Microsoft

Exam Questions 70-764

Administering a SQL Database Infrastructure (beta)

NEW QUESTION 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 have an on-premises server that runs Microsoft SQL Server 2016 Standard Edition. You need to identify missing indexes.
What should you use?

A. Activity Monitor
B. Sp_who3
C. SQL Server Management Studio (SSMS) Object Explorer
D. SQL Server Data Collector
E. SQL Server Data Tools (SSDT)
F. SQL Server Configuration Manager

Answer: D

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

NEW QUESTION 2
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 3
You are a database administrator for a Microsoft SQL Server 2016 environment.
You want to deploy a new application that will scale out the workload to at least five different SQL Server instances.
You need to ensure that for each copy of the database, users are able to read and write data that will then be synchronized between all of the database instances.
Which feature should you use?

A. Database Mirroring
B. Peer-to-Peer Replication
C. Log Shipping
D. Availability Groups

Answer: B

NEW QUESTION 4
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 5
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 6
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 7
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 8
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 9
You administer a Microsoft SQL Server 2016 database named Orders.
Users report that during peak usage periods, certain operations are taking more time than expected. Your initial analysis suggests that blocking is the cause.
You need to gather more data to be able to determine which processes are being blocked and to identify the root cause.
What should you do?

A. Start a trace using SQL Server Profiler to catch the Lock: Deadlock event.
B. Use sp_configure to set the blocked process threshol
C. Start a trace using SQL Server Profiler to catch the Blocked Process Report event.
D. Schedule a SQL Agent job to run every 60 seconds and insert the results of executing the sys.dm_os_wait_stats DMV into a table.
E. Use System Monitor to catch the Lock Waits/sec event.

