



**Vendor: Microsoft**

**Exam Code: 70-764**

**Exam Name: Administering a SQL Database Infrastructure**

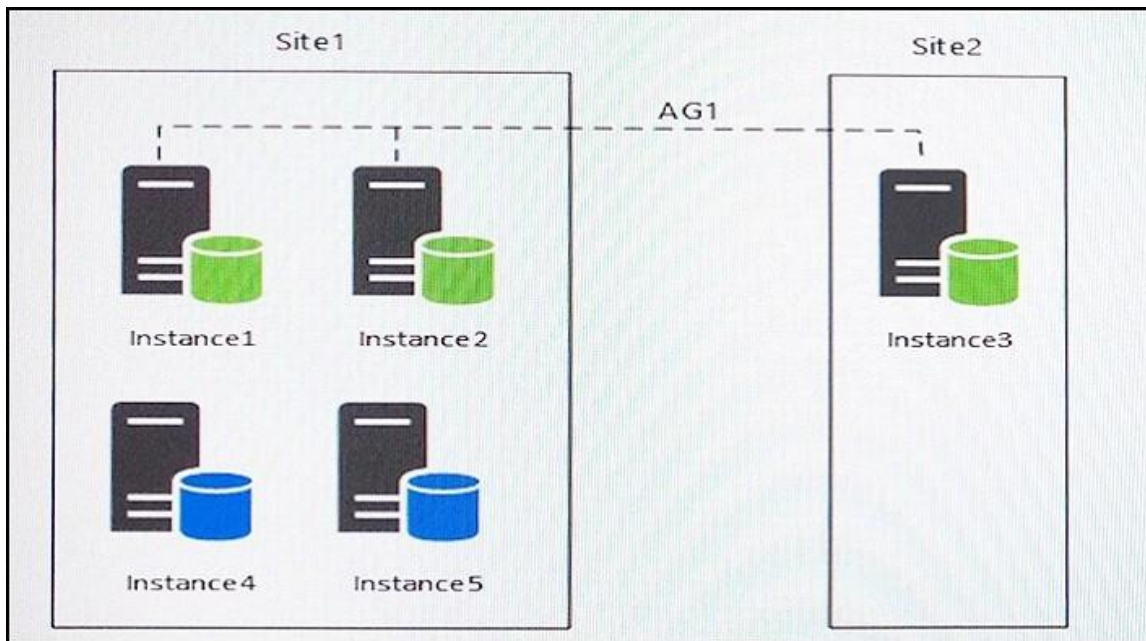
**Version: Demo**

**QUESTION 1**

**HOTSPOT**

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

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



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

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

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

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

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

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

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

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

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

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

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

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

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

You need to 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.	<input type="radio"/>	<input type="radio"/>
	Connect to the current primary replica SQL instance using ApplicationIntent=ReadOnly.	<input type="radio"/>	<input type="radio"/>
	Connect to any current read-only replica SQL instance.	<input type="radio"/>	<input type="radio"/>
	Connect to a Listener.	<input type="radio"/>	<input type="radio"/>
	Connect to the current primary replica SQL instance.	<input type="radio"/>	<input type="radio"/>

**Correct Answer:**

Answer Area	Option	Reporting system	Operations system
	Connect to a Listener using ApplicationIntent=ReadOnly.	<input type="radio"/>	<input type="radio"/>
	Connect to the current primary replica SQL instance using ApplicationIntent=ReadOnly.	<input type="radio"/>	<input type="radio"/>
	Connect to any current read-only replica SQL instance.	<input checked="" type="radio"/>	<input type="radio"/>
	Connect to a Listener.	<input type="radio"/>	<input type="radio"/>
	Connect to the current primary replica SQL instance.	<input type="radio"/>	<input checked="" type="radio"/>

**QUESTION 2**

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

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

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

**Correct 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\\_Administrati on/DBA\\_20.shtml](http://www.idevelopment.info/data/SQLServer/DBA_tips/Database_Administrati on/DBA_20.shtml)

**QUESTION 3**

**DRAG DROP**

You plan to deploy a database by using SQL Server 2016.

