Exam Questions 70-764

Administering a SQL Database Infrastructure (beta)

https://www.2passeasy.com/dumps/70-764/

NEW QUESTION 1
- (Exam Topic 1)
You administer a Microsoft SQL Server 2016 database that includes a table named Application.Events. Application.Events contains millions of records about user activity in an application.
Records in Application.Events that are more than 90 days old are purged nightly.
When records are purged, table locks are causing contention with inserts.
You need to be able to modify Application.Events without requiring any changes to the applications that utilize Application.Events.
Which type of solution should you use?

A. Partitioned tables
B. Online index rebuild
C. Change data capture
D. Change tracking

Answer: A

NEW QUESTION 2
- (Exam Topic 1)
You administer a single server that contains a Microsoft SQL Server 2016 default instance on which several production databases have been deployed.
You plan to install a new ticketing application that requires the deployment of a database on the server. The SQL login for this application requires sysadmin permissions. You need to ensure that the login for the
ticketing application cannot access other production databases.
What should you do?

A. Use the SQL Server default instance and enable Contained Databases.
B. Use the SQL Server default instance and configure a user-defined server rol
C. Add the login for the ticketing application to this role.
D. Install a new named SQL Server instance on the server.
E. Install a new default SQL Server instance on the server.

Answer: C

Explanation:
SQL Server supports multiple instances of SQL Server on a single server or processor, but only one instance can be the default instance. All others must be named instances. A computer can run multiple instances of SQL Server concurrently, and each instance runs independently of other instances.
References: https://msdn.microsoft.com/en-us/library/ms143531(v=SQL.105).aspx

NEW QUESTION 3
- (Exam Topic 1)
You administer a Microsoft SQL Server 2016 server that hosts a transactional database and a reporting database.
The transactional database is updated through a web application and is operational throughout the day. The reporting database is only updated from the transactional database.
The recovery model and backup schedule are configured as shown in the following table:

Database	Description
Transactional database	Recovery model: • Full Backup schedule: • Full database backup: midnight, daily • Differential database backup: on the hour, every two hours starting at 02:00 hours except at 00:00 hours • Log backup: every half hour, except at the times of full and differential backups
Reporting database	Recovery model: • Simple Backup schedule: • Full database backup: 01:00 hours daily • Differential database backup: 13:00 hours daily Data updates: • Changes in data are updated from the transactional database to the reporting database at 00:30 hours and at 12:30 hours • The update takes 15 minutes

At 16:20 hours, you discover that pages 17, 137, and 205 on one of the database files are corrupted on the transactional database. You need to ensure that the transactional database is restored. You also need to ensure that data loss is minimal.
What should you do?

A. Perform a partial restore.
B. Restore the latest full backup, and restore the latest differential backu
C. Then, restore each log backup taken before the time of failure from the most recent differential backup.
D. Perform a point-in-time restore.
E. Restore the latest full backup.
F. Restore the latest full backup, and restore the latest differential backu
G. Then, restore the latest log backup.
H. Perform a page restore.
I. Restore the latest full backu
J. Then, restore each differential backup taken before the time of failure from the most recent full backup.
K. Restore the latest full backu
L. Then, restore the latest differential backup.

Answer: F

Explanation:
The goal of a page restore is to restore one or more damaged pages without restoring the whole database. Typically, pages that are candidates for restore have been marked as "suspect" because of an error that is encountered when accessing the page.
Note: Requirements for Restoring Pages
A page restore is subject to the following requirements:
The databases must be using the full or bulk-logged recovery model. Etc.
References: https://docs.microsoft.com/en-us/sql/relational-databases/backup-restore/restore-pages-sql-server

NEW QUESTION 4
- (Exam Topic 1)
You administer a Microsoft SQL Server 2016 database.
You want to make a full backup of the database to a file on disk. In doing so, you need to output the progress of the backup. Which backup option should you use?

A. STATS
B. COMPRESSION
C. CHECKSUM
D. IN IT

Answer: A

NEW QUESTION 5
- (Exam Topic 1)
You administer two Microsoft SQL Server 2016 servers named ProdSrv1 and ProdSrv2. ProdSrv1 is configured as a Distributor.
Both servers are configured to use the Windows NT Service virtual accounts for all SQL Services.
You are configuring snapshot replication from ProdSrv1 to ProdSrv2 by using ProdSrv2 as a pull subscriber.
The distribution agent on ProdSrv2 regularly fails, displaying the following error message: "Cannot access the file. Operating system error code 5 (Access is denied.)."
You need to configure the distribution agent by granting only the minimum required access to all accounts. What should you do?

A. Configure the Subscriber to use the Local System account.
B. Configure the SQL Server Agent service to run under the Local System accoun
C. Configure the Subscriber to use the SQL Server Agent service account.
D. Configure the SQL Server Agent service to run under a Windows domain accoun
E. Configure the Subscriber to use the SQL Server Agent service accoun
F. Grant FULL CONTROL access for the domain account to the ReplData share on ProdSrv1.
G. Configure the Subscriber to use a Windows domain accoun
H. Grant READ access for the domain account to the ReplData share on ProdSrv1.

Answer: D

NEW QUESTION 6
- (Exam Topic 1)
Note: This question is part of a series of questions that use the same or similar answer choices. An answer choice may be correct for more than one question in the series. Each question is independent of the other questions in this series. Information and details provided in a question apply only to that question.
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

NEW QUESTION 7
- (Exam Topic 1)
You have a database that stores information for a shipping company. You create a table named Customers by running the following Transact-SQL 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-SQL 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

NEW QUESTION 8
- (Exam Topic 1)
Note: This question is part of a series of questions that use the same scenario. For your convenience, the scenario is repeated in each question. Each question presents a different goal and answer choices, but the text of the scenario is exactly the same in each question in this series.
You have five servers that run Microsoft Windows 2012 R2. Each server hosts a Microsoft SQL Server instance. The topology for the environment is shown in the following diagram.

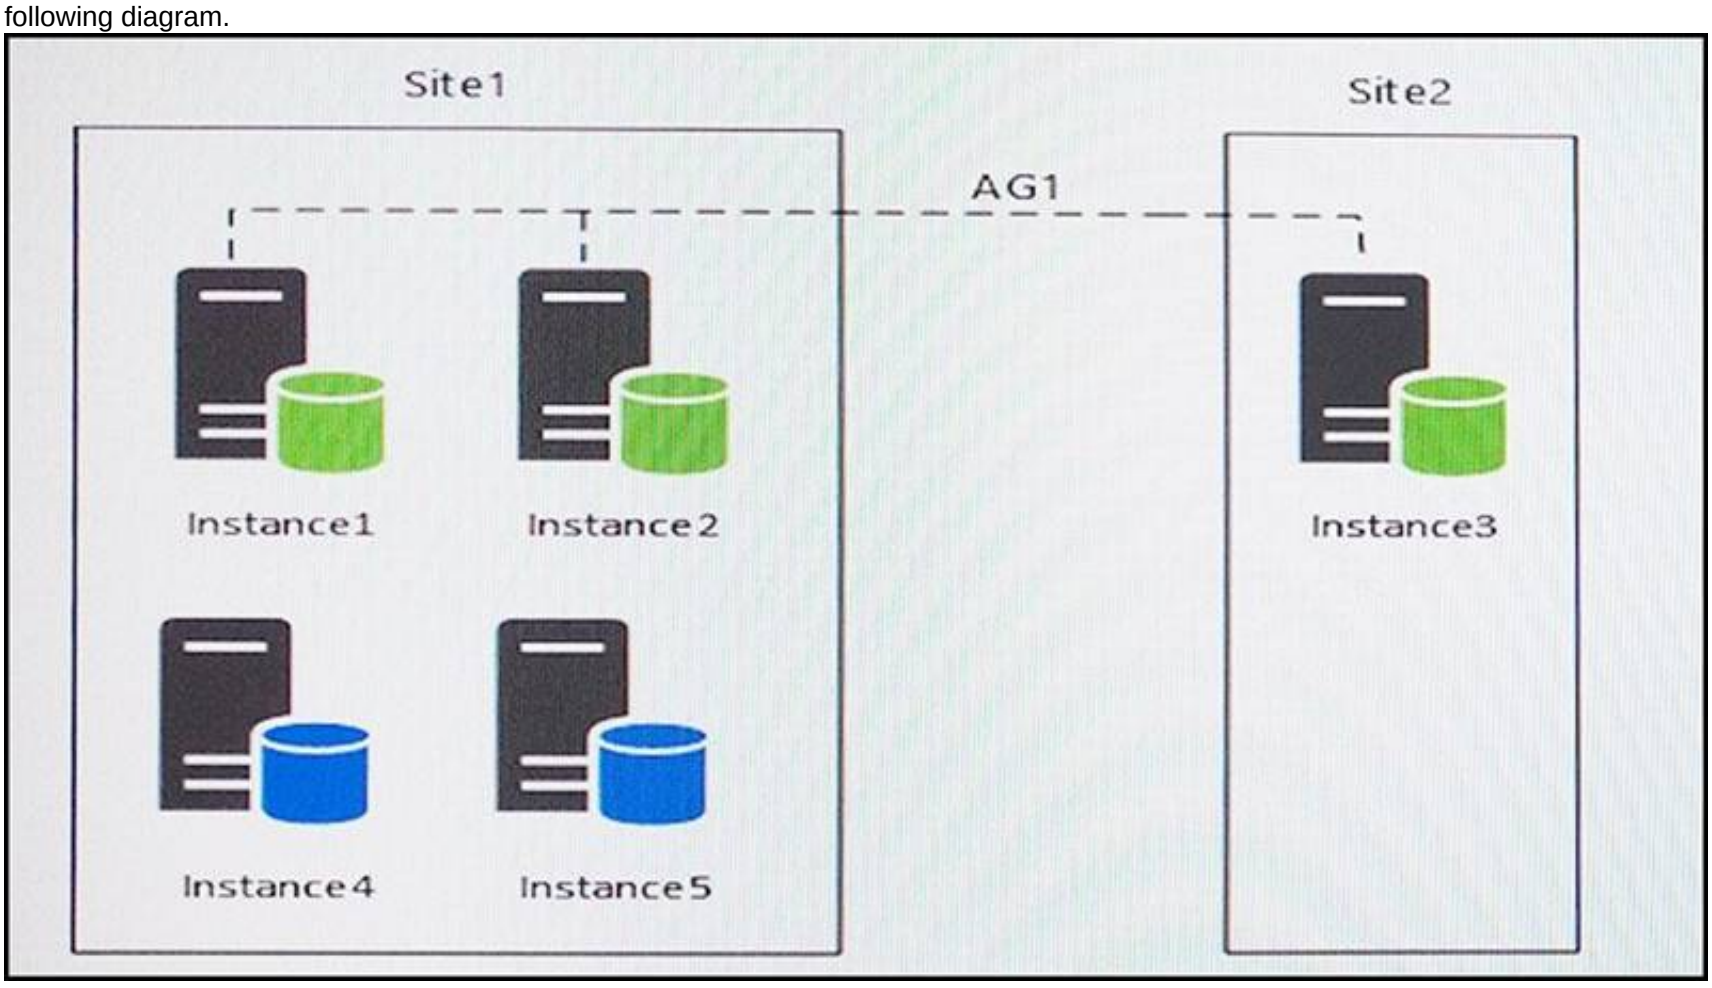

You have an Always On Availability group named AG1. The details for AG1 are shown in the following table.

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

Instance1 experiences heavy read-write traffic. The instance hosts a database named OperationsMain that is four terabytes (TB) in size. The database has multiple data files and filegroups. One of the filegroups is read_only and is half of the total database size.
Instance4 and Instance5 are not part of AG1. Instance4 is engaged in heavy read-write I/O.
Instance5 hosts a database named StagedExternal. A nightly BULK INSERT process loads data into an empty table that has a rowstore clustered index and two nonclustered rowstore indexes.
You must minimize the growth of the StagedExternal database log file during the BULK INSERT operations and perform point-in-time recovery after the BULK INSERT transaction. Changes made must not interrupt the log backup chain.
You plan to add a new instance named Instance6 to a datacenter that is geographically distant from Site1 and Site2. You must minimize latency between the nodes in AG1.
All databases use the full recovery model. All backups are written to the network location \\SQLBackup\. A separate process copies backups to an offsite location. You should minimize both the time required to restore the databases and the space required to store backups. The recovery point objective (RPO) for each instance is shown in the following table.

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

Full backups of OperationsMain take longer than six hours to complete. All SQL Server backups use the keyword COMPRESSION.
You plan to deploy the following solutions to the environment. The solutions will access a database named DB1 that is part of AG1.
Reporting system: This solution accesses data inDB1with a login that is mapped to a database user that is a member of the db_datareader role. The user has EXECUTE permissions on the database. Queries make no changes to the data. The queries must be load balanced over variable read-only replicas.
Operations system: This solution accesses data inDB1with a login that is mapped to a database user that is a member of the db_datareader and db_datawriter roles. The user has EXECUTE permissions on the database. Queries from the operations system will perform both DDL and DML operations.
The wait statistics monitoring requirements for the instances are described in the following table.

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

You need to propose a new process for the StagedExternal database.
Which five actions should you recommended be performed in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

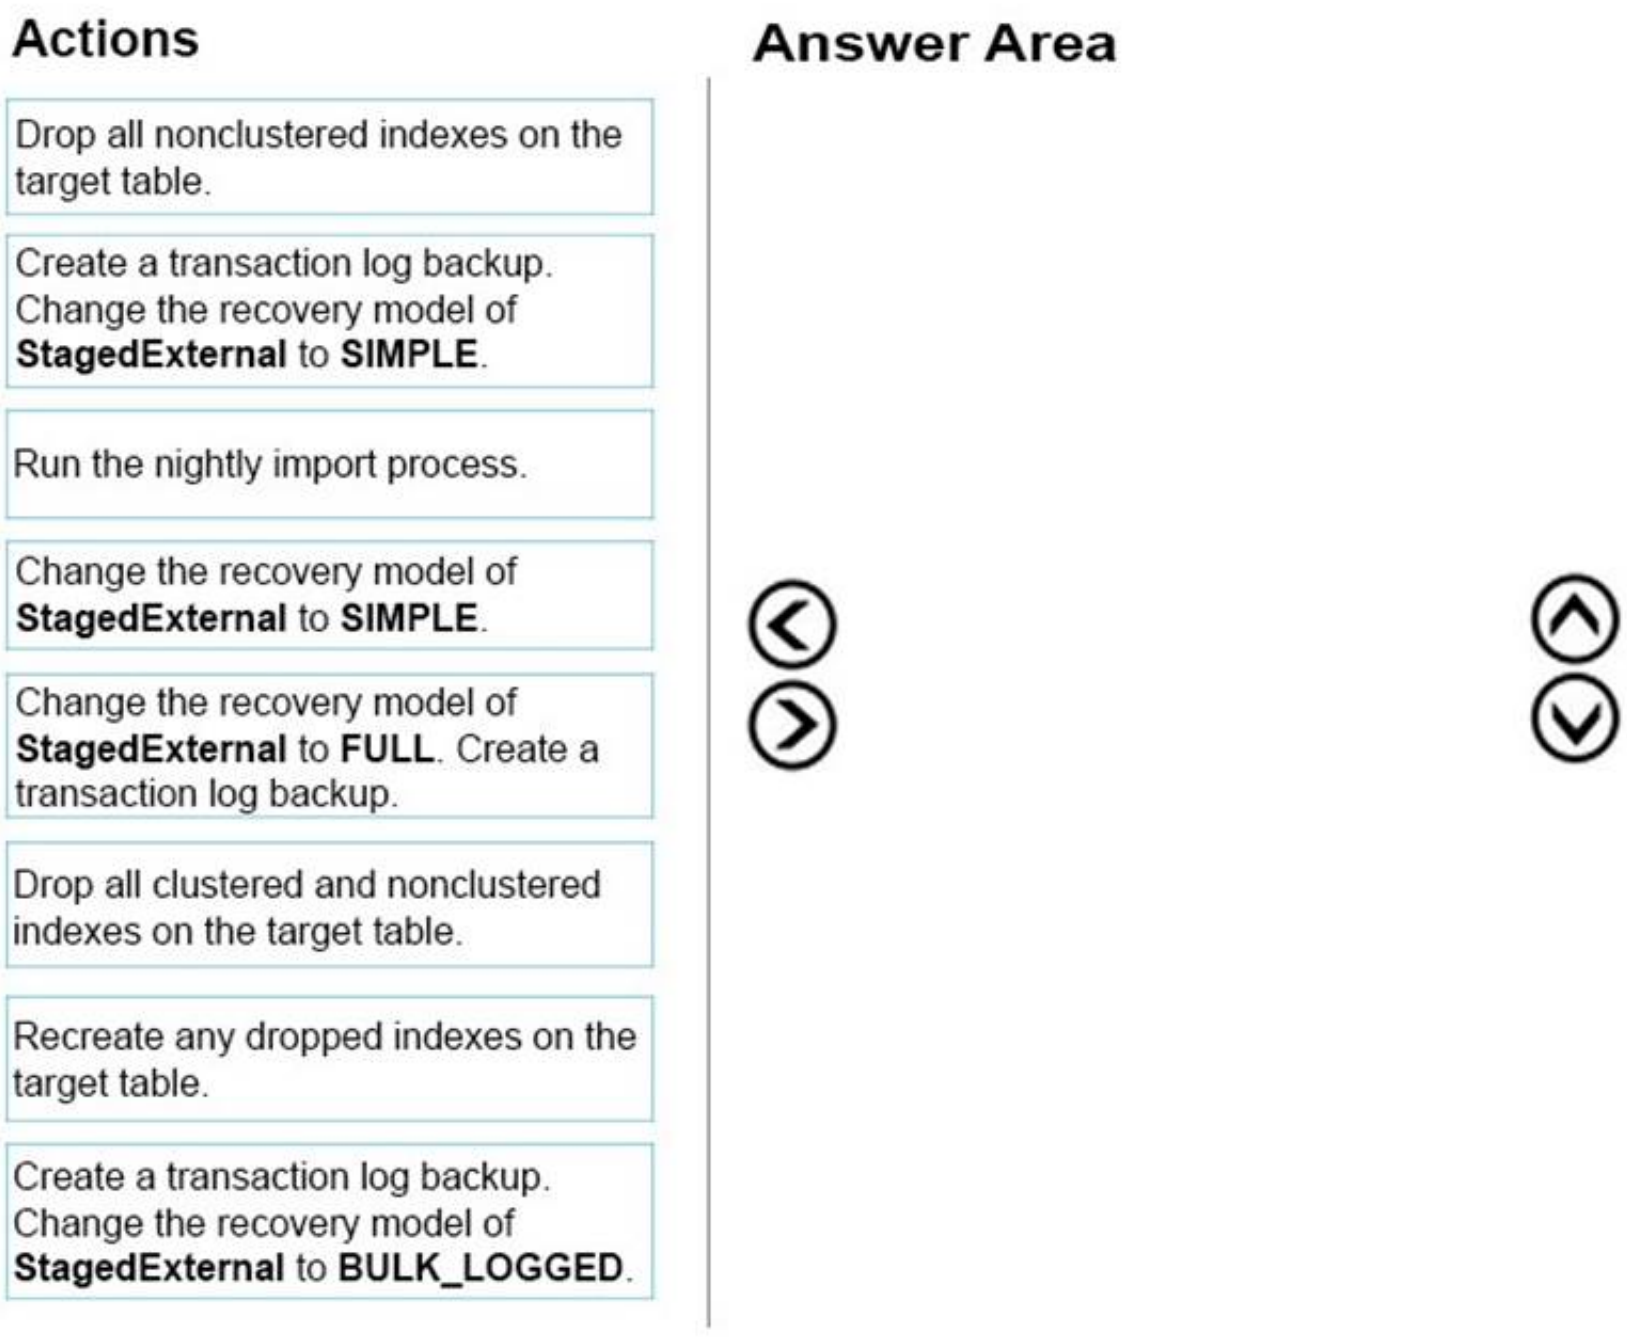

A. Mastered
B. Not Mastered

Answer: A

Explanation:
From scenario: Instance5 hosts a database named StagedExternal. A nightly BULK INSERT process loads data into an empty table that has a rowstore clustered index and two nonclustered rowstore indexes.
You must minimize the growth of the StagedExternaldatabase log file during the BULK INSERT operations and perform point-in-time recovery after the BULK INSERT transaction. Changes made must not interrupt the log backup chain.
All databases use the full recovery model.
References: https://technet.microsoft.com/en-us/library/ms190421(v=sql.105).aspx

NEW QUESTION 9
- (Exam Topic 1)
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You 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 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

NEW QUESTION 10
- (Exam Topic 1)
You are creating an application that will connect to the AgentPortal database by using a SQL login named AgentPortalUser. Stored procedures in the database will use sp_send_dbmail to send email messages.
You create a user account in the msdb database for the AgentPortalUser login.
You use the Database Mail Configuration Wizard to create a Database Mail profile. Security has not been configured for the Database Mail profile.
You need to ensure that AgentPortalUser can send email messages. What should you do?