Answer: B

NEW QUESTION 10
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 10
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 12
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 15
Note: This question is part of a series of question that present the same scenario. Each question in the series contains I 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 SQL Saver instance. Each instance hosts many database. You observe I/O corruption on some of the instance.
You need to perform the following actions:
• identify databases where the PAGE verity 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 != 'CHECKSUM'
GO
```

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

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

Does the solution meet the goal?

A. Yes

B. No

Answer: A

NEW QUESTION 19
You have configured Resource Governor with three resource pools.
You have assigned the first resource pool a minimum CPU and memory value of 20%. You have assigned the second resource pool a minimum CPU and memory value of 30%. You want to assign maximum CPU and memory values to the third resource pool.
What is the maximum CPU and memory value you can assign to this resource pool?

A. 30%
B. 50%
C. 70%
D. 100%

Answer: B

Explanation: The maximum resource value assigned to the third pool is 100%; the sum of the minimum resource values assigned to the other pools is 50%.

NEW QUESTION 20
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 23
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 26
You use Microsoft SQL Server 2016 to write code for a transaction that contains several statements.
There is high contention between readers and writers on several tables used by your transaction. You need to minimize the use of the tempdb space.
You also need to prevent reading queries from blocking writing queries. Which isolation level should you use?

A. SERIALIZABLE
B. SNAPSHOT
C. READ COMMITTED SNAPSHOT
D. REPEATABLE READ

Answer: C

Explanation: For most applications, read committed isolation using row versioning is recommended over snapshot isolation for the following reasons:
It consumes less tempdb space than snapshot isolation. Etc.
References: https://msdn.microsoft.com/en-us/library/ms188277.aspx

NEW QUESTION 31
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.

Answer Area

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

Answer:

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 32
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 36
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 41
You administer a Microsoft SQL Server 2016 database.
You have a SQL Server Agent job instance that runs using the service account. You have a job step within the job that requires elevated privileges.
You need to ensure that the job step can run using a different user account. What should you use?

A. a schedule
B. an alert
C. an operator
D. a proxy

Answer: D

NEW QUESTION 43

You administer a Microsoft SQL Server 2016 database named Contoso on a server named Server01. You need to collect data for a long period of time to troubleshoot wait statistics when querying Contoso. You also need to ensure minimum impact to the server.
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

Answer: C

NEW QUESTION 44
You plan to migrate the db to azure.
You verify that all objects are valid for azure sql database. You need to ensure that users and logins are migrated to azure.
What should you do?

A. Use the Copy Database wizard
B. Use the Database Transfer wizard
C. Use the SQL Management Studio to deploy the db to azure
D. Back up the databasae from the local server and restore it to azure

Answer: CD

NEW QUESTION 45
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.

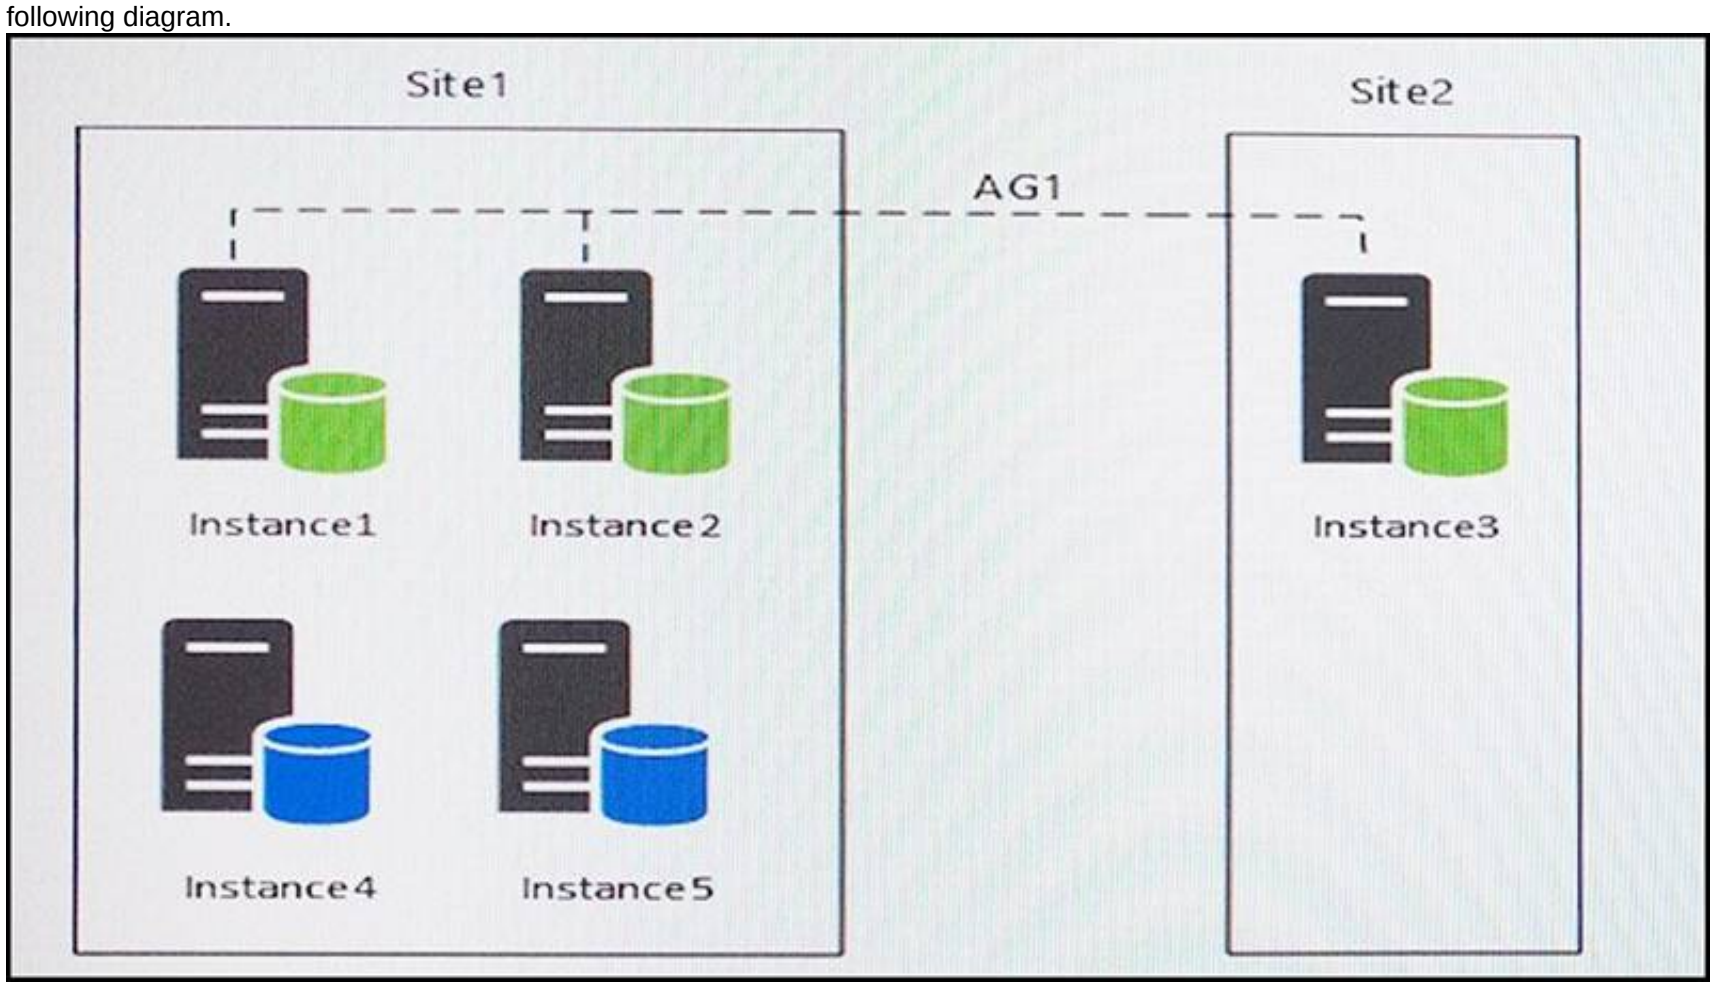

You have an Always On Availability group named AG1. The details for AG1 are shown in the following table.

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 48
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 a server that runs Microsoft SQL Server 2016 Web edition. The server has a default instance that hosts a database named DB1.
You need to ensure that you can perform auditing at the database level for DB1.
Solution: You migrate DB1 to the default instance on a server that runs Microsoft SQL Server 2016 Standard edition.
Does the solution meet the goal?

A. Yes
B. No

Answer: B

Explanation: All editions of SQL Server support server level audits. All editions support database level audits beginning with SQL Server 2016 SP1. Prior to that, database level auditing was limited to Enterprise, Developer, and Evaluation editions.
References:
https://docs.microsoft.com/en-us/sql/relational-databases/security/auditing/sql-server-audit-database-engine

NEW QUESTION 50
You administer a Microsoft SQL Server 2016 environment.
One of the SQL Server 2016 instances contains a database named Sales. You plan to migrate Sales to Windows Azure SQL Database.
To do so, you need to implement a contained database.
What should you do? (Each correct answer presents part of the solution. Choose all that apply.)

A. Set database containment to AZURE.
B. Enable server property contained database authentication.
C. Disable server property cross db ownership chaining.
D. Set database containment to PARTIAL.
E. Disable server property contained database authentication.
F. Set database containment to FULL.

Answer: BD

NEW QUESTION 55
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

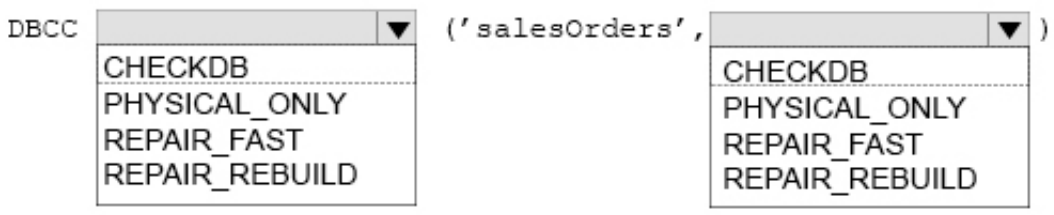

Answer:

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
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

The differential backup of the reporting database fails. Then, the reporting database fails at 14:00 hours.
You need to ensure that the reporting database is restored. You also need to ensure that data loss is minimal.
What should you do?

A. Restore the latest full backup, and restore the latest differential backu

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

Answer: C

Explanation: The differential backup of the reporting database has failed, so it can't be used.

NEW QUESTION 61
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 62
You are the administrator of a Microsoft SQL Server 2016 server. Some applications consume significant resources.
You need to manage the server workload by restricting resource-intensive applications. You need to dynamically limit resource consumption.
What should you do?

A. Set up Service Broker to ensure that applications are not allowed to consume more than the specified amount of resources.
B. Configure Resource Pools, Workload Groups, and Classifier Function, and then enable the Resource Governor.
C. Configure Extended Events to monitor and restrict resource limits allowed by each application type.
D. Create a new Plan Guide with a Scope Type of sql and define the resource limits for each application.

Answer: B

NEW QUESTION 66
You are configuring log shipping for a Microsoft SQL Server database named salesOrders. You run the following Transact-SQL script:

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

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

Answer Area

A dedicated file share **[answer choice]** used to store the backups.

- is
- is not

A SQL Server monitor instance **[answer choice]** on a server named ADATUM-SQL11.

- runs
- does not run

Backup files will be deleted after **[answer choice]**.

- 24 hours
- 48 hours
- 72 hours

The backup job will run every **[answer choice]**.

