Test Guide 70-764

Volume: 55 Questions

Question No: 1
You need to configure a Microsoft SQL Server instance to ensure that a user named Mail1 can send mail by using Database Mail.
Solution: You add the DatabaseMailUserRole to Mail1 in the tempdb database. Does the solution meet the goal?

A. Yes

B. No

Answer: B

Explanation:
Database Mail is guarded by the database role DatabaseMailUserRole in the msdb database, not the tempdb database, in order to prevent anyone from sending arbitrary emails. Database users or roles must be created in the msdb database and must also be a member of DatabaseMailUserRole in order to send emails with the exception of sysadmin who has all privileges.
Note: Database Mail was first introduced as a new feature in SQLServer 2005 and replaces the SQL Mail feature found in previous versions. References:
http://www.idevelopment.info/data/SQLServer/DBA_tips/Database_Administration/DBA_20.shtml

Question No: 2
You need to configure a Microsoft SQL Server instance to ensure that a user named Mail1 can send mail by using Database Mail.
Solution: You add the DatabaseMailUserRole to Ma ill in the msdb database. Does the solution meet the goal?

A. Yes

B. No

Answer: A

Explanation:
Database Mail is guarded by the database role DatabaseMailUserRole in the msdb database in order to prevent anyone from sending arbitrary emails. Database users or roles must be created

in the msdb database and must also be a member of DatabaseMailUserRole in order to send emails with the exception of sysadmin who has all privileges.
Note: Database Mail was first introduced as a new feature in SQL Server 2005 and replaces the SQL Mail feature found in previous versions. References:
http://www.idevelopment.info/data/SQLServer/DBA_tips/Database_Administration/DBA_20.shtml

Question No: 3
You need to configure a Microsoft SQL Server instance to ensure that a user named Mail1 can send mail by using Database Mail.
Solution: You add the DatabaseMailUserRole to Ma ill in the master database. Does the solution meet the goal?

A. Yes

B. No

Answer: B

Explanation:
Database Mail is guarded by the database role DatabaseMailUserRole in the msdb database, not the master database, in order to prevent anyone from sending arbitrary emails. Database users or roles must be created in the msdb database and must also be a member of DatabaseMailUserRole in order to send emails with the exception of sysadmin who has all privileges.
Note: Database Mail was first introduced as a new feature in SQL Server 2005 and replaces the SQL Mail feature found in previous versions. References:
http://www.idevelopment.info/data/SQLServer/DBA_tips/Database_Administration/DBA_20.shtml

Question No: 4
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 the default instance on a server that 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-serveraudit -database-engine

Question No: 5
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 that runs Microsoft SQL Server 2016 Enterprise edition.
Does the solution meet the goal?

A. Yes

B. No

Answer: A

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-serveraudit -database-engine

Question No: 6
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-serveraudit -database-engine

Question No: 7
A company has an on-premises Microsoft SQL Server environment and Microsoft Azure SQL Database instances. The environment hosts several customer databases.
One customer reports that their database is not responding as quickly as the service level agreements dictate. You observe that the database is fragmented.
You need to optimize query performance.
Solution: You reorganize all indexes.
Does the solution meet the goal?

A. Yes

B. No

Answer: A

Explanation:
You can remedy index fragmentation by either reorganizing an index or by rebuilding an index.
References:
https://msdn.microsoft.com/en-us/library/ms189858(v=sql.105).aspx

Question No: 8
A company has an on-premises Microsoft SQL Server environment and Microsoft Azure SQL Database instances. The environment hosts several customer databases.
One customer reports that their database is not responding as quickly as the service level agreements dictate. You observe that the database is fragmented.
You need to optimize query performance.
Solution: You rebuild all indexes.
Does the solution meet the goal?

A. Yes

B. No

Answer: A

Explanation:
You can remedy index fragmentation by either reorganizing an index or by rebuilding an index.
References:
https://msdn.microsoft.com/en-us/library/ms189858(v=sql.105).aspx

Question No: 9
A company has an on-premises Microsoft SQL Server environment and Microsoft Azure SQL Database instances. The environment hosts several customer databases.
One customer reports that their database is not responding as quickly as the service level agreements dictate. You observe that the database is fragmented.
You need to optimize query performance.
Solution: You run the DBCC CHECKDB command. Does
the solution meet the goal?

A. Yes

B. No

Answer: B

Explanation:
DBCC CHECKDB only checks the logical and physical integrity of all the objects in the specified database. It does not update any indexes, and does not improve query performance.
References: https://docs.microsoft.com/en-us/sql/t-sql/database-console-commands/dbcc-checkdbtrans act-sql

Question No: 10
You are the database administrator for a company that hosts Microsoft SQL Server. You manage both on-premises and Microsoft Azure SQL Database environments.
One instance hosts a user database named HRDB. The database contains sensitive human resources data.
You need to grant an auditor permission to view the SQL Server audit logs while following the principle of least privilege.
Which permission should you grant?

A. DDLAdmin

B. db_datawriter

C. dbcreator

D. dbo

E. View Database State

F. View Server State

G. View Definition

H. sysadmin

Answer: F

Explanation:
Unless otherwise specified, viewing catalog views requires a principal to have one of the following: References:
https://technet.microsoft.com/en-us/library/cc280386(v=sql.110).aspx

Question No: 11
You manage a Microsoft SQL Server environment. You implement Transparent Data Encryption (TDE).
A user will assist in managing TDE.
You need to ensure that the user can view the TDE metadata while following the principle of lease privilege.
Which permission should you grant?

A. DDLAdmin

B. db_datawriter

C. dbcreator

D. dbo
E. View Database State

F. View Server State

G. View Definition

H. sysadmin

Answer: G

Explanation:
Viewing the metadata involved with TDE requires the VIEW DEFINITION permission on the certificate. References:
https://docs.microsoft.com/en-us/sql/relationaldatabases/security/encryption/transparentdata-encryption-tde

Question No: 12
You are the database administrator for a company that hosts Microsoft SQL Server. You manage both on-premises and Microsoft Azure SQL Database environments.
You have a user database named HRDB that contains sensitive human resources data. The HRDB backup files must be encrypted.
You need to grant the correct permission to the service account that backs up the HRDB database.
Which permission should you grant?

A. DDLAdmin

B. db_datawriter

C. dbcreator

D. dbo

E. View Database State

F. View Server State

G. View Definition

H. sysadmin

Answer: G

Explanation:
Restoring the encrypted backup: SQL Server restore does not require any encryption parameters to be specified during restores. It does require that the certificate or the asymmetric key used to encrypt the backup file be available on the instance that you are restoring to. The user account performing the restore must have VIEW DEFINITION permissions on the certificate or key.

References: https://docs.microsoft.com/en-us/sql/relational-databases/backup-restore/backupencryptio n

Question No: 13
You are the database administrator for a company that hosts Microsoft SQL Server. You manage both on-premises and Microsoft Azure SQL Database environments.
You plan to delegate encryption operations to a user.
You need to grant the user permission to implement cell-level encryption while following the principle of least privilege.
Which permission should you grant?

A. DDLAdmin

B. db_datawriter

C. dbcreator

D. dbo

E. View Database State

F. View ServerState

G. View Definition

H. sysadmin

Answer: G

Explanation:
The following permissions are necessary to perform column-level encryption, or cell-level encryption. References: https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/encrypta-col umn-of-data

Question No: 14
A company has an on-premises Microsoft SQL Server environment and Microsoft Azure SQL Database instances. The environment hosts a customer database named DB1.
Customers connect to hosted database instances by using line-of-business applications. Developers connect by using SQL Server Management Studio (SSMS).
You need to grant the developers permission to alter views for DB1 while following the principle of least privilege.

Which permission should you grant?

A. DDLAdmin

B. db_datawriter

C. dbcreator

D. dbo

E. View Database State

F. View Server State

G. View Definition

H. sysadmin

Answer: A

Explanation:
To execute ALTER VIEW, at a minimum, ALTER permission on OBJECT is required.
Members of the db_ddladmin fixed database role can run any Data Definition Language (DDL) command in a database.
References:
https://technet.microsoft.com/en-us/library/ms190667(v=sql.90).aspx

Question No: 15
You are the database administrator for a company that hosts Microsoft SQL Server. You manage both on-premises and Microsoft Azure SQL Database environments.
Clients connect to databases by using line-of-business applications. Developers connect by using SQL Server Management Studio (SSMS).
You need to provide permissions to a service account that will be used to provision a new database for a client.
Which permission should you grant?

A. DDLAdmin

B. db_datawriter

C. dbcreator

D. dbo

E. View Database State

F. View Server State

G. View Definition

H. sysadmin

Answer: C

Explanation:
Members of the dbcreator fixed server role can create, alter, drop, and restore any database.
References:
https://docs.microsoft.com/en-us/sql/relational-databases/security/authenticationaccess/server-level-roles

Question No: 16
You are examining information about users, sessions, and processed in an on-premises Microsoft SQL Server Database Engine instance.
You need to return information about processes that are not idle, that belong to a specific user, or that belong to a specific session.
What should you use?

A. Activity Monitor

B. sp_who3

C. SQL Server Management Studio (SSMS) Object Explorer

D. SQL Server Data Collector
E. SQL Server Data Tools (SSDT)

F. SQL Server Configuration Manager

Answer: B

Explanation:
Use sp_who3 to first view the current system load and to identify a session of interest. You should execute the query several times to identify which session id is most consuming the system resources. Parameters

sp_who3 null - who is active;
sp_who3 1 or 'memory' - who is consuming the memory;
sp_who3 2 or 'cpu' - who has cached plans that consumed the most cumulative CPU (top 10);
sp_who3 3 or 'count' - who is connected and how many sessions it has; sp_who3 4 or 'idle' - who is idle that has open transactions;
sp_who3 5 or 'tempdb' - who is running tasks that use tempdb (top 5); and, sp_who3 6 or 'block' - who is blocking.

Question No: 17
You observe that several indexes are fragmented.
You need to rebuild the indexes.
What should you use?

A. Activity Monitor

B. Sp_who3

C. Object Explorer in the SQL Server Management Studio (SSMS)

D. SQL Server Data Collector

E. SQL Server Data Tools (SSDT)

F. SQL Server Configuration Manager

Answer: C

Explanation:
How to: Rebuild an Index (SQL Server Management Studio)
To rebuild an index References:
https://technet.microsoft.com/en-us/library/ms187874(v=sql.105).aspx

Question No: 18
You need to examine information about logins, CPU times, and Disk I/O on a particular database in Microsoft Azure.
What should you use?

A. Activity Monitor

B. Sp_who3

C. SQL Server Management Studio (SSMS) Object Explorer

D. SQL Server Data Collector

E. SQL Server Data Tools (SSDT)

F. SQL Server Configuration Manager

Answer: A

Explanation:
Activity Monitor displays information about SQL Server processes and how these processes affect the current instance of SQL Server.
Activity Monitor is a tabbed document window with the following expandable and collapsible panes:
Overview, Active User Tasks, Resource Waits, Data File I/O, and Recent Expensive Queries. The Activity User Tasks Pane shows information for active user connections to the instance, including the following column:
* Login: The SQL Server login name under which the session is currently executing.
The Recent Expensive Queries Pane shows information about the most expensive queries that have been run on the instance over the last 30 seconds, including the following column: * CPU (ms/sec): The rate of CPU use by the query References:
https://technet.microsoft.com/en-us/library/cc879320(v=sql.105).aspx

Question No: 19
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. Activity Monitor

B. Sp_who3 stored procedure

C. Object Explorer in the Microsoft SQL Server Management Studio (SSMS)

D. SQL Server Data Collector

E. SQL Server Data Tools (SSDT)

F. SQL Server Configuration Manager

Answer: D

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

Question No: 20
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: What should you use?

A. Activity Monitor

B. Sp_who3

C. SQL Server Management Studio (SSMS) Object Explorer

D. SQL Server Data Collector

E. SQL Server Data Tools (SSDT)

F. SQL Server Configuration Manager

Answer: E

Explanation:
SQL Server Data Tools (SSDT) is a Microsoft Visual Studio environment for creating business intelligence solutions. SSDT features the Report Designer authoring environment, where you can open, modify, preview, save, and deploy Reporting Services paginated report definitions, shared data sources, shared datasets, and report parts.
References:
https://msdn.microsoft.com/en-us/library/hh272686(v=vs.103).aspx

Question No: 21
You have an on-premises server that runs Microsoft SQL Server 2016 Standard Edition. You need to identify missing indexes.

