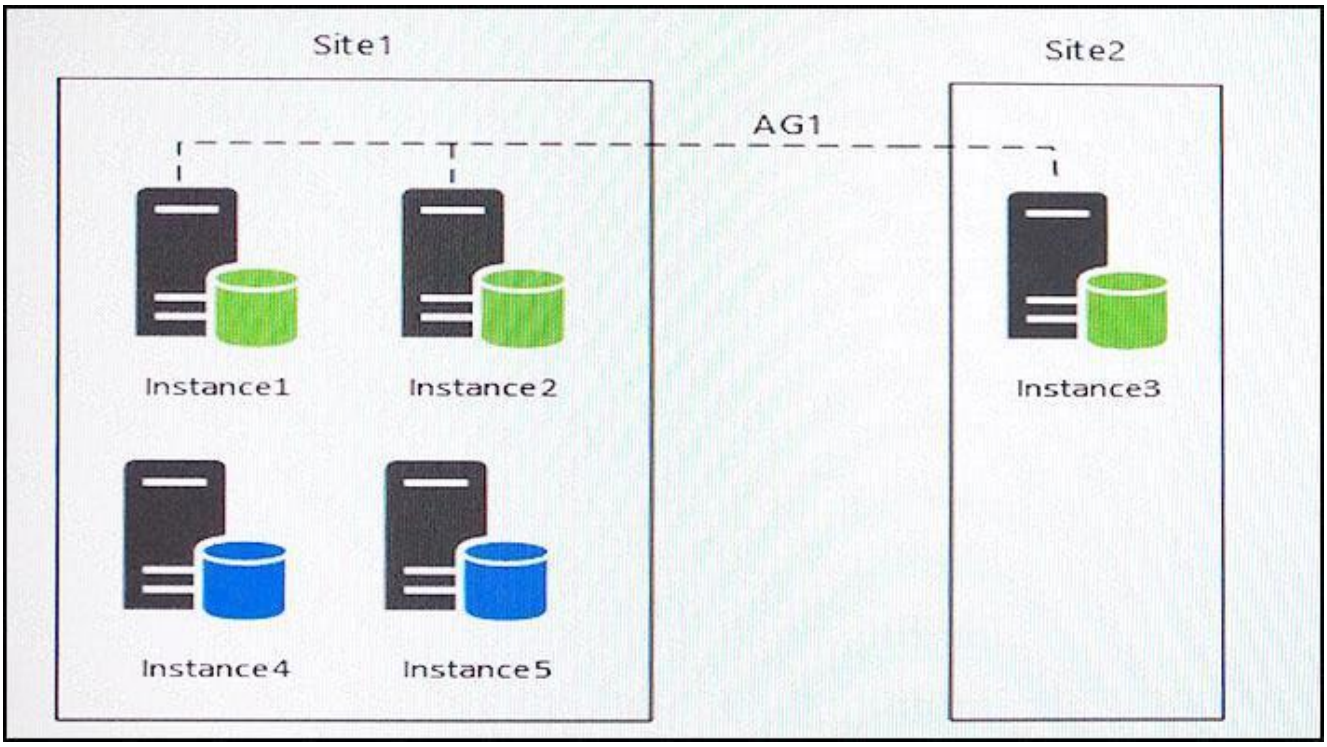

You have an Always On Availability group named AG1. The details for AG1 are shown in the following table.

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.
* Reporting system: This solution accesses data inDB1with a login that is mapped to a database user that is a member of the db_datareader role. The user has EXECUTE permissions on the database. Queries make no changes to the data. The queries must be load balanced over variable read-only replicas.
* Operations system: This solution accesses data inDB1with a login that is mapped to a database user that is a member of the db_datareader and db_datawriter roles. The user has EXECUTE permissions on the database. Queries from the operations system will perform both DDL and DML operations.
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.

End of repeated scenario.
You need to create a backup plan for Instance4.
Which backup plan should you create?
A. Weekly full backups, nightly differential backups, transaction log backups every 12 hours.
B. Full backups every 60 minutes, transaction log backups every 30 minutes.
C. Weekly full backups, nightly differential. No transaction log backups are necessary.
D. Weekly full backups, nightly differential backups, transaction log backups every 30 minutes.
Answer: D
Explanation
Scenario: Instance4 is engaged in heavy read-write I/O. The Recovery Point Objective of Instance4 is 60 minutes.