- 15 minutes
- 60 minutes
- 24 hours

Answer:

Explanation: Box 1: is
The dedicated backup file share is \\localhost\Backup Box 2: does not run
The only thing with a name related to ADATM-SQL11 is the schedule name. Box 3: 72 hours
4320 minutes equals 72 hours.
Note: @backup_retention_period=] backup_retention_period
Isthe length of time, in minutes, to retain the log backup file in the backup directory on the primary server. backup_retention_period is int, with no default, and cannot be NULL.
Box 4: 15 minutes.
[@freq_subday_type =] freq_subday_type
Specifies the units for freq_subday_interval. freq_subday_typeis int, with a default of 0, and can be one of these values.
Here it is 4, which means minutes.
[@freq_subday_interval =] freq_subday_interval
The number of freq_subday_type periods to occur between eachexecution of a job. freq_subday_intervalis int, with a default of 0.
Note: Interval should be longer than 10 seconds. freq_subday_interval is ignored in those cases where freq_subday_type is equal to 1.
Here it is 15. References:
https://docs.microsoft.com/en-us/sql/relational-databases/system-stored-procedures/sp-add-schedule-transact-sql https://docs.microsoft.com/en-us/sql/relational-databases/system-stored-procedures/sp-add-log-shipping-primary

NEW QUESTION 70
You are planning to deploy log shipping for Microsoft SQL Server and store all backups on a dedicated fileshare.
You need to configure the servers to perform each log shipping step.
Which server instance should you configure to perform each action? To answer, select the appropriate server instances in the dialog box in the answer area.

Answer Area

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

Answer:

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

NEW QUESTION 73
You plan to integrate an on-premises Microsoft SQL Server environment with Microsoft Azure. You need to create the authentication object so that you can connect to Azure.
Which Windows PowerShell command or commands should you run?

A. Invoke-Sqlcmd “CREATE EXTERNAL DATA SOURCE MyAzureStorage WITH (LOCATION = ‘wasbs://Azure@myaccount.blob.core.windows.net/”, CREDENTIAL = Pa$$w0rd)”
B. New-SqlAzureKeyVaultColumnMasterKeySettings-KeyUrlhttps://myvault.vault.contoso.net:443/keys/C
C. Invoke-Sqlcmd “CREATE CREDENTIAL AzureCred WITH IDENTITY = ‘AzureKey’, SECRET = ‘Pa$$w0rd”
D. Invoke-Sqlcmd “CREATE LOGIN AzureCred WITH CREDENTIAL = ‘AzureKey’, PASSWORD = ‘Pa$$w0rd”

Answer: C

Explanation: Invoke-Sqlcmd runs a script containing statements supported by the SQL Server SQLCMD utility.
The following example creates a SQL Server credential for the Database Engine to use when accessing the Azure Key Vault using the SQL Server Connector for Microsoft Azure Key Vault.
CREATE CREDENTIAL Azure_EKM_TDE_cred WITH IDENTITY = 'ContosoKeyVault',
SECRET = 'EF5C8E094D2A4A769998D93440D8115DSECRET_DBEngine'

FOR CRYPTOGRAPHIC PROVIDER AzureKeyVault_EKM_Prov ;
References: https://docs.microsoft.com/en-us/sql/t-sql/statements/create-credential-transact-sql

NEW QUESTION 78
You administer all the deployments of Microsoft SQL Server 2016 in your company.
You need to ensure that data changes are sent to a non-SQL Server database server in near real time. You also need to ensure that data on the primary server is unaffected.
Which configuration should you use?

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

Answer: A

Explanation: SQL Server supports the following heterogeneous scenarios for transactional and snapshot replication:
Publishing data from SQL Server to non- SQL Server Subscribers.
Publishing data to and from Oracle has some restrictions.
References:
https://docs.microsoft.com/en-us/sql/relational-databases/replication/non-sql/heterogeneous-database-replication

NEW QUESTION 82
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 implement a process for importing data into WingDB.
Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

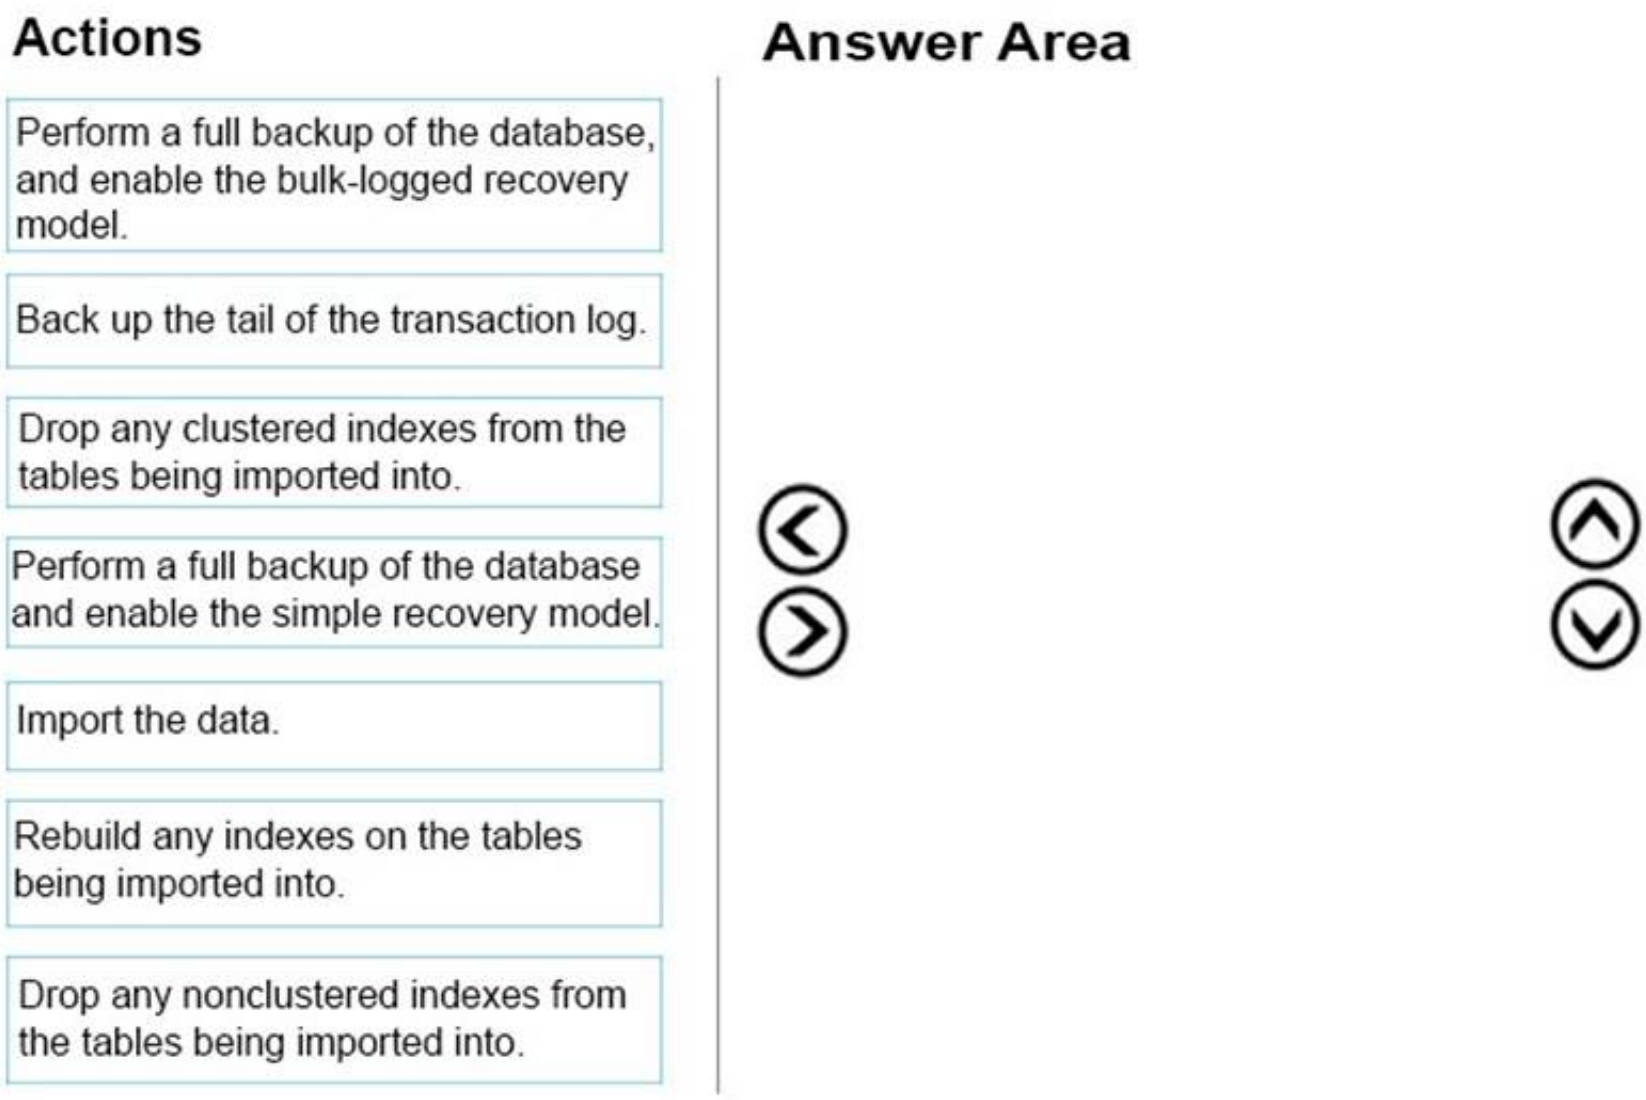

Answer:

Explanation: Step 1: Perform a full backup of the database and enable the bulk-logged recovery model. Not: Simple recovery model.
With the Simple recovery model we cannot minimize data loss. Step 2: Import the data
Step 3: Backup the tail of the transaction log.
For databases that use full and bulk-logged recovery, database backups are necessary but not sufficient. Transaction log backups are also required.
Note: Three recovery models exist: simple, full, and bulk-logged. Typically, a database uses the full recovery model or simple recovery model. A database can be switched to another recovery model at any time.
References:
https://docs.microsoft.com/en-us/sql/relational-databases/backup-restore/recovery-models-sql-server

NEW QUESTION 86
You administer a Microsoft SQL Server 2016 instance that has several SQL Server Agent jobs configured. SQL Server Agent jobs fail, the error messages returned by the job steps are truncated.
The following error message is an example of the truncated error message:
"Executed as user CONTOSO\ServiceAccount. ...0.4035.00 for 64-bit Copyright (C) Microsoft Corp
1984-2011. All rights reserved. Started 63513 PM Error 2012-06-23 183536.87 Code 0XC001000E Source UserImport Description Code 0x00000000 Source Log Import Activity Descript... The package execution fa... The step failed."
You need to ensure that all the details of the job step failures are retained for SQL Server Agent jobs. What should you do?

A. Expand agent logging to include information from all events.
B. Disable the Limit size of job history log feature.
C. Configure event forwarding.
D. Configure output files.

Answer: D

Explanation: When you have a multiple-step job, then log all steps against a single file. Check the 'Append output to existing file' checkbox for all steps in the job that execute after the initial step. This results in a log file with all of the job steps from the last job execution. Each time the first step executes (each time the job is kicked-off) the file will be overwritten, so we have a record of the last set of output.
References: https://www.mssqltips.com/sqlservertip/1411/verbose-sql-server-agent-logging/

NEW QUESTION 88
You administer a Microsoft SQL Server database named Contoso. You create a stored procedure named Sales.ReviewInvoice by running the following Transact-SQL statement:

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

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

Answer:

Explanation:

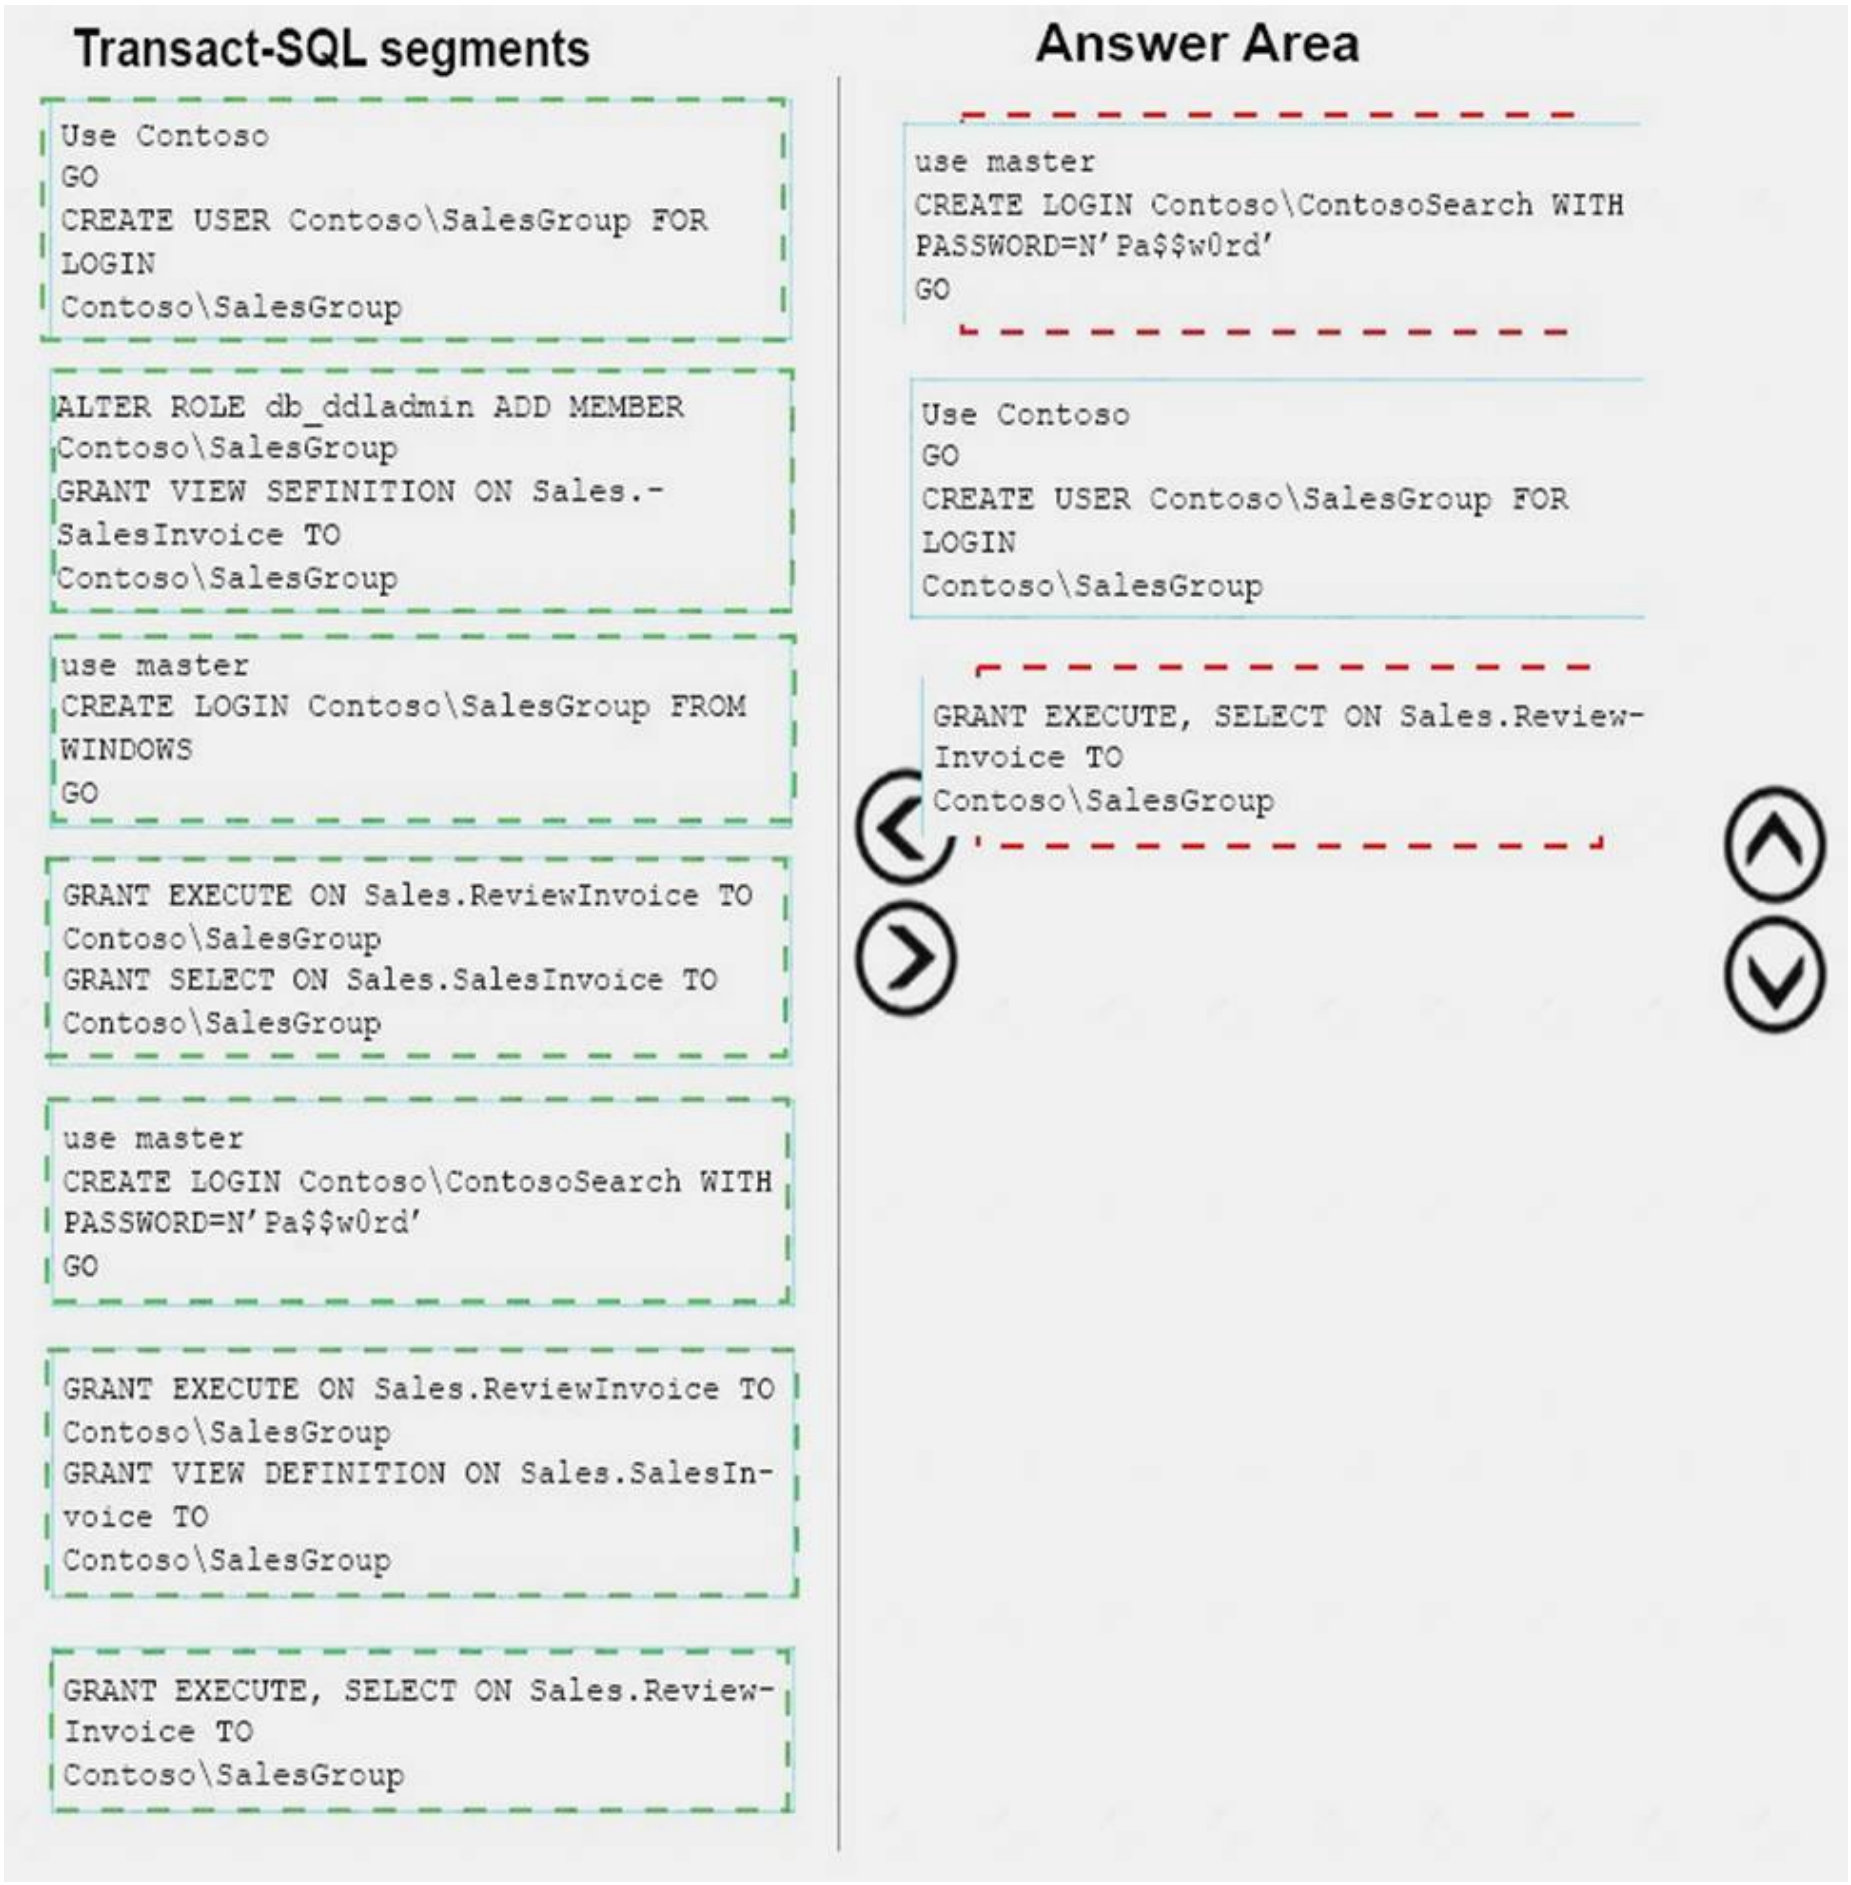

NEW QUESTION 93
Note: This question is part of a series of questions that use the same or similar answer choices. An answer choice may be correct for more than one question in the series. Each question is independent of the other questions in this series. Information and details provided in a question 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.
One instance hosts a user database named HRDB. The database contains sensitive human resources data. You need to grant an auditor permission to view the SQL Server audit logs while following the principle of
least privilege.
Which permission should you grant?

A. DDLAdmin
B. db_datawriter
C. dbcreator
D. dbo
E. View Database State
F. View Server State
G. View Definition
H. sysadmin

Answer: F

Explanation: Unless otherwise specified, viewing catalog views requires a principal to have one of the following:
Membership in the sysadmin fixed server role.
The CONTROL SERVER permission.
The VIEW SERVER STATE permission.