What should you use?

A. Activity Monitor

B. Sp_who3

C. SQL Server Management Studio (SSMS) Object Explorer

D. SQL Server Data Collector

E. SQL Server Data Tools (SSDT)

F. SQL Server Configuration Manager

Answer: D

Explanation:
Data Collector can gather performance information from multiple SQL Server instances and store it in a single repository. It has three built-in data collecting specifications (data collectors) designed to collect the most important performance metrics. The information collected by default is about disk usage, query statistics, and server activity.
The Query Statistics data collection set collects information about query statistics, activity, execution plans and text on the SQL Server instance.
Missing indexes can be found with the execution plans.
References: https ://www.sqlshack.com/sql-server-performance-monitoring-data-collector/

Question No: 22
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 monitoring for Tailspin Toys.
In the table below, identify the monitoring tool that you must use for each activity. NOTE: Make only one selection in each column.

Answer Area

Monitoring option	Monitoring from application	Trend analysis
Error logs	○	○
Transact-SQL	○	○
System Monitor	○	○
Distributed Replay	○	○

Answer:

Answer Area

Monitoring option	Monitoring from application	Trend analysis
Error logs	○	○
Transact-SQL	○	○
System Monitor	○	○
Distributed Replay	○	○

Test Guide 70-764

Question No: 23 DRAG DROP
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 implement a process for importing data into WingDB.
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.

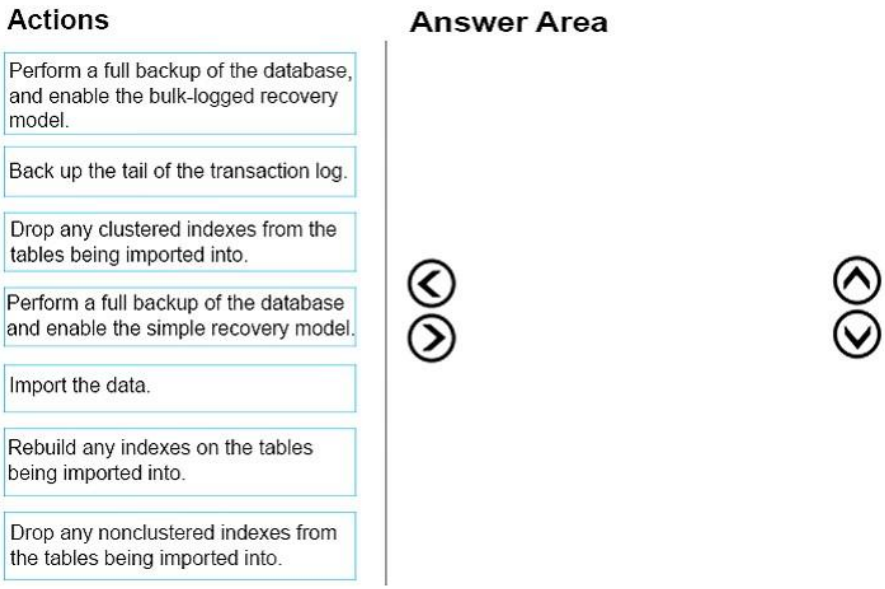

Answer:

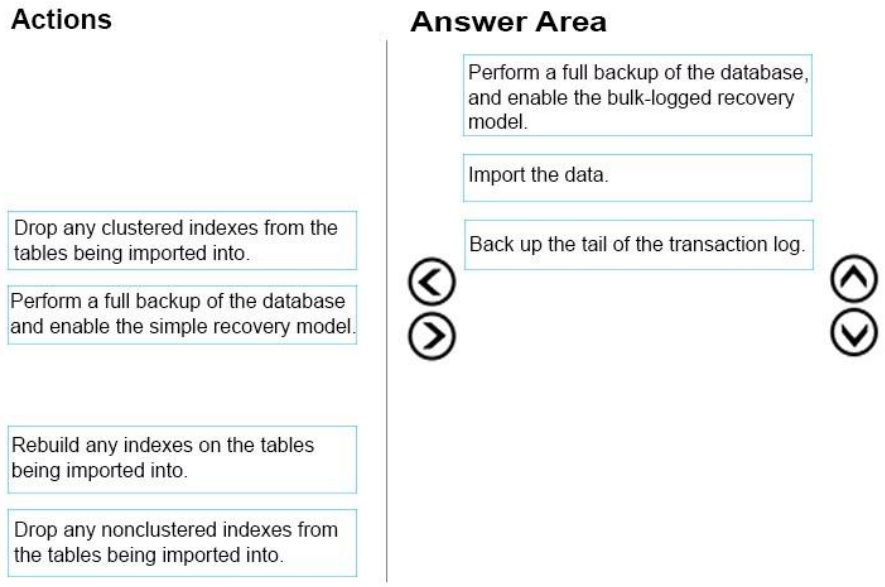

Explanation:
Step 1: Perform a full backup of the database and enable the bulk-logged recovery model.
Not: Simple recovery model.
With the Simple recovery model we cannot minimize data loss.
Step 2: Import the data
Step 3: Backup the tail of the transaction log.
For databases that use full and bulk-logged recovery, database backups are necessary but not sufficient. Transaction log backups are also required.

Note: Three recovery models exist: simple, full, and bulk-logged. Typically, a database uses the full recovery model or simple recovery model. A database can be switched to another recovery model at any time.

Question No: 24 HOTSPOT
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 the Adventure Works environment.
How should you complete the Transact-SOL statement? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Test Guide 70-764

Answer Area

```
USE master
GO
[ CREATE DATABASE AUDIT | ALTER DATABASE AUDIT | CREATE SERVER AUDIT | ALTER SERVER AUDIT ] AuditADVAccess
    TO FILE ( FILEPATH = 'C:\ADVAudit\' )
    WHERE object_name = 'SensitiveData'
GO
[ CREATE DATABASE AUDIT | ALTER DATABASE AUDIT | CREATE SERVER AUDIT | ALTER SERVER AUDIT ] AuditADVAccess WITH (STATE = ON)
GO
Use Adventureworks
[ CREATE DATABASE AUDIT | ALTER DATABASE AUDIT | CREATE SERVER AUDIT | ALTER SERVER AUDIT ] SPECIFICATION [FilterForSensitiveData]
[ FOR SERVER AUDIT | FOR DATABASE AUDIT | USE [AuditDataAcces] | SELECT ID ] [AuditADVAccess]
ADD (SELECT ON SCHEMA::[ADVSchema] BY [public])
WITH (STATE = ON)
GO
```

Answer:

Answer Area

```
USE master
GO
[ CREATE DATABASE AUDIT | ALTER DATABASE AUDIT | CREATE SERVER AUDIT | ALTER SERVER AUDIT ] AuditADVAccess
    TO FILE ( FILEPATH = 'C:\ADVAudit\' )
    WHERE object_name = 'SensitiveData'
GO
[ CREATE DATABASE AUDIT | ALTER DATABASE AUDIT | CREATE SERVER AUDIT | ALTER SERVER AUDIT ] AuditADVAccess WITH (STATE = ON)
GO
Use Adventureworks
[ CREATE DATABASE AUDIT | ALTER DATABASE AUDIT | CREATE SERVER AUDIT | ALTER SERVER AUDIT ] SPECIFICATION [FilterForSensitiveData]
[ FOR SERVER AUDIT | FOR DATABASE AUDIT | USE [AuditDataAcces] | SELECT ID ] [AuditADVAccess]
ADD (SELECT ON SCHEMA::[ADVSchema] BY [public])
WITH (STATE = ON)
GO
```

Explanation:
Box 1: CREATE SERVER AUDIT
Create the server audit.
You must implement auditing to record access to data that is considered sensitive by the company.
Create database audit Box 2:
ALTER SERVER AUDIT
Enable the server audit.
Box 3: CREATE DATABASE AUDIT
Create the database audit specification.
Box 4: FOR SERVER AUDIIT
You must implement auditing for all objects in the ADVSchema.

Question No: 25 HOTSPOT
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 the Contoso instance.

How should you complete the Transact-SOL statement? To answer, select the appropriate TransactSQL segments in the answer area.

Answer Area

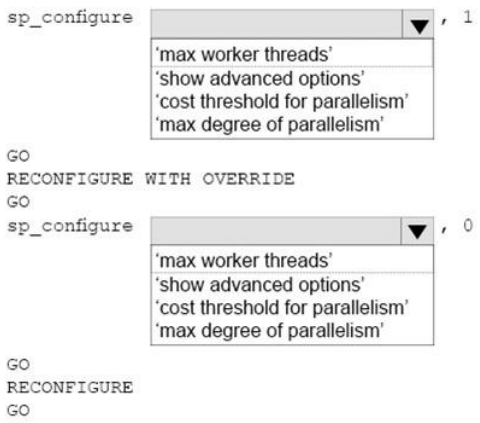

Answer:

Answer Area

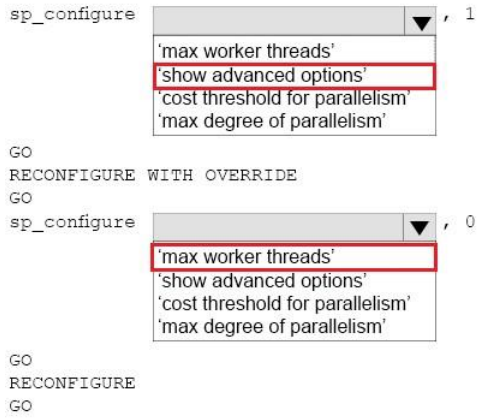

Explanation:
Box 1: show advanced options
Advanced configuration options are displayed by first setting show advanced option to 1.
Box 2: max worker threads
SQL Server uses the native thread services of the operating systems so that one or more threads support each network that SQL Server supports simultaneously, another thread handles database checkpoints, and a pool of threads handles all users. The default value for max worker threads is 0.
This enables SQL Server to automatically configure the number of worker threads at startup. The default setting is best for most systems. References:
https://docs.microsoft.com/en-us/sql/database-engine/configure-windows/configurethe-max-worker-threads-server-configuration-option
Question No: 26

Test Guide 70-764

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 monitor WingDB and gather information for troubleshooting issues. What should you use?

A. sp_updatestats

B. sp_lock

C. sys.dm_os_waiting_tasks

D. sys.dm_tran_active_snapshot_database_transactions

Answer: B

Explanation:
The sp_lock system stored procedure is packaged with SQL Server and will give you insight into the locks that are happening on your system. This procedure returns much of its information

from the syslock info in the master database, which is a system table that contains information on all granted, converting, and waiting lock requests.
Note: sp_lock will be removed in a future version of Microsoft SQL Server. Avoid using this feature in new development work, and plan to modify applications that currently use this feature. To obtain information about locks in the SQL Server Database Engine, use the sys.dm_tran_locks dynamic management view. sys.dm_tran_locks returns information about currently active lock manager resources in SQL Server 2008 and later. Each row represents a currently active request to the lock manager for a lock that has been granted or is waiting to be granted.
References: https://docs.microsoft.com/en-us/sql/relational-databases/system-storedprocedures/sp-lock -transact-sql

Question No: 27 HOTSPOT
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	○	○

Question No: 28 HOTSPOT
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.

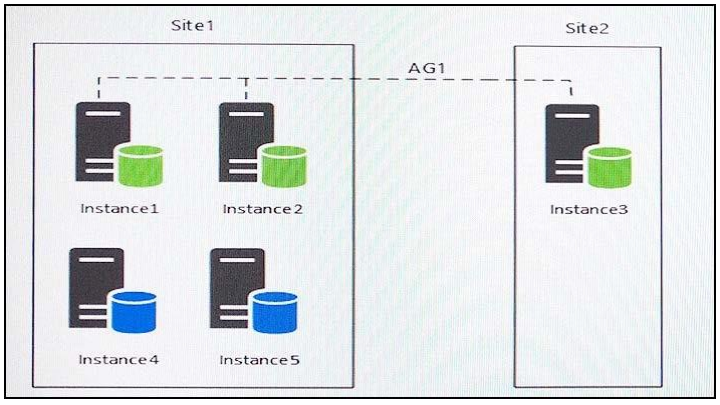

You have an Always On Availability group named AG1. The details for AG1 are shown in the following table.

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

lnstance1 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.
lnstance4 and Instances are not part of AG1. lnstance4 is engaged in heavy read-write I/O. Instances 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 lnstance6 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 create the connection strings for the operations and reporting systems.