A. In the Database Mail Configuration Wizard, configure the Database Mail profile as a private profile for the AgentPortalUser account.
B. Disable the guest user in the msdb database.
C. Use the sysmail_help_profileaccount_sp stored procedure to add accounts to the Database Mail profile.
D. In the Database Mail Configuration Wizard, create an email account for each recipient's email address in the Database Mail profile.

Answer: A

Explanation:
You enable and configure Database Mail using the Database Mail Configuration Wizard. Profiles are either public or private. A private profile is accessible only to specific users or roles.
References: https://docs.microsoft.com/en-us/sql/relational-databases/database-mail/configure-database-mail

NEW QUESTION 10
- (Exam Topic 1)
You plan to install a Microsoft SQL Server 2016 instance.
The instance will support a database that has the following requirements:
Store Excel workbooks on the file system.
Access the workbooks through Transact-SQL.
Include the workbooks in database backups.
During installation, you need to ensure that the requirements will be met.
Which feature should you use?

A. Excel Services
B. FILESTREAM
C. SQL Server Integration Services (SSIS)
D. OpenXML

Answer: B

NEW QUESTION 11
- (Exam Topic 1)
You plan to install Microsoft SQL Server 2016 for a web hosting company.
The company plans to host multiple web sites, each supported by a SQL Server database.
You need to select an edition of SQL Server that features backup compression of databases, basic data integration features, and low total cost of ownership.
Which edition should you choose?

A. Express Edition with Tools
B. Standard Edition
C. Web Edition
D. Express Edition with Advanced Services

Answer: B

Explanation:
Backup compression is supported on SQL Server 2016 editions: Enterprise, Standard, and Developer. References: https://docs.microsoft.com/en-us/sql/sql-server/editions-and-components-of-sql-server-2016

NEW QUESTION 15
- (Exam Topic 1)
Note: This question is part of a series of questions that use the same or similar answer choices. An answer choice may be correct for more than one question in the 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 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 permissionon the certificate. References:
https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/transparent-data-encryption-tde

NEW QUESTION 16
- (Exam Topic 1)
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
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-checkdb-transact-sql

NEW QUESTION 18
- (Exam Topic 1)
You administer a Microsoft SQL Server 2016 server.
When transaction logs grow, SQL Server must send an email message to the database administrators. You need to configure SQL Server to send the email messages.
What should you configure?

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

Answer: C

Explanation:
Operators are aliases for people or groups that can receive electronic notification when jobs have completed or alerts have been raised. The SQL Server Agent service supports the notification of administrators through operators. Operators enable notification and monitoring capabilities of SQL Server Agent.
References:
https://docs.microsoft.com/en-us/sql/relational-databases/database-mail/configure-sql-server-agent-mail-to-use-d

NEW QUESTION 19
- (Exam Topic 1)
You administer a Microsoft SQL Server 2016 default instance. The instance is hosted by a server that has a local firewall configured.
The firewall only allows inbound connections on port 1433. The server only hosts a single instance of SQL Server.
You need to ensure that the instance is configured to allow remote connections even if the SQL Server is unresponsive to client connections.
What should you do?

A. Enable inbound connections on TCP port 1434 in the Windows Firewall on the server.
B. Execute the following Transact-SQL command: sp_configure 'remote admin connections',
C. Execute the Reconfigure command.
D. Execute the following Transact-SQL command: sp_configure 'remote access', 1
E. Restart the SQL Server Agent Service.
F. Enable inbound connections on TCP port 135 in the Windows Firewall on the server.

Answer: ABC

Explanation:
SQL Server provides a dedicated administrator connection (DAC). The DAC lets an administrator access a running server to execute diagnostic functions or Transact-SQL statements, or to troubleshoot problems on the server, even when the server is locked or running in an abnormal state and not responding to a SQL Server Database Engine connection. By default, the DAC is only available from a client on the server. To enable client applications on remote computers to use the DAC, use the remote admin connections option of sp_configure.
By default, the DAC only listens on the loop-back IP address (127.0.0.1), port 1434 The following example enables the DAC from a remote computer.
sp_configure 'remote admin connections', 1; GO
RECONFIGURE; GO
References:
https://docs.microsoft.com/en-us/sql/database-engine/configure-windows/remote-admin-connections-server-con

NEW QUESTION 20
- (Exam Topic 1)
You install Microsoft SQL Server 2016 on a new server.
After setup is complete, you attempt to start the SQL Server service.
After being in a starting state for a few moments, the service goes back to a stopped state. You need to determine the cause of the failure. Which file should you

use?

A. %programfiles%\Microsoft SQLServer\MSSQL11.MSSQLSERVER|MSSQL\Log\Errorlog
B. %programfiles%\Microsoft SQL Server\110\setupBootstrap\Log\Summary.txt
C. %programfiles%\Microsoft SQL Server\MSSQL11.MSSQLSERVER\MSSQL\DATA\mastlog.idf
D. %programfiles%\Microsoft SQLServer\110\Shared\ErrorDmpr[XXXX] .mdmp

Answer: A

NEW QUESTION 22
- (Exam Topic 1)
Note: This question is part of a series of questions that use the same scenario. For your convenience, the scenario is repeated in each question. Each question presents a different goal and answer choices, but the text of the scenario is exactly the same in each question in this series.
You are a database administrator for a company that has an on-premises Microsoft SQL Server environment and Microsoft Azure SQL Database instances. The environment hosts several customer databases, and each customer uses a dedicated instance. The environments that you manage are shown in the following table.

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

You need to configure 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	○	○

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Monitoring from application: Transact-SQL
Transact-SQL can be used to monitor a customized application. Trend analysis: System Monitor
System Monitor can provide trend analysis. From question:

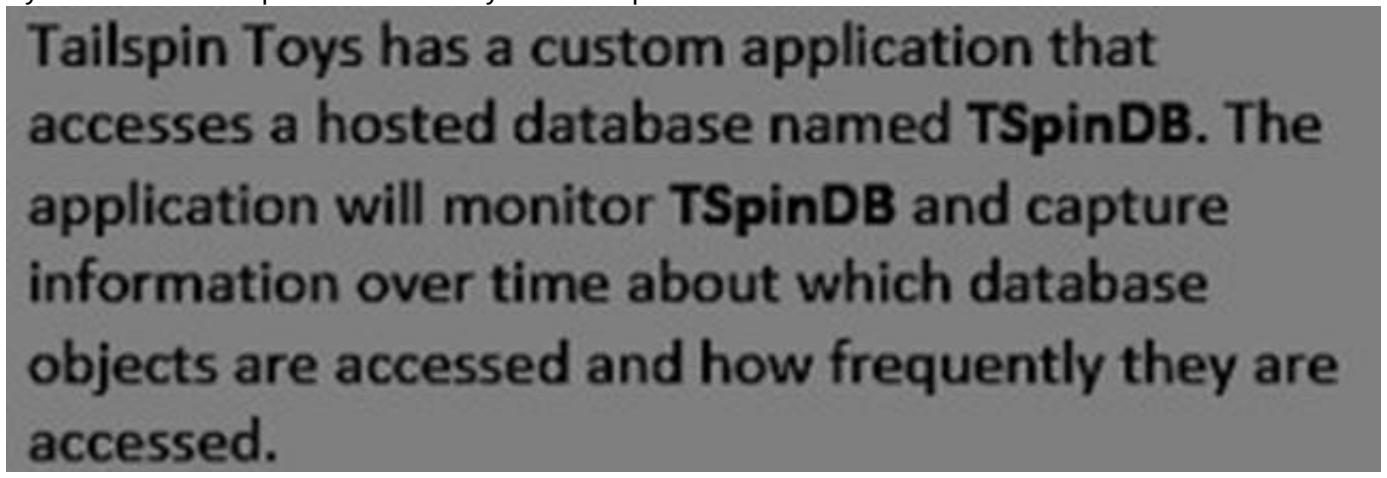

Tailspin Toys has a custom application that accesses a hosted database named TSpinDB. The application will monitor TSpinDB and capture information over time about which database objects are accessed and how frequently they are accessed.
References:
https://docs.microsoft.com/en-us/sql/relational-databases/performance/performance-monitoring-and-tuning-tools

NEW QUESTION 23
- (Exam Topic 1)
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:

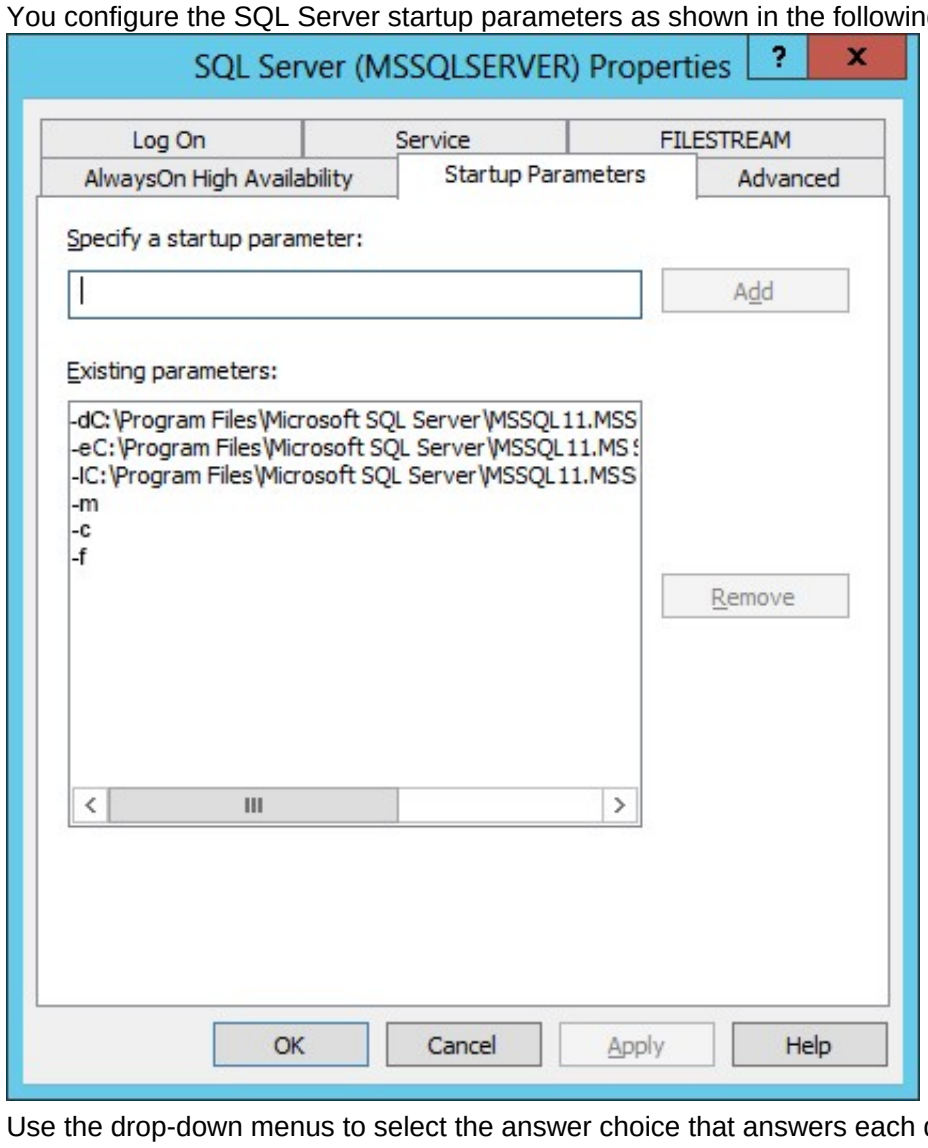

Use the drop-down menus to select the answer choice that answers each question. NOTE: Each correct selection is worth one point.

Answer Area

In which user mode will the SQL Server instance start?

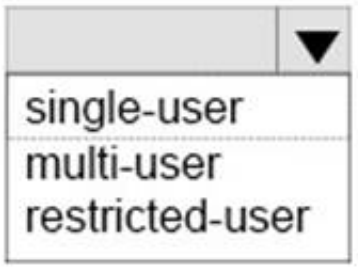

With which server role can a local WIndows administrator connect to the database?

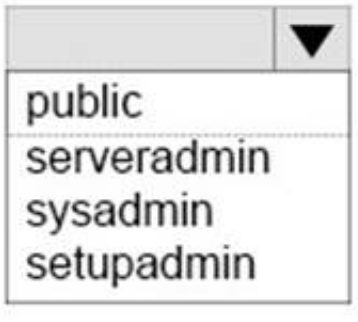

A. Mastered
B. Not Mastered

Answer: A

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 anymember 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/database-engine-service-startup-option

NEW QUESTION 24
- (Exam Topic 1)
You administer a Windows Azure SQL Database database named Inventory that contains a stored procedure named p_AddInventory.
Users need to be able to SELECT from all tables in the database and execute the stored procedure. You need to grant only the necessary permissions.
What should you do?

A. Grant EXECUTE permission on p_AddInventory to all user
B. Grant VIEW DEFINITION to all users.
C. Grant EXECUTE permission on p_AddInventory to all user
D. Add all users to the db_datawriter role.
E. Add all users to the db_owner role.
F. Grant EXECUTE permission on p_Add!nventory to all user
G. Add all users to the db_datareader role.

Answer: D

NEW QUESTION 26
- (Exam Topic 1)
You administer a Microsoft SQL Server 2016 database named Contoso on a server named Server01.
You need to track all SELECT statements issued in the Contoso database only by users in a role named Sales. What should you create?

A. An Alert
B. A Resource Pool
C. An Extended Event session
D. A Server Audit Specification
E. A SQL Profiler Trace
F. A Database Audit Specification
G. A Policy
H. A Data Collector Set

Answer: F

NEW QUESTION 31
- (Exam Topic 1)
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
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

NEW QUESTION 33
- (Exam Topic 1)
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:
Account SMTP1_Account must only be usable by logins that have been given explicit permissions to use the SMTP1_profile.
Account SMTP2_Account must only be usable by logins who are a member of the [DatabaseMailUserRole] role in msdb.
In the table below. identify the profile type that must be used for each account. NOTE: Make only one selection in each column.

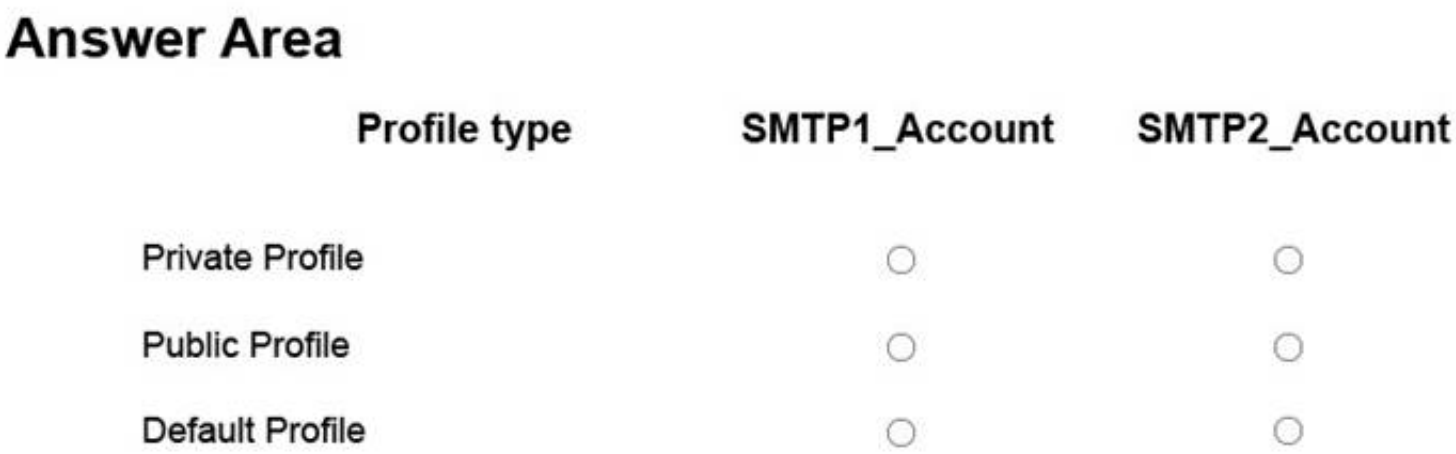

A. Mastered
B. Not Mastered

Answer: A

Explanation:
SMTP1_Account1: Private Profile
When no profile_name is specified, sp_send_dbmail uses the default private profile for the current user. I 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-stored-procedures/sp-send-dbmail-transact-sql

NEW QUESTION 38
- (Exam Topic 1)
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You 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 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

NEW QUESTION 41
- (Exam Topic 1)
Note: This question is part of a series of questions that use the same or similar answer choices. An answer choice may be correct for more than one question in the 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 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/backup-encryption

NEW QUESTION 43
- (Exam Topic 1)
You administer a Microsoft SQL Server 2016 database.
Users report that a billing application becomes unresponsive during busy times of the day. While investigating, you notice large number of processes taking or waiting for table locks. You suspect that SQL Server is assigning stronger locks to queries.
You start a SQL Profiler trace. Which event should you select?

A. Deadlock graph
B. Lock: Escalation
C. Lock: Timeout
D. Lock: Deadlock

Answer: B

NEW QUESTION 45
- (Exam Topic 1)
Note: This question is part of a series of questions that use the same scenario. For your convenience, the scenario is repeated in each question. Each question presents a different goal and answer choices, but the text of the scenario is exactly the same in each question in this series.
You have five servers that run Microsoft Windows 2012 R2. Each server hosts a Microsoft SQL Server instance. The topology for the environment is shown in the following diagram.