Your company identifies the following requirements for the database:

The name of all stored procedures must start with "usp\_"always. All distribution statistics must be updated daily.

You need to identify which feature must be used to meet each database requirement.

Which features should you identify?

To answer, drag the appropriate feature to the correct database requirement in the answer area.

Features	Answer Area
change data capture	The name of all stored procedures must start with "usp_"always. Feature
the CHECK constraint	All distribution statistics must be updated daily. Feature
Extended Event	
a maintenance plan	
Policy-Based Management	

**Correct Answer:**

Features	Answer Area
change data capture	The name of all stored procedures must start with "usp_"always. Policy-Based Management
the CHECK constraint	All distribution statistics must be updated daily. a maintenance plan
Extended Event	
a maintenance plan	
Policy-Based Management	

**QUESTION 4**

**DRAG DROP**

You use SQL Server 2016. You need to create a single object that inserts a provided value into Table1, and then returns a count of the records in Table1.

Develop the solution by selecting and arranging the required code blocks in the correct order. You may not need all of the code blocks.

Code Blocks	Answer Area
END	
CREATE FUNCTION dbo.Func_Table1 @InsertWord Varchar (10), @Var1 int OUTPUT AS BEGIN	
CREATE FUNCTION dbo.Func_Table1 (@InsertWord Varchar (10)) Returns INT AS BEGIN	
CREATE PROCEDURE dbo.Spr_Table1 @InsertWord Varchar (10), @Var1 INT AS BEGIN	
Select @Var1 = count(*) from TABLE1	
Declare @Var1 INT	
CREATE PROCEDURE dbo.Spr_Table1 @InsertWord Varchar (10), @Var1 int OUTPUT AS BEGIN	
INSERT INTO TABLE1 (FIELD1) values (@InsertWord)	

**Correct Answer:**

Code Blocks	Answer Area
END	CREATE PROCEDURE dbo.Spr_Table1 @InsertWord Varchar (10), @Var1 int OUTPUT AS BEGIN
CREATE FUNCTION dbo.Func_Table1 @InsertWord Varchar (10), @Var1 int OUTPUT AS BEGIN	INSERT INTO TABLE1 (FIELD1) values (@InsertWord)
CREATE FUNCTION dbo.Func_Table1 (@InsertWord Varchar (10)) Returns INT AS BEGIN	Select @Var1 = count(*) from TABLE1
CREATE PROCEDURE dbo.Spr_Table1 @InsertWord Varchar (10), @Var1 INT AS BEGIN	END
Select @Var1 = count(*) from TABLE1	
Declare @Var1 INT	
CREATE PROCEDURE dbo.Spr_Table1 @InsertWord Varchar (10), @Var1 int OUTPUT AS BEGIN	
INSERT INTO TABLE1 (FIELD1) values (@InsertWord)	

### QUESTION 5

You have a query that is used by a reporting dashboard.

Users report that the query sometimes takes a long time to run.

You need to recommend a solution to identify what is causing the issue.

What should you recommend?

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

- A. Set the blocked process threshold, and then run SQL Server Profiler.
- B. Set the blocked process threshold, and then create an alert.
- C. Enable trace flag 1204, and then create an alert.
- D. Create a job that queries the sys.dm\_os\_waiting\_tasks dynamic management view.

**Correct Answer: B**

#### **Explanation:**

Step 1: Turn on the blocked process report. This will look for any blocking taking 20 seconds or longer.

Make sure you don't have any pending changes

```
SELECT *  
FROM sys.configurations  
WHERE value <> value_in_use;  
GO
```

```
exec sp_configure 'show advanced options', 1;  
GO  
RECONFIGURE  
GO
```

```
exec sp_configure 'blocked process threshold (s)', 20; GO  
RECONFIGURE  
GO
```

Step 2: Set up a trace to capture the blocked process report. Run it as a server side trace.

Reference: blocked process threshold Server Configuration Option

### QUESTION 6

You have a server named Server1 that has 16 processors.

You plan to deploy multiple instances of SQL Server 2016 to Server1.

You need to recommend a method to allocate processors to each instance.