In the table below, identify the option that must be specified in each connection string. NOTE: Make only one selection in each column.

Answer Area

Option	Reporting system	Operations system
Connect to a Listener using ApplicationIntent=ReadOnly.	○	○
Connect to the current primary replica SQL instance using ApplicationIntent=ReadOnly.	○	○
Connect to any current read-only replica SQL instance.	○	○
Connect to a Listener.	○	○
Connect to the current primary replica SQL instance.	○	○

Answer:

Answer Area

Option	Reporting system	Operations system
Connect to a Listener using ApplicationIntent=ReadOnly.	○	○
Connect to the current primary replica SQL instance using ApplicationIntent=ReadOnly.	○	○
Connect to any current read-only replica SQL instance.	●	○
Connect to a Listener.	○	○
Connect to the current primary replica SQL instance.	○	●

Explanation:
Reporting system: Connect to any current read-only replica instance
We configure Read-OnlyAccess on an Availability Replica. We select Read-intent only. Only read-only connections are allowed to secondary databases of this replica. The secondary database(s) are all available for read access.
From Scenario: 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.
Operating system: Connect to the current primary replica SQL instance
By default both read-write and read-intent access are allowed to the primary replica and no connections are allowed to secondary replicas of an Always On availability group.
From scenario: 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. References:
https://docs.microsoft.com/en-us/sql/database-engine/availabilitygroups/windows/configure-read-only-access-on-an-availability-replica-sql-server

Question No: 29 DRAG DROP
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.

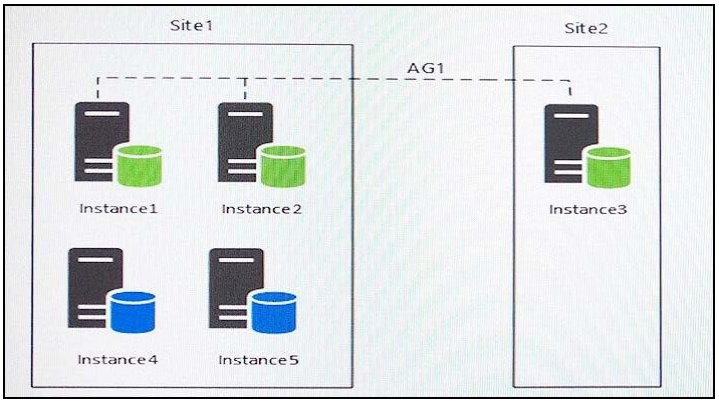

You have an Always On Availability group named AG1. The details for AG1 are shown in the following table.

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

lnstance1 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. lnstance4 and Instances are not part of AG1. lnstance4 is engaged in heavy read-write I/O. Instances 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 lnstance6 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 Staged External 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.

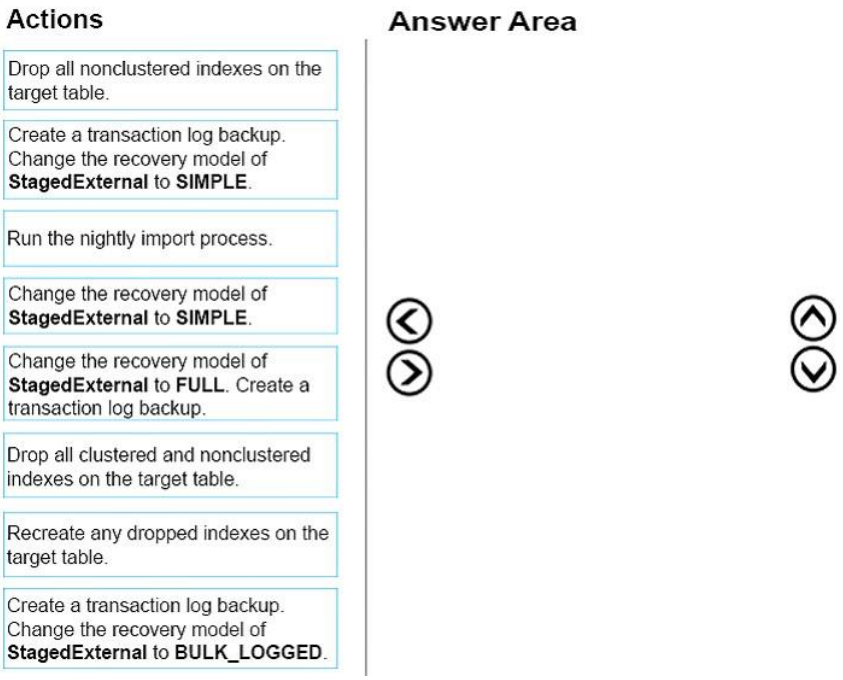

Answer:

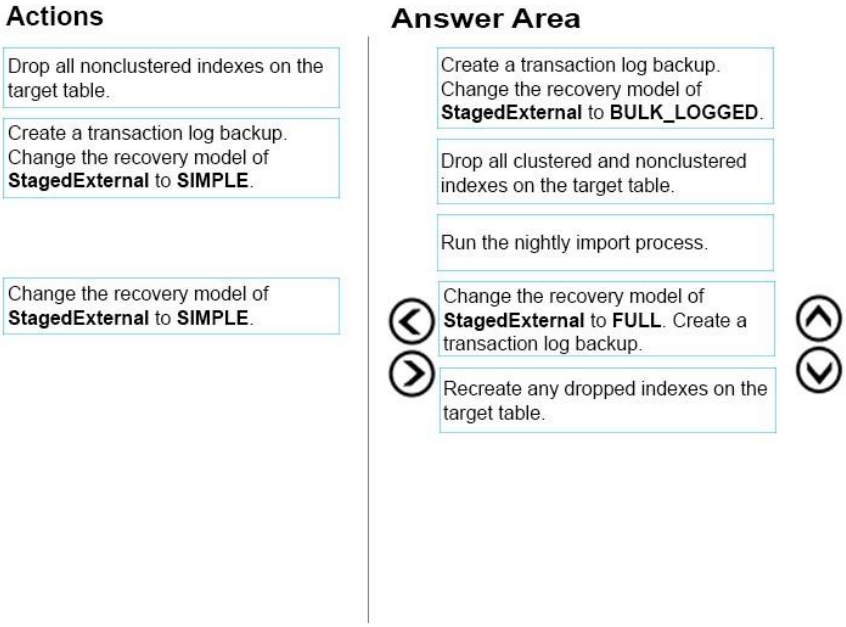

Explanation:
From scenario: Instances hosts a database named StagedExternal. A nightly BULK INSERT process loads data into an empty table that has a rowstore clustered index and two nonclustered rowstore indexes.

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.

Question No: 30 DRAG DROP
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.

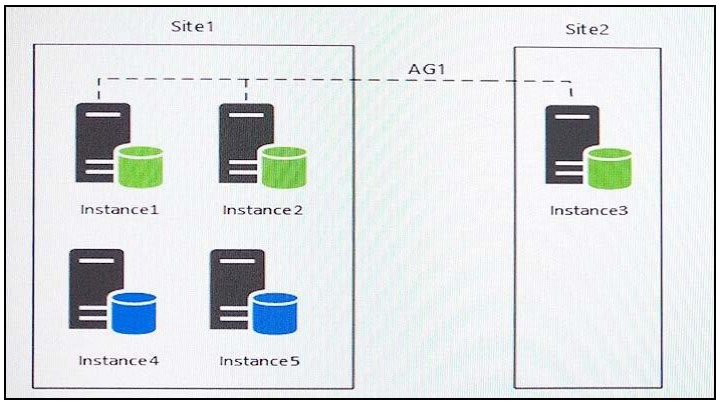

You have an Always On Availability group named AG1. The details for AG1 are shown in the following table.

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

lnstance1 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.
lnstance4 and Instances are not part of AG1. lnstance4 is engaged in heavy read-write I/O. Instances 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 lnstance6 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 analyze the wait type and statistics for specific instanced in the environment.

Which object should you use to gather information about each instance? To answer, drag the appropriate objects to the correct instances. Each object 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. NOTE: Each correct selection is worth one point.

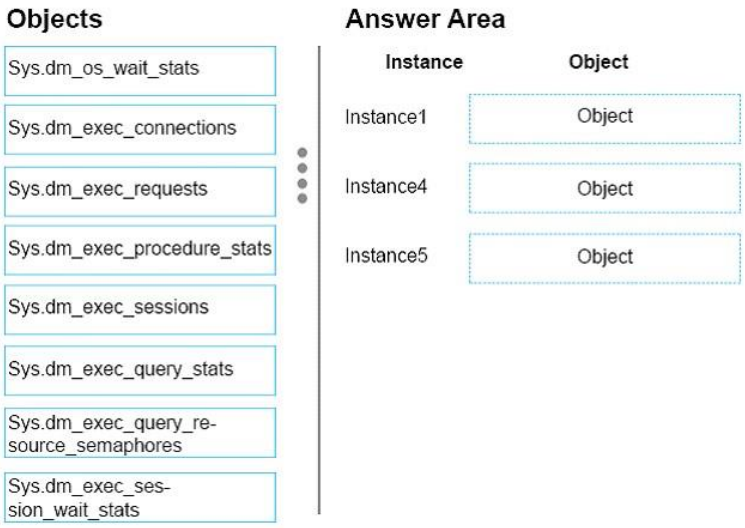

Answer:

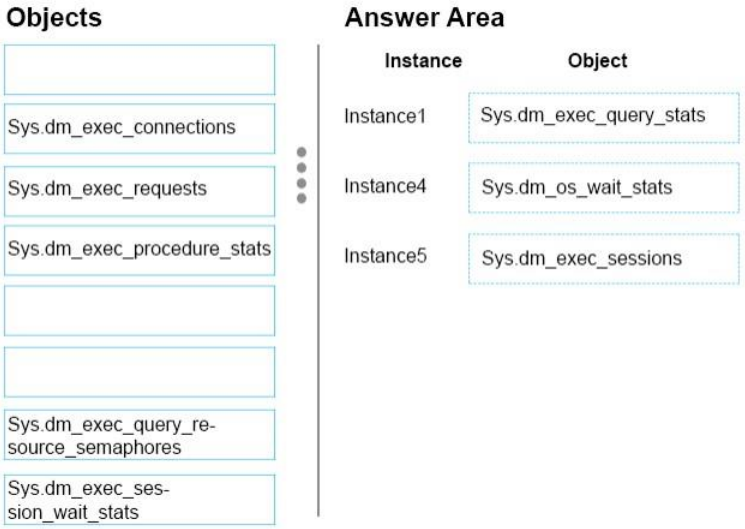

Explanation:
Instance 1: sys.dm_exec_query_stats
From Scenario: Instance! requirement: Aggregate statistics since last server restart. sys.dm_exec_query_stats returns aggregate performance statistics for cachedquery plans in SQL Server.
Instance 4: sys.dm_os_wait_stats
sys.dm_os_wait_statsreturns information about all the waits encountered by threads that executed.
From Scenario: lnstance4 requirement: Identify the most prominent wait types.

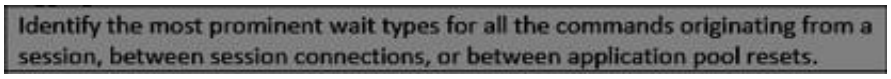

Instance 5:sys.dm_exec_session_wait_stats
From Scenario: Instances requirement: Identify all wait types for queries currently running on the server.
sys.dm_exec_session_wait_stats returns information about all the waits encountered by threads that executed for each session.

Question No: 31 DRAG DROP
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.

Test Guide 70-764

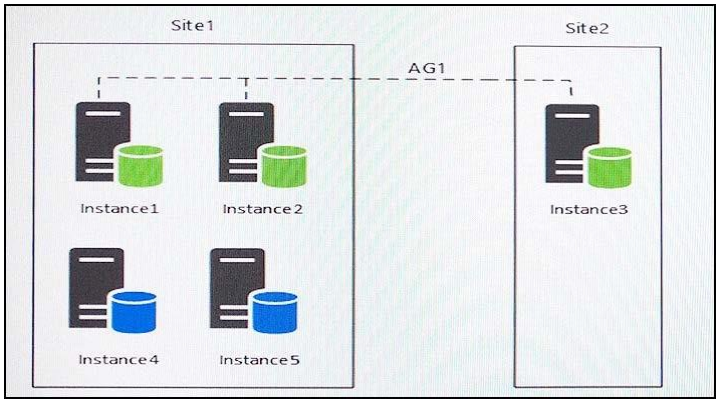

You have an Always On Availability group named AG1. The details for AG1 are shown in the following table.