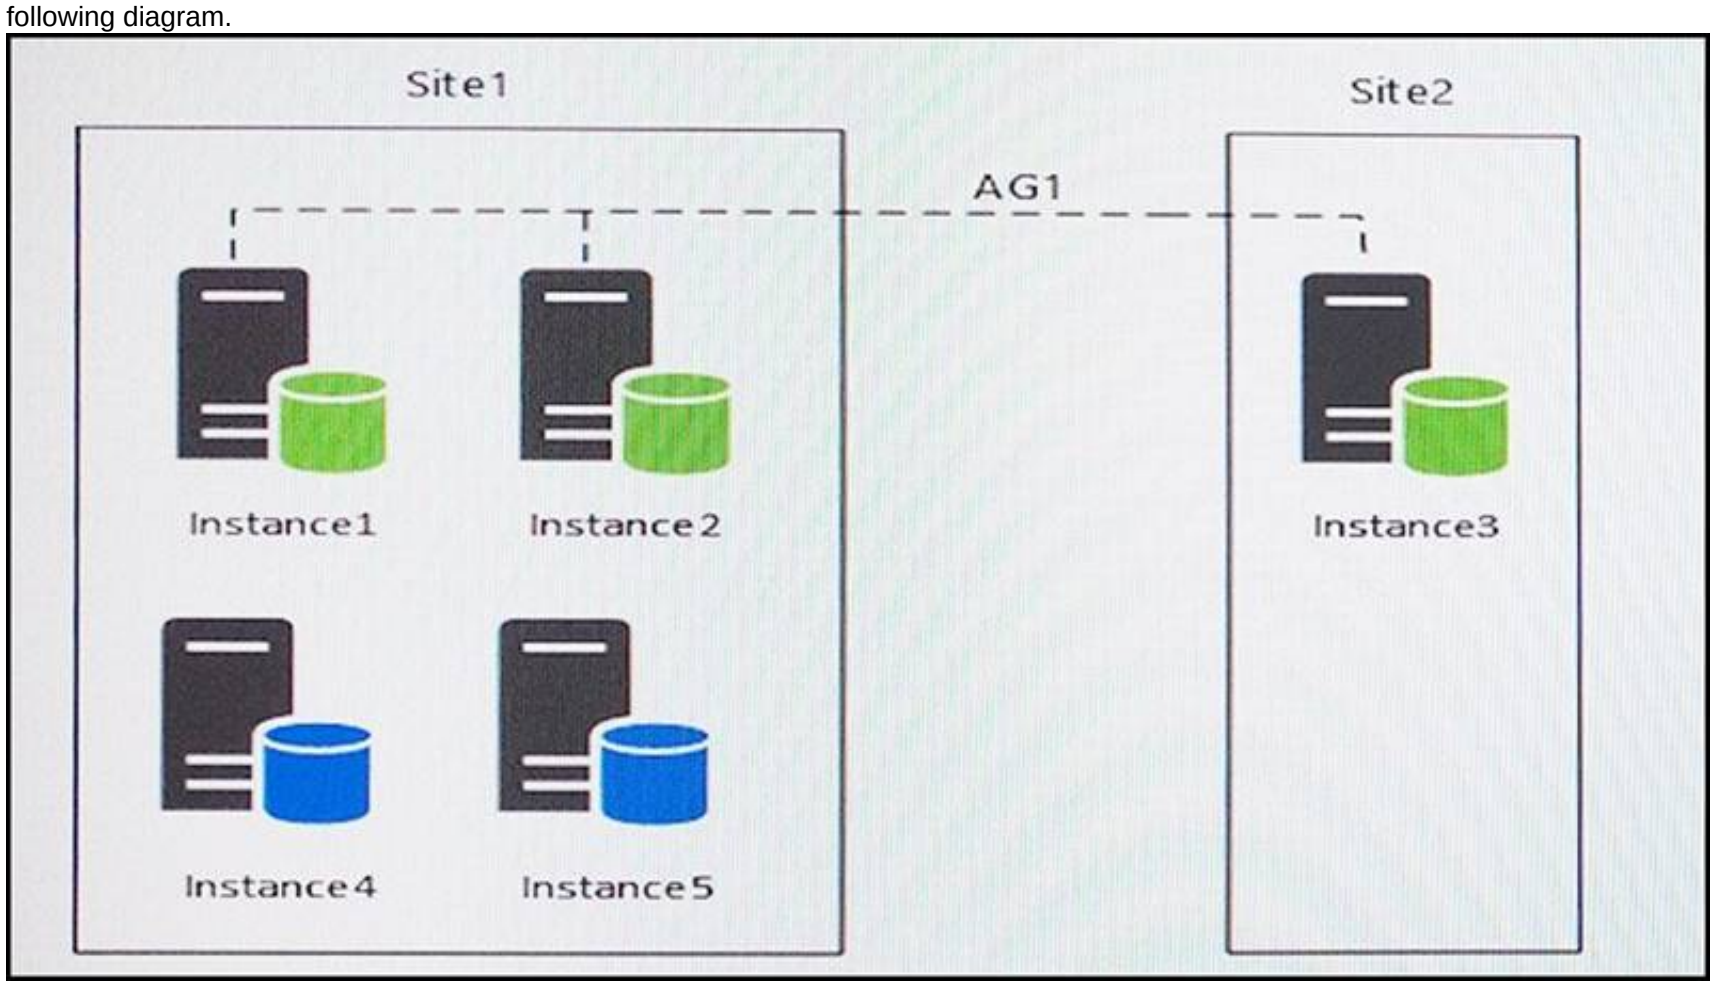

You have an Always On Availability group named AG1. The details for AG1 are shown in the following table.

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

Instance1 experiences heavy read-write traffic. The instance hosts a database named OperationsMain that is four terabytes (TB) in size. The database has multiple data files and filegroups. One of the filegroups is read_only and is half of the total database size.
Instance4 and Instance5 are not part of AG1. Instance4 is engaged in heavy read-write I/O.
Instance5 hosts a database named StagedExternal. A nightly BULK INSERT process loads data into an empty table that has a rowstore clustered index and two nonclustered rowstore indexes.

You must minimize the growth of the StagedExternal database log file during the BULK INSERT operations and perform point-in-time recovery after the BULK INSERT transaction. Changes made must not interrupt the log backup chain.
You plan to add a new instance named Instance6 to a datacenter that is geographically distant from Site1 and Site2. You must minimize latency between the nodes in AG1.
All databases use the full recovery model. All backups are written to the network location \\SQLBackup\. A separate process copies backups to an offsite location. You should minimize both the time required to restore the databases and the space required to store backups. The recovery point objective (RPO) for each instance is shown in the following table.

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

Full backups of OperationsMain take longer than six hours to complete. All SQL Server backups use the keyword COMPRESSION.
You plan to deploy the following solutions to the environment. The solutions will access a database named DB1 that is part of AG1.
Reporting system: This solution accesses data inDB1with a login that is mapped to a database user that is a member of the db_datareader role. The user has EXECUTE permissions on the database. Queries make no changes to the data. The queries must be load balanced over variable read-only replicas.
Operations system: This solution accesses data inDB1with a login that is mapped to a database user that is a member of the db_datareader and db_datawriter roles. The user has EXECUTE permissions on the database. Queries from the operations system will perform both DDL and DML operations.
The wait statistics monitoring requirements for the instances are described in the following table.

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

You need to configure a new replica of AG1 on Instance6.
How should you complete the Transact-SQL statement? To answer, drag the appropriate Transact-SQL statements to the correct locations. Each Transact-SQL segment may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

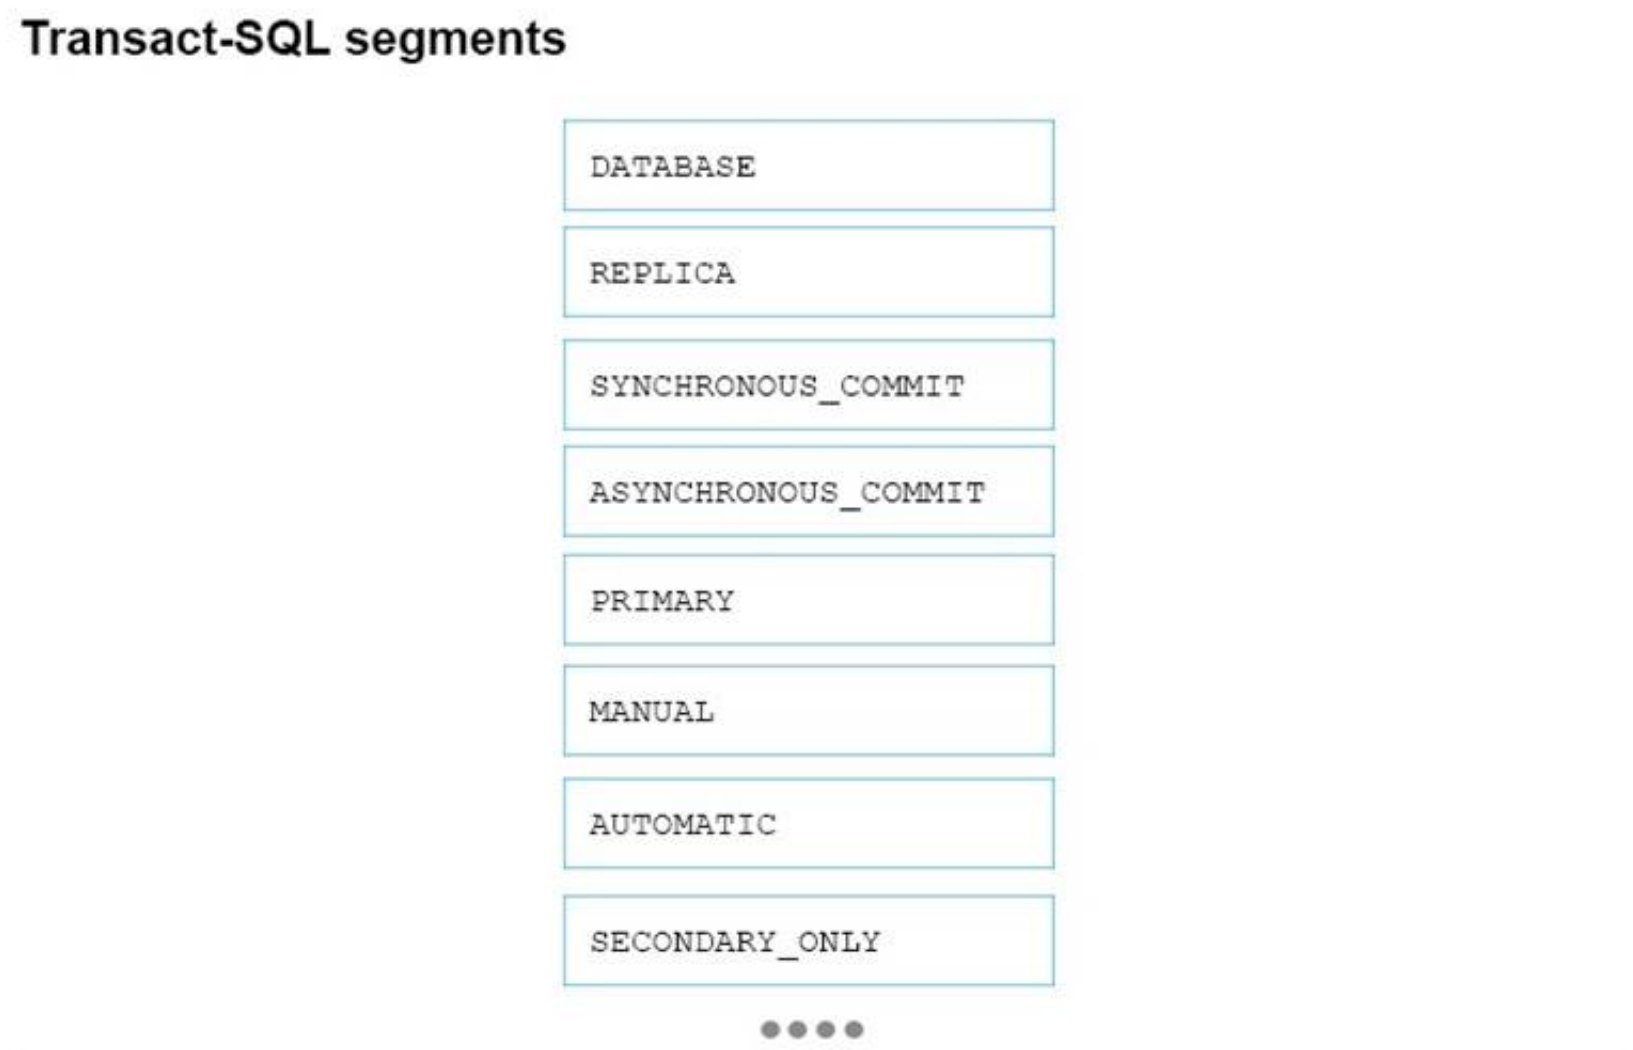

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Scenario: You plan to add a new instance named Instance6 to a datacenter that is geographically distant from Site1 and Site2. You must minimize latency between the nodes in AG1.
Box 1: REPLICA
MODIFY REPLICA ON modifies any of the replicas ofthe availability group. Box 2: SYNCHRONOUS_COMMIT
You must minimize latency between the nodes in AG1
AVAILABILITY_MODE = { SYNCHRONOUS_COMMIT | ASYNCHRONOUS_COMMIT }
Specifies whether the primary replica has to wait for the secondary availability group to acknowledge the hardening (writing) of the log records to disk before the primary replica can commit the transaction on a given primary database.
FAILOVER AUTOMATIC (box 4) requires SYNCHRONOUS_COMMIT Box 3: REPLICA
MODIFY REPLICA ON modifies any of the replicas of the availability group. Box 4: AUTOMATIC
You must minimize latency between the nodes in AG1 FAILOVER_MODE = { AUTOMATIC | MANUAL }
Specifies the failover mode of the availability replica that you are defining.
FAILOVER_MODE is required in the ADD REPLICA ON clause and optional in the MODIFY REPLICA ON clause.
AUTOMATIC enables automatic failover. AUTOMATIC is supported only if you also specify
AVAILABILITY_MODE = SYNCHRONOUS_COMMIT.
References: https://docs.microsoft.com/en-us/sql/t-sql/statements/alter-availability-group-transact-sql

NEW QUESTION 48
- (Exam Topic 1)
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 MD5 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/backup-encryption

NEW QUESTION 52
- (Exam Topic 1)
Note: This question is part of a series of questions that use the same scenario. For your convenience, the scenario is repeated in each question. Each question presents a different goal and answer choices, but the text of the scenario is exactly the same in each question in this series.
You have five servers that run Microsoft Windows 2012 R2. Each server hosts a Microsoft SQL Server instance. The topology for the environment is shown in the following diagram.

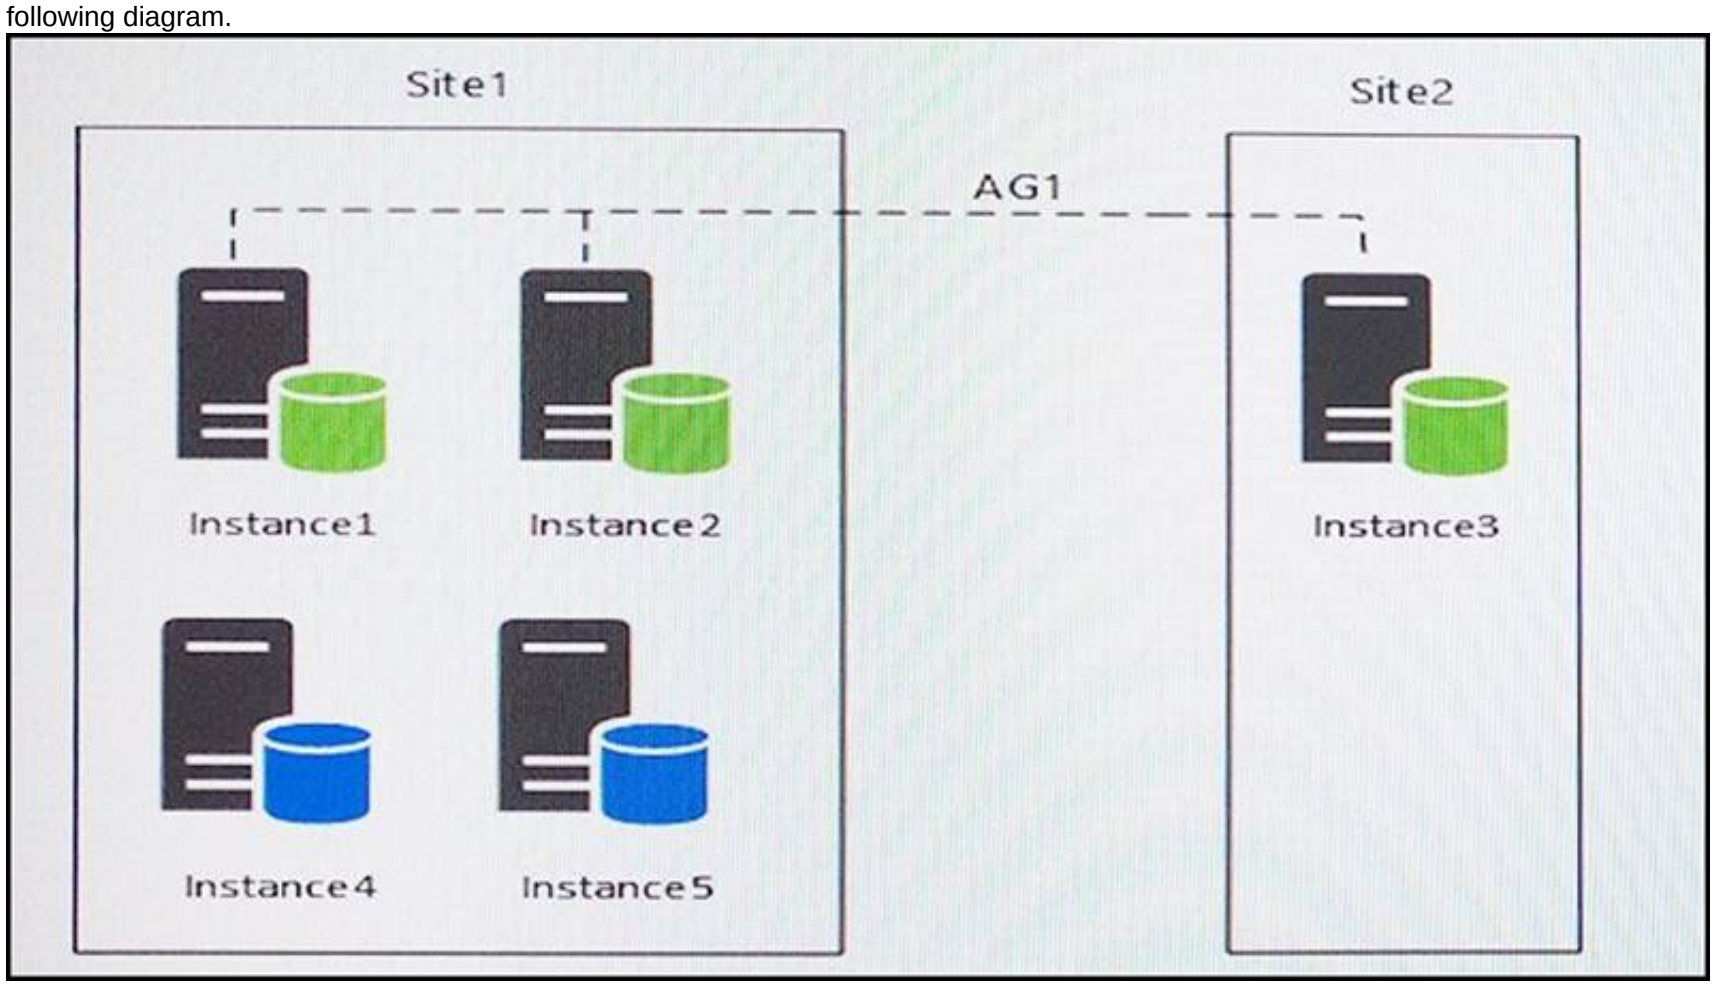

You have an Always On Availability group named AG1. The details for AG1 are shown in the following table.

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

Instance1 experiences heavy read-write traffic. The instance hosts a database named OperationsMain that is four terabytes (TB) in size. The database has multiple data files and filegroups. One of the filegroups is read_only and is half of the total database size.
Instance4 and Instance5 are not part of AG1. Instance4 is engaged in heavy read-write I/O.
Instance5 hosts a database named StagedExternal. A nightly BULK INSERT process loads data into an empty table that has a rowstore clustered index and two nonclustered rowstore indexes.
You must minimize the growth of the StagedExternal database log file during the BULK INSERT operations and perform point-in-time recovery after the BULK INSERT transaction. Changes made must not interrupt the log backup chain.
You plan to add a new instance named Instance6 to a datacenter that is geographically distant from Site1 and Site2. You must minimize latency between the nodes in AG1.
All databases use the full recovery model. All backups are written to the network location \\SQLBackup\. A separate process copies backups to an offsite location.
You should minimize both the time required to restore
the databases and the space required to store backups. The recovery point objective (RPO) for each instance is shown in the following table.

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

Full backups of OperationsMain take longer than six hours to complete. All SQL Server backups use the keyword COMPRESSION.
You plan to deploy the following solutions to the environment. The solutions will access a database named DB1 that is part of AG1.
Reporting system: This solution accesses data inDB1with a login that is mapped to a database user that is a member of the db_datareader role. The user has EXECUTE permissions on the database. Queries make no changes to the data. The queries must be load balanced over variable read-only replicas.
Operations system: This solution accesses data inDB1with a login that is mapped to a database user that is a member of the db_datareader and db_datawriter roles. The user has EXECUTE permissions on the database. Queries from the operations system will perform both DDL and DML operations.
The wait statistics monitoring requirements for the instances are described in the following table.