The ALTER ANY AUDIT permission.

The VIEW AUDIT STATE permission (gives only the principal access to the sys.server_audits catalog view).
References: https://technet.microsoft.com/en-us/library/cc280386(v=sql.110).aspx

NEW QUESTION 94
You create an availability group named HaContoso that has replicas named Server01/HA, Server02/HA, and Server03/HA.
Currently, Server01l/HA is the primary replica.
You need to ensure that the following requirements are met:
Backup operations occur on Server02/HA.
If Server02/HA is unavailable, backup operations occur on Server03/HA.
Backup operations do not occur on Server01/HA.
How should you configure HaContoso?

A. Set the backup preference of HaContoso to Prefer Secondar
B. Set the backup priority of Server02/HA to20. Set the backup priority of Server03/HA to 10.
C. Set the backup preference of HaContoso to Secondary onl
D. Set the backup priority of Server02/HA to20. Set the backup priority of Server03/HA to 10.
E. Set the backup preference of HaContoso to Secondary onl
F. Set the backup priority of Server02/HA to10. Set the backup priority of Server03/HA to 20.
G. set the exclude replica of Server01/HA to tru
H. Set the backup priority of Server02/HA to 10. Set the backup priority of Server03/HA to 20.

Answer: B

Explanation: Secondary only: Specifies that backups should never be performed on the primary replica. If the primary replica is the only replica online, the backup should not occur.
Backup Priority (Lowest=1, Highest=100)
Specifies your priority for performing backups on this replica relative to the other replicas in the same availability group. The value is an integer in the range of 0..100. 1 indicates the lowest priority, and 100 indicates the highest priority. If Backup Priority = 1, the availability replica would be chosen for performing backups only if no higher priority availability replicas are currently available.
References:
https://docs.microsoft.com/en-us/sql/database-engine/availability-groups/windows/configure-backup-on-availab

NEW QUESTION 96
You administer two instances of Microsoft SQL Server 2016.
You deploy an application that uses a database on the named instance.
The application is unable to connect to the database on the named instance. You need to ensure that the application can connect to the named instance. What should you do?

A. Use the Data Quality Client to configure the application.
B. Start the SQL Server Browser Service.
C. Use the Master Data Services Configuration Manager to configure the application.
D. Start the SQL Server Integration Services Service.

Answer: B

Explanation: The SQL ServerBrowser program runs as a Windows service. SQL Server Browser listens for incoming requests for Microsoft SQL Server resources and provides information about SQL Server instances installed on the computer. SQL Server Browser contributes to the following actions:
Browsing a list of available servers Connecting to the correct server instance Etc.
References: https://docs.microsoft.com/en-us/sql/tools/configuration-manager/sql-server-browser-service

NEW QUESTION 101
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.

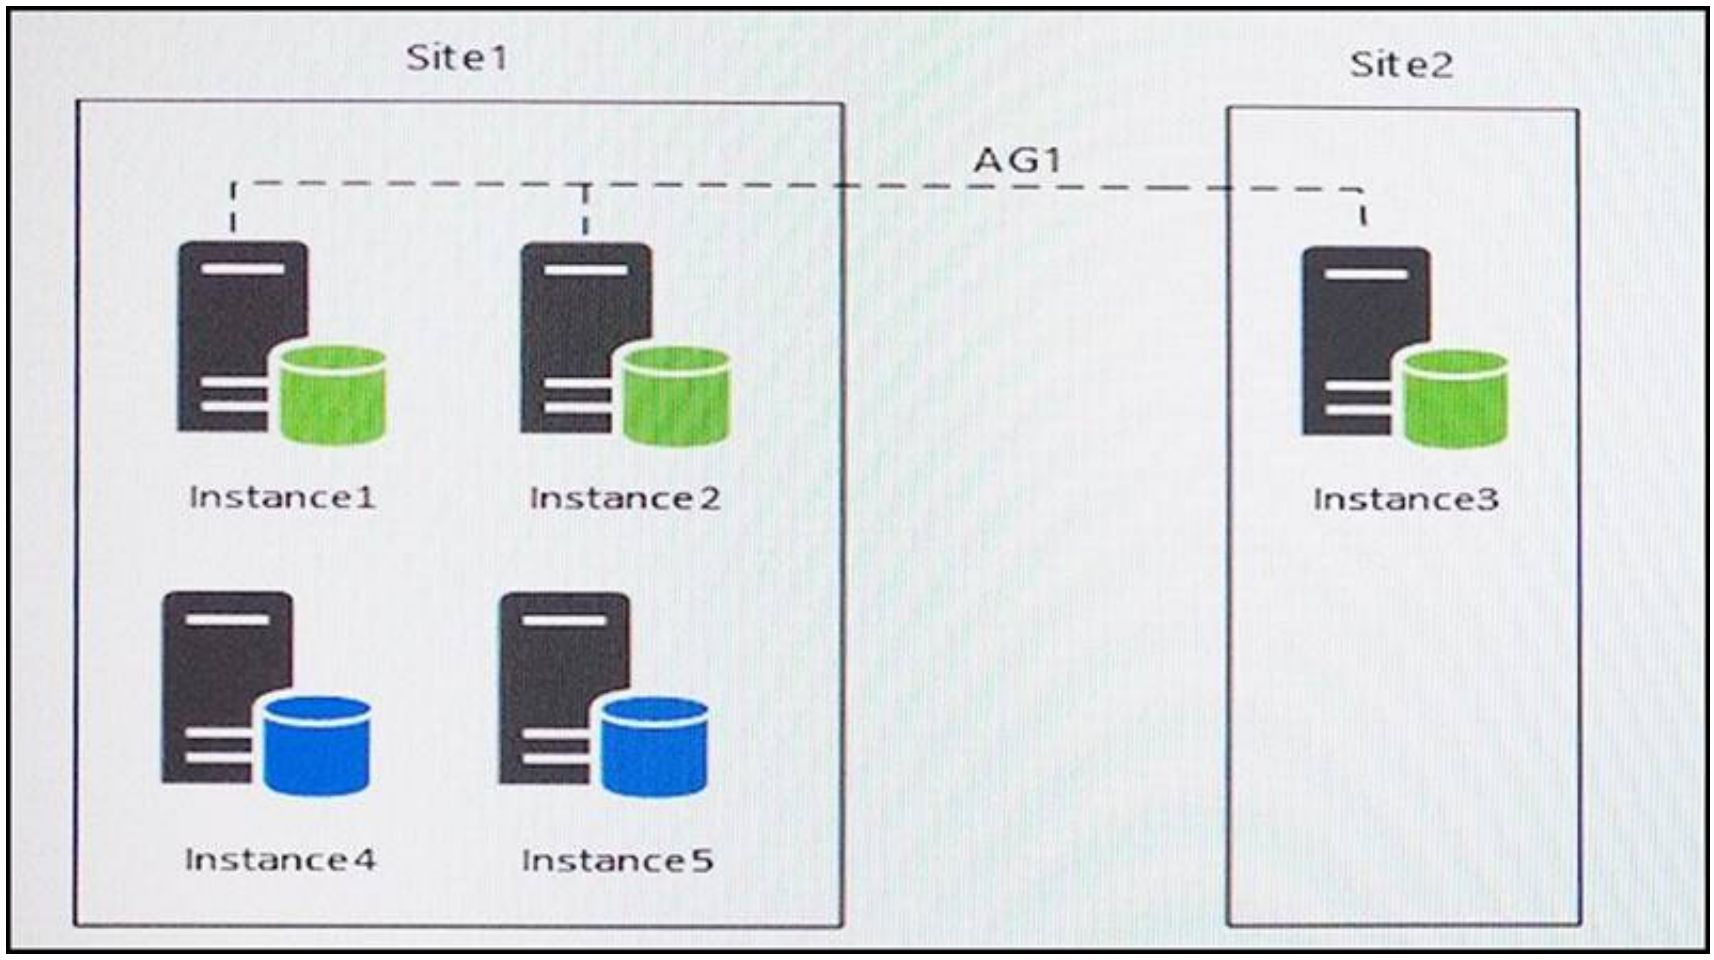

You have an Always On Availability group named AG1. The details for AG1 are shown in the following table.

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.	○	○

Answer:

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 106
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 111
You are configuring a new Microsoft SQL Server Always On Availability Group. You plan to configure a shared network location at \\DATA-CI1\SQL.
You need to create an availability group listener named AGL1 on port 1433.
In which order should you perform the actions? To answer, move all actions from the list of actions to the answer area and arrange them in the correct order.

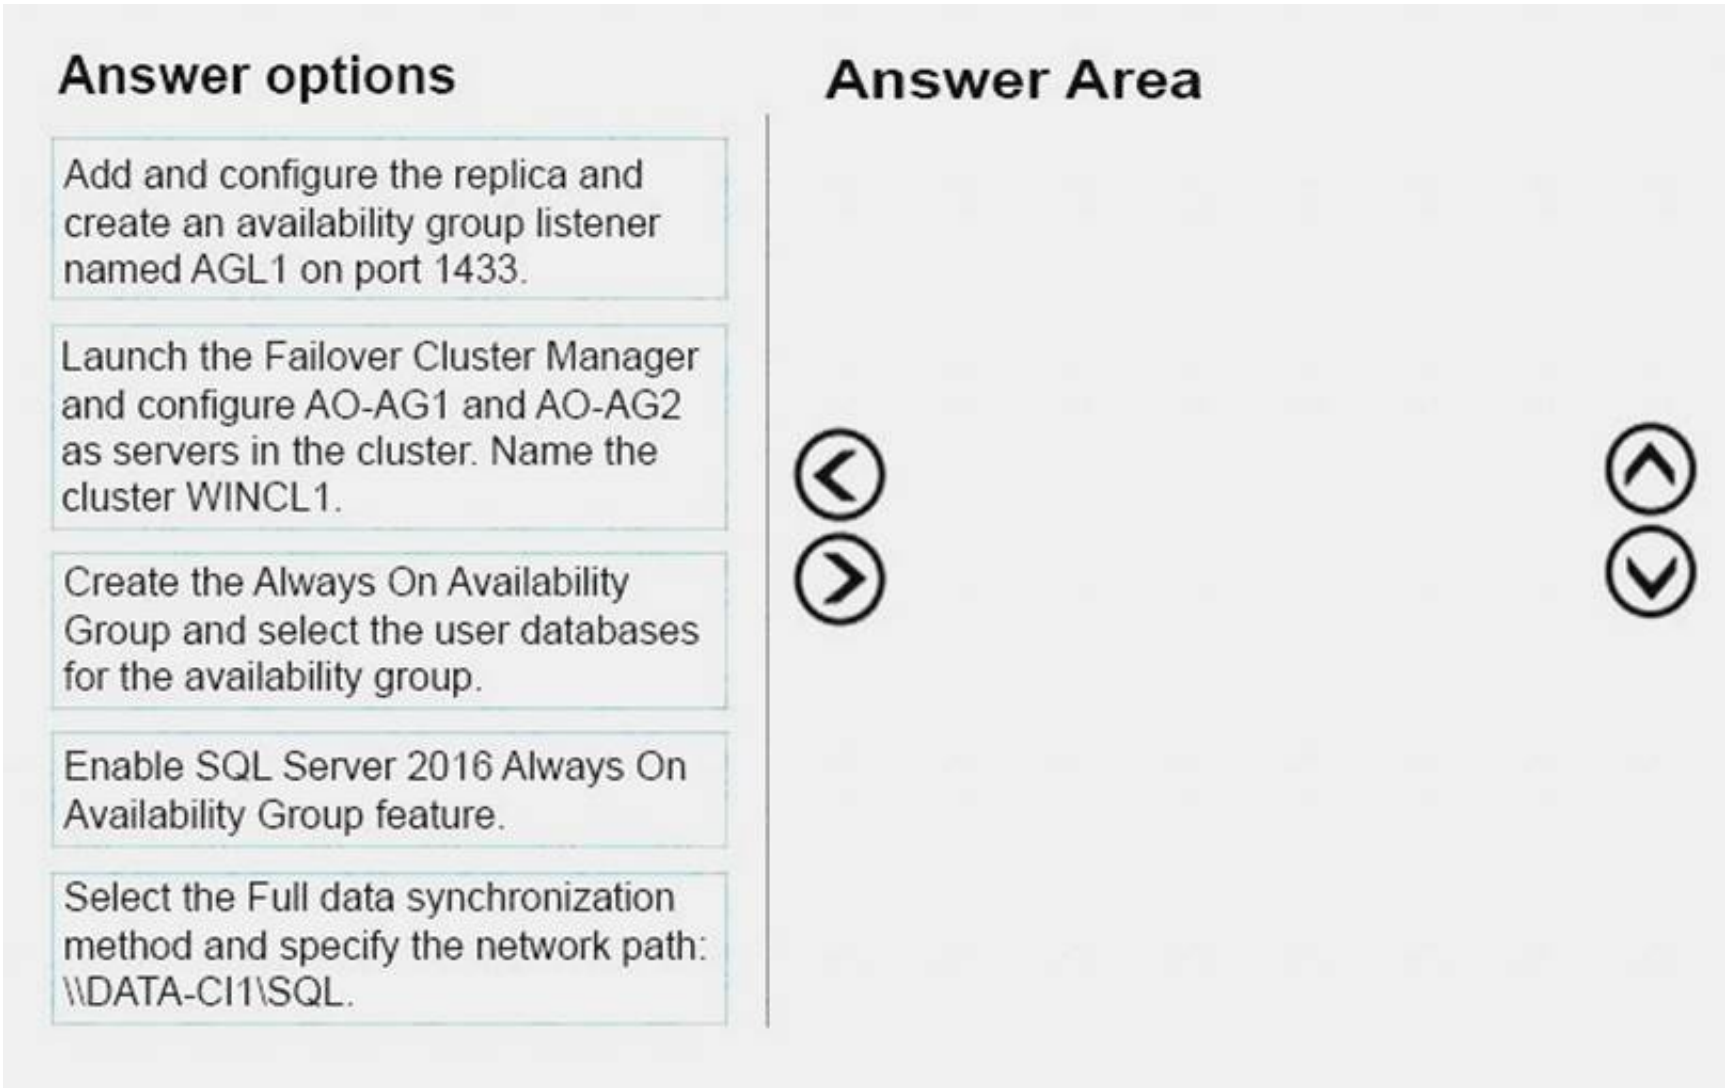

Answer:

Explanation: Step 1: Launch the Failover Cluster Manager and..
To support theAlways On availability groups feature, ensure that every computer that is to participate in one or more availability groups meets requirements including:
* Ensure that each computer is a node in a WSFC (Windows Server Failover Clustering). Step 2: Add andconfigure the replica and…
All the server instances that host availability replicas for an availability group must use the same SQL Server collation.
Step 3: Enable the SQL Server 2016 Always On Availability Group feature.
Enable the Always On availability groups feature on each server instance that will host an availability replica for any availability group. On a given computer, you can enable as many server instances for Always On availability groups as your SQL Server installation supports.
Step 4: Create the Always On Availability Group and..
Using Transact-SQL to create or configure an availability group listener Step 5: Select the Full data synchronization method and…
References: https://technet.microsoft.com/en-us/library/jj899851(v=sc.12).aspx
https://docs.microsoft.com/en-us/sql/database-engine/availability-groups/windows/create-or-configure-an-availa

NEW QUESTION 114
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 processes in an on-premises Microsoft SQL Server 2016 Standard Edition server.
You need to identify waits for resources and return only the following information:
a list of all databases on the SQL Server instance, along with information about the database files, their paths, and names
a list of the queries recently executed that use most of memory, disk, and network resources
What should you use?

A. Activity Monitor
B. Sp_who3
C. SQL Server Management Studio (SSMS) Object Explorer
D. SQL Server Data Collector
E. SQL Server Data Tools (SSDT)
F. SQL Server Configuration Manager

Answer: E

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

NEW QUESTION 116
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 120
You are the database administrator for a Microsoft SQL Server instance. You develop an Extended Events package to look for events related to application performance.
You need to change the event session to include SQL Server errors that are greater than error severity 15. Which five Transact-SQL segments should you use to develop the solution? To answer, move the appropriate
Transact-SQL segments from the list of Transact-SQL segments to the answer area and arrange them in the correct order.

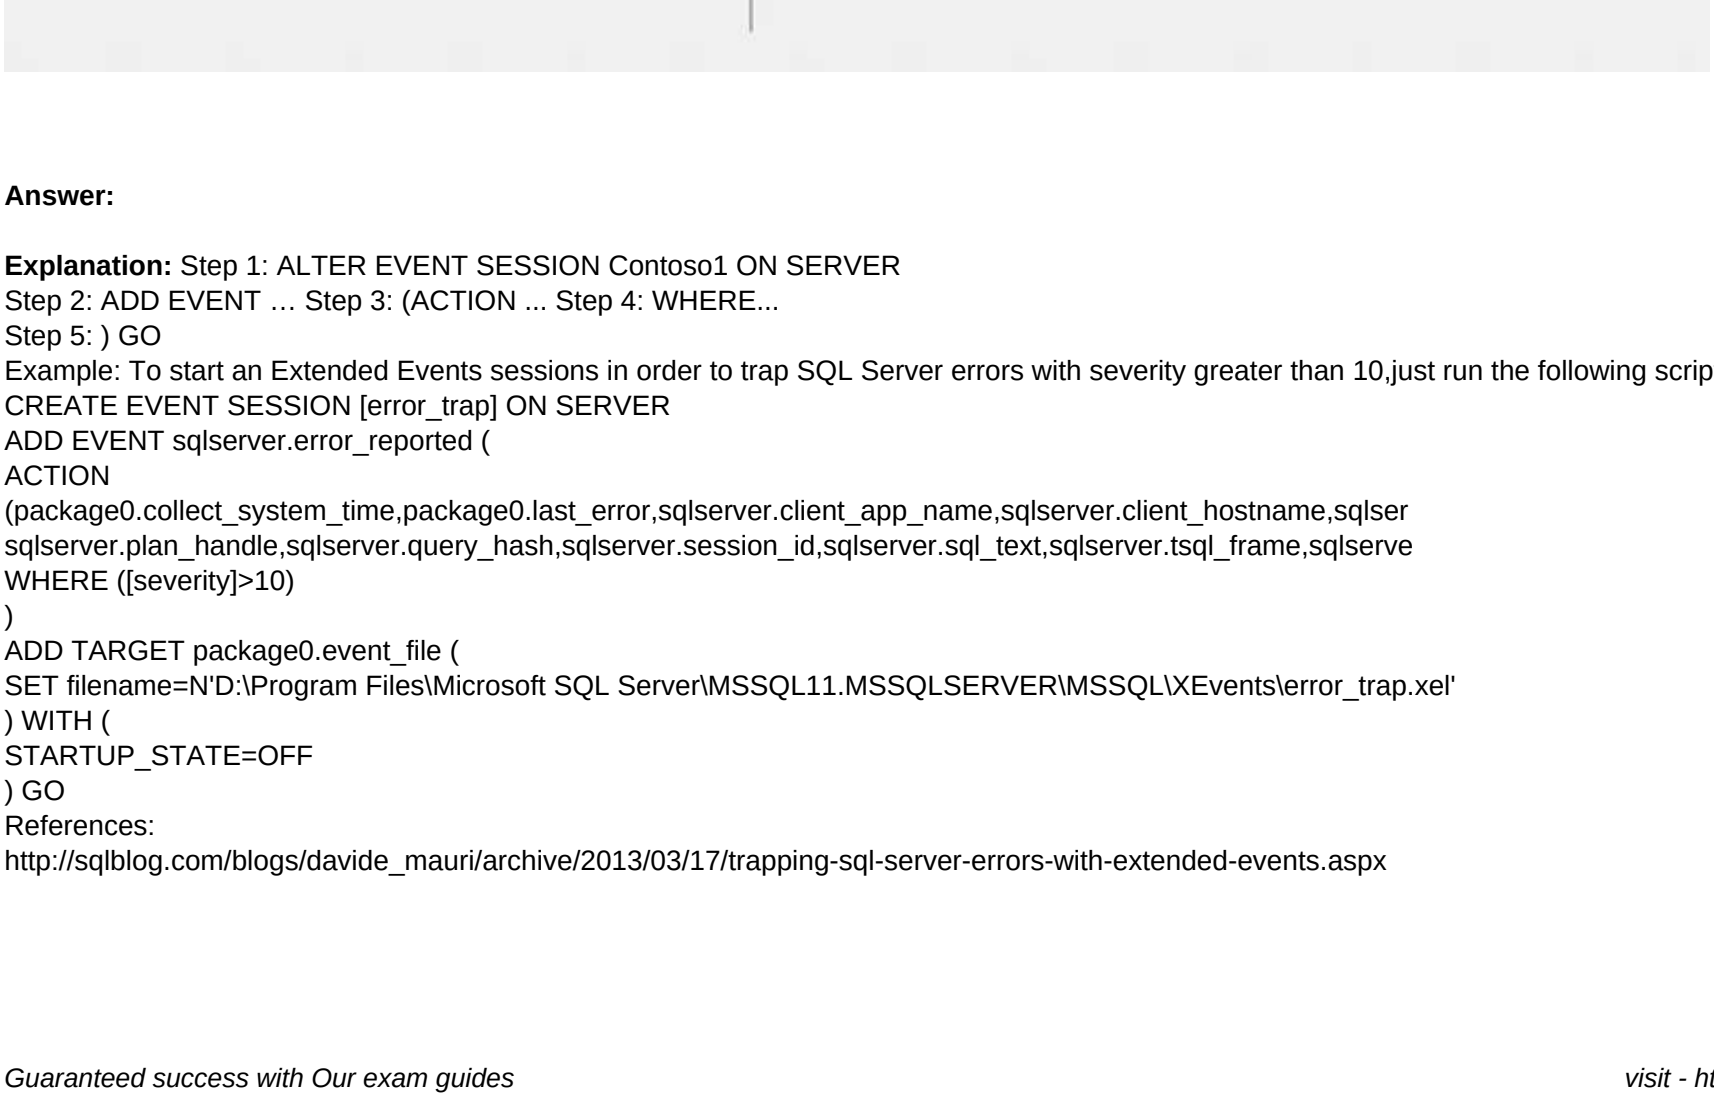

Answer:

Explanation: Step 1: ALTER EVENT SESSION Contoso1 ON SERVER
Step 2: ADD EVENT … Step 3: (ACTION ... Step 4: WHERE...
Step 5:) GO
Example: To start an Extended Events sessions in order to trap SQL Server errors with severity greater than 10,just run the following script:
CREATE EVENT SESSION [error_trap] ON SERVER
ADD EVENT sqlserver.error_reported (
ACTION
(package0.collect_system_time,package0.last_error,sqlserver.client_app_name,sqlserver.client_hostname,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 121
You have a Microsoft SQL Server instance that hosts a database named DB1 that contains 800 gigabyte (GB) of data. The database is used 24 hours each day.
You implement indexes and set the value of the Auto Update Statistics option set to True.
Users report that queries take a long time to complete.
You need to identify statistics that have not been updated for a week for tables where more than 1,000 rows changed.
How should you complete the Transact-SQL statement? To answer, configure the appropriate Transact-SQL segments in the answer area.

Answer Area