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

lnstance1 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. lnstance4 and Instances are not part of AG1. lnstance4 is engaged in heavy read-write I/O. Instances 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 lnstance6 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 configure a new replica of AG1 on lnstance6.

How should you complete the Transact-SOL statement? To answer, drag the appropriate Transact-SOL statements to the correct locations. Each Transact-SOL 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:

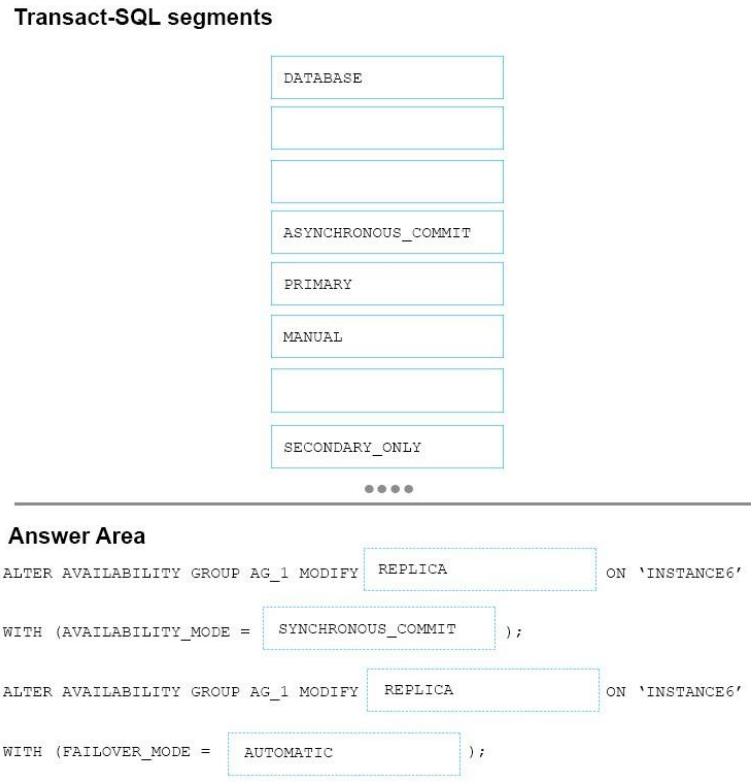

Explanation:
Scenario: You plan to add a new instance named lnstance6 to a datacenter that is geographically distant from Site1 and Site2. You must minimize latency between the nodes in AG1.
Box 1: REPLICA
MODIFY REPLICA ON modifies any of the replicas of the availability group.
Box 2: SYNCHRONOUS_COMMIT
You must minimize latency between the nodes in AG1
AVAILABILITY_MODE = { SYNCHRONOUS_COMMIT I 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 I 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.

Test Guide 70-764

Question No: 32
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.

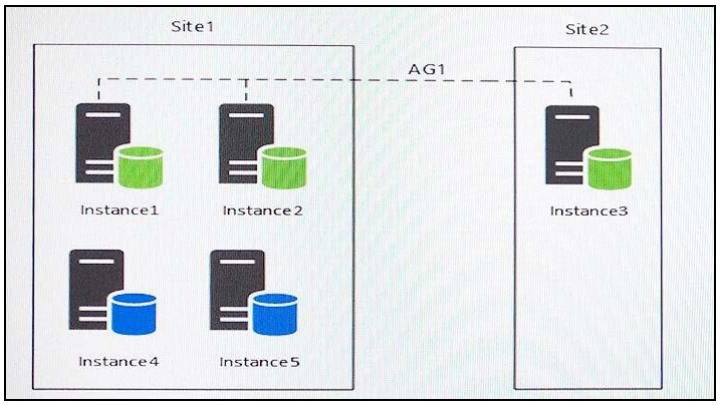

You have an Always On Availability group named AG1. The details for AG1 are shown in the following table.

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

lnstance1 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. lnstance4 and Instances are not part of AG1. lnstance4 is engaged in heavy read-write I/O. Instances 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 lnstance6 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 create a backup plan for lnstance4.
Which backup plan should you create?

A. Weekly full backups, nightly differential. No transaction log backups are necessary.

B. Weekly full backups, nightly differential backups, transaction log backups every 5 minutes.

C. Weekly full backups, nightly differential backups, transaction log backups every 12 hours.

D. Weekly full backups, nightly differential backups, nightly transaction log backups.

Answer: B

Explanation:
From scenario: lnstance4 and Instances are not part of AG1. lnstance4 is engaged in heavy read-write I/O. The recovery point objective of lnstancse4 is 60 minutes. RecoveryPoint Objectives are commonly described as the amount of data that was lost during the outage and recovery period.
References: http://sqlmag.com/blog/sql-server-recovery-time-objectives-and-recovery-pointobjectives

Question No: 33
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.

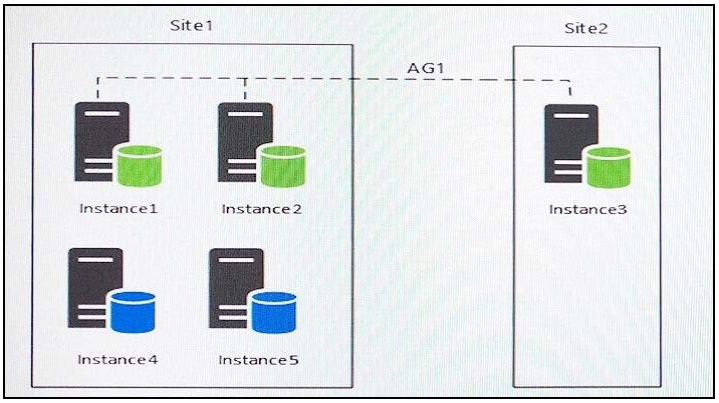

You have an Always On Availability group named AG1. The details for AG1 are shown in the following table.

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

lnstance1 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.
lnstance4 and Instances are not part of AG1. lnstance4 is engaged in heavy read-write I/O. Instances 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 lnstance6 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/partialbackups-sql-server

Question No: 34 DRAG DROP
You have a database. The existing backups for the database and their corresponding files are listed in the following table.

Backup type	Backup date/time	File name
Full	05/02/2016 21:00	Full_20160502_2100.bak
Transaction log	05/03/2016 6:00	Log_20160503_0600.trn
Transaction log	05/03/2016 9:00	Log_20160503_0900.trn
Differential	05/03/2016 12:00	Diff_20160503_1200.bak
Transaction log	05/03/2016 15:00	Log_20160503_1500.trn
Differential	05/03/2016 17:00	Diff_20160503_1700.bak
Transaction log	05/03/2016 19:00	Log_20160503_1900.trn

You purchase a new server. You must restore the database to the new server.
You need to restore the data to the most recent time possible.
Which three files should you restore in sequence? To answer, move the appropriate files from the list of files to the answer area and arrange them in the correct order.

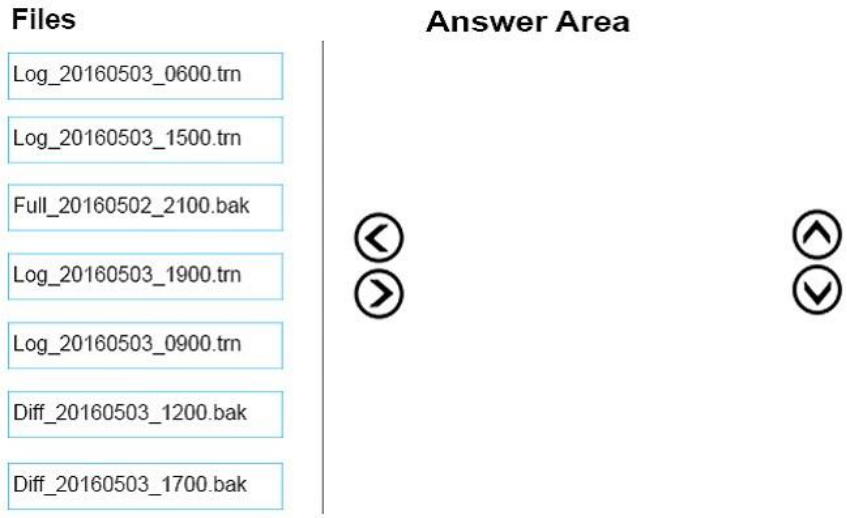

Answer:

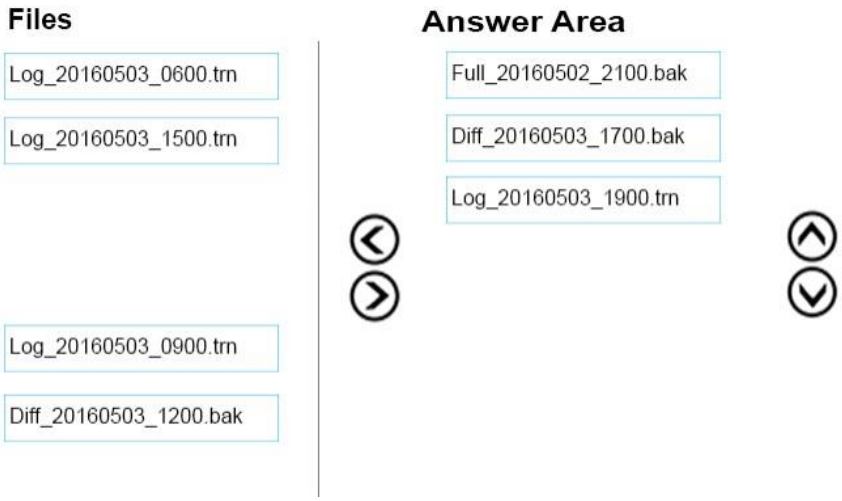

Explanation:
Step 1: Full..
Start with the full backup.
Step 2: Diff_20160503_1700.bak
Followed by the most recent differential backup.
Step 3: Log_20160503_1900.bak
And finally the most recent log backup (the only log backup done after the most recent differential backup).
References: https://docs.microsoft.com/en-us/sql/relational-databases/backup-restore/differentialbacku ps-sql-server

Question No: 35
A Microsoft SQL Server database named DB1 has two filegroups named FG1 and FG2. You implement a backup strategy that creates backups for the filegroups.
DB1 experiences a failure. You must restore FG1 and then FG2.
You need to ensure that the database remains in the RECOVERING state until the restoration of FG2 completes. After the restoration of FG2 completes, the database must be online. What should you specify when you run the recovery command?

A. the WITH NORECOVERY clause for FG1 and the WITH RECOVERY clause for FG2

B. the WITH RECOVERY clause for FG1 and the WITH RECOVERY clause for FG2

C. the WITH RECOVERY clause for both FG1 and FG2

D. the WITH NORECOVERY clause for both FG1 and FG2

Answer: A

Question No: 36 DRAG DROP
You have a test server that contains a database named DB1. Backups of the database are written to a single backup device. The backup device has a full, differential, and transaction log backup. You discover that the database is damaged. You restore the database to the point at which the differential backup was taken.
You need to rebuild the database with data stored in the latest transaction logs.
How should you complete the Transact-SOL statement? To answer. drag the appropriate Transact-SOL segments to the correct locations. Each Transact-SOL 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.