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

You need to create a backup plan for Instance4. Which backup plan should you create?

A. Weekly full backups, nightly differentia
B. No transaction log backups are necessary.
C. Weekly full backups, nightly differential backups, transaction log backups every 5 minutes.
D. Weekly full backups, nightly differential backups, transaction log backups every 12 hours.
E. Full backups every 60 minutes, transaction log backups every 30 minutes.

Answer: B

Explanation:
From scenario: Instance4 and Instance5 are not part of AG1. Instance4 is engaged in heavy read-write I/O. The recovery point objective of Instancse4 is 60 minutes. RecoveryPoint Objectives are commonly described as the amount of data that was lost during the outage and recovery period. You should minimize both the time required to restore the databases and the space required to store backups.
References:
http://sqlmag.com/blog/sql-server-recovery-time-objectives-and-recovery-point-objectives

NEW QUESTION 53
- (Exam Topic 1)
Note: This question is part of a series of questions that use the same scenario. For your convenience, the scenario is repeated in each question. Each question presents a different goal and answer choices, but the text of the scenario is exactly the same in each question in this series.
You have five servers that run Microsoft Windows 2012 R2. Each server hosts a Microsoft SQL Server instance. The topology for the environment is shown in the following diagram.

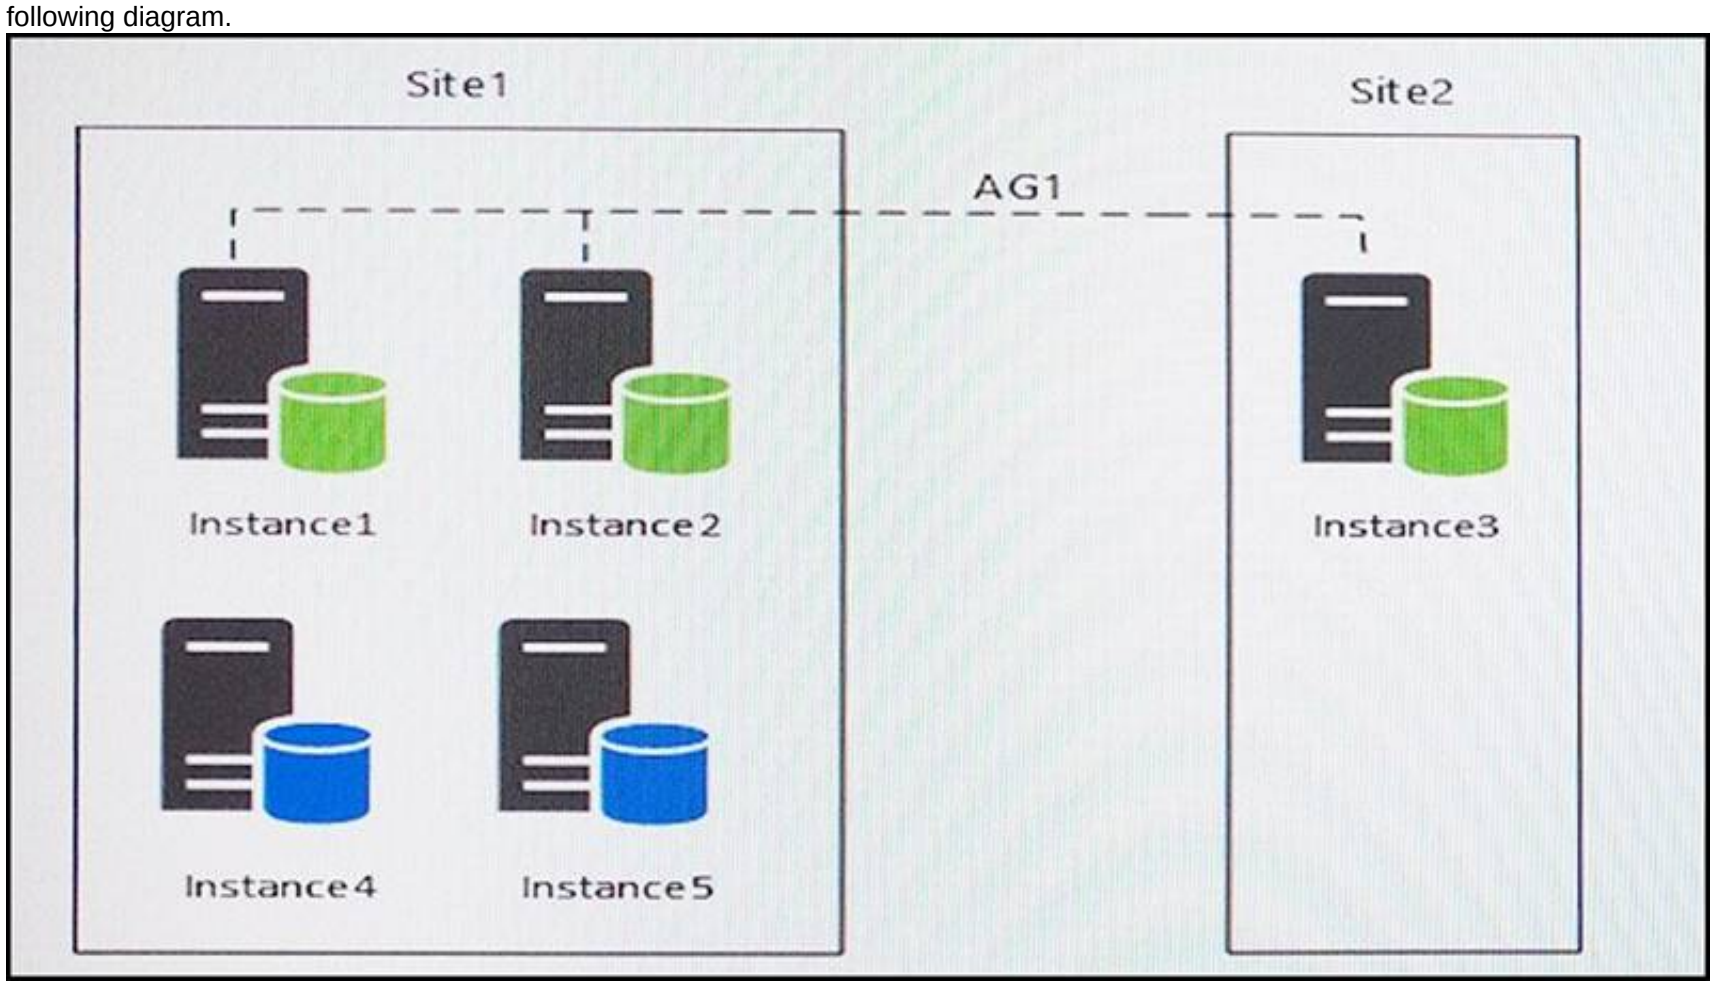

You have an Always On Availability group named AG1. The details for AG1 are shown in the following table.

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

Instance1 experiences heavy read-write traffic. The instance hosts a database named OperationsMain that is four terabytes (TB) in size. The database has multiple data files and filegroups. One of the filegroups is read_only and is half of the total database size.
Instance4 and Instance5 are not part of AG1. Instance4 is engaged in heavy read-write I/O.
Instance5 hosts a database named StagedExternal. A nightly BULK INSERT process loads data into an empty table that has a rowstore clustered index and two nonclustered rowstore indexes.
You must minimize the growth of the StagedExternal database log file during the BULK INSERT operations and perform point-in-time recovery after the BULK INSERT transaction. Changes made must not interrupt the log backup chain.
You plan to add a new instance named Instance6 to a datacenter that is geographically distant from Site1 and Site2. You must minimize latency between the nodes in AG1.
All databases use the full recovery model. All backups are written to the network location \\SQLBackup\. A separate process copies backups to an offsite location. You should minimize both the time required to restore
the databases and the space required to store backups. The recovery point objective (RPO) for each instance is shown in the following table.

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

Full backups of OperationsMain take longer than six hours to complete. All SQL Server backups use the keyword COMPRESSION.
You plan to deploy the following solutions to the environment. The solutions will access a database named DB1 that is part of AG1.
Reporting system: This solution accesses data inDB1with a login that is mapped to a database user that is a member of the db_datareader role. The user has EXECUTE permissions on the database. Queries make no changes to the data. The queries must be load balanced over variable read-only replicas.
Operations system: This solution accesses data inDB1with a login that is mapped to a database user that is a member of the db_datareader and db_datawriter roles. The user has EXECUTE permissions on the database. Queries from the operations system will perform both DDL and DML operations.
The wait statistics monitoring requirements for the instances are described in the following table.

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

You need to reduce the amount of time it takes to backup OperationsMain. What should you do?

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

Answer: D

Explanation:
One of the filegroup is read_only should be as it only need to be backup up once. Partial backups are useful whenever you want to exclude read-only filegroups. A partial backup resembles a full database backup, but a partial backup does not contain all the filegroups. Instead, for a read-write database, a partial backup contains the data in the primary filegroup, every read-write filegroup, and, optionally, one or more read-only files. A partial backup of a read-only database contains only the primary filegroup.
From scenario: Instance1 experiences heavy read-write traffic. The instance hosts a database named OperationsMainthat is four terabytes (TB) in size. The database has multiple data files and filegroups. One of the filegroups is read_only and is half of the total database size.
References:
https://docs.microsoft.com/en-us/sql/relational-databases/backup-restore/partial-backups-sql-server

NEW QUESTION 55
- (Exam Topic 1)
You manage a Microsoft-SQL 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

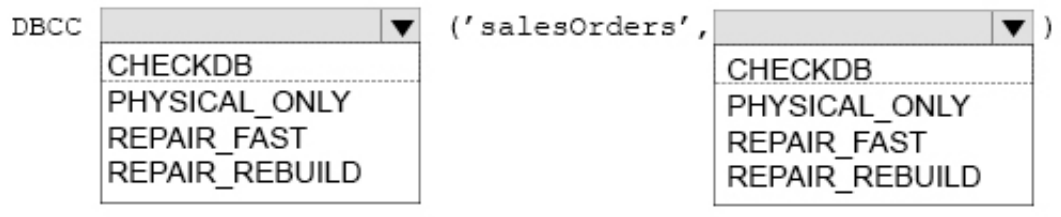

A. Mastered
B. Not Mastered

Answer: A

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 | database_id | 0 [, NOINDEX
| , { REPAIR_ALLOW_DATA_LOSS | REPAIR_FAST | REPAIR_REBUILD }]
….
Box 2: REPAIR_REBUILD
DBCC CHECKDB …REPAIR_ALLOW_DATA_LOSS | REPAIR_FAST |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-checkdb-transact-sql

NEW QUESTION 59
- (Exam Topic 1)
Note: This question is part of a series of questions that use the same scenario. For your convenience, the scenario is repeated in each question. Each question presents a different goal and answer choices, but the text of the scenario is exactly the same in each question in this series.
You are a database administrator for a company that has an on-premises Microsoft SQL Server environment and Microsoft Azure SQL Database instances. The environment hosts several customer databases, and each customer uses a dedicated instance. The environments that you manage are shown in the following table.

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

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

Answer Area

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

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

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

NEW QUESTION 60
- (Exam Topic 1)
Note: This question is part of a series of questions that use the same or similar answer choices. An answer choice may be correct for more than one question in the 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.
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/authentication-access/server-level-roles

NEW QUESTION 64
- (Exam Topic 1)
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 TableB.
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.

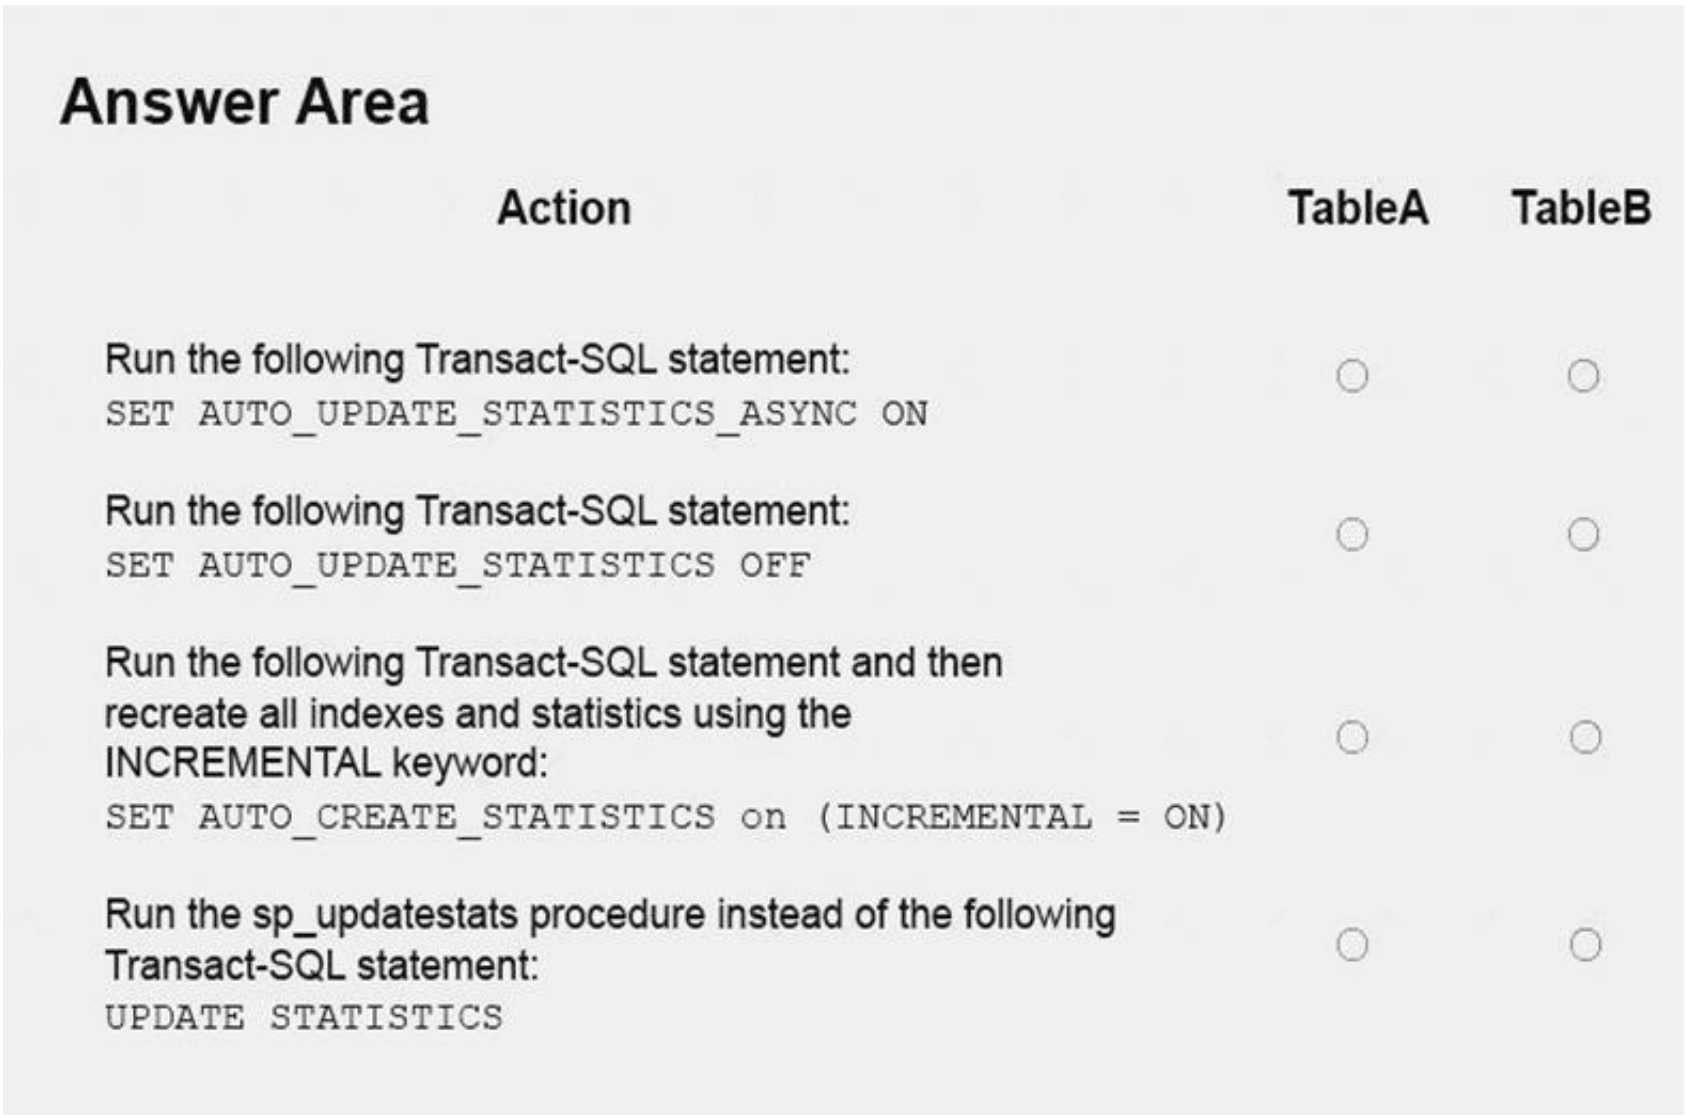

A. Mastered
B. Not Mastered

Answer: A

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-create-statistics-options/

NEW QUESTION 66
- (Exam Topic 1)
You administer a Microsoft SQL Server 2016 database instance.
You plan to migrate the database to Windows Azure SQL Database.
You verify that all objects contained in the database are compatible with Windows Azure SQL Database. You need to ensure that database users and required server logins are migrated to Windows Azure SQL Database.
What should you do?

A. Use the Copy Database wizard.
B. Back up the database from the local server and restore it to Windows Azure SQL Database.
C. Use the Database Transfer wizard.
D. Use SQL Server Management Studio to deploy the database to Windows Azure SQL Database.

Answer: D

NEW QUESTION 67
- (Exam Topic 1)
You administer a Microsoft SQL Server 2016 database.
Users report that an application that accesses the database displays an error, but the error does not provide meaningful information.
No entries are found in the SQL Server log or Windows event logs related to the error. You need to identify the root cause of the issue by retrieving the error message.
What should you do?

A. Create an Extended Events session by using the sqlserver.error_reported event.
B. Create a SQL Profiler session to capture all ErrorLog and EventLog events.
C. Flag all stored procedures for recompilation by using sp_recompile.
D. Execute sp_who.

Answer: A

Explanation:

Trapping SQL Server Errors with Extended Events
One very useful usage of Extended Events is the ability to trap SQL Server error without the need to have a server trace running (which, btw, is deprecated), with the additional feature of being able to query the data as soon as it comes in. This means that we a solution to monitor and trap errors as soon as they happen can be easily created, in order to help developers to fix problems as soon as they are detected. This is really, really, really helpful especially in very big applications, where the code base is quite old and there is no-one really knowing everything of the solution.
To start a 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 Etc.
References:
http://sqlblog.com/blogs/davide_mauri/archive/2013/03/17/trapping-sql-server-errors-with-extended-events.aspx

NEW QUESTION 69
- (Exam Topic 1)
You administer a Microsoft SQL Server 2016 database that includes a table named dbo.Log. This table contains millions of records about user activity in an application.
Records in dbo.Log that are more than 90 days old are purged nightly. When records are purged, table locks are causing contention withinserts.
You need to be able to modify dbo.Log without requiring any changes to the applications that utilize dbo.Log. Which type of solution should you use?

A. Extended events
B. Columnstore index
C. Partitioned tables
D. Read committed snapshot

Answer: C

NEW QUESTION 70
- (Exam Topic 1)
You administer a Microsoft SQL Server 2016 server that has SQL Server Integration Services (SSIS) installed. You plan to deploy new SSIS packages to the server.
The SSIS packages use the Project Deployment Model together with parameters and Integration Services environment variables.
You need to configure the SQL Server environment to support these packages. What should you do?