```
SELECT OBJECT_NAME(id), name, [ rowcnt | stats_date | rowmodctr | stats_collect ] (id, indid), [ rowcnt | stats_date | rowmodctr | stats_collect ]
FROM sys.sysindexes
WHERE [ rowmodctr | stats_collect | stats_date | rowcnt ] (id, indid) <= DATEADD(DAY, -7, GETDATE())
AND [ stats_collect | rowmodctr | stats_date | rowcnt ] > 1000
AND id IN (SELECT object_id FROM sys.tables)
```

Answer:

Explanation: Box 1: stats_date See example below. Box 2: rowmodctr See examplebelow. Box 3: stats_date
You need to identify statistics that have not been updated for a week. Box 4: rowmodctr
You need to identify that more than 1,000 rows changed.
Rowmodctr counts the total number of inserted, deleted, or updated rows since the last time statistics were updated for the table.
Example: We will query every statistics object which was not updated in the last day and has rows modified since the last update. We will use the rowmodctr field of sys.sysindexes because it shows how many rows were inserted, updated or deleted since the last update occurred. Please note that it is not always 100% accurate in SQL Server 2005 and later, but it can be used to check if any rows were modified.
--Get the list of outdated statistics
SELECTOBJECT_NAME(id),name,STATS_DATE(id, indid),rowmodctr FROM sys.sysindexes
WHERE STATS_DATE (id, indid)<=DATEADD(DAY,-1,GETDATE())
AND rowmodctr>0
AND id IN (SELECT object_id FROM sys.tables) GO
After collecting this information, we can decide which statistics require an update.
References:
https://docs.microsoft.com/en-us/sql/relational-databases/system-compatibility-views/sys-sysindexes-transact-sq
https://www.mssqltips.com/sqlservertip/2628/how-to-find-outdated-statistics-in-sql-server-2008/

NEW QUESTION 123
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 126
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 [          ] ('Table1' [          ] )
```