What should you include in the recommendation?

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

- A. Processor affinity
- B. Windows System Resource Manager (WSRM)
- C. Max Degree of Parallelism

D. Resource Governor

**Correct Answer:** A

**Explanation:**

CPU affinity management through Windows System Resource Manager is not recommended for SQL Server multi-instance management. Instead, use the processor affinity settings in SQL Server.

Reference: Server Properties (Processors Page)

#### **QUESTION 7**

You are building a stored procedure for a Windows Azure SQL Database. The procedure will add multiple rows to a table.

You need to design the stored procedure to meet the following requirements:

If any of the new rows violates a table constraint, then no further additions must be attempted and all changes made by the stored procedure must be discarded.

If any errors occur, a row must be added to an audit table, and the original error must be returned to the caller of the stored procedure.

What should you include in the design?

- A. An explicit transaction that has XACT\_ABORT disabled
- B. An implicit transaction that has error handling enabled
- C. An explicit transaction that has error handling enabled
- D. An implicit transaction that has XACT.ABORT enabled

**Correct Answer:** C

**Explanation:**

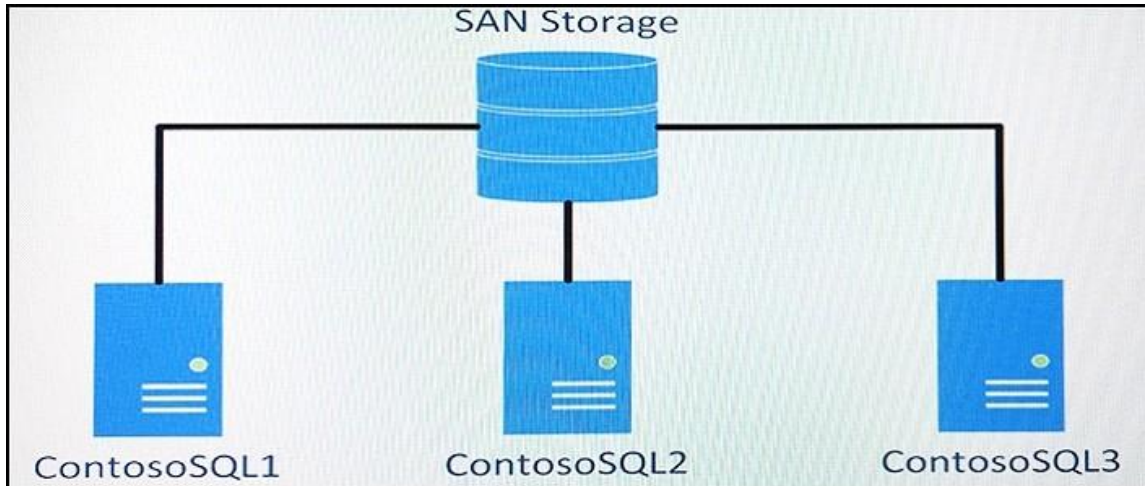
[http://technet.microsoft.com/en-us/library/ms175127\(v=SQL.105\).aspx](http://technet.microsoft.com/en-us/library/ms175127(v=SQL.105).aspx)