NO.13 You have multiple Microsoft SQL Server databases that are deployed in an Always On availability group. You configure the SQL Server Agent service to start automatically.
You need to automate backups for all user databases.
What should you create?
A. SQL Agent job
B. SQL Agent operator

C. SQL Server message

D. SQL script

Answer: A

Explanation

To schedule backups using a SQL Server Agent job

To automate and schedule a backup with SQL Server Agent:

* In the Object Explorer panel, under the SQL Server Agent node, right click Jobs and select New job from the context menu
* In the New Job dialog enter a job's name
* Under the Steps tab click on the New button and create a backup step by inserting a T-SQL statement. In this case the CHECKSUM clause has to be included in T-SQL code.
* Click ok to add a step, and click OK to create a job
* To schedule a job, in the New Job dialog, under the Schedule tab click New.
* In the Job Schedule select an occurring frequency, duration and a start date and click OK:
* To check a created job in the Object Explorer pane and under the SQL Server Agent Jobs node right click the job create above and select the Start job at step option References:

https://sqlbackupandftp.com/blog/how-to-automate-sql-server-database-backups

NO.14 You manage a database named DB1 that uses the following filegroups:

Filegroup name	Size	Description	Backup logical device
PRIMARY	100 MB	Primary filegroup. Contains no business-related data. All files are stored on the D drive.	Backup1
FG1	10 GB	Filegroup has read/write access. Contains mission-critical data. All files are stored on the E drive.	Backup2
FG2	5 TB	Filegroup has read-only access. Contains historical data but no critical reporting data. All files are stored on the F drive. You create a backup after you configure the filegroup for read-only access.	Backup3
FG3	100 GB	Filegroup has read-only access. Contains mission-critical data. All files are stored on the E drive.	Backup2

The database is configured to use full recovery model. Transaction logs are backed up to a backup set named TLogBackup.

The PRIMARY and FG2 for DB1

You need to design a piecemeal restore plan that meets all the above requirements. You need to bring critical filegroups online as soon as possible while minimizing restoration time. All damaged filegroups must be online after the restore operation completes.

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

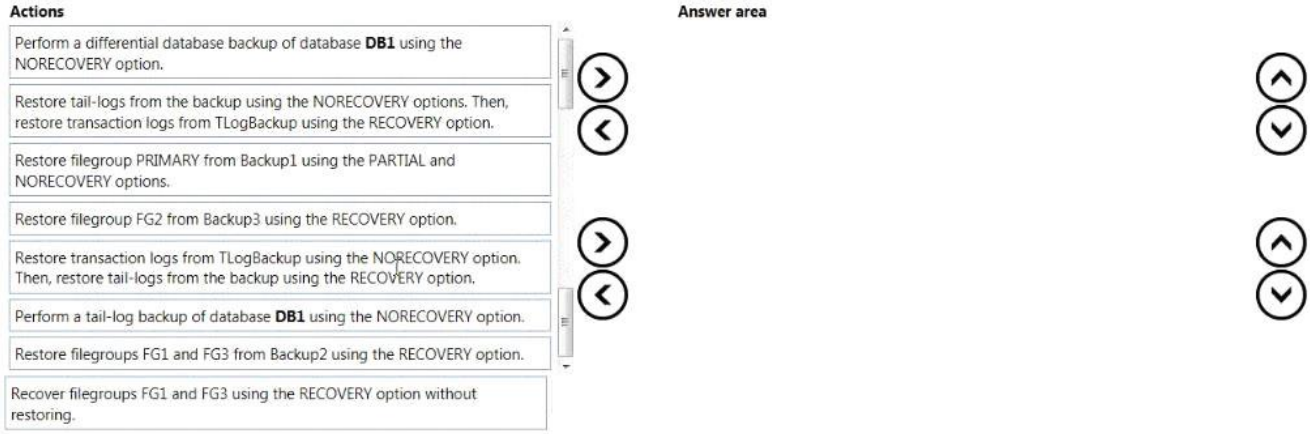

Answer:

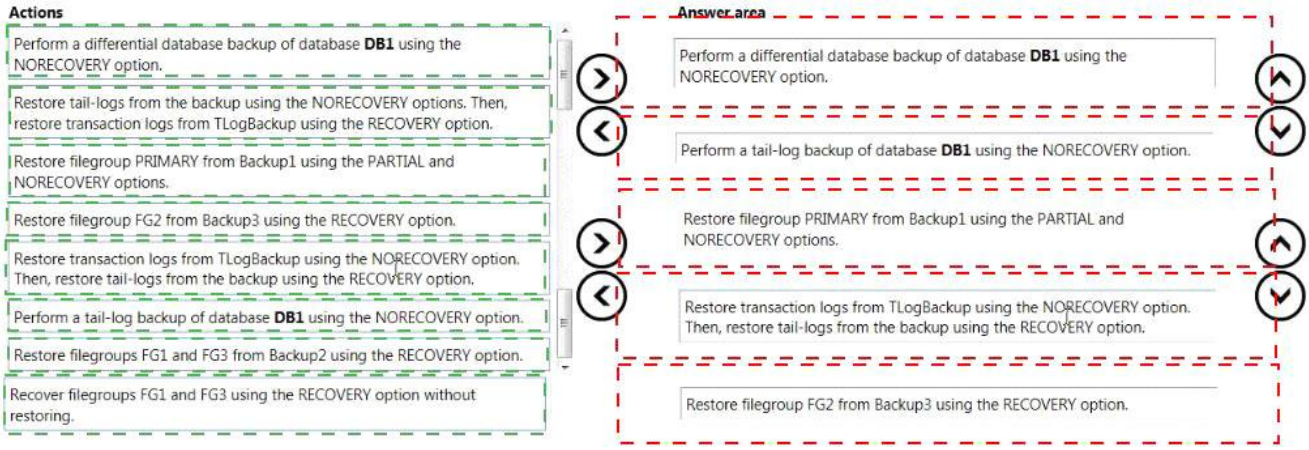

Explanation

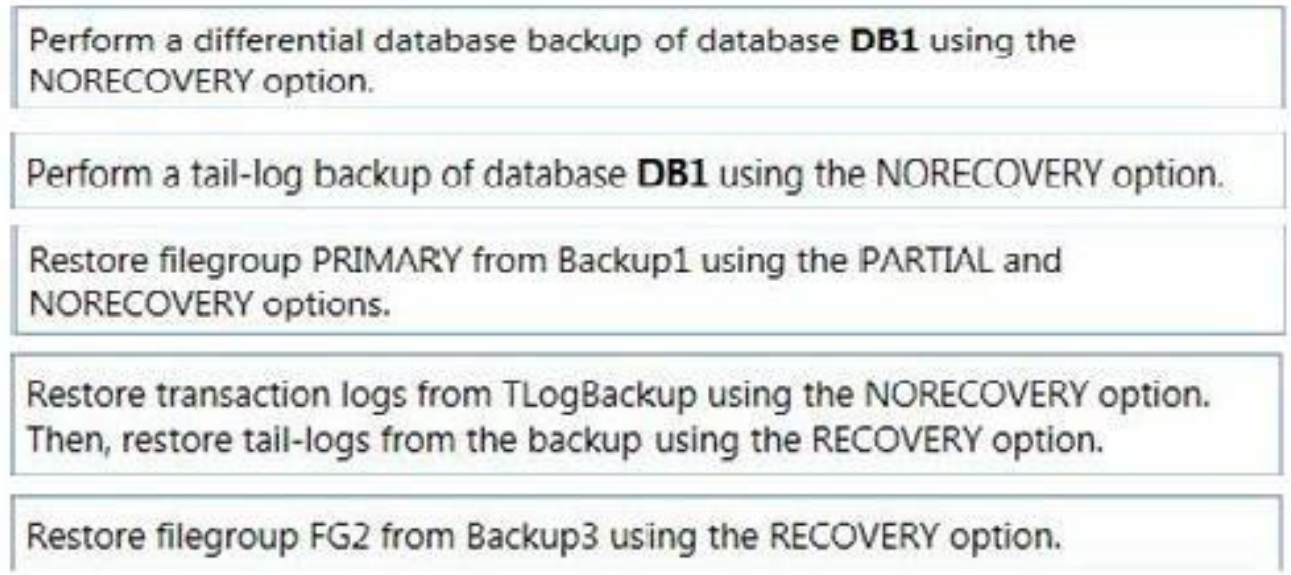

Step 1: Perform a differential database backup...
Step 2: Perform a tail-log backup...
A tail-log backup captures any log records that have not yet been backed up (the tail of the log) to prevent work loss and to keep the log chain intact. Before you can recover a SQL Server database to its latest point in time, you must back up the tail of its transaction log. The tail-log backup will be the last backup of interest in the recovery plan for the database.
Step 3:
The PRIMARY and FG2 for DB1 are damaged. FG1 and FG3 are intact.
Step 4:
Transaction logs are backed up to a backup set named TLogBackup.
Step 5:
The PRIMARY and FG2 for DB1 are damaged.
References:
https://docs.microsoft.com/en-us/sql/relational-databases/backup-restore/restore-files-and-filegroups-sql-server?v
https://docs.microsoft.com/en-us/sql/relational-databases/backup-restore/tail-log-backups-sql-server?view=sql-se

NO.15 You administer a Microsoft SQL Server 2016 database instance.
You create a new user named UserA. You need to ensure that UserA is able to create SQL Server Agent jobs and to execute SQL Server Agent jobs.
To which role should you add UserA?
A. RSExecRole
B. Securityadmin
C. DatabaseMailUserRole
D. SQLAgentUserRole
Answer: D

NO.16 Overview
Application Overview
Contoso, Ltd., is the developer of an enterprise resource planning (ERP) application.
Contoso is designing a new version of the ERP application. The previous version of the ERP application used SQL Server 2008 R2.
The new version will use SQL Server 2014.
The ERP application relies on an import process to load supplier data. The import process updates thousands of rows simultaneously, requires exclusive access to the database, and runs daily.
You receive several support calls reporting unexpected behavior in the ERP application. After analyzing the calls, you conclude that users made changes directly to the tables in the database.
Tables
The current database schema contains a table named OrderDetails.
The OrderDetails table contains information about the items sold for each purchase order.
OrderDetails stores the product ID, quantities, and discounts applied to each product in a purchase order.
The product price is stored in a table named Products. The Products table was defined by using the SQL_Latin1_General_CP1_CI_AS collation.
A column named ProductName was created by using the varchar data type. The database contains a table named Orders.
Orders contains all of the purchase orders from the last 12 months. Purchase orders that are older than 12 months are stored in a table named OrdersOld.
The previous version of the ERP application relied on table-level security.
Stored Procedures
The current version of the database contains stored procedures that change two tables. The following shows the relevant portions of the two stored procedures:

```
CREATE PROC Sales.Proc1
AS
BEGIN TRAN
UPDATE Sales.Table1 …
UPDATE Sales.Table2 …
COMMIT TRAN
GO

CREATE PROC Sales.Proc2
AS
BEGIN TRAN
UPDATE Sales.Table2 …
UPDATE Sales.Table1 …
COMMIT TRAN
GO
```