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

Answer: B

Explanation:
Use can use Project Deployment Model for a project, containing packages and parameters, which is deployed to the SSISDB catalog on an instance of SQL Server.
References:
https://docs.microsoft.com/en-us/sql/integration-services/packages/deploy-integration-services-ssis-projects-and

NEW QUESTION 72
- (Exam Topic 1)
Note: This question is part of a series of questions that use the same scenario. For your convenience, the scenario is repeated in each question. Each question presents a different goal and answer choices, but the text of the scenario is exactly the same in each question in this series.
You have five servers that run Microsoft Windows 2012 R2. Each server hosts a Microsoft SQL Server instance. The topology for the environment is shown in the following diagram.

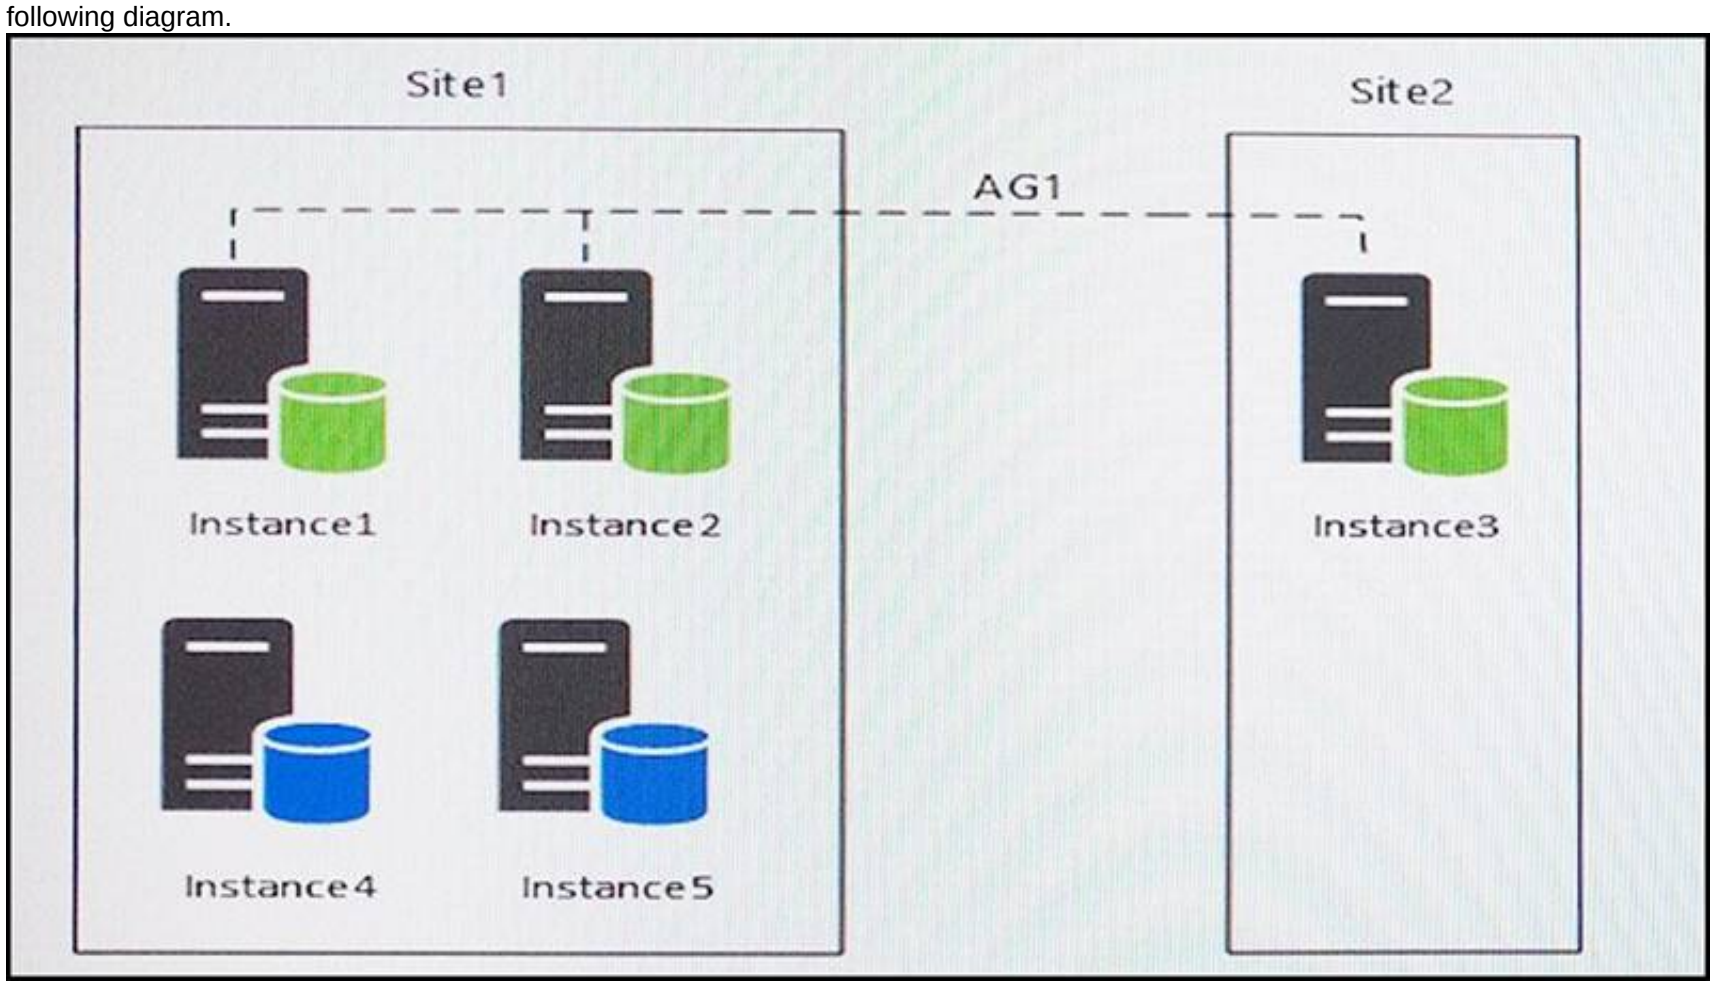

You have an Always On Availability group named AG1. The details for AG1 are shown in the following table.

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

Instance1 experiences heavy read-write traffic. The instance hosts a database named OperationsMain that is four terabytes (TB) in size. The database has multiple data files and filegroups. One of the filegroups is read_only and is half of the total database size.
Instance4 and Instance5 are not part of AG1. Instance4 is engaged in heavy read-write I/O.
Instance5 hosts a database named StagedExternal. A nightly BULK INSERT process loads data into an empty table that has a rowstore clustered index and two nonclustered rowstore indexes.
You must minimize the growth of the StagedExternal database log file during the BULK INSERT operations and perform point-in-time recovery after the BULK INSERT transaction. Changes made must not interrupt the log backup chain.
You plan to add a new instance named Instance6 to a datacenter that is geographically distant from Site1 and Site2. You must minimize latency between the nodes in AG1.
All databases use the full recovery model. All backups are written to the network location \\SQLBackup\. A separate process copies backups to an offsite location. You should minimize both the time required to restore the databases and the space required to store backups. The recovery point objective (RPO) for each instance is shown in the following table.

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

Full backups of OperationsMain take longer than six hours to complete. All SQL Server backups use the keyword COMPRESSION.
You plan to deploy the following solutions to the environment. The solutions will access a database named DB1 that is part of AG1.
Reporting system: This solution accesses data inDB1with a login that is mapped to a database user that is a member of the db_datareader role. The user has EXECUTE permissions on the database. Queries make no changes to the data. The queries must be load balanced over variable read-only replicas.
Operations system: This solution accesses data inDB1with a login that is mapped to a database user that is a member of the db_datareader and db_datawriter roles. The user has EXECUTE permissions on the database. Queries from the operations system will perform both DDL and DML operations.
The wait statistics monitoring requirements for the instances are described in the following table.

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

You need to 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.	○	○

A. Mastered
B. Not Mastered

Answer: A

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/availability-groups/windows/configure-read-only-access-o

NEW QUESTION 77
- (Exam Topic 1)
Note: This question is part of a series of questions that use the same or similar answer choices. An answer choice may be correct for more than one question in the 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 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 teh 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.

NEW QUESTION 80
- (Exam Topic 1)
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 goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it As a result these questions will not appear in the review screen.
Your company has several Microsoft SQL Server instances, Each instance hosts many databases. You observe I/O corruption on some of the instances
You need to perform the following actions:
• Identify databases where the PAGE VERIFY option is not set
• Configure full page protection for the identified databases. Solution: You run the following Transact-SQL statement:

```
SELECT NAME, page_verify_option_desc
FROM master.sys.databases
WHERE page_verify_option_desc = 'NONE'
GO
```

For each database that you identify, you run the following Transact-SQL statement:

```
ALTER DATABASE <database_name>
SET PAGE_VERIFY TORN_PAGE_DETECTION
```

Does the solution meet the goal?

A. Yes
B. No

Answer: B

NEW QUESTION 85
- (Exam Topic 1)
You administer a Microsoft SQL Server 2016 server that hosts a transactional database and a reporting database. The transactional database is updated through a web application and is operational throughout the day.
The reporting database is only updated from the transactional database.
The recovery model and backup schedule are configured as shown in the following table:

Database	Description
Transactional database	Recovery model: • Full Backup schedule: • Full database backup: midnight, daily • Differential database backup: on the hour, every two hours starting at 02:00 hours except at 00:00 hours • Log backup: every half hour, except at the times of full and differential backups
Reporting database	Recovery model: • Simple Backup schedule: • Full database backup: 01:00 hours daily • Differential database backup: 13:00 hours daily Data updates: • Changes in data are updated from the transactional database to the reporting database at 00:30 hours and at 12:30 hours • The update takes 15 minutes

At 14:00 hours, you discover that pages 71, 520, and 713 on one of the database files are corrupted on the reporting database.
You need to ensure that the databases are restored.
You also need to ensure that data loss is minimal. What should you do?

A. Perform a partial restore.
B. Restore the latest full backup, and restore the latest differential backu
C. Then, restore each log backup taken before the time of failure from the most recent differential backup.
D. Restore the latest full backup.
E. Restore the latest full backup, and restore the latest differential backu
F. Then, restore the latest log backup.
G. Perform a page restore.
H. Restore the latest full backu
I. Then, restore each differential backup taken before the time of failure from the most recent full backup.
J. Perform a point-in-time restore.
K. Restore the latest full backu
L. Then, restore the latest differential backup.

Answer: H

Explanation:
At restore time, before you restore a differential backup, you must restore its base. Then, restore only the most recent differential backup to bring the database forward to the time when that differential backup was created. Typically, you would restore the most recent full backup followed by the most recent differential backup that is based on that full backup.
References: https://technet.microsoft.com/en-us/library/ms345448(v=sql.105).aspx

NEW QUESTION 90
- (Exam Topic 1)
Note: This question is part of a series of questions that use the same or similar answer choices. An answer choice may be correct for more than one question in the 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 observe that several indexes are fragmented. You need to rebuild the indexes.
What should you use?

A. Activity Monitor
B. Sp_who3 stored procedure
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
In Object Explorer, connect to an instance of the SQL Server Database Engine and then expand that instance.
Expand Databases, expand the database that contains the table with the specified index, and then expand Tables.
Expand the table in which the index belongs and then expand Indexes.
Right-click the index to rebuild and then click Rebuild.
To start the rebuild operation, click OK.
References: https://technet.microsoft.com/en-us/library/ms187874(v=sql.105).aspx

NEW QUESTION 95
- (Exam Topic 1)
You are the database administrator for a Microsoft SQL Server instance. You develop an Extended Events package to look for events related to application performance.
You need to change the event session to include SQL Server errors that are greater than error severity 15. Which five Transact-SQL segments should you use to develop the solution? To answer, move the appropriate
Transact-SQL segments from the list of Transact-SQL segments to the answer area and arrange them in the correct order.

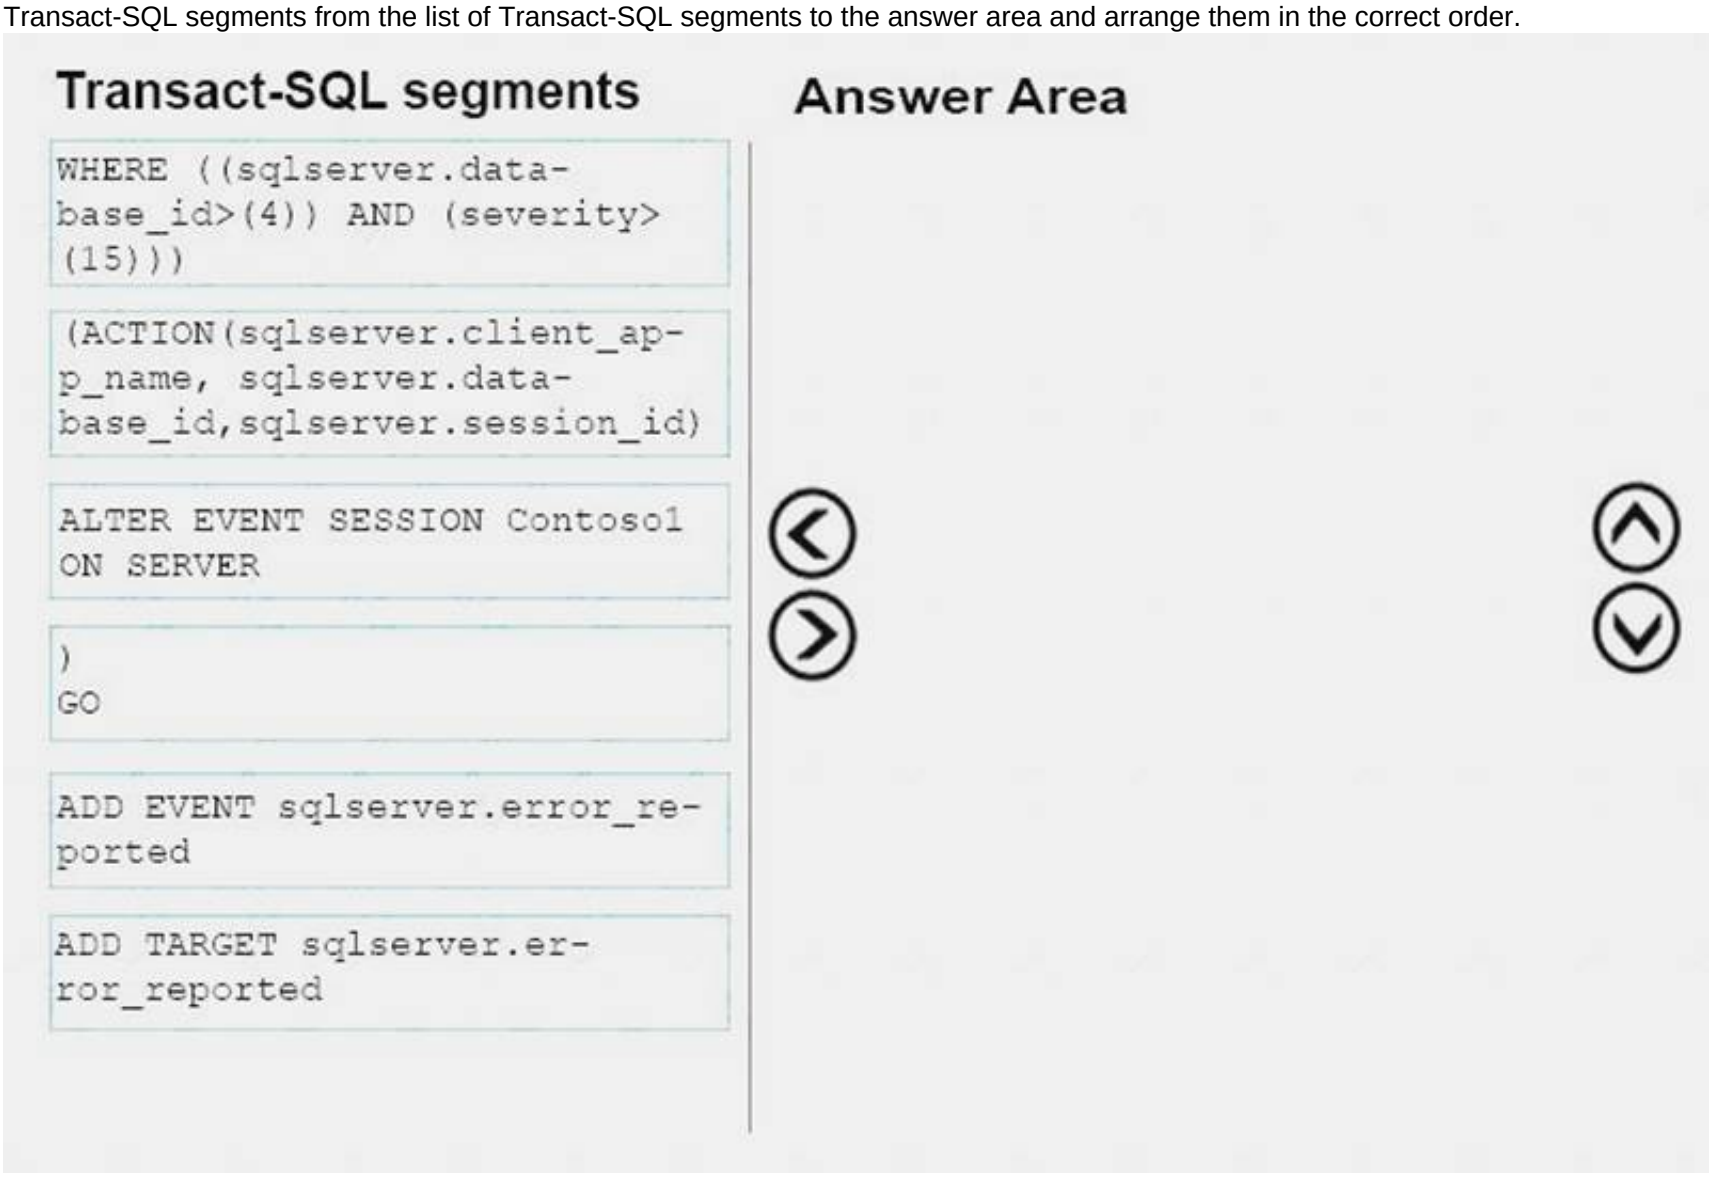

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Step 1: ALTER EVENT SESSION Contoso1 ON SERVER
Step 2: ADD EVENT … Step 3: (ACTION ... Step 4: WHERE...
Step 5:) GO
Example: To start an Extended Events sessions in order to trap SQL Server errors with severity greater than 10,just run the following script:
CREATE EVENT SESSION [error_trap] ON SERVER
ADD EVENT sqlserver.error_reported (
ACTION
(package0.collect_system_time,package0.last_error,sqlserver.client_app_name,sqlserver.client_hostname,sqlser
sqlserver.plan_handle,sqlserver.query_hash,sqlserver.session_id,sqlserver.sql_text,sqlserver.tsql_frame,sqlserve
WHERE ([severity]>10)
)
ADD TARGET package0.event_file (
SET filename=N'D:\Program Files\Microsoft SQL Server\MSSQL11.MSSQLSERVER\MSSQL\XEvents\error_trap.xel'
) WITH (
STARTUP_STATE=OFF
) GO
References:
http://sqlblog.com/blogs/davide_mauri/archive/2013/03/17/trapping-sql-server-errors-with-extended-events.aspx

NEW QUESTION 96
- (Exam Topic 1)
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the

stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
Your company has several Microsoft SOI Server instances. Each instance hosts many databases. You observe I/O corruption on some of the instances.
You need to perform the following actions:
• Identify databases where the PAGE VERIFY option is not set.
• Configure full page protection for the identified databases. Solution: You run the following Transact-SQL statement:

```
SELECT NAME, page_verity_option_desc
FROM master.sys.databases
WHERE page_verify_option_desc != 'TORN_PAGE_DETECTION'
GO
```

For each database that you identify, you run the following Transact-SQL statement:

```
ALTER DATABASE <database_name>
SET PAGE_VERIFY TORN_PAGE_DETECTION
```

Does the solution meet the goal?