#### **QUESTION 8**

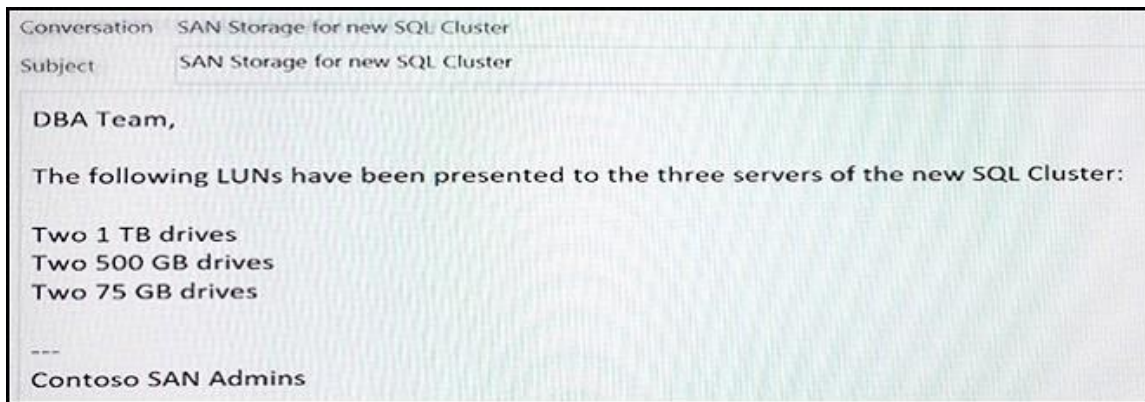
**HOTSPOT**

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.	<input type="radio"/>	<input type="radio"/>
One virtual network name for each SQL Server instance must be configured in the cluster.	<input type="radio"/>	<input type="radio"/>
The shared storage has been formatted and configured on ContosoSQL1.	<input type="radio"/>	<input type="radio"/>

Correct Answer:

## Answer Area

Statements	Yes	No
The Tempdb database for each cluster instance can be placed on the D: drive for the instance.	<input checked="" type="radio"/>	<input type="radio"/>
One virtual network name for each SQL Server instance must be configured in the cluster.	<input checked="" type="radio"/>	<input type="radio"/>
The shared storage has been formatted and configured on ContosoSQL1.	<input type="radio"/>	<input checked="" type="radio"/>

### QUESTION 9

#### HOTSPOT

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

You are a database administrator for a company that has an on-premises Microsoft SQL Server environment and Microsoft Azure SQL Database instances. The environment hosts several customer databases, and each customer uses a dedicated instance. The environments that you manage are shown in the following table.

Customer	Cloud Type	Description
AdventureWorks Cycles	Private	The environment includes a database named <b>Adventureworks</b> 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 <b>TSpinDB</b> . The application will monitor <b>TSpinDB</b> 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 <b>ConDB</b> that was recently upgraded to Microsoft SQL Server 2016. Contoso reports that <b>ConDB</b> is slow to return results when the server is busy. You must modify the startup parameters to <b>ConDB</b> to optimize performance.
Wingtip Toys	Private	Wingtip Toys has a database named <b>WingDB</b> . 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 <b>WingDB</b> . You must use minimal logging and minimized data loss during import process.
Wide World Importers	Public	The environment includes a database named <b>WDWDB</b> . Neither auditing nor statistics are configured for <b>WDWDB</b> . 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-SQL statement? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

## Answer Area

```
USE master
GO
```

	▼ AuditADUAccess
CREATE DATABASE AUDIT	
ALTER DATABASE AUDIT	
CREATE SERVER AUDIT	
ALTER SERVER AUDIT	

```
    TO FILE ( FILEPATH = 'C:\ADUAudit\' )
    WHERE object_name = 'SensitiveData'
```

```
GO
```

	▼ AuditADUAccess WITH (STATE = ON)
CREATE DATABASE AUDIT	
ALTER DATABASE AUDIT	
CREATE SERVER AUDIT	
ALTER SERVER AUDIT	

```
GO
```

```
Use Adventureworks
```

	▼ SPECIFICATION [FilterForSensitiveData]
CREATE DATABASE AUDIT	
ALTER DATABASE AUDIT	
CREATE SERVER AUDIT	
ALTER SERVER AUDIT	

	▼ [AuditADUAccess]
FOR SERVER AUDIT	
FOR DATABASE AUDIT	
USE [AuditDataAcces]	
SELECT ID	

```
ADD (SELECT ON SCHEMA::[ADUSchema] BY [public])
WITH (STATE = ON)
GO
```

Correct Answer:

## Answer Area

```
USE master
GO
```

```

AuditADUAccess
CREATE DATABASE AUDIT
ALTER DATABASE AUDIT
CREATE SERVER AUDIT
ALTER SERVER AUDIT

```

```

    TO FILE ( FILEPATH = 'C:\ADUAudit\' )
    WHERE object_name = 'SensitiveData'