Customer Problems

Installation Issues

Column	Data type
id	uniquedentifier
lastModified	datetime
modifiedBy	Varchar(200)

The current version of the ERP application requires that several SQL Server logins be set up to function correctly. Most customers set up the ERP application in multiple locations and must create logins multiple times.

Index Fragmentation Issues

Customers discover that clustered indexes often are fragmented. To resolve this issue, the customers defragment the indexes more frequently. All of the tables affected by fragmentation have the following columns that are used as the clustered index key:

Backup Issues

Customers who have large amounts of historical purchase order data report that backup time is unacceptable.

Search Issues

Users report that when they search product names, the search results exclude product names that contain accents, unless the search string includes the accent.

Missing Data Issues

Customers report that when they make a price change in the Products table, they cannot retrieve the price that the item was sold for in previous orders.

Query Performance Issues

Customers report that query performance degrades very quickly. Additionally, the customers report that users cannot run queries when SQL Server runs maintenance tasks. Import Issues During the

monthly import process, database administrators receive many supports call from users who report that they cannot access the supplier data. The database administrators want to reduce the amount of time required to import the data.
Design Requirements
File Storage Requirements
The ERP database stores scanned documents that are larger than 2 MB. These files must only be accessed through the ERP application. File access must have the best possible read and write performance.
Data Recovery Requirements
If the import process fails, the database must be returned to its prior state immediately.
Security Requirements
You must provide users with the ability to execute functions within the ERP application, without having direct access to the underlying tables.
Concurrency Requirements
You must reduce the likelihood of deadlocks occurring when Sales.Prod and Sales.Proc2 execute.
You need to recommend a solution that addresses the index fragmentation and index width issue.
What should you include in the recommendation? (Each correct answer presents part of the solution. Choose all that apply.)

A. Change the data type of the lastModified column to smalldatetime.

B. Remove the modifiedBy column from the clustered index.

C. Change the data type of the modifiedBy column to tinyint.

D. Remove the id column from the clustered index.

E. Remove the lastModified column from the clustered index.

F. Change the data type of the id column to bigint.

Answer: B,E
Explanation
Scenario: Index Fragmentation Issues Customers discover that clustered indexes often are fragmented. To resolve this issue, the customers defragment the indexes more frequently. All of the tables affected by fragmentation have the following columns that are used as the clustered index key:

Column	Data type
id	uniquedentifier
lastModified	datetime
modifiedBy	Varchar(200)

NO.17 Overview
General Overview
ADatum Corporation has offices in Miami and Montreal.
The network contains a single Active Directory forest named adatum.com. The offices connect to each other by using a WAN link that has 5-ms latency.
A: Datum standardizes its database platform by using SQL Server
2014 Enterprise edition.
Databases

Each office contains databases named Sales, Inventory, Customers, Products, Personnel, and Dev.
Servers and databases are managed by a team of database administrators. Currently, all of the database administrators have the same level of permissions on all of the servers and all of the databases.
The Customers database contains two tables named Customers and Classifications.
The following graphic shows the relevant portions of the tables:

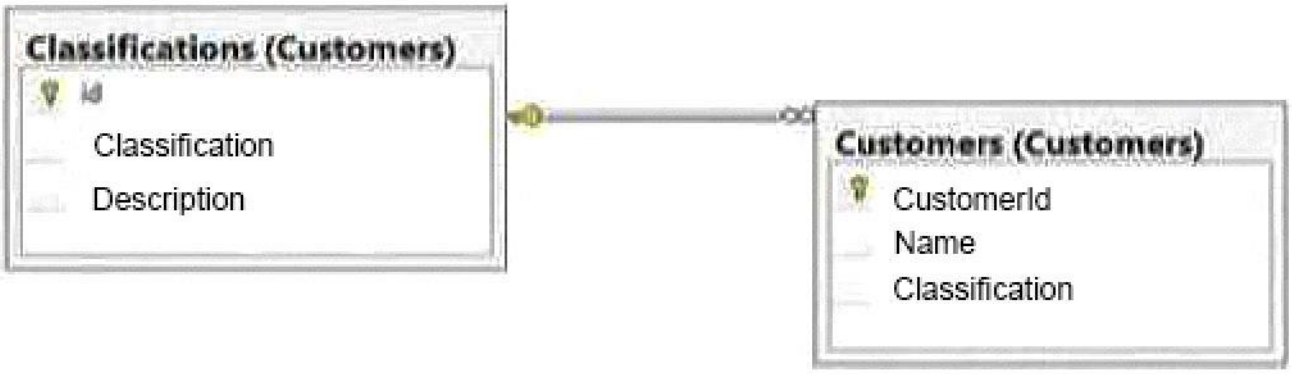

The following table shows the current data in the Classifications table:

ID	Classification	Description
1	Platinum	Yearly sales over 1,000,000
2	Gold	Yearly sales over 500,000
3	Silver	Yearly sales over 100,000

The Inventory database is updated frequently.
The database is often used for reporting.
A full backup of the database currently takes three hours to complete.
Stored Procedures
A stored procedure named USP_1 generates millions of rows of data for multiple reports. USP_1 combines data from five different tables from the Sales and Customers databases in a table named Table1.
After Table1 is created, the reporting process reads data from Table1 sequentially several times.
After the process is complete, Table1 is deleted.
A stored procedure named USP_2 is used to generate a product list. The product list contains the names of products grouped by category.
USP_2 takes several minutes to run due to locks on the tables the procedure accesses. The locks are caused by USP_1 and USP_3.
A stored procedure named USP_3 is used to update prices. USP_3 is composed of several UPDATE statements called in sequence from within a transaction.
Currently, if one of the UPDATE statements fails, the stored procedure fails. A stored procedure named USP_4 calls stored procedures in the Sales, Customers, and Inventory databases.
The nested stored procedures read tables from the Sales, Customers, and Inventory databases.
USP_4 uses an EXECUTE AS clause.
All nested stored procedures handle errors by using structured exception handling. A stored procedure named USP_5 calls several stored procedures in the same database. Security checks are performed each time USP_5 calls a stored procedure.

You suspect that the security checks are slowing down the performance of USP_5. All stored procedures accessed by user applications call nested stored procedures.
The nested stored procedures are never called directly.
Design Requirements
Data Recovery
You must be able to recover data from the Inventory database if a storage failure occurs. You have a Recovery Time Objective (RTO) of 5 minutes.
You must be able to recover data from the Dev database if data is lost accidentally. You have a Recovery Point Objective (RPO) of one day.
Classification Changes
You plan to change the way customers are classified. The new classifications will have four levels based on the number of orders. Classifications may be removed or added in the future. Management requests that historical data be maintained for the previous classifications. Security A group of junior database administrators must be able to manage security for the Sales database. The junior database administrators will not have any other administrative rights.
A: Datum wants to track which users run each stored procedure.
Storage
ADatum has limited storage. Whenever possible, all storage space should be minimized for all databases and all backups.
Error Handling
There is currently no error handling code in any stored procedure.
You plan to log errors in called stored procedures and nested stored procedures. Nested stored procedures are never called directly.
You need to recommend a solution to minimize the amount of time it takes to execute USP_1. With what should you recommend replacing Table1?

A. A table variable

B. A temporary table

C. A function

D. An indexed view

Answer: B

Explanation
Scenario:
A stored procedure named USP_1 generates millions of rows of data for multiple reports. USP_1 combines data from five different tables from the Sales and Customers databases in a table named Table1.
After Table1 is created, the reporting process reads data from a table in the Products database and searches for information in Table1 based on input from the Products table. After the process is complete, Table1 is deleted.

NO.18 Note: This question is part of a series of questions that use the same or similar answer choices. An answer choice may be correct for more than one question in the series. Each question is independent of the other questions in this series. Information and details provided in a question apply only to that question.
You collect performance metrics on multiple Microsoft SQL Server instances and store the data in a single repository.
You need to examine disk usage, query statistics, and server activity without building custom

counters.
What should you use?
A. SQL Server Data Collector
B. SQL Server Configuration Manager
C. Sp_who3 stored procedure
D. Object Explorer in the Microsoft SQL Server Management Studio (SSMS)
E. Activity Monitor
F. SQL Server Data Tools (SSDT)
Answer: A
Explanation
The data collector is a core component of the data collection platform for SQL Server 2017 and the tools that are provided by SQL Server. The data collector provides one central point for data collection across your database servers and applications. This collection point can obtain data from a variety of sources and is not limited to performance data

NO.19 Background
Corporate Information
Fabrikam, Inc. is a retailer that sells electronics products on the Internet. The company has a headquarters site and one satellite sales office. You have been hired as the database administrator, and the company wants you to change the architecture of the Fabrikam ecommerce site to optimize performance and reduce downtime while keeping capital expenditures to a minimum. To help with the solution, Fabrikam has decided to use cloud resources as well as on-premise servers.
Physical Locations
All of the corporate executives, product managers, and support staff are stationed at the headquarters office.
Half of the sales force works at this location. There is also a satellite sales office. The other half of the sales force works at the satellite office in order to have sales people closer to clients in that area. Only sales people work at the satellite location.
Problem Statement
To be successful, Fabrikam needs a website that is fast and has a high degree of system uptime. The current system operates on a single server and the company is not happy with the single point of failure this presents.
The current nightly backups have been failing due to insufficient space on the available drives and manual drive cleanup often needing to happen to get past the errors. Additional space will not be made available for backups on the HQ or satellite servers. During your investigation, you discover that the sales force reports are causing significant contention.
Configuration
Windows Logins
The network administrators have set up Windows groups to make it easier to manage security. Users may belong to more than one group depending on their role. The groups have been set up as shown in the following table:

Group	Members
OurDomain\Management	All corporate executives
OurDomain\SalesStaff	All sales people
OurDomain\ProductionStaff	All product managers and support staff
OurDomain\AllUsers	Everyone
OurDomain\CustomerSupport	Customer support representatives

Server Configuration The IT department has configured two physical servers with Microsoft Windows Server
2012 R2 and SQL Server 2014 Enterprise Edition and one Windows Azure Server. There are two tiers of storage available for use by database files only a fast tier and a slower tier. Currently the data and log files are stored on the fast tier of storage only. If a possible use case exists, management would like to utilize the slower tier storage for data files. The servers are configured as shown in the following table:

Location	Server
Company headquarters	HQ_Server
Satellite sales office	Satellite_Server
Microsoft Windows Azure (cloud)	Cloud_File Server

Database
Currently all information is stored in a single database called ProdDB, created with the following script:

```
CREATE DATABASE ProdDB
GO
ALTER DATABASE ProdDB SET RECOVERY SIMPLE
GO
```

The Product table is in the Production schema owned by the ProductionStaff Windows group. It is the main table in the system so access to information in the Product table should be as fast as possible. The columns in the Product table are defined as shown in the following table:

Column	Data type
ProductID	INT
ProductName	VARCHAR(100)
ProductDescription	VARCHAR(MAX)
ProductPrice	SMALLMONEY
QuantityOnHand	INT
ProductCost	SMALLMONEY
ProductSupplierID	INT

The SalesOrderDetail table holds the details about each sale. It is in the Sales schema owned by the SalesStaff Windows group. This table is constantly being updated, inserted into, and read. The columns in the SalesOrderDetail table are defined as shown in the following table:

Column	Data type
SalesOrderDetailID	INT
ProductID	INT
SalePrice	SMALLMONEY
SaleQuantity	INT

Database Issues

The current database does not perform well. Additionally, a recent disk problem caused the system to go down, resulting in lost sales revenue. In reviewing the current system, you found that there are no automated maintenance procedures. The database is severely fragmented, and everyone has read and write access.

Requirements

Database

The database should be configured to maximize uptime and to ensure that very little data is lost in the event of a server failure. To help with performance, the database needs to be modified so that it can support in-memory data, specifically for the Product table, which the CIO has indicated should be a memoryoptimized table. The auto-update statistics option is set off on this database. Only product managers are allowed to add products or to make changes to the name, description, price, cost, and supplier. The changes are made in an internal database and pushed to the Product table in ProdDB during system maintenance time. Product managers and others working at the headquarters location also should be able to generate reports that include supplier and cost information.

Customer data access

Customers access the company's website to order products, so they must be able to read product information such asname, description, and price from the Product table. When customers place orders, stored procedures calledby the website update product quantityon-hand values. This means the product table is constantly updated at randomtimes.

Customer support data access

Customer support representatives need to be able to view and not update or change product information.

Management does not want the customer support representatives to be able to see the product cost

or any supplier information.
Sales force data access
Sales people at both the headquarters office and the satellite office must generate reports that read from the Product and SalesOrderDetail tables. No updates or inserts are ever made by sales people. These reports are run at random times and there can be no reporting downtime to refresh the data set except during the monthly maintenance window. The reports that run from the satellite office are process intensive queries with large data sets. Regardless of which office runs a sales force report, the SalesOrderDetail table should only return valid, committed order data; any orders not yet committed should be ignored.
Historical Data
The system should keep historical information about customers who access the site so that sales people can see how frequently customers log in and how long they stay on the site.
The information should be stored in a table called Customer Access. Supporting this requirement should have minimal impact on production website performance.
Backups
The recovery strategy for Fabrikam needs to include the ability to do point in time restores and minimize the risk of data loss by performing transaction log backups every 15 minutes.
Database Maintenance
The company has defined a maintenance window every month when the server can be unavailable. Any maintenance functions that require exclusive access should be accomplished during that window.
Project milestones completed
* Revoked all existing read and write access to the database, leaving the schema ownership in place.
* Configured an Azure storage container secured with the storage account name MyStorageAccount with the primary access key StorageAccountKey on the cloud file server.
* SQL Server 2014 has been configured on the satellite server and is ready for use.
* On each database server, the fast storage has been assigned to drive letter F:, and the slow storage has been assigned to drive letter D:.
You need to create a job to automate some database maintenance tasks. Which code fragment should you use in each location in the command to complete one of the commands you will need to include in the job? To answer, drag the appropriate lines of code to the correct locations in the command. Each line of code may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