A. Yes
B. NO

Answer: B

NEW QUESTION 99
- (Exam Topic 1)
Note: This question is part of a series of questions that use the same or similar answer choices. An answer choice may be correct for more than one question in the 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

Answer: G

Explanation:
The following permissions are necessary to perform column-level encryption, or cell-level encryption.
CONTROL permission on the database.
CREATE CERTIFICATE permission on the database. Only Windows logins, SQL Server logins, and application roles can own certificates. Groups and roles cannot own certificates.
ALTER permission on the table.
Some permission on the key and must not have been denied VIEW DEFINITION permission. References:
https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/encrypt-a-column-of-data

NEW QUESTION 103
- (Exam Topic 1)
Note: This question is part of a series of questions that use the same or similar answer choices. An answer choice may be correct for more than one question in the series. Each question is independent of the other questions in this series. Information and details provided in a question apply only to that question.
You collect performance metrics on multiple Microsoft SQL Server instances and store the data in a single repository.
You need to examine disk usage, query statistics, and server activity without building custom counters.
What should you use?

A. 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

NEW QUESTION 106
- (Exam Topic 1)

You administer a Microsoft SQL Server 2016 instance that contains a financial database hosted on a storage area network (SAN). The financial database has the following characteristics:
A data file of 2 terabytes is located on a dedicated LUN (drive D).
A transaction log of 10 GB is located on a dedicated LUN (drive E).
Drive D has 1 terabyte of free disk space.
Drive E has 5 GB of free disk space.
The database is continually modified by users during business hours from Monday through Friday between 09:00 hours and 17:00 hours.
Five percent of the existing data is modified each day.
The Finance department loads large CSV files into a number of tables each business day at 11:15 hours and 15:15 hours by using the BCP or BULK INSERT commands.
Each data load adds 3 GB of data to the database.
These data load operations must occur in the minimum amount of time. A full database backup is performed every Sunday at 10:00 hours.
Backup operations will be performed every two hours (11:00, 13:00, 15:00, and 17:00) during business hours. You implement log shipping of the financial database to another SQL Server 2016 instance. You decide to
failover to this secondary database.
You need to ensure that all transactions will be replicated to the secondary database. Which backup option should you use?

A. Differential
B. Transaction Log
C. FULL
D. SIMPLE
E. SKIP
F. RESTART
G. STANDBY
H. CHECKSUM
I. DBO_ONLY
J. COPY_ONLY
K. NORECOVERY
L. NO_CHECKSUM
M. CONTINUE_AFTER_ERROR
N. BULK_LOGGED

Answer: K

Explanation:
Roll back is controlled by the RESTORE statement through the [RECOVERY | NORECOVERY] options: NORECOVERY specifies that roll back not occur. This allows roll forward to continue with the next
statement in the sequence.
References: https://docs.microsoft.com/en-us/sql/t-sql/statements/restore-statements-transact-sql

NEW QUESTION 107
- (Exam Topic 1)
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
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

NEW QUESTION 108
- (Exam Topic 1)
You administer all the deployments of Microsoft SQL Server 2016 in your company.
You need to ensure that an OLTP database that uses a storage area network (SAN) remains available if any of the servers fail.
You also need to minimize the amount of storage used by the database. Which configuration should you use?

A. Two servers configured in different data centers SQL Server Availability Group configured in Synchronous-Commit Availability Mode One server configured as an Active Secondary
B. SQL Server that includes an application database configured to perform transactional replication
C. Two servers configured in the same data center SQL Server Availability Group configured in AsynchronousCommit Availability Mode One server configured as an Active Secondary
D. Two servers configured in different data centers SQL Server Availability Group configured in AsynchronousCommit Availability Mode
E. Two servers configured in the same data center A primary server configured to perform log-shippingevery 10 minutes A backup server configured as a warm standby
F. Two servers configured on the same subnet SQL Server Availability Group configured in Synchronous-Commit Availability Mode
G. SQL Server that includes an application database configured to perform snapshot replication
H. Two servers configured in a Windows Failover Cluster in the same data center SQL Server configured as a clustered instance

Answer: H

Explanation:
A Windows Server Failover Cluster (WSFC) is a group of independent servers that work together to increase the availability of applications and services. SQL

Server takes advantage of WSFC services and capabilities to support Always On availability groups and SQL Server Failover Cluster Instances.
References:
https://docs.micosoft.com/en-us/sql/sql-server/failover-clusters/windows/windows-server-failover-clustering-ws

NEW QUESTION 111
- (Exam Topic 1)
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-SQL code statement? To answer, select the appropriate Transact-SQL code segments in the dialog box in the answer area.

Answer Area

```
USE salesOrders
DBCC [ CHECKDB | CHECKCATALOG | CHECKCONSTRAINTS ] ('Table1' [ ALL_CONSTRAINTS | ALL_ERRORMSGS | NO_INFOMSGS ] )
```

A. Mastered
B. Not Mastered

Answer: A

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 | table_id | constraint_name | constraint_id
)
]
[WITH
[{ ALL_CONSTRAINTS | ALL_ERRORMSGS }] [,] [NO_INFOMSGS]
]
References:
https://docs.microsoft.com/en-us/sql/t-sql/database-console-commands/dbcc-checkconstraints-transact-sql

NEW QUESTION 112
- (Exam Topic 1)
You administer a Microsoft SQL Server 2016 failover cluster that contains two nodes named Node A and Node B.
A single instance of SQL Server is installed on the cluster.
An additional node named Node C has been added to the existing cluster.
You need to ensure that the SQL Server instance can use all nodes of the cluster. What should you do?
A. Create a ConfigurationFile.ini file from Node B, and then run the AddNode command-line tool on Node A.
B. Use Node A to install SQL Server on Node C.
C. Run the Add Node to SQL Server Failover Cluster Wizard on Node C.
D. Use Cluster Administrator to add a new Resource Group to Node B.

Answer: C

Explanation:
To add a node to an existing SQL Server failover cluster
Insert the SQL Server installation media, and from the root folder, double-click Setup.exe. To install from a network share, navigate to the root folder on the share, and then double-click Setup.exe.
The Installation Wizard will launch the SQL Server Installation Center. To add a node to an existing failover cluster instance, click Installation in the left-hand pane. Then, select Add node to a SQL Server failover cluster.
Etc.
References:
https://docs.microsoft.com/en-us/sql/sql-server/failover-clusters/install/add-or-remove-nodes-in-a-sql-server-fail

NEW QUESTION 117
- (Exam Topic 1)
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 goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it As a result these questions will not appear in the review screen.
You have a database named DB1 that is 640 GB and is updated frequently.
You enabled log shipping for DB1 and configure backup and restore to occur every 30 minutes. You discover that the disks on the data server are almost full.
You need to reduce the amount of disk space used by the log shipping process. Solution: You enable compression for the transaction log backups:

Does this meet the goal?

A. Yes
B. No

Answer: A

NEW QUESTION 118
- (Exam Topic 1)
You administer several Microsoft SQL Server 2016 database servers.
Merge replication has been configured for an application that is distributed across offices throughout a wide area network (WAN). Many of the tables involved in replication use the XML and varchar (max) data types.
Occasionally, merge replication fails due to timeout errors. You need to reduce the occurrence of these timeout errors. What should you do?

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

Answer: A

Explanation:
You might have different profiles for different instances of an agent. For example, a Merge Agent that connects to the Publisher and Distributor over a dialup connection could use a set of parameters that are better suited to the slower communications link by using the slow link profile.
Note: When replication is configured, a set of agent profiles is installed on the Distributor. An agent profile contains a set of parameters that are used each time an agent runs: each agent logs in to the Distributor during its startup process and queries for the parameters in its profile.
References:
https://docs.microsoft.com/en-us/sql/relational-databases/replication/agents/replication-agent-profiles

NEW QUESTION 123
- (Exam Topic 1)
Note: This question is part of a series of questions that use the same or similar answer choices. An answer choice may be correct for more than one question in the 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 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

NEW QUESTION 124
- (Exam Topic 2)
You use SQL Server 2014 Enterprise Edition.
Your database contains a partitioned table named AuditData. AuditData is partitioned by year. Partition 1 contains data from the year 2010 and prior.
Management has decided to archive all AUDITDATA records from 2010 and prior.
Management wants the records to be removed from the database entirely and provided to the backup team as a zipped text file. The data must no longer reside in the database.
There is very little tolerance for performance degradation in your environment. You need to remove all 2010 and prior data from the AuditData table by using the least amount of system resources possible. Develop the solution by selecting and arranging the required SQL actions in the correct order.
You may not need all of the actions.

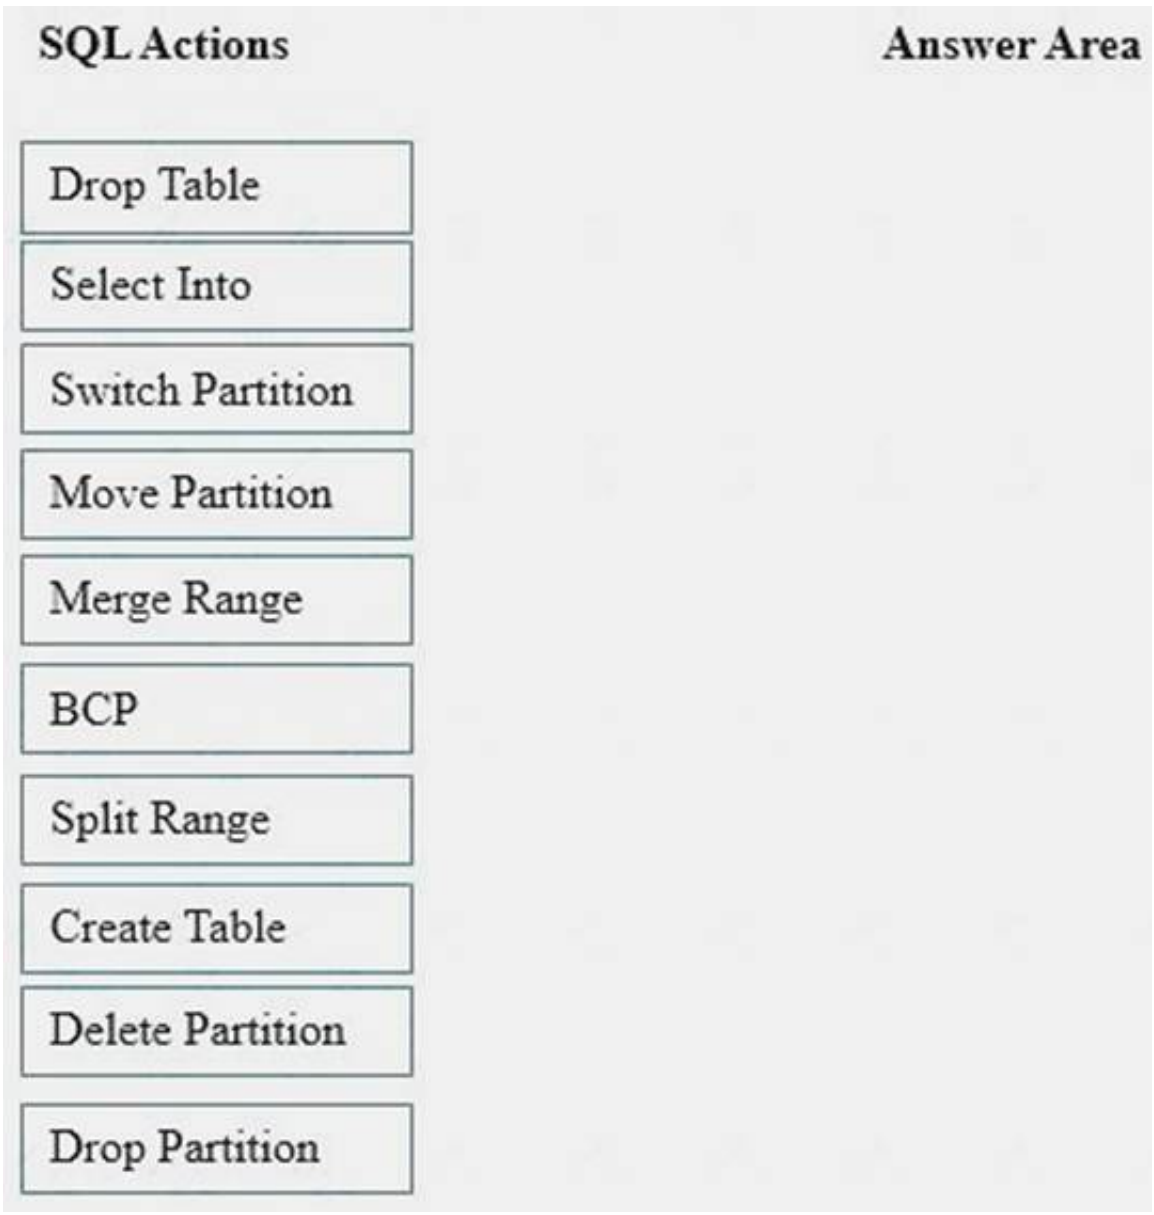

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Note:
- Create a new partitioned table with the partition function you want, and then insert the data from the old table into the new table by using an INSERT INTO...SELECT FROM statement.
- SPLIT RANGE (boundary_value)
Adds one partition to the partition function. boundary_value determines the range of the new partition, and must differ from the existing boundary ranges of the partition function. Based on boundary_value, the Database Engine splits one of the existing ranges into two.
Of these two, the one where the new boundary_value resides is considered the new partition.
- BCP can be used top produce the zipped text file.
- Example: plitting a partition of a partitioned table or index into two partitions
The following example creates a partition function to partition a table or index into four partitions. ALTER PARTITION FUNCTION splits one of the partitions into two to create a total of five partitions. CREATE PARTITION FUNCTION myRangePF1 (int)
AS RANGE LEFT FOR VALUES (1, 100, 1000); GO
-Split the partition between boundary_values 100 and 1000
-to create two partitions between boundary_values 100 and 500
--and between boundary_values 500 and 1000. ALTER PARTITION FUNCTION myRangePF1 () SPLIT RANGE (500);

NEW QUESTION 126
- (Exam Topic 2)
Overview
General Overview
ADatum Corporation has offices in Miami and Montreal.
The network contains a single Active Directory forest named adatum.com. The offices connect to each other by using a WAN link that has 5-ms latency. A. Datum standardizes its database platform by using SQL Server 2014 Enterprise edition.
Databases
Each office contains databases named Sales, Inventory, Customers, Products, Personnel, and Dev. Servers and databases are managed by a team of database administrators. Currently, all of the database
administrators have the same level of permissions on all of the servers and all of the databases.
The Customers database contains two tables named Customers and Classifications. The following graphic shows the relevant portions of the tables:

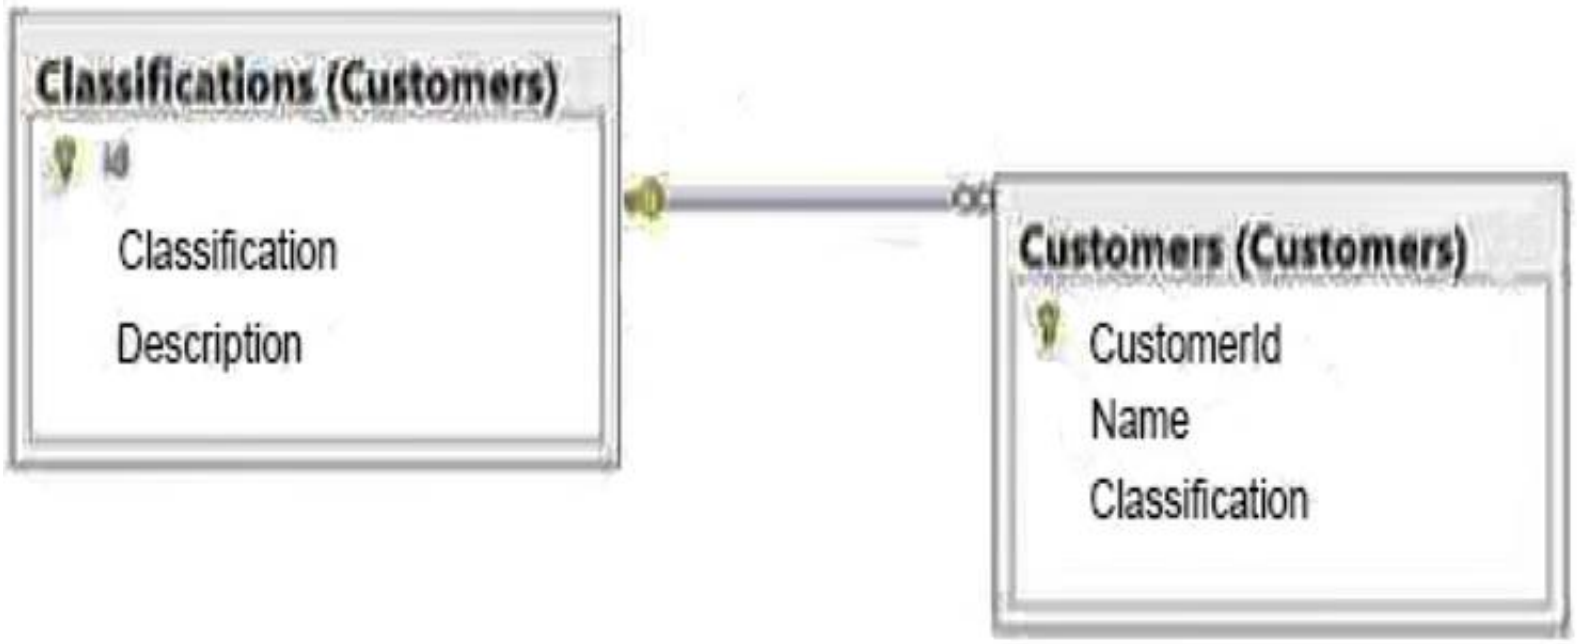

The following table shows the current data in the Classifications table:

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

The Inventory database is updated frequently. The database is often used for reporting.
A full backup of the database currently takes three hours to complete. Stored Procedures
A stored procedure named USP_1 generates millions of rows of data for multiple reports. USP_1 combines data from five different tables from the Sales and Customers databases in a table named Table1.
After Table1 is created, the reporting process reads data from Table1 sequentially several times. After the process is complete, Table1 is deleted.
A stored procedure named USP_2 is used to generate a product list. The product list contains the names of products grouped by category.
USP_2 takes several minutes to run due to locks on the tables the procedure accesses. The locks are caused by USP_1 and USP_3.
A stored procedure named USP_3 is used to update prices. USP_3 is composed of several UPDATE statements called in sequence from within a transaction.
Currently, if one of the UPDATE statements fails, the stored procedure fails. A stored procedure named USP_4 calls stored procedures in the Sales, Customers, and Inventory databases.
The nested stored procedures read tables from the Sales, Customers, and Inventory databases. USP_4 uses an EXECUTE AS clause.
All nested stored procedures handle errors by using structured exception handling. A stored procedure named USP_5 calls several stored procedures in the same database. Security checks are performed each time USP_5 calls a stored procedure.
You suspect that the security checks are slowing down the performance of USP_5. All stored procedures accessed by user applications call nested stored procedures.
The nested stored procedures are never called directly. Design Requirements
Data Recovery
You must be able to recover data from the Inventory database if a storage failure occurs. You have a Recovery Time Objective (RTO) of 5 minutes.
You must be able to recover data from the Dev database if data is lost accidentally. You have a Recovery Point Objective (RPO) of one day.
Classification Changes
You plan to change the way customers are classified. The new classifications will have four levels based on the number of orders. Classifications may be removed or added in the future. Management requests that historical data be maintained for the previous classifications. Security A group of junior database administrators must be able to manage security for the Sales database. The junior database administrators will not have any other administrative rights. A. Datum wants to track which users run each stored procedure.
Storage
ADatum has limited storage. Whenever possible, all storage space should be minimized for all databases and all backups.
Error Handling
There is currently no error handling code in any stored procedure.
You plan to log errors in called stored procedures and nested stored procedures. Nested stored procedures are never called directly.
You need to recommend a solution to minimize the amount of time it takes to execute USP_5. What should you include in the recommendation?

A. Enable cross-database chaining.
B. Use a server role to group all logins.
C. Use the EXECUTE AS clause in USP_5.
D. Copy USP.5 to each database.