```

```
GO
```

```

AuditADUAccess WITH (STATE = ON)
CREATE DATABASE AUDIT
ALTER DATABASE AUDIT
CREATE SERVER AUDIT
ALTER SERVER AUDIT

```

```
GO
```

```
Use Adventureworks
```

```

SPECIFICATION [FilterForSensitiveData]
CREATE DATABASE AUDIT
ALTER DATABASE AUDIT
CREATE SERVER AUDIT
ALTER SERVER AUDIT

```

```

[AuditADUAccess]
FOR SERVER AUDIT
FOR DATABASE AUDIT
USE [AuditDataAcces]
SELECT ID

```

```

ADD (SELECT ON SCHEMA::[ADUSchema] BY [public])
WITH (STATE = ON)
GO

```

**QUESTION 10**

**HOTSPOT**

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

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.	<div style="border: 1px solid black; padding: 2px;"><div style="background-color: #cccccc; padding: 2px; text-align: right;">▼</div><div style="padding: 2px;">Primary server instance Secondary server instance Monitor server instance Backup share file server</div></div>
Copy the backup job.	<div style="border: 1px solid black; padding: 2px;"><div style="background-color: #cccccc; padding: 2px; text-align: right;">▼</div><div style="padding: 2px;">Primary server instance Secondary server instance Monitor server instance Backup share file server</div></div>
Restore the backup.	<div style="border: 1px solid black; padding: 2px;"><div style="background-color: #cccccc; padding: 2px; text-align: right;">▼</div><div style="padding: 2px;">Primary server instance Secondary server instance Monitor server instance Backup share file server</div></div>

**Correct Answer:**

## Answer Area

Action	Server instance
Complete the backup job.	<div style="border: 1px solid gray; padding: 5px;"> <div style="background-color: #f0f0f0; padding: 2px; text-align: right;">▼</div> <ul style="list-style-type: none"> <li>Primary server instance</li> <li>Secondary server instance</li> <li>Monitor server instance</li> <li>Backup share file server</li> </ul> </div>
Copy the backup job.	<div style="border: 1px solid gray; padding: 5px;"> <div style="background-color: #f0f0f0; padding: 2px; text-align: right;">▼</div> <ul style="list-style-type: none"> <li>Primary server instance</li> <li>Secondary server instance</li> <li>Monitor server instance</li> <li>Backup share file server</li> </ul> </div>
Restore the backup.	<div style="border: 1px solid gray; padding: 5px;"> <div style="background-color: #f0f0f0; padding: 2px; text-align: right;">▼</div> <ul style="list-style-type: none"> <li>Primary server instance</li> <li>Secondary server instance</li> <li>Monitor server instance</li> <li>Backup share file server</li> </ul> </div>

### QUESTION 11

You are troubleshooting an application that runs a query. The application frequently causes deadlocks.

You need to identify which transaction causes the deadlock.

What should you do?

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

- A. Query the sys.dm\_exec\_sessions dynamic management view.
- B. Query the sys.dm\_exec\_requests dynamic management view.
- C. Create a trace in SQL Server Profiler that contains the Deadlock graph event
- D. Create an extended events session to capture deadlock information.

**Correct Answer:** D

#### Explanation:

Troubleshooting deadlocks

You have been receiving reports from users indicating that certain applications are returning deadlock errors. To maximize the effectiveness of troubleshooting these problems, you decide to focus on the deadlocks that are hit most frequently. You create an Extended Events session that:

Configures deadlock event tracking for the session.

Specifies a target that aggregates based on an identifier for the deadlock.

You run the Extended Events session, and after additional deadlocks are reported you are able to obtain aggregated deadlock information, along with the complete XML deadlock graph for each source. Using this information you are able to pin point the most common deadlocks and start working on a solution.