CHECKDB	ALL_CONSTRAINTS
CHECKCATALOG	ALL_ERRORMSGS
CHECKCONSTRAINTS	NO_INFOMSGS

Answer:

Explanation: Box 1: CHECKCONSTRAINTS
DBCC CHECKCONSTRAINTS checks the integrity of a specified constraint or all constraints on a specified table in the current database.
Box 2: ALL_CONSTRAINTS
ALL_CONSTRAINTS checks all enabled and disabled constraints on the table if the table name is specified or if all tables are checked;otherwise, checks only the enabled constraint.
Note: Syntax: DBCC CHECKCONSTRAINTS [
(
table_name | 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 127
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 130
A Microsoft SQL Server database named DB1 has two filegroups named FG1 and FG2. You implement a backup strategy that creates backups for the filegroups. DB1 experiences a failure. You must restore FG1 and then FG2.
You need to ensure that the database remains in the RECOVERING state until the restoration of FG2 completes. After the restoration of FG2 completes, the database must be online.
What should you specify when you run the recovery command?

A. the WITH NORECOVERY clause for FG1 and the WITH RECOVERY clause for FG2
B. the WITH RECOVERY clause for FG1 and the WITH RECOVERY clause for FG2
C. the WITH RECOVERY clause for both FG1 and FG2
D. the WITH NORECOVERY clause for both FG1 and FG2

Answer: A

NEW QUESTION 131
You are a database administrator for a Microsoft SQL Server 2016 instance.
You need to ensure that data can be replicated from a production server to two reporting servers in real time. You also need to ensure that data on the reporting server is always accessible.
Which solution should you use?

A. Availability Groups
B. Extended Events
C. Snapshot Replication
D. Policy Based Management

Answer: A

NEW QUESTION 136
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.

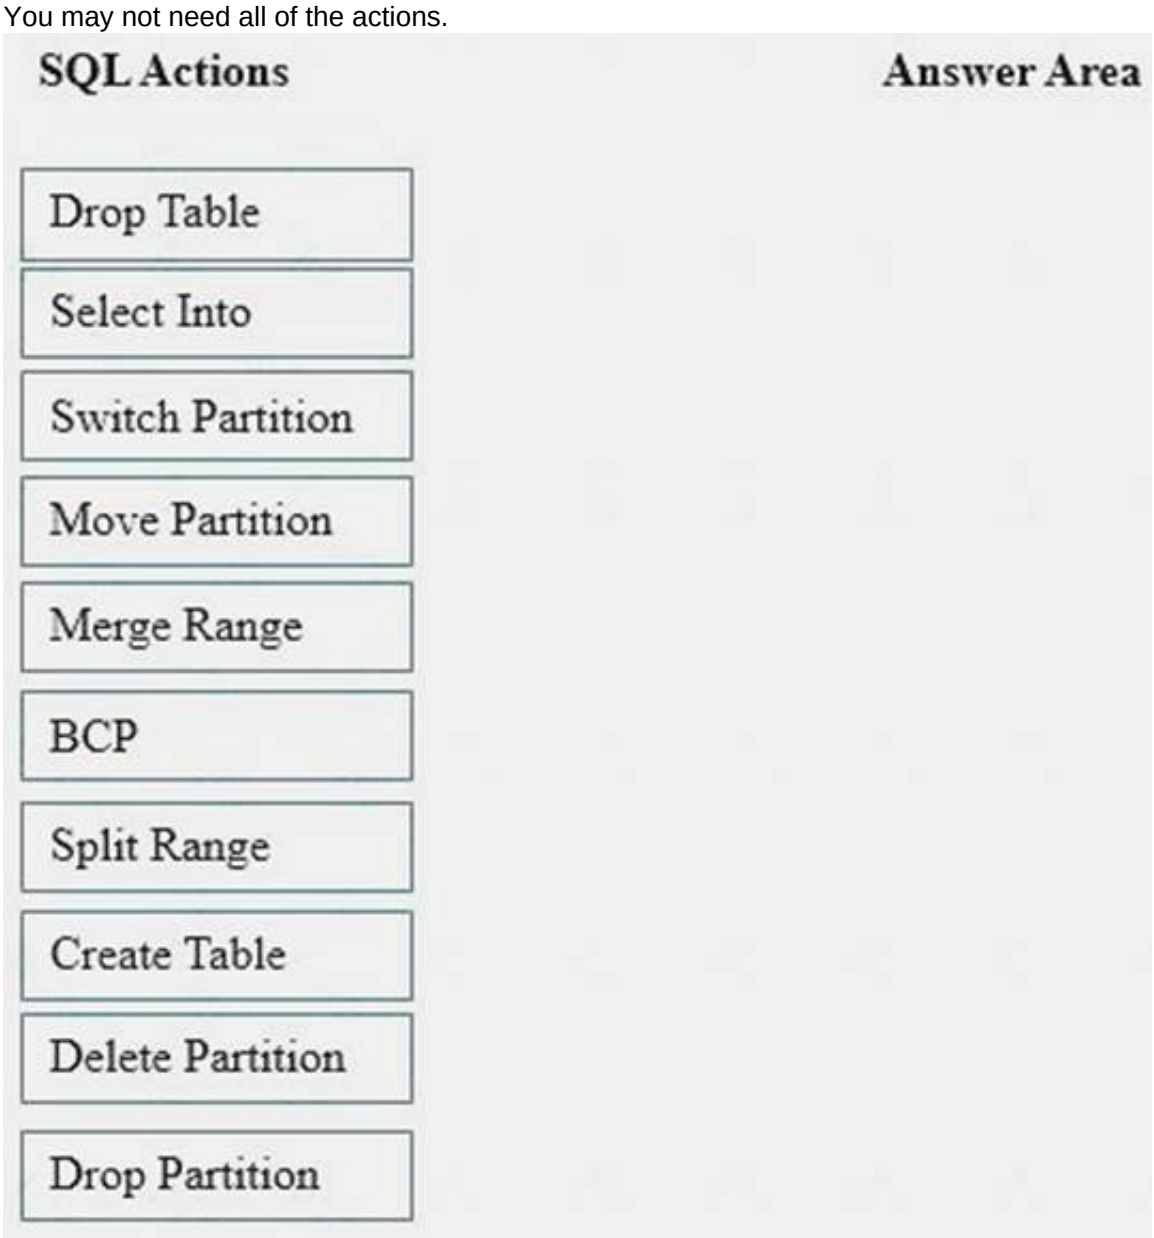

Answer:

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 137
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 141
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 144
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 148
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.