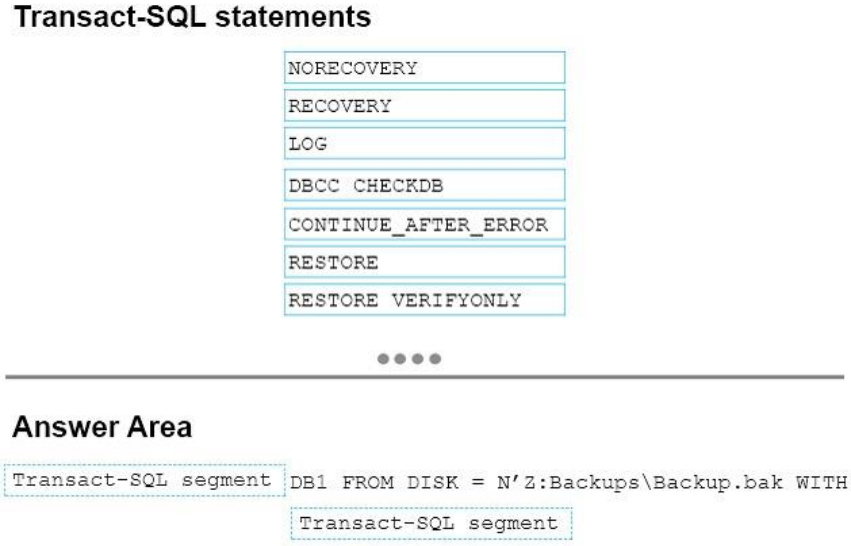

Answer:

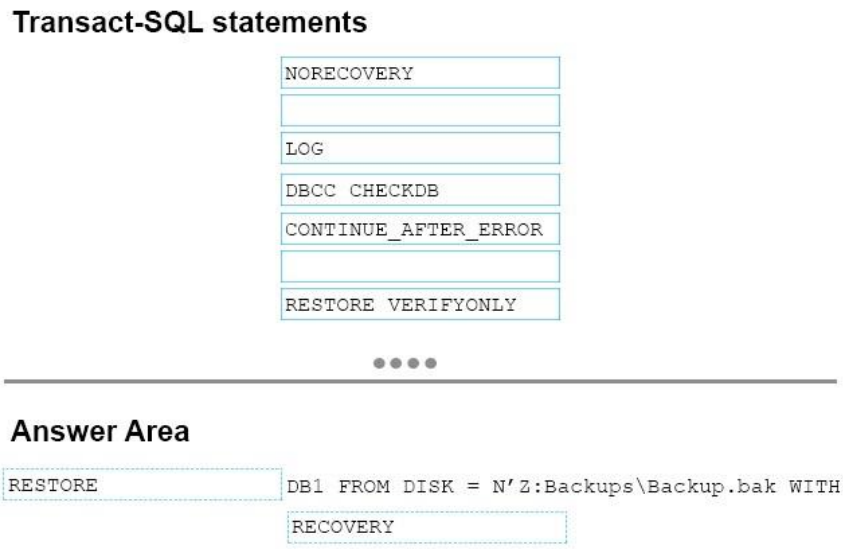

Explanation:
Box 1: RESTORE
Box 2: RECOVERY
The RESTORE ... WITH RECOVERY option puts the database into a useable state, so users can access a restored database.
References: https://www.mssqltips.com/sqlservertutorial/112/recovering-a-database-that-is-in-therestori ng-state/

Question No: 37
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 D81 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-SOL 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
```

Answer: B

Question No: 38 HOTSPOT
You manage a Microsoft-SOL Server database named sales Orders.
You need to verify the integrity of the database and attempt to repair any errors that are found. Repair must not cause any data to be lost in the database.
How should you complete the DBCC command? To answer, select the appropriate options in the answer area.

Answer Area

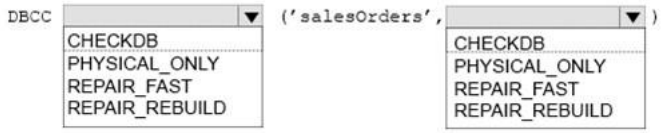

Answer:

Answer Area

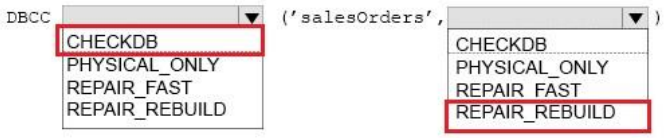

Explanation:
Box 1: CHECKDB
DBCC CHECKDB checks the logical and physical integrity of all the objects in the specified database. Partial syntax:
DBCC CHECKDB
[(database_name I database_id I 0
[, NOINDEX I, { REPAIR_ALLOW_DATA_LOSS I REPAIR_FAST I REPAIR_REBUILD}]
Box 2: REPAIR_REBUILD
DBCC CHECKDB ... REPAIR_ALLOW_DATA_LOSS I REPAIR_FAST I REPAIR_REBUILD specifies that DBCC CHECKDB repair the found errors.
REPAIR_REBUILD performs repairs that have no possibility of data loss. This can include quick repairs, such as repairing missing rows in non-clustered indexes, and more time-consuming repairs, such as rebuilding an index.
References:
https://docs.microsoft.com/en-us/sql/t-sql/database-console-commands/dbcc-checkdbtransact-sql

Question No: 39

You manage a Microsoft SQL Server environment. You have a database named salesOrders that includes a table named Table1.
Table1 becomes corrupt. You repair the table.
You need to verify that all the data in Table1 complies with the schema.
How should you complete the Transact-SOL code statement? To answer, select the appropriate Transact-SOL code segments in the dialog box in the answer area.

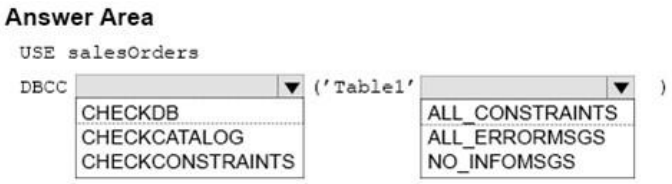

Answer:

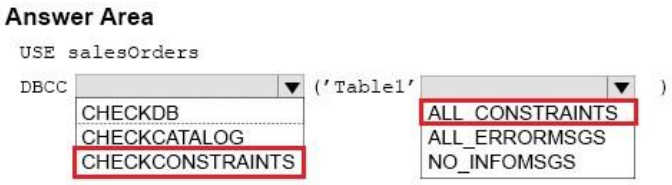

Explanation:
Box 1: CHECKCONSTRAINTS
DBCC CHECKCONSTRAINTS checks the integrity of a specified constraint or all constraints on a specified table in the current database.
Box 2: ALL_CONSTRAINTS
ALL_CONSTRAINTS checks all enabled and disabled constraints on the table if the table name is specified or if all tables are checked;otherwise, checks only the enabled constraint. Note: Syntax: DBCC CHECKCONSTRAINTS
[
(
table_name I table_id I constraint_name I constraint_id
) 1
[WITH
[{ ALL_CONSTRAINTS I ALL_ERRORMSGS}]
[,] [NO_INFOMSGS]
]
References:
https://docs.microsoft.com/en-us/sql/t-sql/database-console-commands/dbccheckconstraints-transact-sql

Question No: 40

Test Guide 70-764

You are configuring a new Microsoft SQL Server Always On Availability Group. You plan to configure a shared network location at \\DATA-CI1\SQL.
You need to create an availability group listener named AGL1 on port 1433.
In which order should you perform the actions? To answer, move all actions from the list of actions to the answer area and arrange them in the correct order.

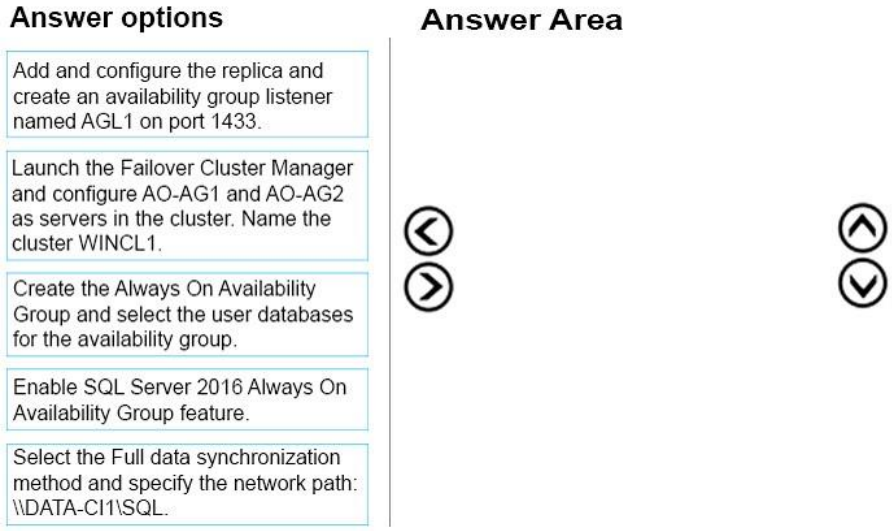

Answer:

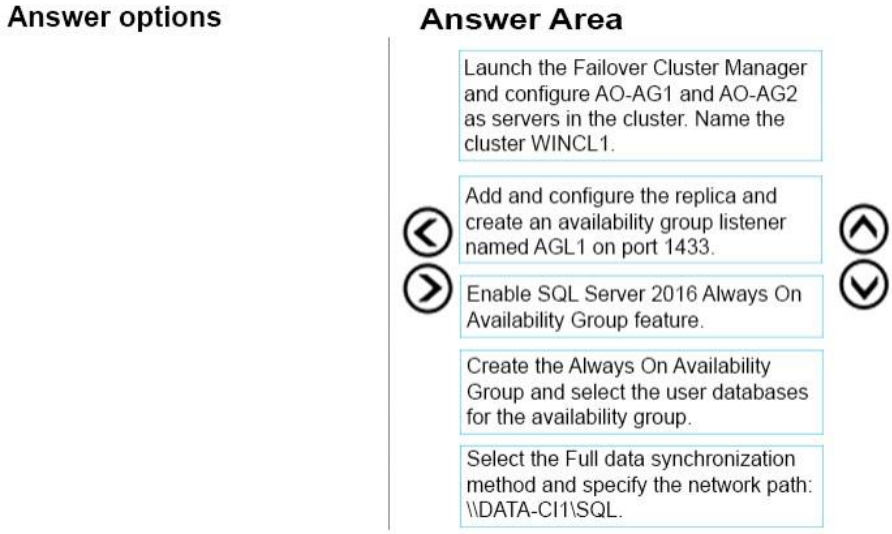

Explanation:
Step 1: Launch the Failover Cluster Manager and ..
To support the Always On availability groups feature, ensure that every computer that is to participate in one or more availability groups meets requirements including:
* Ensure that each computer is a node in a WSFC (Windows Server Failover Clustering).
Step 2: Add and configure the replica and ...
All the server instances that host availability replicas for an availability group must use the same SQL Server collation.
Step 3: Enable the SQL Server 2016 Always On Availability Group feature.

Test Guide 70-764

Enable the Always On availability groups feature on each server instance that will host an availability replica for any availability group. On a given computer, you can enable as many server instances for Always On availability groups as your SQL Server installation supports.
Step 4: Create the Always On Availability Group and ..
Using Transact-SOL to create or configure an availability group listener Step
5: Select the Full data synchronization method and ...

Question No: 41
You are configuring log shipping for a Microsoft SQL Server database named salesOrders.
You run the following Transact-SOL script:

```
DECLARE @LS_BackupJobId AS uniqueidentifier
DECLARE @LS_PrimaryId AS uniqueidentifier
DECLARE @SP_Add_RetCode As int
EXEC @SP_Add_RetCode = master.dbo.sp_add_log_shipping_primary_database
     @database = N'salesOrders'
     ,@backup_directory = N'C:\Backup'
     ,@backup_share = N'\\localhost\Backup'
     ,@backup_job_name = N'LSBackup_salesOrders'
     ,@backup_retention_period = 4320
     ,@backup_compression = 1
     ,@backup_threshold = 60
     ,@threshold_alert_enabled = 1
     ,@history_retention_period = 5760
     ,@backup_job_id = @LS_BackupJobId OUTPUT
     ,@primary_id = @LAS_PrimaryId OUTPUT
     ,@overwrite = 1
IF (@@ERROR = 0 AND @SP_Add_RetCode = 0)
     BEGIN
          DECLARE @LS_BackUpScheduleUID As uniqueidentifier
          DECLARE @LA_BackUpScheduleID AS int
          EXEC msdb.dbo.sp_add_schedule
               @schedule_name = N'LSBackupSchedule_ADATUM-SQL11'
               ,@enabled = 1
               ,@freq_type = 4
               ,@freq_interval = 1
               ,@freq_subday_type = 4
               ,@freq_subday_interval = 15
               ,@freq_recurrence_factor = 0
               ,@active_start_date = 20160720
               ,@active_end_date = 99991231
               ,@active_start_time = 0
               ,@active_end_time = 235900
               ,@schedule_uid = @LS_BackUpScheduleUID OUTPUT
               ,@schedule_id = @LS_BackupScheduleID OUTPUT
          EXEC msdb.dbo.sp_attach_schedule
               @job_id = @LS_BackupJobId
               ,@schedule_id = @LS_BackupScheduleID
          EXEC msdb.dbo.sp_update_job
               @job_id = @LS_BackupJobId
               ,@enabled = 1
     END
EXEC master.dbo.sp_add_log_shipping_alert_job
```

You need to determine the changes that the script has on the environment.
How does the script affect the environment? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Answer Area

A dedicated file share **[answer choice]** used to store the backups.
- is
- is not

A SQL Server monitor instance **[answer choice]** on a server named ADATUM-SQL11.
- runs
- does not run

Backup files will be deleted after **[answer choice]**.
- 24 hours
- 48 hours
- 72 hours

The backup job will run every **[answer choice]**.
- 15 minutes
- 60 minutes
- 24 hours

Answer:

Answer Area

A dedicated file share **[answer choice]** used to store the backups.
- **is**
- is not

A SQL Server monitor instance **[answer choice]** on a server named ADATUM-SQL11.
- runs
- **does not run**

Backup files will be deleted after **[answer choice]**.
- 24 hours
- 48 hours
- **72 hours**

The backup job will run every **[answer choice]**.
- **15 minutes**
- 60 minutes
- 24 hours

Explanation:
Box 1: is
The dedicated backup file share is \ \localhost\Backup
Box 2: does not run
The only thing with a name related to ADATM-SQL11 is the schedule name.
Box 3: 72 hours
4320 minutes equals 72 hours.
Note: @backup_retention_period=] backup_retention_period
Ls the length of time, in minutes, to retain the log backup file in the backup directory on the primary server. backup_retention_period is int, with no default, and cannot be NULL.
Box 4: 15 minutes.
[@freq_subday_type =] freq_subday_type
Specifies the units for freq_subday_interval. freq_subday_typeis int, with a default of 0, and can be one of these values.
Here it is 4, which means minutes.
[@freq_subday_interval =] freq_subday_interval

The number of freq_subday_type periods to occur between eachexecution of a job.
freq_subday_intervalis int, with a default of 0. Note: Interval should be longer than 10 seconds.
freq_subday_interval is ignored in those cases where freq_subday_type is equal to 1.
Here it is 15.
References:
https://docs.microsoft.com/en-us/sql/relational-databases/system-storedprocedures/sp-add-schedule-transact-sql
https://docs.microsoft.com/en-us/sql/relational-databases/system-stored-procedures/sp-add-logshipping-primary-database-transact-sql

Question No: 42
You are planning the deployment of two new Always On Failover Cluster Instances (FCIs) of Microsoft SQL Server to a single Windows Server Cluster with three nodes. The planned configuration for the cluster is shown in the Server Layout exhibit. (Click the Exhibit button.)

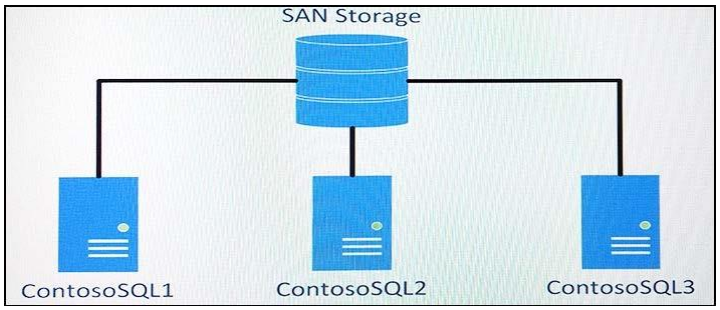

The SAN team has configured storage for the cluster and sent the configuration to you in the email shown in the SAN Team Email exhibit. (Click the Exhibit button.)

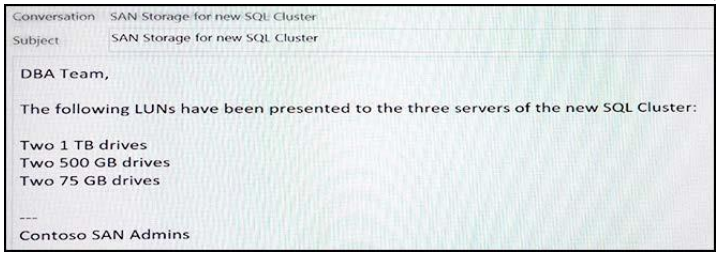

Each node of the cluster has identical local storage available as shown in the Local Storage exhibit. (Click the Exhibit button.)

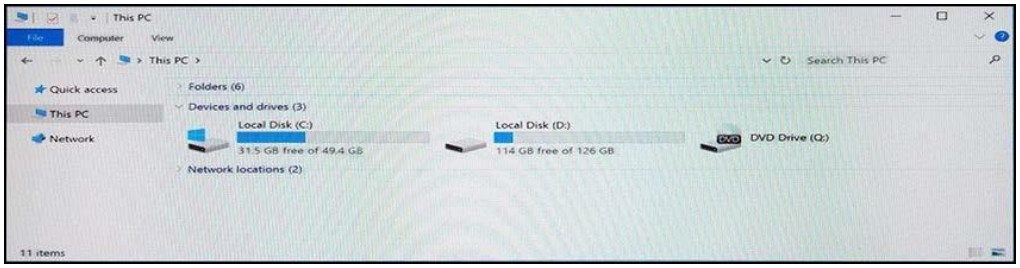

All local storage is on SSD.
You need to plan specific configurations for the new cluster.
For each of the following statement, select Yes if the statement is true. Otherwise, select No.

Answer Area

Statements	Yes	No
The **Tempdb** database for each cluster instance can be placed on the D: drive for the instance.	○	○
One virtual network name for each SQL Server instance must be configured in the cluster.	○	○
The shared storage has been formatted and configured on ContosoSQL1.	○	○

Answer:

Answer Area

Statements	Yes	No
The **Tempdb** database for each cluster instance can be placed on the D: drive for the instance.	○	○
One virtual network name for each SQL Server instance must be configured in the cluster.	○	○
The shared storage has been formatted and configured on ContosoSQL1.	○	○

Question No: 43

You are planning to deploy log shipping for Microsoft SQL Server and store all backups on a dedicated file share.

You need to configure the servers to perform each log shipping step.

Which server instance should you configure to perform each action? To answer, select the appropriate server instances in the dialog box in the answer area.

Answer Area

Action	Server instance
Complete the backup job.	Primary server instance Secondary server instance Monitor server instance Backup share file server
Copy the backup job.	Primary server instance Secondary server instance Monitor server instance Backup share file server
Restore the backup.	Primary server instance Secondary server instance Monitor server instance Backup share file server

Answer:

Answer Area

Action	Server instance
Complete the backup job.	Primary server instance
Copy the backup job.	Secondary server instance
Restore the backup.	Secondary server instance

Explanation:
Note: Before you configure log shipping, you must create a share to make the transaction log backups available to the secondary server.
SQL Server Log shipping allows you to automatically send transaction log backups from a primary database on a primary server instance to one or more secondary databases on separate secondary server instances. The transaction log backups are applied to each of the secondary databases individually. An optional third server instance, known as the monitor server, records the history and status of backup and restore operations and, optionally, raises alerts if these operations fail to occur as scheduled.
Box 1: Primary server instance.
The primary server instance runs the backup job to back up the transaction log on the primary database.
backup job: A SQL Server Agent job that performs the backup operation, logs history to the local server and the monitor server, and deletes old backup files and history information. When log shipping is enabled, the job category "Log Shipping Backup" is created on the primary server instance.
Box 2: Secondary server instance
Each of the three secondary server instances runs its own copy job to copy the primary log-backup file to its own local destination folder.
copy job: A SQL Server Agent job that copies the backup files from the primary server to a configurable destination on the secondary server and logs history on the secondary server and the monitor server. When log shipping is enabled on a database, the job category "Log Shipping Copy" is created on each secondary server in a log shipping configuration.
Box 3: Secondary server instance.
Each secondary server instance runs its own restore job to restore the log backup from the local destination folder onto the local secondary database.
restore job: A SQL Server Agent job that restores the copied backup files to the secondary databases.

It logs history on the local server and the monitor server, and deletes old files and old history information. When log shipping is enabled on a database, the job category "Log Shipping Restore" is created on the secondary server instance.
References: https://docs.microsoft.com/en-us/sql/database-engine/log-shipping/about-log-shippingsql-s erver

Question No: 44 HOTSPOT
You manage a Microsoft SQL Server instance. You have a user named User1.
You need to grant the minimum permissions necessary to allow User1 to review audit logs. For each action, which option should you use? To answer, select the appropriate options in the answer area.

Answer Area

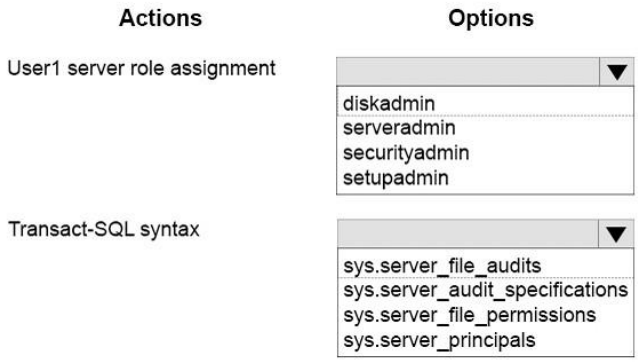

Answer:

Answer Area

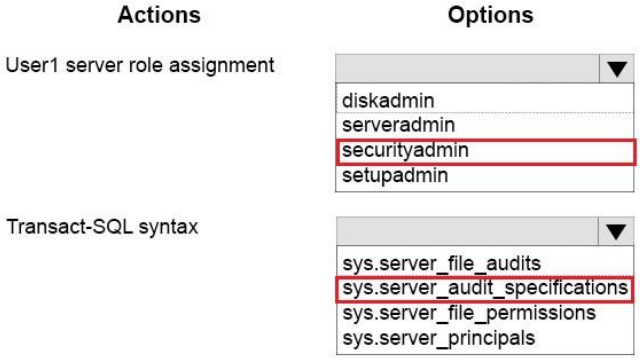

Explanation:
Box 1: securityadmin
To access log files for instances of SQL Server that are online, this requires membership in the securityadmin fixed server role. Box 2: sys.server audit_specifications
sys.server_audit_specifications contains information about the server audit specifications in a SQL Server audit on a server instance.

Test Guide 70-764

Question No: 45
You administer a Microsoft SQL Server database named Contoso. You create a stored procedure named Sales.Reviewlnvoice by running the following Transact-SOL statement:

```
CREATE PROCEDURE Sales.ReviewInvoice (@SaleID int)
AS
     DECLARE @tsql nvarchar(4000) = 'SELECT SaleID, CustomerID, TotalAmount FROM Sales.SalesIn-
voice WHERE SaleID = '
     SET @tsql = @tsql + CAST(@saleID AS varchar(20))
     EXEC sp_executesql @TSQL
```

You need to create a Windows-authenticated login named ContosoSearch and ensure that ContosoSearch can run the Sales.Reviewlnvoices stored procedure.
Which three Transact-SOL segments should you use to develop the solution? To answer, move the appropriate Transact-SOL segments from the list of Transact-SOL segments to the answer area and arrange them in the correct order.

Transact-SQL segments

```
Use Contoso
GO
CREATE USER Contoso\SalesGroup FOR
LOGIN
Contoso\SalesGroup
```

```
ALTER ROLE db_ddladmin ADD MEMBER
Contoso\SalesGroup
GRANT VIEW SEFINITION ON Sales.-
SalesInvoice TO
Contoso\SalesGroup
```

```
use master
CREATE LOGIN Contoso\SalesGroup FROM
WINDOWS
GO
```

```
GRANT EXECUTE ON Sales.ReviewInvoice TO
Contoso\SalesGroup
GRANT SELECT ON Sales.SalesInvoice TO
Contoso\SalesGroup
```

```
use master
CREATE LOGIN Contoso\ContosoSearch WITH
PASSWORD=N'Pa$$w0rd'
GO
```

```
GRANT EXECUTE ON Sales.ReviewInvoice TO
Contoso\SalesGroup
GRANT VIEW DEFINITION ON Sales.SalesIn-
voice TO
Contoso\SalesGroup
```