Create an Extended Events Session

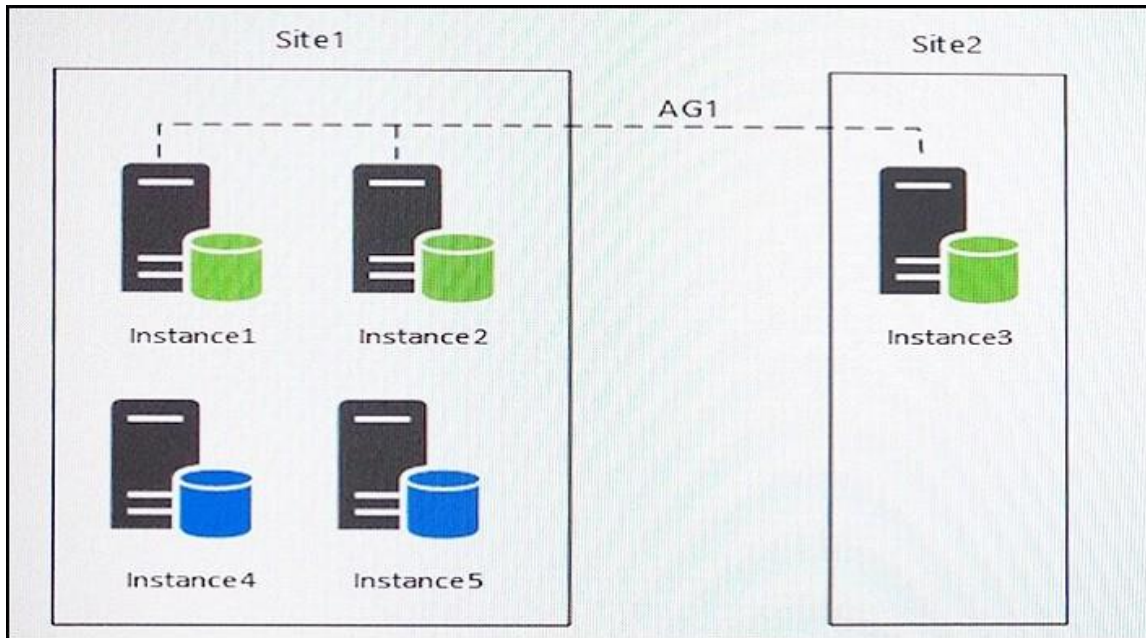
View Event Session Data

**QUESTION 12**

DRAG DROP

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

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



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

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

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

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



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

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

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

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

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

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

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

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

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

You need to configure a new replica of AG1 on Instance6.

How should you complete the Transact-SQL statement?

To answer, drag the appropriate Transact-SQL statements to the correct locations. Each Transact-SQL segment may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

## Transact-SQL segments

DATABASE
REPLICA
SYNCHRONOUS_COMMIT
ASYNCHRONOUS_COMMIT
PRIMARY
MANUAL
AUTOMATIC
SECONDARY_ONLY



---

### 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] );
```

**Correct Answer:**

## Transact-SQL segments

● ● ● ●

### Answer Area

```

ALTER AVAILABILITY GROUP AG_1 MODIFY  ON 'INSTANCE6'
WITH (AVAILABILITY_MODE =  );

ALTER AVAILABILITY GROUP AG_1 MODIFY  ON 'INSTANCE6'
WITH (FAILOVER_MODE =  );
    
```

### QUESTION 13

You are designing a monitoring application for a new SQL Server 2016 instance.

You need to recommend a solution to generate a report that displays the 10 most frequent wait types that occur for the instance.

What should you include in the recommendation?

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

- A. The SQL Server error log
- B. The sys.dm\_os\_wait\_stats dynamic management view
- C. The DBCC SQLPERF(WAITSTATS) command
- D. SQL Server Profiler

**Correct Answer:** B

**Explanation:**

sys.dm\_os\_wait\_stats

Returns information about all the waits encountered by threads that executed. You can use this aggregated view to diagnose performance issues with SQL Server and also with specific queries and batches.

Columns include:

waiting\_tasks\_count

Number of waits on this wait type. This counter is incremented at the start of each wait.

Reference:

sys.dm\_db\_wait\_stats (Windows Azure SQL Database)

#### QUESTION 14

Note: This question is part of a series of questions that use the same or similar answer choices. An answer choice may be correct for more than one question in the series. Each question is independent of the other questions in this series. Information and details provided in a question apply only to that question.

You 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

**Correct 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/encrypt-a-column-of-data>