Answer: A

Explanation:
Scenario:
A stored procedure named USP_5 changes data in multiple databases. Security checks are performed each time USP_5 accesses a database.
- Cross-database ownership chaining occurs when a procedure in one database depends on objects in another database. A cross-database ownership chain works in the same way as ownership chaining within a single database, except that an unbroken ownership chain requires that all the object owners are mapped to the same login account. If the source object in the source database and the target objects in the target databases are owned by the same login account, SQL Server does not check permissions on the target objects.

NEW QUESTION 130
- (Exam Topic 2)
You are designing a monitoring application for a new SQL Server 2014 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

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.

NEW QUESTION 133
- (Exam Topic 2)
You are designing a SQL Server database for an order fulfillment system. You create a table named Sales.Orders by using the following script:

```
CREATE TABLE Sales.Orders
 (
      OrderID int IDENTITY (1,1) NOT NULL PRIMARY KEY,
      OrderDate date NOT NULL,
      CustomerID int NOT NULL
 );
```

Each order is tracked by using one of the following statuses:
Fulfilled
Shipped
Ordered
Received
You need to design the database to ensure that you can retrieve the status of an order on a given date. The solution must ensure that new statuses can be added in the future.
What should you do? More than one answer choice may achieve the goal. Select the BEST answer.

A. To the Sales.Orders table, add a column named Status that will store the order statu
B. Update the Status column as the order status changes.
C. Create a new table named Sales.OrderStatus that contains three columns named OrderID, StatusDate, and Statu
D. Insert new rows into the table as the order status changes.
E. Implement change data capture on the Sales.Orders table.
F. To the Sales.Orders table, add three columns named FulfilledDate, ShippedDate, and ReceivedDate.Update the value of each column from null to the appropriate date as the order status changes.

Answer: A

NEW QUESTION 136
- (Exam Topic 2)
Overview
You are a database administrator for a company named Litware, Inc.
Litware is a book publishing house. Litware has a main office and a branch office.
You are designing the database infrastructure to support a new web-based application that is being developed. The web application will be accessed at www.litwareinc.com. Both internal employees and external partners
will use the application.
You have an existing desktop application that uses a SQL Server 2008 database named App1_DB. App1_DB will remain in production.
Requirements Planned Changes
You plan to deploy a SQL Server 2014 instance that will contain two databases named Database1 and
Database2.
All database files will be stored in a highly available SAN. Database1 will contain two tables named Orders and OrderDetails.
Database1 will also contain a stored procedure named usp_UpdateOrderDetails.
The stored procedure is used to update order information. The stored procedure queries the Orders table twice each time the procedure executes.
The rows returned from the first query must be returned on the second query unchanged along with any rows added to the table between the two read operations.
Database1 will contain several queries that access data in the Database2 tables. Database2 will contain a table named Inventory.
Inventory will contain over 100 GB of data.
The Inventory table will have two indexes: a clustered index on the primary key and a nonclustered index. The column that is used as the primary key will use the identity property.
Database2 wilt contains a stored procedure named usp_UpdateInventory. usp_UpdateInventory will manipulate a table that contains a self-join that has an unlimited number of hierarchies. All data in Database2 is recreated each day ad does not change until the next data creation process. Data from Database2 will be accessed periodically by an external application named Application1. The data from Database2 will be sent to a database named Appl_Dbl as soon as changes occur to the data in Database2. Litware plans to use offsite storage for all SQL Server 2014 backups.
Business Requirements
You have the following requirements:
Costs for new licenses must be minimized.
Private information that is accessed by Application must be stored in a secure format.
Development effort must be minimized whenever possible.
The storage requirements for databases must be minimized.
System administrators must be able to run real-time reports on disk usage.
The databases must be available if the SQL Server service fails.
Database administrators must receive a detailed report that contains allocation errors and data corruption.
Application developers must be denied direct access to the database tables. Applications must be denied

direct access to the tables.
You must encrypt the backup files to meet regulatory compliance requirements.
The encryption strategy must minimize changes to the databases and to the applications. You need to recommend a solution to improve the performance of usp.UpdateInventory.
The solution must minimize the amount of development effort. What should you include in the recommendation?

A. A table variable
B. A common table expression
C. A subquery
D. A cursor

Answer: A

Explanation:
- Scenario: Database2 will contain a stored procedure named usp_UpdateInventory. Usp_UpdateInventory will manipulate a table that contains a self-join that has an unlimited number of hierarchies.
- A table variable can be very useful to store temporary data and return the data in the table format.
- Example: The following example uses a self-join to find the products that are supplied by more than one vendor. Because this query involves a join of the ProductVendor table with itself, the ProductVendor table appears in two roles. To distinguish these roles, you must give the ProductVendor table two different aliases (pv1 and pv2) in the FROM clause. These aliases are used to qualify the column names in the rest of the query. This is an example of the self-join Transact-SQL statement:

```
USE AdventureWorks2008R2;
GO
SELECT DISTINCT pv1.ProductID, pv1.VendorID
FROM Purchasing.ProductVendor pv1
INNER JOIN Purchasing.ProductVendor pv2
ON pv1.ProductID = pv2.ProductID
AND pv1.VendorID <> pv2.VendorID
ORDER BY pv1.ProductID
```

NEW QUESTION 140
- (Exam Topic 2)
You plan to create a database.
The database will be used by a Microsoft .NET application for a special event that will last for two days. During the event, data must be highly available. After the event, the database will be deleted. You need to
recommend a solution to implement the database while minimizing costs. The solution must not affect any existing applications.
What should you recommend? More than one answer choice may achieve the goal. Select the BEST answer.

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

Answer: D

NEW QUESTION 142
- (Exam Topic 2)
You plan to deploy SQL Server 2014.
You identify the following security requirements for the deployment:
Users must be prevented from intercepting and reading the T-SQL statements sent from the clients to the database engine.
All database files and log files must be encrypted if the files are moved to another disk on another server.
You need to identify which feature meets each security requirement. The solution must minimize processor overhead.
Which features should you identify? To answer, drag the appropriate feature to the correct requirement in the answer area.

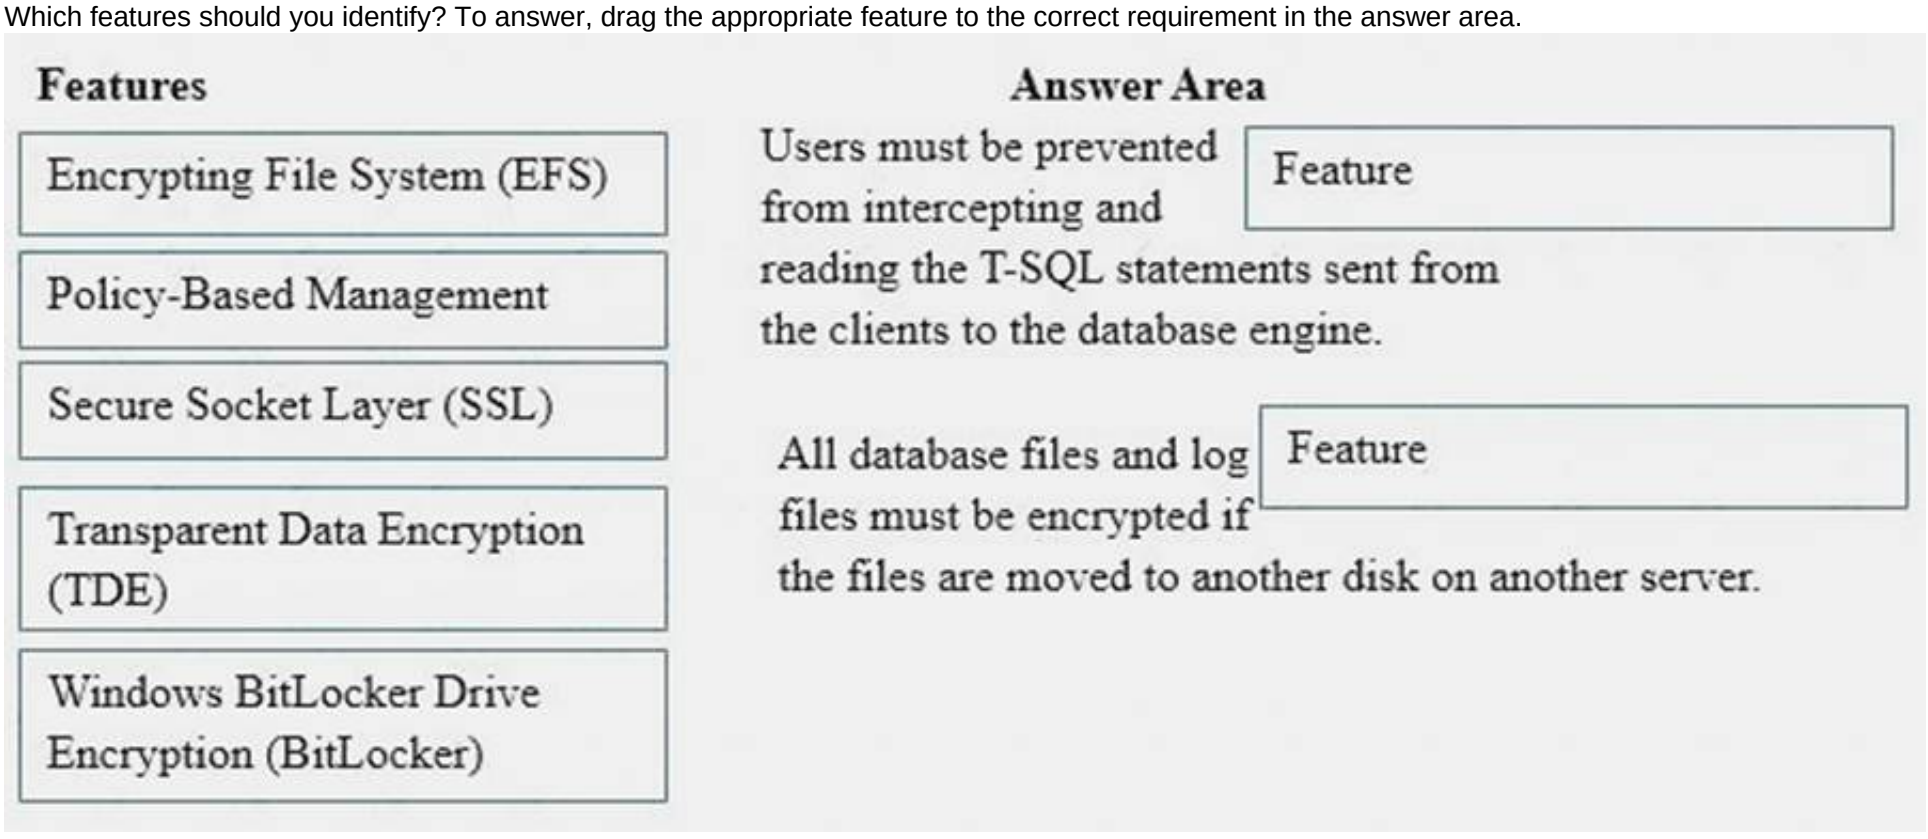

A. Mastered
B. Not Mastered

Answer: A

Explanation:
- Secure Sockets Layer (SSL) encryption enables transmitting encrypted data across the network between an instance of SQL Server and a client application.
- Transparent data encryption (TDE) performs real-time I/O encryption and decryption of the data and log files.

NEW QUESTION 146
- (Exam Topic 2)
Overview
General Overview
ADatum Corporation has offices in Miami and Montreal.
The network contains a single Active Directory forest named adatum.com. The offices connect to each other by using a WAN link that has 5-ms latency. A. Datum standardizes its database platform by using SQL Server 2014 Enterprise edition.
Databases
Each office contains databases named Sales, Inventory, Customers, Products, Personnel, and Dev. Servers and databases are managed by a team of database administrators. Currently, all of the database
administrators have the same level of permissions on all of the servers and all of the databases.
The Customers database contains two tables named Customers and Classifications. The following graphic shows the relevant portions of the tables:

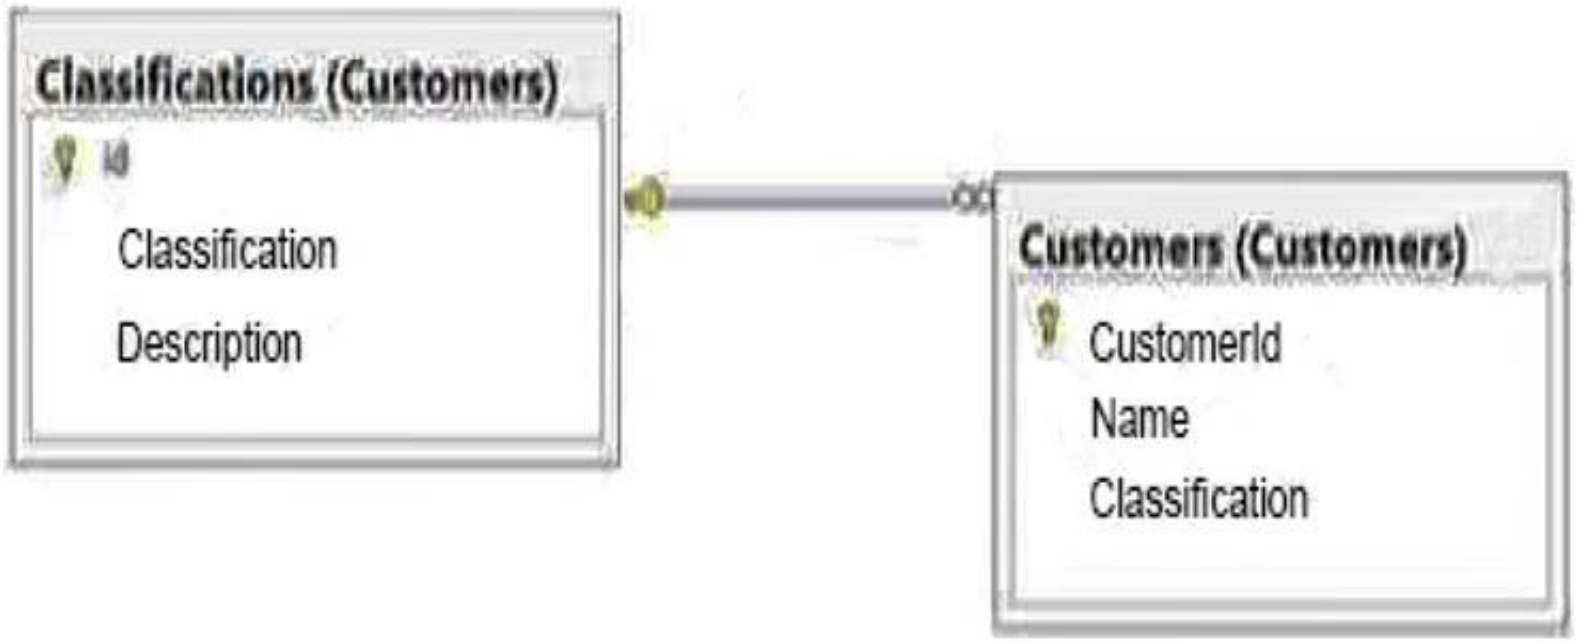

The following table shows the current data in the Classifications table:

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

The Inventory database is updated frequently. The database is often used for reporting.
A full backup of the database currently takes three hours to complete. Stored Procedures
A stored procedure named USP_1 generates millions of rows of data for multiple reports. USP_1 combines data from five different tables from the Sales and Customers databases in a table named Table1.
After Table1 is created, the reporting process reads data from Table1 sequentially several times. After the process is complete, Table1 is deleted.
A stored procedure named USP_2 is used to generate a product list. The product list contains the names of products grouped by category.
USP_2 takes several minutes to run due to locks on the tables the procedure accesses. The locks are caused by USP_1 and USP_3.
A stored procedure named USP_3 is used to update prices. USP_3 is composed of several UPDATE statements called in sequence from within a transaction.
Currently, if one of the UPDATE statements fails, the stored procedure fails. A stored procedure named USP_4 calls stored procedures in the Sales, Customers, and Inventory databases.
The nested stored procedures read tables from the Sales, Customers, and Inventory databases. USP_4 uses an EXECUTE AS clause.
All nested stored procedures handle errors by using structured exception handling. A stored procedure named USP_5 calls several stored procedures in the same database. Security checks are performed each time USP_5 calls a stored procedure.
You suspect that the security checks are slowing down the performance of USP_5. All stored procedures accessed by user applications call nested stored procedures.
The nested stored procedures are never called directly. Design Requirements
Data Recovery
You must be able to recover data from the Inventory database if a storage failure occurs. You have a Recovery Time Objective (RTO) of 5 minutes.
You must be able to recover data from the Dev database if data is lost accidentally. You have a Recovery Point Objective (RPO) of one day.
Classification Changes
You plan to change the way customers are classified. The new classifications will have four levels based on the number of orders. Classifications may be removed or added in the future. Management requests that historical data be maintained for the previous classifications. Security A group of junior database administrators must be able to manage security for the Sales database. The junior database administrators will not have any other administrative rights. A. Datum wants to track which users run each stored procedure.
Storage
ADatum has limited storage. Whenever possible, all storage space should be minimized for all databases and all backups. Error Handling
There is currently no error handling code in any stored procedure.
You plan to log errors in called stored procedures and nested stored procedures. Nested stored procedures are never called directly.
You need to recommend a disaster recovery solution for the Dev database. What should you include in the recommendation?

A. The simple recovery model and full backups
B. The full recovery model, full backups, and transaction log backups

C. The full recovery model, full backups, and differential backups
D. The bulk-logged recovery model and full backups

Answer: A

Explanation:
Scenario:
You must be able to recover data from the Dev database if data is lost accidentally. You have a Recovery Point Objective (RPO) of one day.
- The simple recovery model provides the simplest form of backup and restore. This recovery model supports both database backups and file backups, but does not support log backups. Transaction log data is backed up only with the associated user data.
The absence of log backups simplifies managing backup and restore. However, a database can be restored only to the end of the most recent backup.