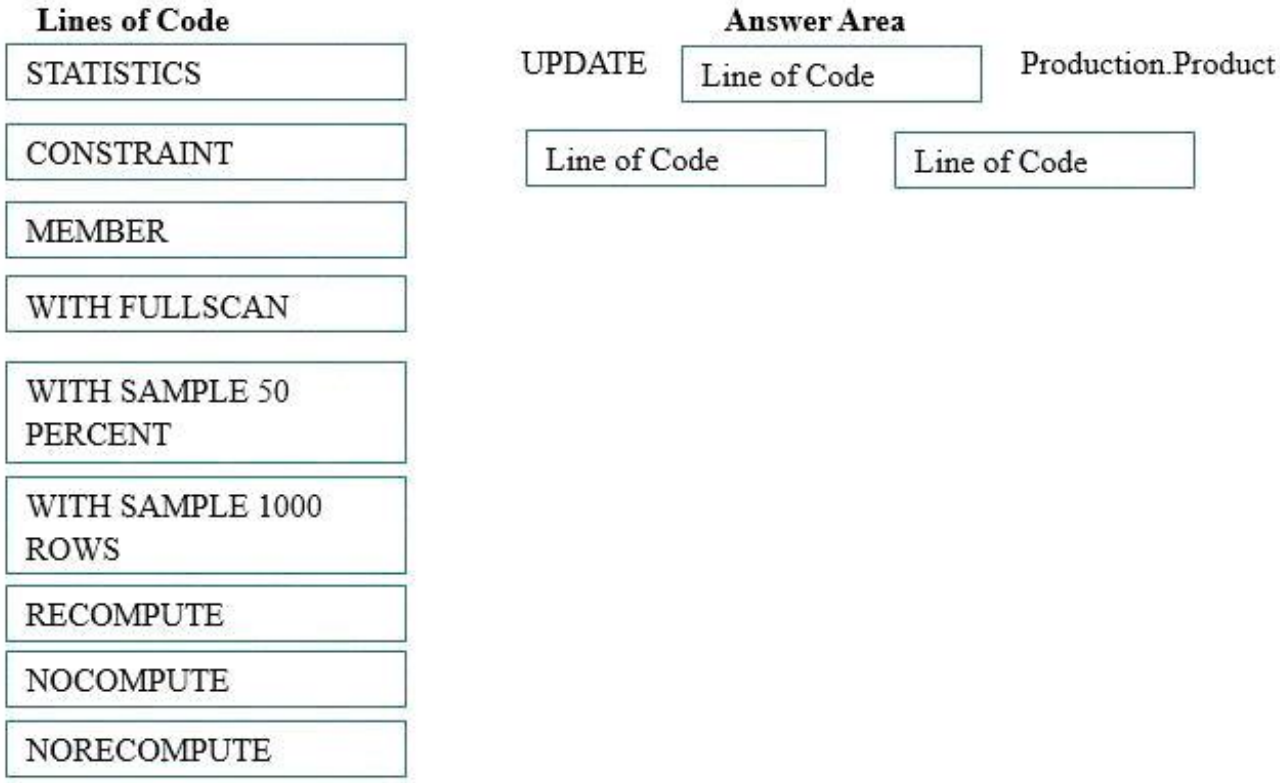

Answer:

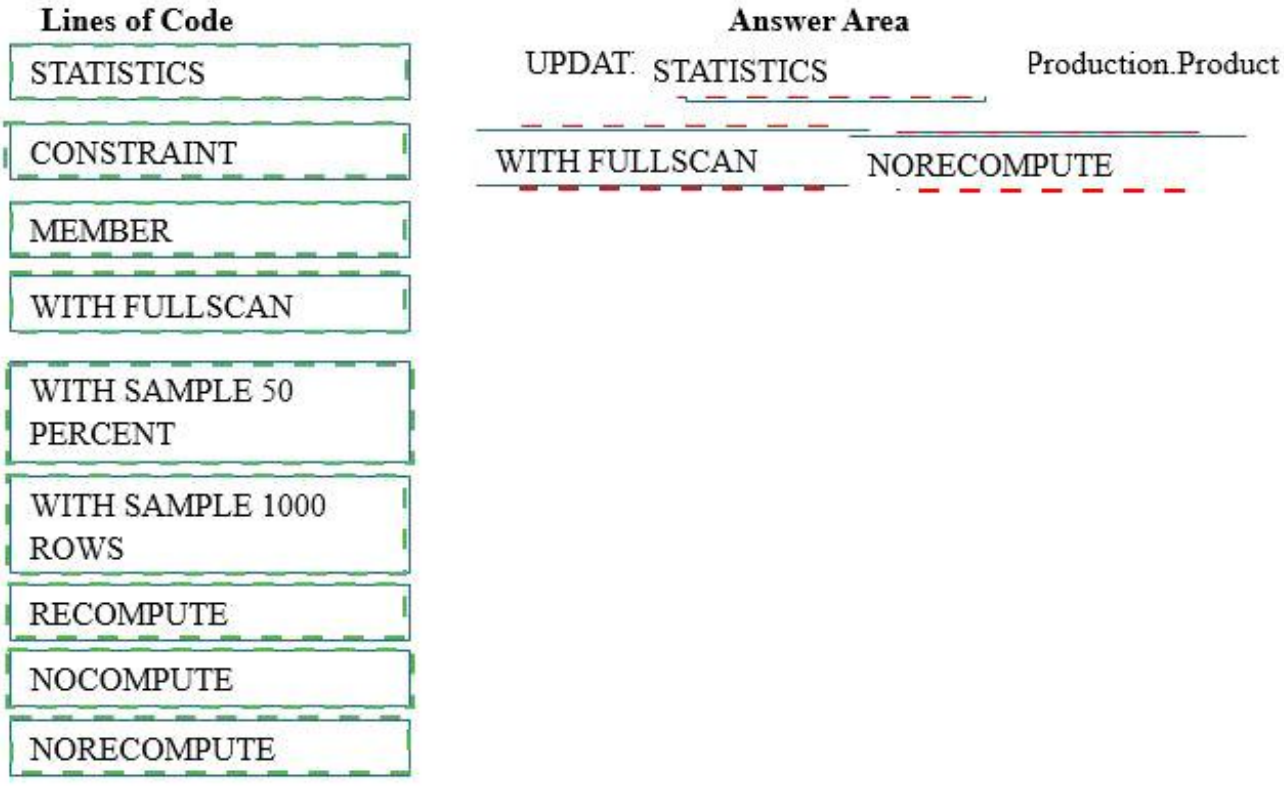

Explanation

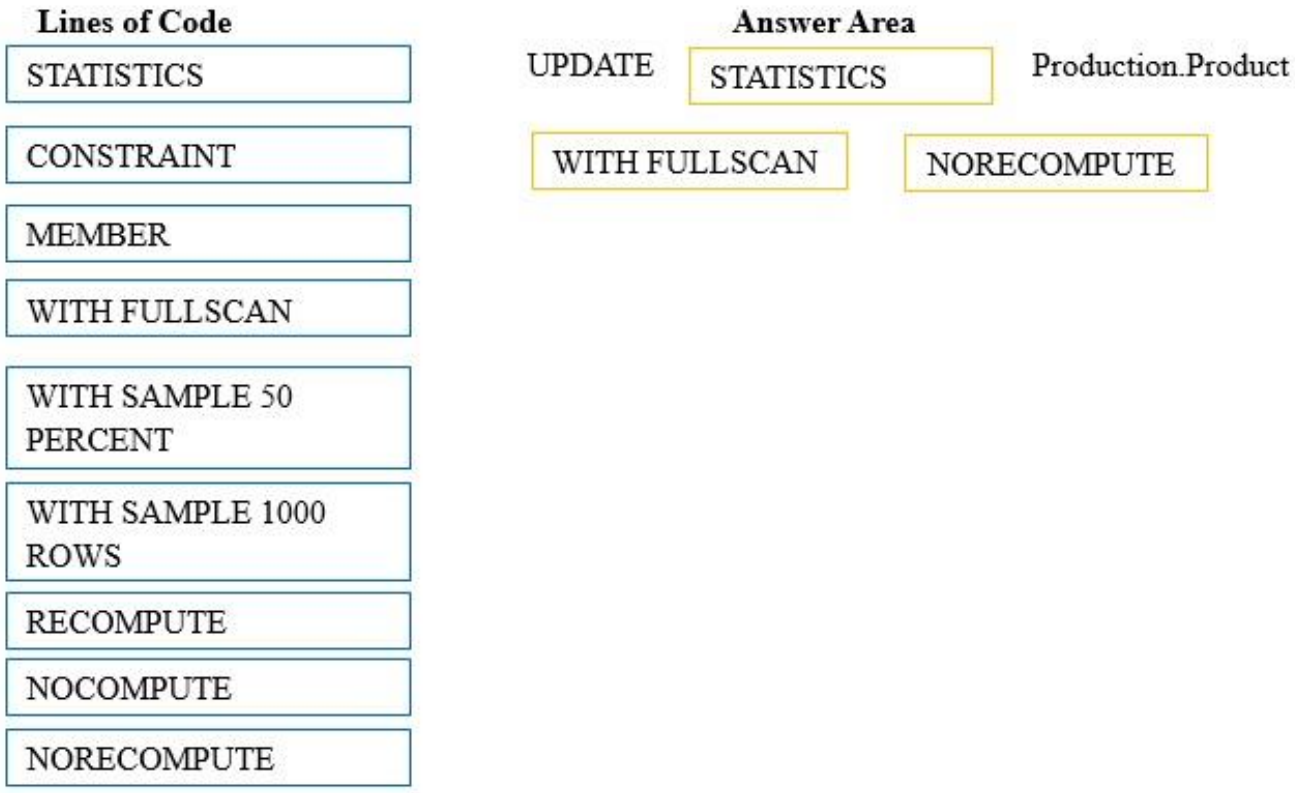

NO.20 You plan to install a Microsoft SQL Server 2016 instance.
The instance will support a database that has the following requirements:
* Store Excel workbooks on the file system.
* Access the workbooks through Transact-SQL.
* Include the workbooks in database backups.
* During installation, you need to ensure that the requirements will be met.
Which feature should you use?

A. FILESTREAM

B. SQL Server Integration Services (SSIS)

C. Excel Services

D. OpenXML

Answer: A