```
GRANT EXECUTE, SELECT ON Sales.Review-
Invoice TO
Contoso\SalesGroup
```

Answer Area

Answer:

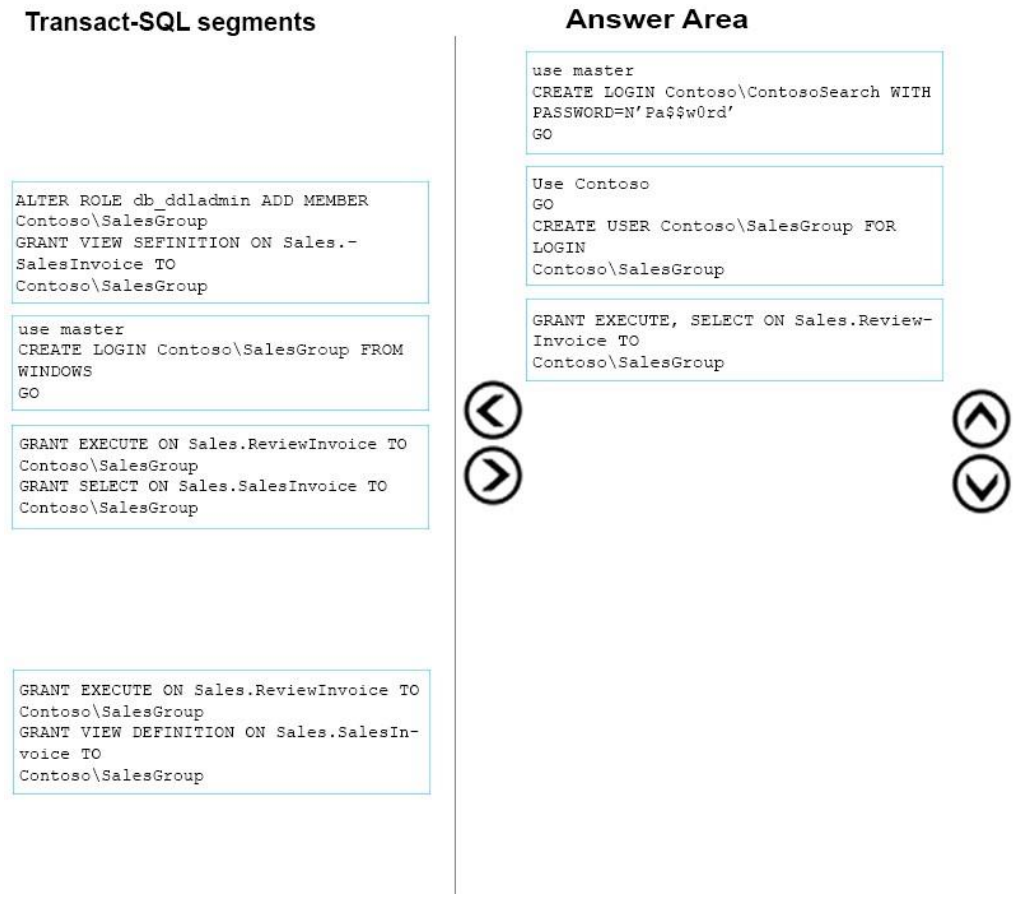

Question No: 46
You have a database that stores information for a shipping company. You create a table named Customers by running the following Transact-SOL statement. (Line numbers are included for reference only.)

```
01  CREATE TABLE dbo.Customers (
02      customerId int,
03      customerName varchar(200),
04      salesPerson varchar(20)
05  )
06  CREATE FUNCTION fn_securitypredicateSalesPerson (@salesPerson sysname)
07
08  AS
09  RETURN SELECT 1 AS [fn_securityPredicateOrder_result]
10  FROM dbo.Customers
11  WHERE @salesPerson = user_name()
```

You need to ensure that salespeople can view data only for the customers that are assigned to them.
Which Transact-SOL segment should you insert at line 07?

A. RETURNS varchar(20)WITH Schemabinding

B. RETURNS dbo.CustomersORDER BY @salesPerson

C. RETURNS tableORDER BY @salesPerson

D. RETURNS tableWITH Schemabinding

Answer: D

Explanation:
The return value can either be a scalar (single) value or a table.
SELECT 1 just selects a 1 for every row, of course. What it's used for in this case is testing whether any rows exist that match the criteria: if a row exists that matches the WHERE clause, then it returns 1, otherwise it returns nothing.
Specify the WITH SCHEMABINDING clause when you are creating the function. This ensures that the objects referenced in the function definition cannot be modified unless the function is also modified.
References: https://docs.microsoft.com/en-us/sql/t-sql/statements/create-function-transact-sql

Question No: 47
You manage a Microsoft SQL Server environment. You plan to encrypt data when you create backups.
You need to configure the encryption options for backups. What
should you configure?

A. a certificate

B. an MDS hash

C. a DES key

D. an AES 256-bit key

Answer: D

Explanation:
To encrypt during backup, you must specify an encryption algorithm, and an encryptor to secure the encryption key. The following are the supported encryption options:
Encryption Algorithm: The supported encryption algorithms are: AES 128, AES 192, AES 256, and Triple DES Encryptor: A certificate or asymmetric Key References:
https://docs.microsoft.com/en-us/sql/relational-databases/backup-restore/backupencryption

Question No: 48

You have a database named DB1 that stores more than 700 gigabyte (GB) of data and serves millions of requests per hour.
Queries on DB1 are taking longer than normal to complete.
You run the following Transact-SOL statement:
SELECT* FROM sys.database_query_store_options
You determine that the Query Store is in Read-Only mode.
You need to maximize the time that the Query Store is in Read-Write mode. Which Transact-SOL statement should you run?

A. ALTER DATABASE DB1SET QUERY_STORE (QUERY_CAPTURE_MODE = ALL)

B. ALTER DATABASE DB1SET QUERY_STORE (MAX_STORAGE_SIZE_MB = SO)

C. ALTER DATABASE DB1SET QUERY_STORE (CLEANUP _POLICY= (STALE_QUERY_THRESHOLD_DAYS = 14));

D. ALTER DATABASE DB1SET QUERY_STORE (QUERY_CAPTURE_MODE = NONE)

Answer: C

Explanation:
Stale Query Threshold (Days): Time-based cleanup policy that controls the retention period of persisted runtime statistics and inactive queries.
By default, Query Store is configured to keep the data for 30 days which may be unnecessarily long for your scenario.
Avoid keeping historical data that you do not plan to use. This will reduce changes to read-only status. The size of Query Store data as well as the time to detect and mitigate the issue will be more predictable. Use Management Studio or the following script to configure time-based cleanup policy:
ALTER DATABASE [QueryStoreDB]
SET QUERY_STORE (CLEANUP _POLICY= (STALE_QUERY_THRESHOLD_DAYS = 14));
References: https://docs.microsoft.com/en-us/sql/relational-databases/performance/best-practicewith-t he-query-store

Question No: 49
You deploy a Microsoft SQL Server instance to support a global sales application. The instance includes the following tables: TableA and TableB.
TableA is a partitioned table that uses an incrementing integer number for partitioning. The table has millions of rows in each partition. Most changes to the data in TableA affect recently added data. The UPDATE STATISTICS for TableA takes longer to complete than the allotted maintenance window.

Thousands of operations are performed against TableB each minute. You observe a large number of Auto Update Statistics events for Table 8.
You need to address the performance issues with each table.
In the table below, identify the action that will resolve the issues for each table. NOTE: Make only one selection in each column.

Answer Area

Action	TableA	TableB
Run the following Transact-SQL statement: `SET AUTO_UPDATE_STATISTICS_ASYNC ON`	○	○
Run the following Transact-SQL statement: `SET AUTO_UPDATE_STATISTICS OFF`	○	○
Run the following Transact-SQL statement and then recreate all indexes and statistics using the INCREMENTAL keyword: `SET AUTO_CREATE_STATISTICS on (INCREMENTAL = ON)`	○	○
Run the sp_updatestats procedure instead of the following Transact-SQL statement: `UPDATE STATISTICS`	○	○

Answer:

Answer Area

Action	TableA	TableB
Run the following Transact-SQL statement: `SET AUTO_UPDATE_STATISTICS_ASYNC ON`	○	◉
Run the following Transact-SQL statement: `SET AUTO_UPDATE_STATISTICS OFF`	◉	○
Run the following Transact-SQL statement and then recreate all indexes and statistics using the INCREMENTAL keyword: `SET AUTO_CREATE_STATISTICS on (INCREMENTAL = ON)`	○	○
Run the sp_updatestats procedure instead of the following Transact-SQL statement: `UPDATE STATISTICS`	○	○

Explanation:
Table A: Auto_update statistics off
Table A does not change much. There is no need to update the statistics on this table. Table B: SET AUTO_UPDATE_STATISTICS_ASYNC ON
You can set the database to update statistics asynchronously:
ALTER DATABASE YourDBName
SET AUTO_UPDATE_STATISTICS_ASYNC ON
If you enable this option then the Query Optimizer will run the query first and update the outdated statistics afterwards. When you set this option to OFF, the Query Optimizer will update the outdated statistics before compiling the query. This option can be useful in OLTP

environments References: https://www.mssqltips.com/sqlservertip/2766/sql-server-auto-update-and-auto-createstatisti cs-options/

Question No: 50
You administer a database that is used for reporting purposes. The database has a large fact table that contains three hundred million rows. The table includes a clustered columnstore index and a nonclustered index on the ProductlD column. New rows are inserted into the table every day.
Performance of queries that filter the Product ID column have degraded significantly.
You need to improve the performance of the queries.
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.

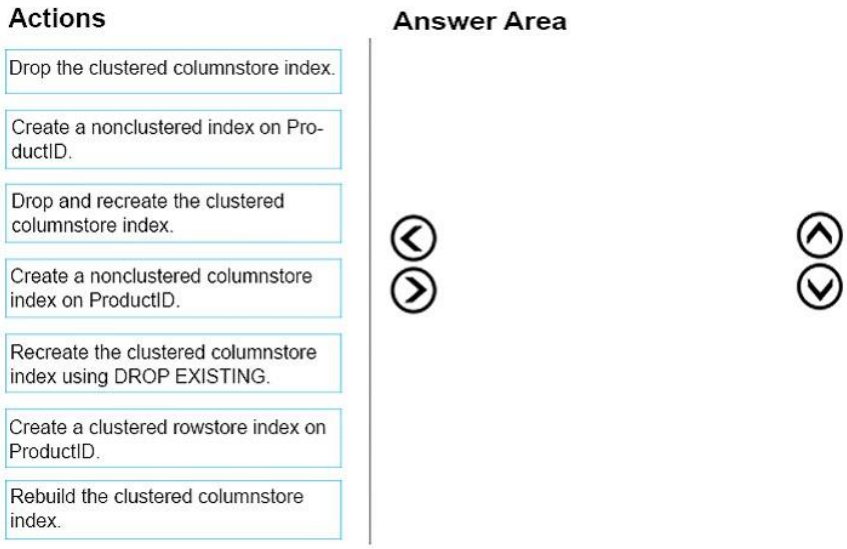

Answer:

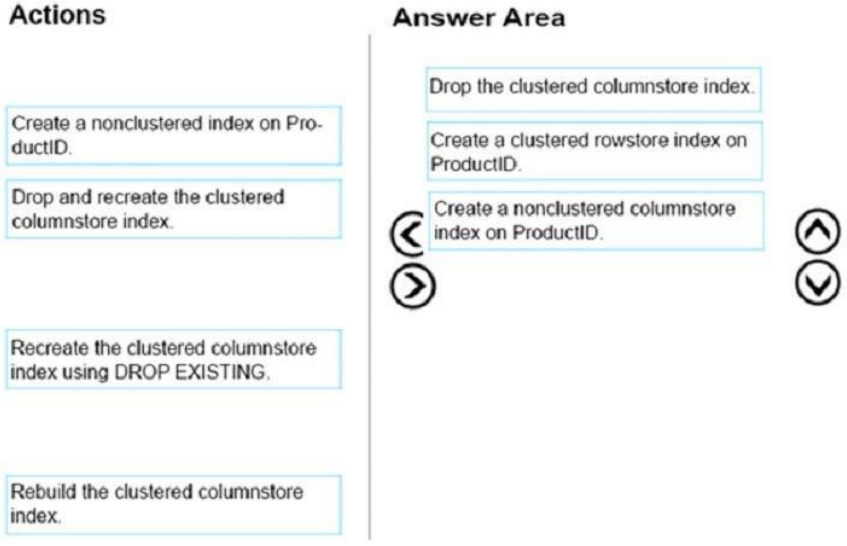

Explanation:
Step 1: Drop the clustered columnstore index
Step 2: Create a clustered rowstore index on ProductlD.
Rowstore indexes perform best on queries that seek into the data, searching for a particular value, or for queries on a small range of values. Use rowstore indexes with transactional workloads since they tend to require mostly table seeks instead of table scans. Step 3: Create a nonclustered index on ProductlD

Question No: 51
You are the database administrator of a Microsoft SQL Server instance. Developers are writing stored procedures to send emails using sp_send_dbmail. Database Mail is enabled.
You need to configure each account's profile security and meet the following requirements:
In the table below. identify the profile type that must be used for each account. NOTE:
Make only one selection in each column.

Answer Area

Profile type	SMTP1_Account	SMTP2_Account
Private Profile	○	○
Public Profile	○	○
Default Profile	○	○

Answer:

Answer Area

Profile type	SMTP1_Account	SMTP2_Account
Private Profile	○	○
Public Profile	○	○
Default Profile	○	○

Explanation:
SMTPl_Accountl: Private Profile
When no profile_name is specified, sp_send_dbmail uses the default private profile for the current user. If the user does not have a default private profile, sp_send_dbmail uses the default public profile for the msdb database.
SMTP1_Account2: Default Profile
Execute permissions forsp_send_dbmail default to all members of the DatabaseMailUser database role in the msdb database.
References: https://docs.microsoft.com/en-us/sql/relational-databases/system-storedprocedures/sp-sen d-dbmail-transact-sql

Test Guide 70-764

Question No: 52
You have a Microsoft SQL Server instance that hosts a database named DB1 that contains 800 gigabyte (GB) of data. The database is used 24 hours each day. You implement indexes and set the value of the Auto Update Statistics option set to True.
Users report that queries take a long time to complete.
You need to identify statistics that have not been updated for a week for tables where more than 1,000 rows changed.
How should you complete the Transact-SOL statement? To answer, configure the appropriate Transact-SOL segments in the answer area.

Answer Area

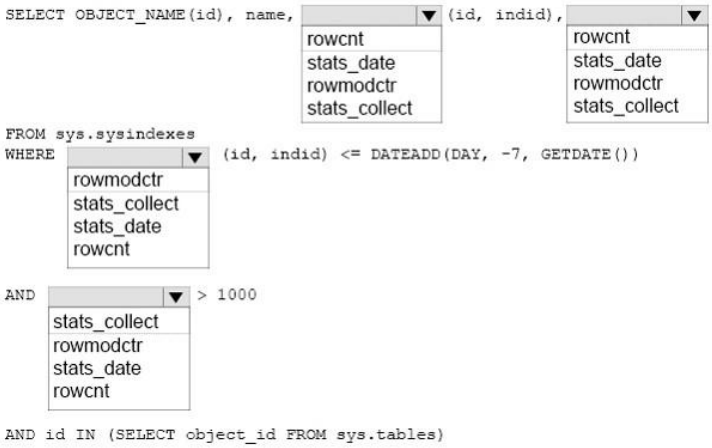

Answer:

Answer Area

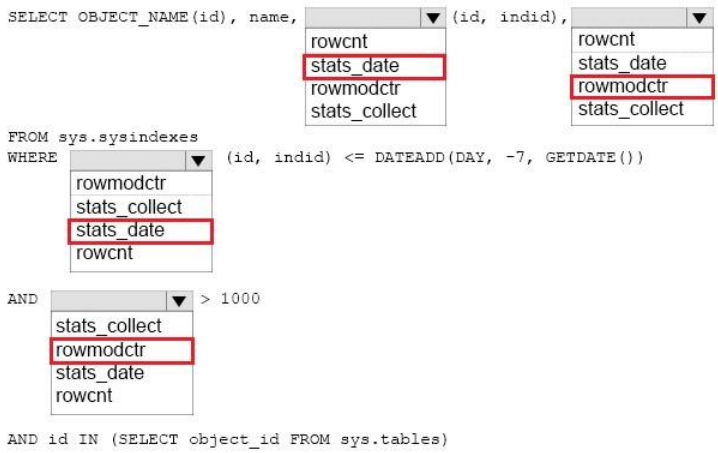

Explanation:
Box 1: stats_date
See example below. Box
2: rowmodctr

See examplebelow.
Box 3: stats_date
You need to identify statistics that have not been updated for a week.
Box 4: rowmodctr
You need to identify that more than 1,000 rows changed.
Rowmodctr counts the total number of inserted, deleted, or updated rows since the last time statistics were updated for the table.
Example: We will query every statistics object which was not updated in the last day and has rows modified since the last update. We will use the rowmodctr field of sys.sysindexes because it shows how many rows were inserted, updated or deleted since the last update occurred. Please note that it is not always 100% accurate in SQL Server 2005 and later, but it can be used to check if any rows were modified. Get the list of outdated statistics
SELECTOBJECT_NAME(id),name,STATS_DATE(id, indid),rowmodctr FROM
sys.sysindexes
WHERE STATS_DATE (id, indid)<=DATEADD(DAY,-1,GETDATE())
AND rowmodctr>O
AND id IN (SELECT object_id FROM sys.tables) GO
After collecting this information, we can decide which statistics require an update.

Question No: 53
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-SOL segments should you use to develop the solution? To answer, move the appropriate Transact-SOL segments from the list of Transact-SOL segments to the answer area and arrange them in the correct order.

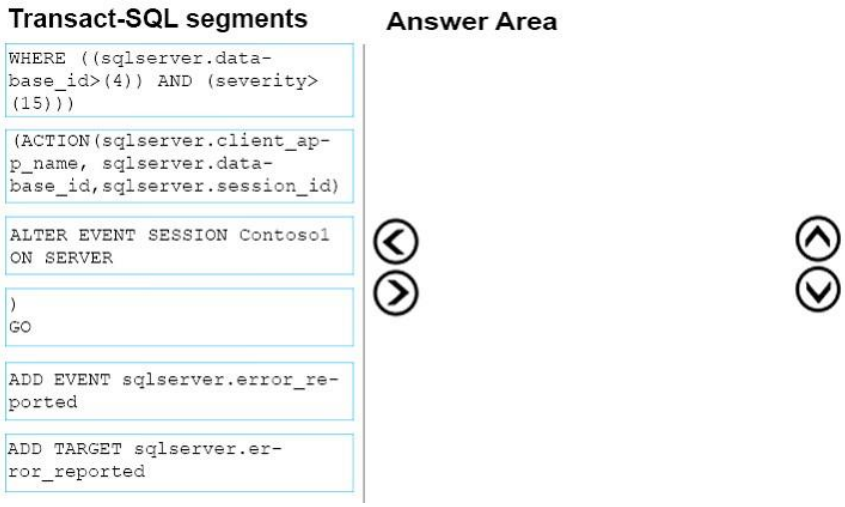

Answer:

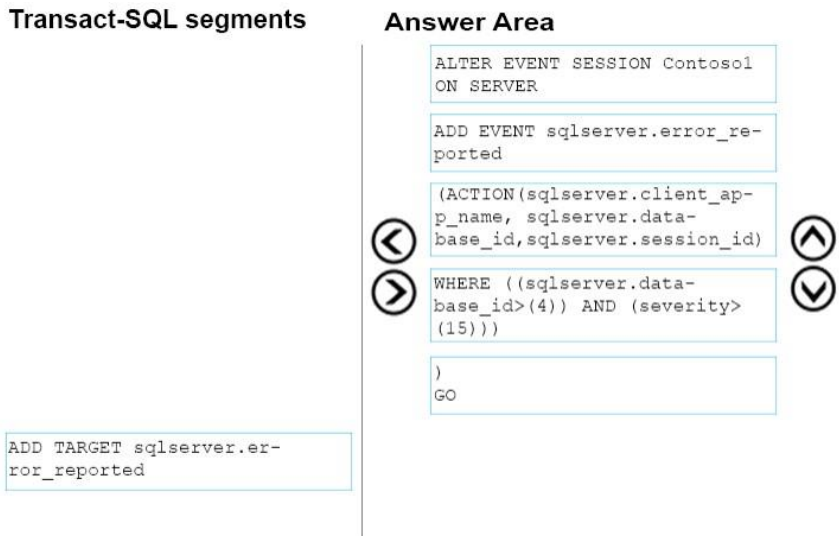

Explanation:
Step 1: ALTER EVENT SESSION Contosol ON SERVER Step
2: ADD EVENT ...
Step 3: (ACTION .. .
Step 4: WHERE .. .
Step 5:) GO
Example: To start an Extended Events sessions in order to trap SQL Server errors with severity greater than 10,just run the following script:

```
CREATE EVENT SESSION [error_trap] ON SERVER
ADD EVENT sqlserver.error _reported
(
ACTION
(packageO.collect_system_time,packageO.last_error,sqlserver.client_app_name,sqlserver.client
_host name,sqlserver.database_id,sqlserver.database_name,sqlserver.nt_username,
sqlserver.plan_handle,sqlserver.query_hash,sqlserver.session_id,sqlserver.sql_text,sqlserver.ts
ql_frame,sqlserver.tsql_stack,sqlserver.username)
WHERE ([severity]>10)
)
ADD TARGET packageO.event_file
(
SET filename=N'D:\Program Files\Microsoft SQL
Server\MSSQL11.MSSQLSERVER\MSSQL\XEvents\error_trap.xel'
)
WITH
(
STARTUP STATE=OFF
)
GO
```

Test Guide 70-764

Question No: 54
You manage a Microsoft SQL Server environment. A server fails and writes the following event to the application event log:
MSG_AUDIT_FORCED_SHUTDOWN
You configure the SQL Server startup parameters as shown in the following graphic:

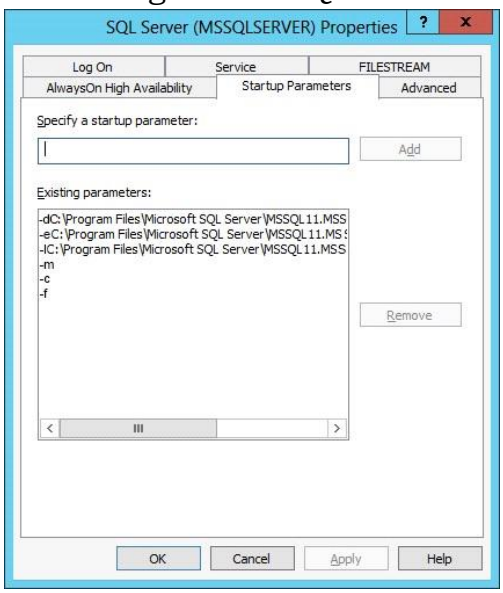

Use the drop-down menus to select the answer choice that answers each question. NOTE: Each correct selection is worth one point.

Answer Area

In which user mode will the SQL Server instance start?
- single-user
- multi-user
- restricted-user

With which server role can a local WIndows administrator connect to the database?
- public
- serveradmin
- sysadmin
- setupadmin

Answer:

Answer Area

In which user mode will the SQL Server instance start?
- single-user
- multi-user
- restricted-user

With which server role can a local WIndows administrator connect to the database?
- public
- serveradmin
- sysadmin
- setupadmin

Explanation:
Box 1: single-user
The startup option -m starts an instance of SQL Server in single-user mode.
Box 2: sysadmin
Starting SQL Server in single-user mode enables any member of the computer's local Administrators group to connect to the instance of SQL Server as a member of the sysadmin fixed server role.
References:
https://docs.microsoft.com/en-us/sql/database-engine/configure-windows/databaseengine-service-startup-options

Question No: 55
A company has an on-premises Microsoft SQL Server environment and Microsoft Azure SQL Database instanced. The environments host several customer databases.
You host a local database and a Stretch database that has a table named Members for one specific customer.
You need to provide the customer with information about the space used in the databases.
In the table below, identify the query that provides the required information for each database.
NOTE: Make only one selection in each column.

Answer Area

Query	Local database	Stretch database
`EXEC sp_spaceused @updateusage = N'TRUE'`	○	○
`EXEC sp_spaceused N'Company.Members'`	○	○
`EXEC sp_spaceused N'Company.Members', @mode = 'REMOTE_ONLY'`	○	○

Answer:

Answer Area

Query	Local database	Stretch database
`EXEC sp_spaceused @updateusage = N'TRUE'`	○	○
`EXEC sp_spaceused N'Company.Members'`	◉	○
`EXEC sp_spaceused N'Company.Members', @mode = 'REMOTE_ONLY'`	○	◉

Explanation:
Local database: EXEC sp_spaceused N'Company.Members'