NEW QUESTION 150
- (Exam Topic 2)
Overview
General Overview
ADatum Corporation has offices in Miami and Montreal.
The network contains a single Active Directory forest named adatum.com. The offices connect to each other by using a WAN link that has 5-ms latency. A. Datum standardizes its database platform by using SQL Server 2014 Enterprise edition.
Databases
Each office contains databases named Sales, Inventory, Customers, Products, Personnel, and Dev.
Servers and databases are managed by a team of database administrators. Currently, all of the database administrators have the same level of permissions on all of the servers and all of the databases.
The Customers database contains two tables named Customers and Classifications. The following graphic shows the relevant portions of the tables:

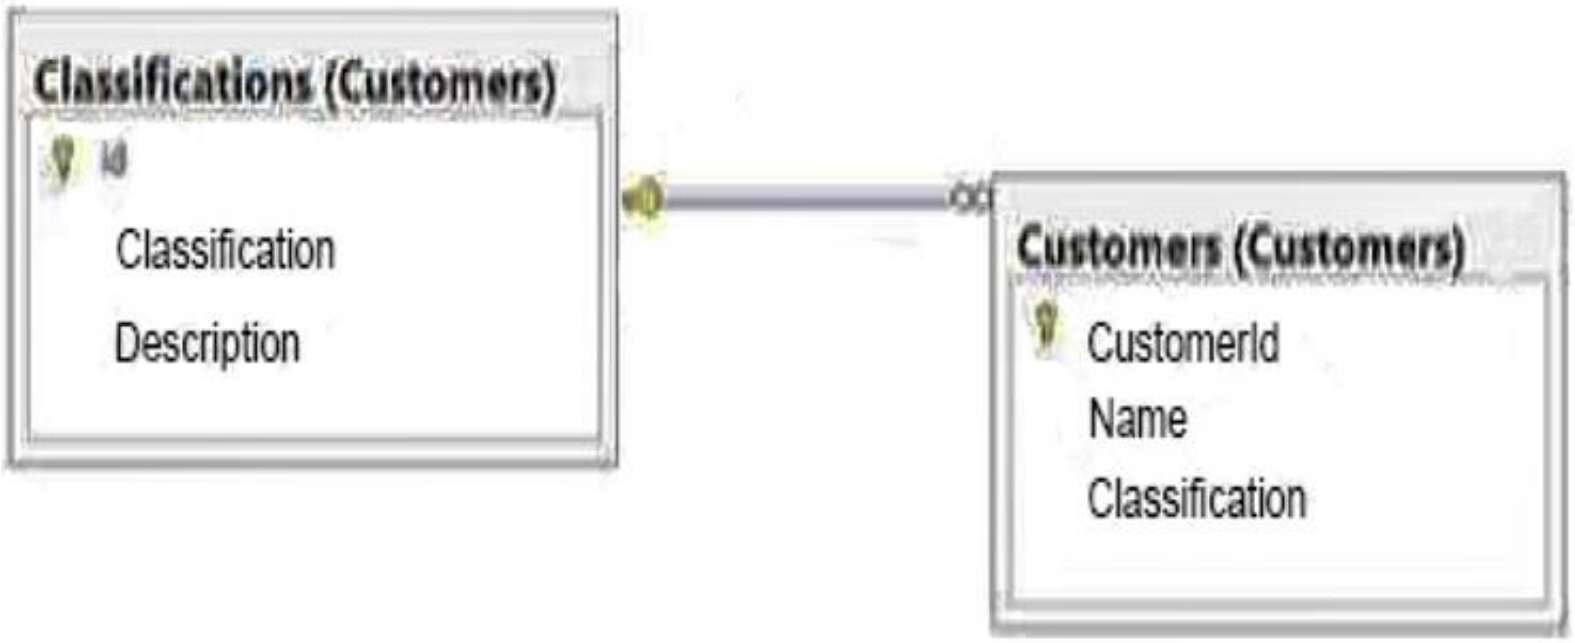

The following table shows the current data in the Classifications table:

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

The Inventory database is updated frequently. The database is often used for reporting.
A full backup of the database currently takes three hours to complete. Stored Procedures
A stored procedure named USP_1 generates millions of rows of data for multiple reports. USP_1 combines data from five different tables from the Sales and Customers databases in a table named Table1.
After Table1 is created, the reporting process reads data from Table1 sequentially several times. After the process is complete, Table1 is deleted.
A stored procedure named USP_2 is used to generate a product list. The product list contains the names of products grouped by category.
USP_2 takes several minutes to run due to locks on the tables the procedure accesses. The locks are caused by USP_1 and USP_3.
A stored procedure named USP_3 is used to update prices. USP_3 is composed of several UPDATE statements called in sequence from within a transaction.
Currently, if one of the UPDATE statements fails, the stored procedure fails. A stored procedure named USP_4 calls stored procedures in the Sales, Customers, and Inventory databases.
The nested stored procedures read tables from the Sales, Customers, and Inventory databases. USP_4 uses an EXECUTE AS clause.
All nested stored procedures handle errors by using structured exception handling. A stored procedure named USP_5 calls several stored procedures in the same database. Security checks are performed each time USP_5 calls a stored procedure.
You suspect that the security checks are slowing down the performance of USP_5. All stored procedures accessed by user applications call nested stored procedures.
The nested stored procedures are never called directly. Design Requirements
Data Recovery
You must be able to recover data from the Inventory database if a storage failure occurs. You have a Recovery Time Objective (RTO) of 5 minutes.
You must be able to recover data from the Dev database if data is lost accidentally. You have a Recovery Point Objective (RPO) of one day.
Classification Changes
You plan to change the way customers are classified. The new classifications will have four levels based on the number of orders. Classifications may be removed or added in the future. Management requests that historical data be maintained for the previous classifications. Security A group of junior database administrators must be able to manage security for the Sales database. The junior database administrators will not have any other administrative rights. A. Datum wants to track which users run each stored procedure.
Storage
ADatum has limited storage. Whenever possible, all storage space should be minimized for all databases and all backups.
Error Handling
There is currently no error handling code in any stored procedure.
You plan to log errors in called stored procedures and nested stored procedures. Nested stored procedures are never called directly.

You need to recommend a change to USP_3 to ensure that the procedure continues to execute even if one of the UPDATE statements fails.
Which change should you recommend?

A. Set the XACT_ABORT option to off.
B. Set the XACT_ABORT option to on.
C. Set the IMPLICIT_TRANSACTIONS option to off.
D. Set the IMPLICIT_TRANSACTIONS option to on.

Answer: A

Explanation:
- Scenario: A stored procedure named USP_3 is used to update prices. USP_3 is composed of several UPDATE statements called in sequence from within a transaction. Currently, if one of the UPDATE statements fails, the stored procedure continues to execute.
- When SET XACT_ABORT is OFF, in some cases only the Transact-SQL statement that raised the error is rolled back and the transaction continues processing.

NEW QUESTION 154
- (Exam Topic 2)
Overview
Application Overview
Contoso, Ltd., is the developer of an enterprise resource planning (ERP) application.
Contoso is designing a new version of the ERP application. The previous version of the ERP application used SQL Server 2008 R2.
The new version will use SQL Server 2014.
The ERP application relies on an import process to load supplier data. The import process updates thousands of rows simultaneously, requires exclusive access to the database, and runs daily.
You receive several support calls reporting unexpected behavior in the ERP application. After analyzing the calls, you conclude that users made changes directly to the tables in the database.
Tables
The current database schema contains a table named OrderDetails.
The OrderDetails table contains information about the items sold for each purchase order. OrderDetails stores the product ID, quantities, and discounts applied to each product in a purchase order.
The product price is stored in a table named Products. The Products table was defined by using the SQL_Latin1_General_CP1_CI_AS collation.
A column named ProductName was created by using the varchar data type. The database contains a table named Orders.
Orders contains all of the purchase orders from the last 12 months. Purchase orders that are older than 12 months are stored in a table named OrdersOld.
The previous version of the ERP application relied on table-level security. Stored Procedures
The current version of the database contains stored procedures that change two tables. The following shows the relevant portions of the two stored procedures:

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

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

Customer Problems Installation Issues
The current version of the ERP application requires that several SQL Server logins be set up to function correctly. Most customers set up the ERP application in multiple locations and must create logins multiple times.
Index Fragmentation Issues
Customers discover that clustered indexes often are fragmented. To resolve this issue, the customers defragment the indexes more frequently. All of the tables affected by fragmentation have the following columns that are used as the clustered index key:

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

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

Users report that when they search product names, the search results exclude product names that contain accents, unless the search string includes the accent.
Missing Data Issues
Customers report that when they make a price change in the Products table, they cannot retrieve the price that the item was sold for in previous orders.
Query Performance Issues
Customers report that query performance degrades very quickly. Additionally, the customers report that users cannot run queries when SQL Server runs maintenance tasks. Import Issues During the monthly import process, database administrators receive many supports call from users who report that they cannot access the supplier data. The database administrators want to reduce the amount of time required to import the data.
Design Requirements
File Storage Requirements
The ERP database stores scanned documents that are larger than 2 MB. These files must only be accessed through the ERP application. File access must have the best possible read and write performance.
Data Recovery Requirements
If the import process fails, the database must be returned to its prior state immediately. Security Requirements
You must provide users with the ability to execute functions within the ERP application, without having direct access to the underlying tables.
Concurrency Requirements
You must reduce the likelihood of deadlocks occurring when Sales.Prod and Sales.Proc2 execute.
You need to recommend a solution that addresses the concurrency requirement. What should you recommend?

A. Call the stored procedures in a Distributed Transaction Coordinator (DTC) transaction.
B. Modify the stored procedures to update tables in the same order for all of the stored procedures.
C. Make calls to Sales.Proc1 and Sales.Proc2 synchronously.
D. Break each stored procedure into two separate procedures, one that changes Sales.Table1 and one that changes Sales.Table2.

Answer: B

Explanation:
- Concurrency Requirements
You must reduce the likelihood of deadlocks occurring when Sales.Proc1 and Sales.Proc2 execute.

NEW QUESTION 159
- (Exam Topic 2)
Overview
Application Overview
Contoso, Ltd., is the developer of an enterprise resource planning (ERP) application.
Contoso is designing a new version of the ERP application. The previous version of the ERP application used SQL Server 2008 R2.
The new version will use SQL Server 2014.
The ERP application relies on an import process to load supplier data. The import process updates thousands of rows simultaneously, requires exclusive access to the database, and runs daily.
You receive several support calls reporting unexpected behavior in the ERP application. After analyzing the calls, you conclude that users made changes directly to the tables in the database.
Tables
The current database schema contains a table named OrderDetails.
The OrderDetails table contains information about the items sold for each purchase order. OrderDetails stores the product ID, quantities, and discounts applied to each product in a purchase order.
The product price is stored in a table named Products. The Products table was defined by using the SQL_Latin1_General_CP1_CI_AS collation.
A column named ProductName was created by using the varchar data type. The database contains a table named Orders.
Orders contains all of the purchase orders from the last 12 months. Purchase orders that are older than 12 months are stored in a table named OrdersOld.
The previous version of the ERP application relied on table-level security. Stored Procedures
The current version of the database contains stored procedures that change two tables. The following shows the relevant portions of the two stored procedures:

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

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

Customer Problems Installation Issues
The current version of the ERP application requires that several SQL Server logins be set up to function correctly. Most customers set up the ERP application in multiple locations and must create logins multiple times.
Index Fragmentation Issues
Customers discover that clustered indexes often are fragmented. To resolve this issue, the customers defragment the indexes more frequently. All of the tables affected by fragmentation have the following columns that are used as the clustered index key:

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

Backup Issues
Customers who have large amounts of historical purchase order data report that backup time is unacceptable. Search Issues
Users report that when they search product names, the search results exclude product names that contain accents, unless the search string includes the accent.
Missing Data Issues
Customers report that when they make a price change in the Products table, they cannot retrieve the price that the item was sold for in previous orders.
Query Performance Issues
Customers report that query performance degrades very quickly. Additionally, the customers report that users cannot run queries when SQL Server runs maintenance tasks. Import Issues During the monthly import process, database administrators receive many supports call from users who report that they cannot access the supplier data. The database administrators want to reduce the amount of time required to import the data.
Design Requirements
File Storage Requirements
The ERP database stores scanned documents that are larger than 2 MB. These files must only be accessed through the ERP application. File access must have the best possible read and write performance.
Data Recovery Requirements
If the import process fails, the database must be returned to its prior state immediately. Security Requirements
You must provide users with the ability to execute functions within the ERP application, without having direct access to the underlying tables.
Concurrency Requirements
You must reduce the likelihood of deadlocks occurring when Sales.Prod and Sales.Proc2 execute. What should you recommend for the updates to Sales.TransactionHistory?

A. a REPEATABLE READ isolation level
B. implicit transactions
C. query hints
D. a SNAPSHOT isolation level

Answer: A

NEW QUESTION 160
- (Exam Topic 2)
You administer a Microsoft SQL Server 2016 database named Contoso on a server named Server01. You need to prevent users from disabling server audits in Server01.
What should you create?

A. an Alert
B. a Resource Pool
C. an Extended Event session
D. a Policy
E. a Database Audit Specification
F. a SQL Profiler Trace
G. a Server Audit Specification

Answer: D

NEW QUESTION 164
- (Exam Topic 2)
You administer a Microsoft SQL Server 2016 server.
The MSSQLSERVER service uses a domain account named CONTOSO\SQLService. You plan to configure Instant File Initialization.
You need to ensure that Data File Autogrow operations use Instant File Initialization. What should you do? Choose all that apply.

A. Restart the SQL Server Agent Service.
B. Disable snapshot isolation.
C. Restart the SQL Server Service.
D. Add the CONTOSO\SQLService account to the Perform Volume Maintenance Tasks local security policy.
E. Add the CONTOSO\SQLService account to the Server Operators fixed server role.
F. Enable snapshot isolation.

Answer: CD

Explanation:
How To Enable Instant File Initialization
Open Local Security Policy and go to Local Policies –> User Rights Assignment.
Double click Perform Volume Maintenance Tasks and add your SQL Server database engine service account.
Restart the SQL Server service using SQL Server Configuration Manager and this setting should now be enabled.
References:
http://msdn.microsoft.com/en-us/library/ms175935.aspx

NEW QUESTION 167
- (Exam Topic 2)

You administer a Microsoft SQL Server 2016 instance.
After a routine shutdown, the drive that contains tempdb fails. You need to be able to start the SQL Server.
What should you do?

A. Modify tempdb location in startup parameters.
B. Start SQL Server in minimal configuration mode.
C. Start SQL Server in single-user mode.
D. Configure SQL Server to bypass Windows application logging.

Answer: B

NEW QUESTION 171
- (Exam Topic 2)
You have two SQL Server instances named SQLDev and SQLProd that have access to various storage media. You plan to synchronize SQLDev and SQLProd.
You need to recommend a solution that meets the following requirements:
The database schemas must be synchronized from SQLDev to SQLProd.
The database on SQLDev must be deployed to SQLProd by using a package.
The package must support being deployed to SQL Azure.
What should you recommend? More than one answer choice may achieve the goal. Select the BEST answer.

A. A database snapshot
B. A data-tier application
C. Change data capture
D. SQL Server Integration Services (SSIS)

Answer: B

Explanation:
*SIS supports connections to SQL Database by using the ADO.NET provider. OLEDB is not supported at this time. You can build the SSIS package connecting to SQL Database and create the data flow tasks the same way as you would against a typical on premise SQL Server.
http://technet.microsoft.com/en-us/library/ee210546.aspx

NEW QUESTION 173
- (Exam Topic 2)
Overview
General Overview
ADatum Corporation has offices in Miami and Montreal.
The network contains a single Active Directory forest named adatum.com. The offices connect to each other by using a WAN link that has 5-ms latency. A. Datum standardizes its database platform by using SQL Server 2014 Enterprise edition.
Databases
Each office contains databases named Sales, Inventory, Customers, Products, Personnel, and Dev. Servers and databases are managed by a team of database administrators. Currently, all of the database
administrators have the same level of permissions on all of the servers and all of the databases.
The Customers database contains two tables named Customers and Classifications. The following graphic shows the relevant portions of the tables:
Classifications (Customers)

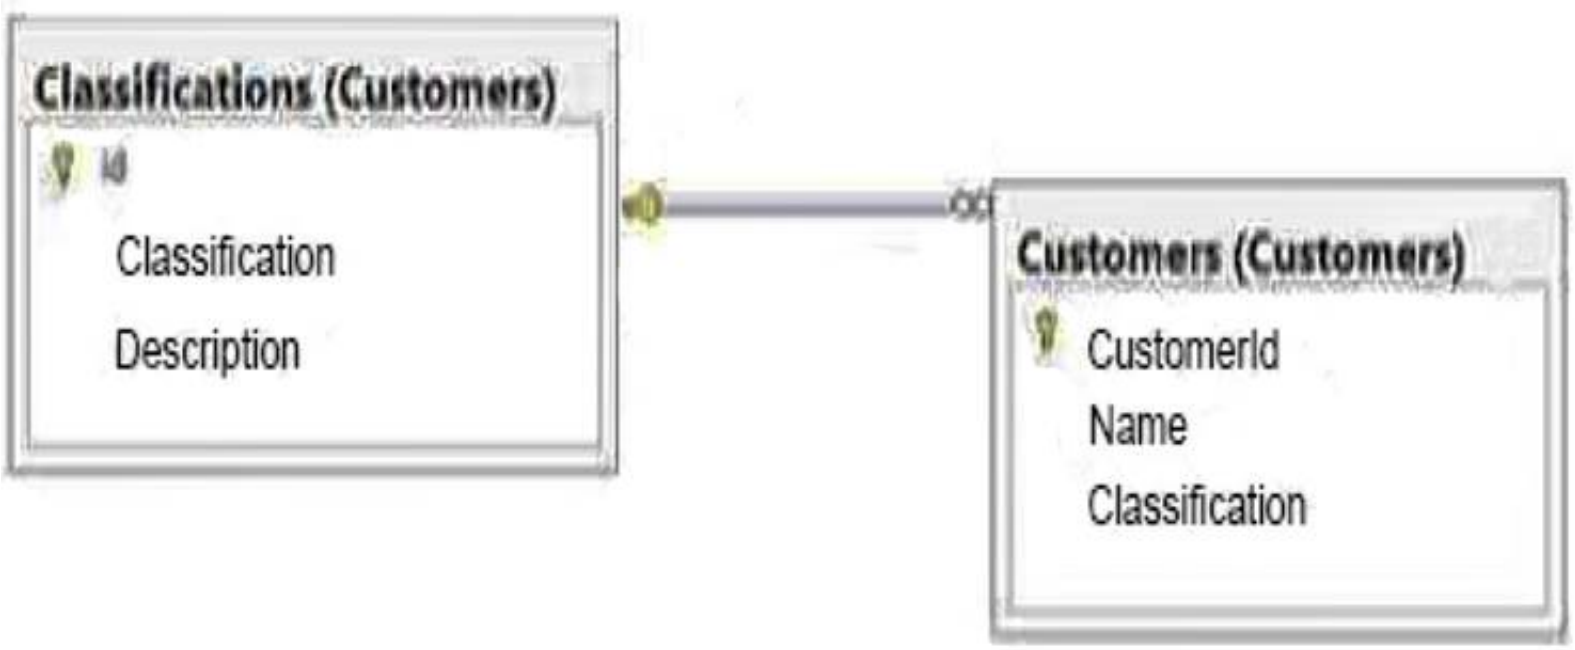

The following table shows the current data in the Classifications table:

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

The Inventory database is updated frequently. The database is often used for reporting.
A full backup of the database currently takes three hours to complete. Stored Procedures
A stored procedure named USP_1 generates millions of rows of data for multiple reports. USP_1 combines data from five different tables from the Sales and Customers databases in a table named Table1.
After Table1 is created, the reporting process reads data from Table1 sequentially several times. After the process is complete, Table1 is deleted.
A stored procedure named USP_2 is used to generate a product list. The product list contains the names of products grouped by category.
USP_2 takes several minutes to run due to locks on the tables the procedure accesses. The locks are caused by USP_1 and USP_3.

A stored procedure named USP_3 is used to update prices. USP_3 is composed of several UPDATE statements called in sequence from within a transaction. Currently, if one of the UPDATE statements fails, the stored procedure fails. A stored procedure named USP_4 calls stored procedures in the Sales, Customers, and Inventory databases.
The nested stored procedures read tables from the Sales, Customers, and Inventory databases. USP_4 uses an EXECUTE AS clause.
All nested stored procedures handle errors by using structured exception handling. A stored procedure named USP_5 calls several stored procedures in the same database. Security checks are performed each time USP_5 calls a stored procedure.
You suspect that the security checks are slowing down the performance of USP_5. All stored procedures accessed by user applications call nested stored procedures.
The nested stored procedures are never called directly. Design Requirements
Data Recovery
You must be able to recover data from the Inventory database if a storage failure occurs. You have a Recovery Time Objective (RTO) of 5 minutes.
You must be able to recover data from the Dev database if data is lost accidentally. You have a Recovery Point Objective (RPO) of one day.
Classification Changes
You plan to change the way customers are classified. The new classifications will have four levels based on the number of orders. Classifications may be removed or added in the future. Management requests that historical data be maintained for the previous classifications. Security A group of junior database administrators must be able to manage security for the Sales database. The junior database administrators will not have any other administrative rights. A. Datum wants to track which users run each stored procedure.
Storage
ADatum has limited storage. Whenever possible, all storage space should be minimized for all databases and all backups.
Error Handling
There is currently no error handling code in any stored procedure.
You plan to log errors in called stored procedures and nested stored procedures. Nested stored procedures are never called directly.
You need to recommend a disaster recovery strategy for the Inventory database. What should you include in the recommendation?

A. Log shipping
B. SQL Server Failover Clustering
C. AlwaysOn availability groups
D. Peer-to-peer replication

Answer: A

Explanation:
Scenario:
- You must be able to recover data from the Inventory database if a storage failure occurs. You have a Recovery Point Objective (RPO) of one hour.
- A. Datum Corporation has offices in Miami and Montreal.
- 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.

NEW QUESTION 177
.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual 70-764 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the 70-764 Product From:

https://www.2passeasy.com/dumps/70-764/

Money Back Guarantee

70-764 Practice Exam Features:

* 70-764 Questions and Answers Updated Frequently

* 70-764 Practice Questions Verified by Expert Senior Certified Staff

* 70-764 Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* 70-764 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year