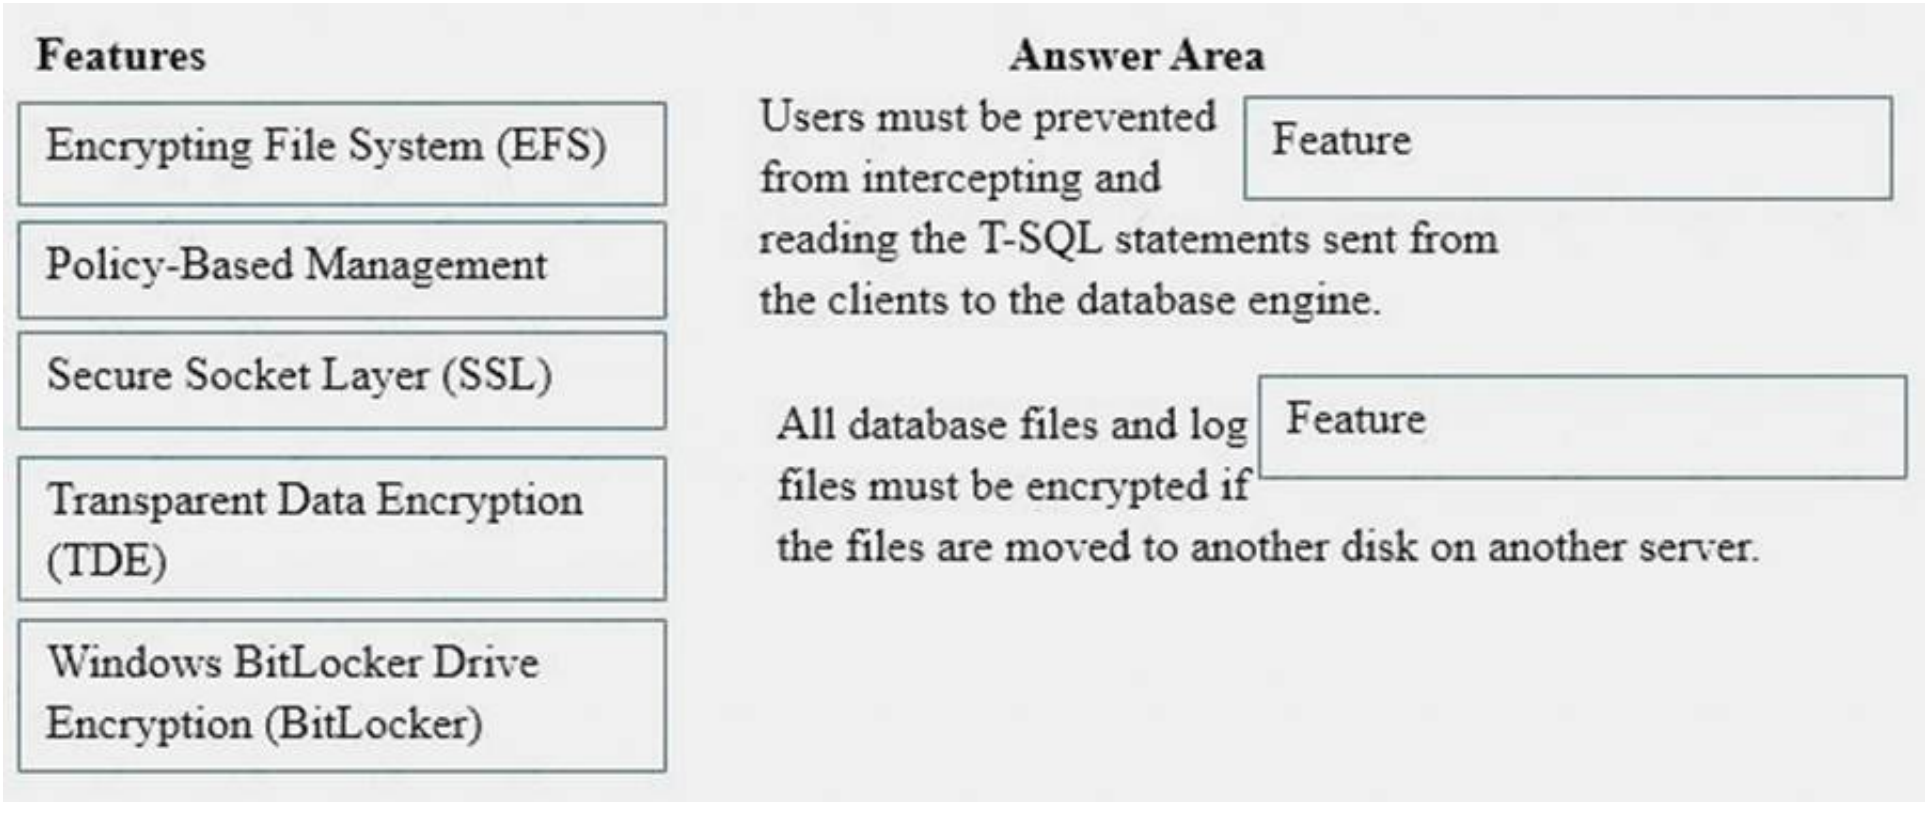

Answer:

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 151
You have a SQL Azure database named Database1.
You need to design the schema for a table named table1. Table1 will have less than one million rows.
Table1 will contain the following information for each row:

Column	Description
ID	An incremental numeric value used to identify the row
Name	A string in English
Code	An alphanumeric code that has five characters
ModifiedDate	The date of the last modification

The solution must minimize the amount of space used to store each row.
Which data types should you recommend for each column? To answer, drag the appropriate data type to the correct column in the answer area.

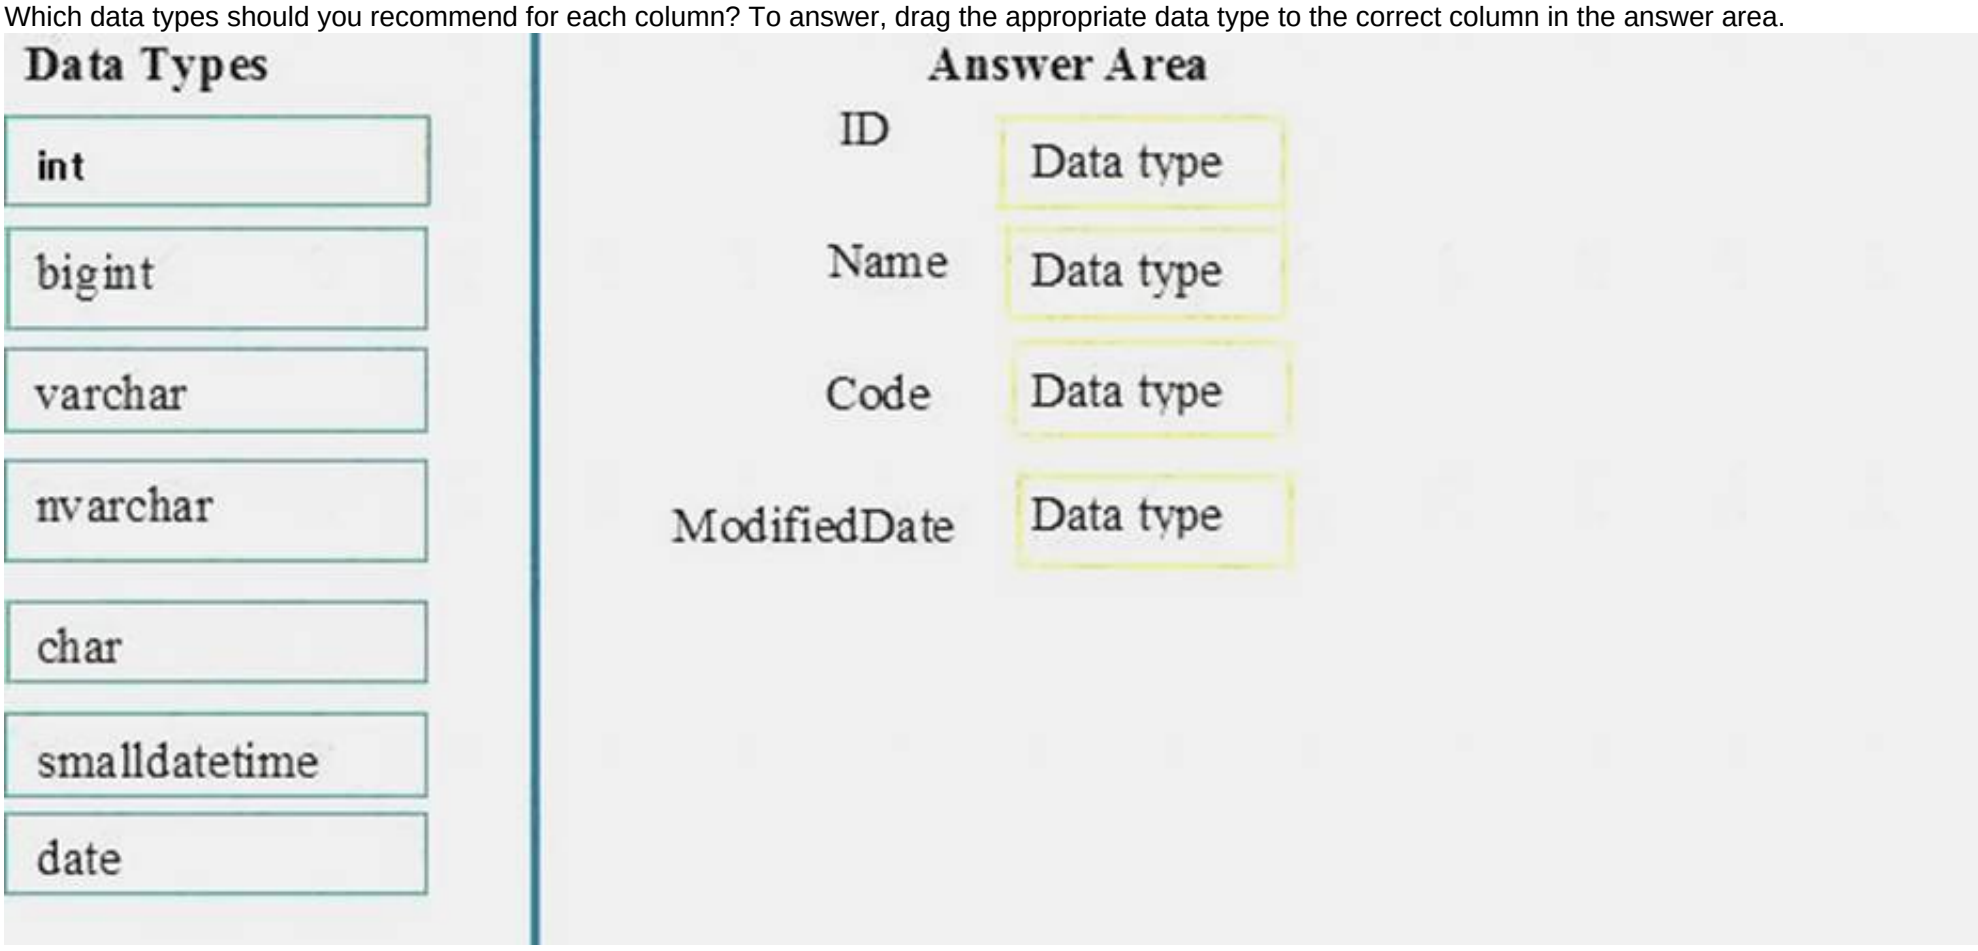

Answer:

Explanation: References:
http://msdn.microsoft.com/en-US/library/ms187752.aspx

NEW QUESTION 156
You deploy a database by using SQL Server 2014. The database contains a table named Table1.
You need to recommend a solution to track all of the deletions executed on Table1. The solution must minimize the amount of custom code required.
What should you recommend?

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

Answer: A

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

NEW QUESTION 159
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 163
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 166
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 database reporting solution that meets the business requirements. What should you include in the recommendation?

A. Data collection
B. Performance Monitor
C. A maintenance plan
D. A dynamic management view

Answer: A

Explanation: 1. Scenario: System administrators must be able to run real-time reports on disk usage.
2. The data collector provides an historical report for each of the System Data collection sets. Each of the following reports use data that is stored in the management data warehouse:
You can use these reports to obtain information for monitoring system capacity and troubleshooting system performance.

NEW QUESTION 169
You are building a stored procedure for a SQL Azure database. The procedure will add multiple rows to a table. You need to design the stored procedure to meet the following requirements:
If any of the new rows violates a table constraint, then no further additions must be attempted and all changes made by the stored procedure must be discarded.
If any errors occur, a row must be added to an audit table, and the original error must be returned to the caller of the stored procedure.
What should you include in the design?

A. An implicit transaction that has XACT_ABORT enabled
B. An explicit transaction that has XACT_ABORT disabled
C. An implicit transaction that has error handling enabled
D. An explicit transaction that has error handling enabled

Answer: D

Explanation: References:
http://technet.microsoft.com/en-us/library/ms175127(v=SQL.105).aspx

NEW QUESTION 174
You have two SQL Server 2012 instances named SQLDev and SQLProd.
You plan to create a new database on SQLProd that will use SQL Server Authentication.
You need to ensure that when the new database is copied from SQLProd to SQLDev, users can connect to the database on SQLDev if they do not have a login on the SQLDev instance.
What should you use? More than one answer choice may achieve the goal. Select the BEST answer.

A. SQL Server Integration Services (SSIS) scripts
B. Extended Events
C. Triggers
D. SQL Server Analysis Services (SSAS) scripts
E. Contained database

Answer: E

Explanation: A fully contained database includes all the settings and metadata required to define the database and has no configuration dependencies on the instance of the SQL Server Database Engine where the database is installed.

NEW QUESTION 176
You need to recommend a backup process for an Online Transaction Processing (OLTP) database. The process must meet the following requirements:
Ensure that if a hardware failure occurs, you can bring the database online with a minimum amount of data loss.
Minimize the amount of administrative effort required to restore any lost data.
What should you include in the recommendation? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Create a database snapshot

Perform a differential backup of the database every night

Ship the logs to a secondary server

Perform a backup of the transaction log every hour

Set the database to the simple recovery model

Set the database to the full recovery model

Perform a weekly full backup of the database

Answer:

Explanation:

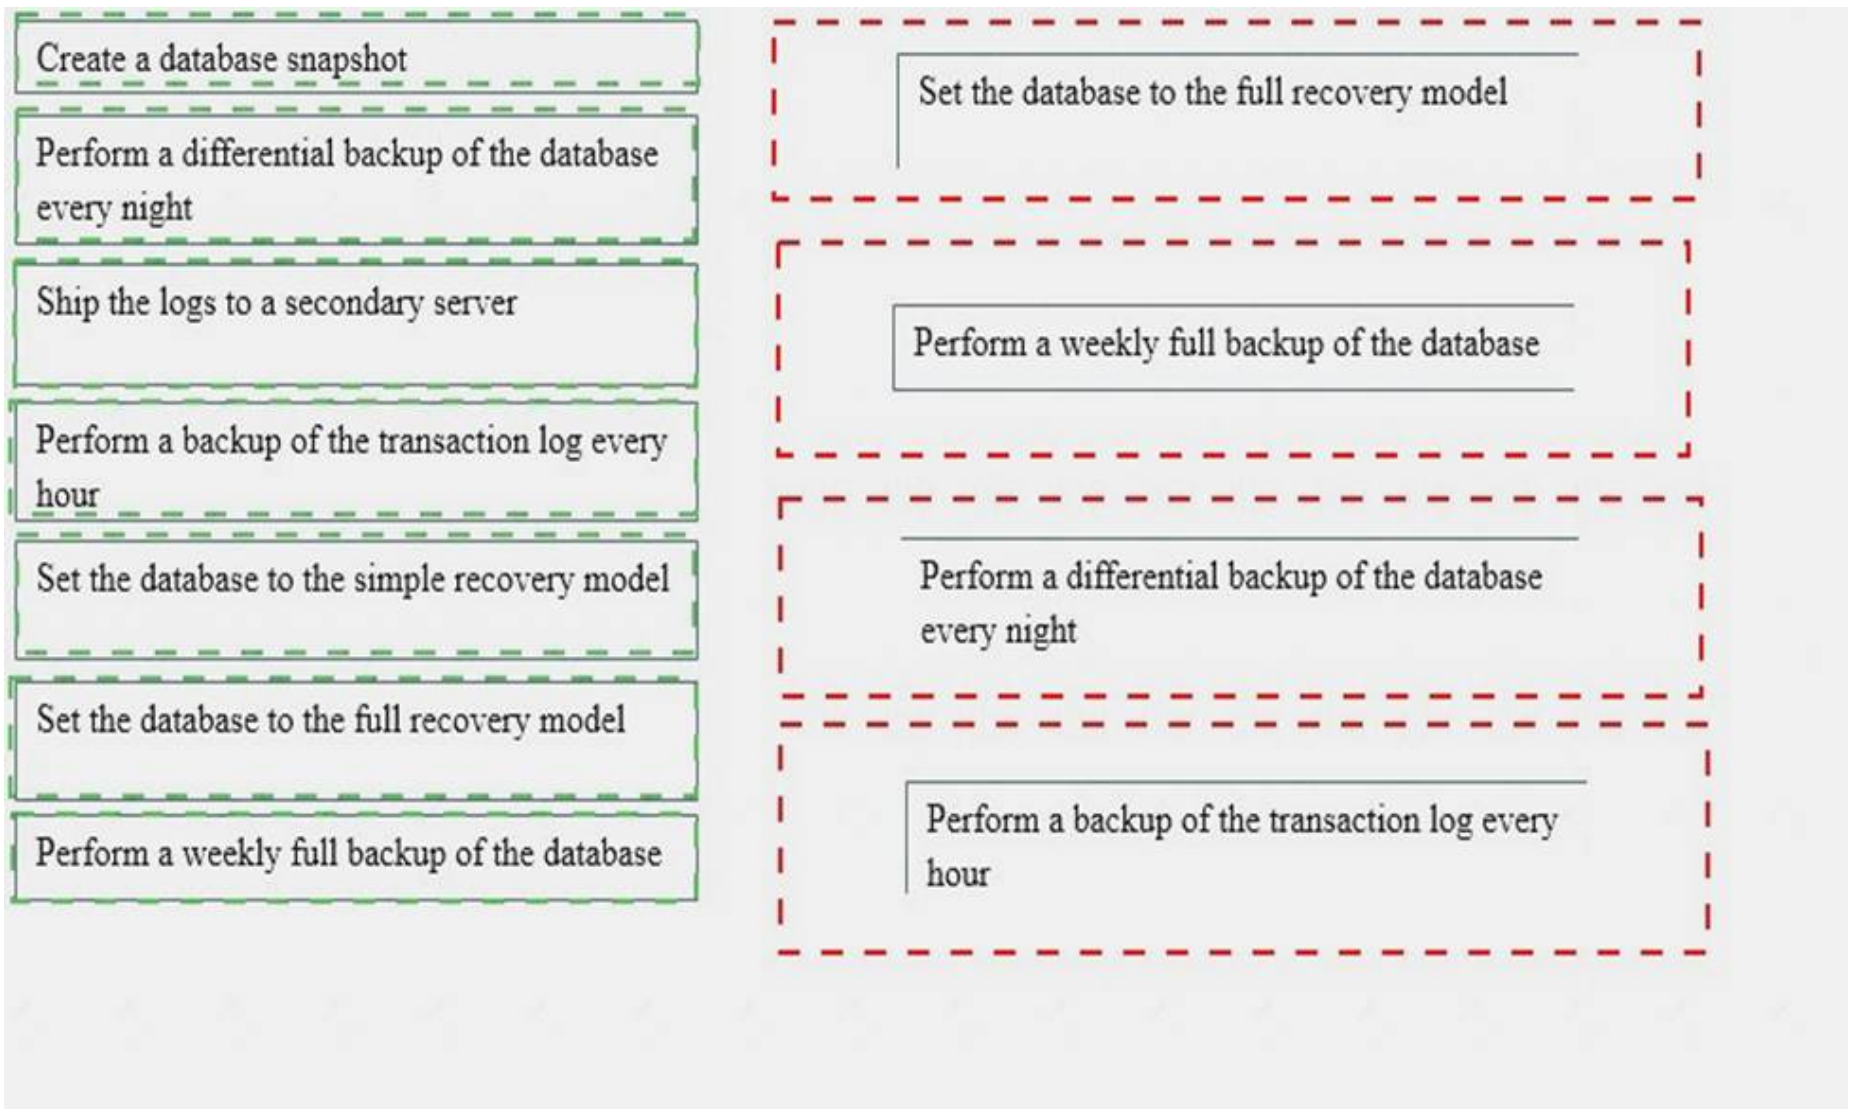

NEW QUESTION 177
Overview
Application Overview
Contoso, Ltd., is the developer of an enterprise resource planning (ERP) application.
Contoso is designing a new version of the ERP application. The previous version of the ERP application used SQL Server 2008 R2.
The new version will use SQL Server 2014.
The ERP application relies on an import process to load supplier data. The import process updates thousands of rows simultaneously, requires exclusive access to the database, and runs daily.

You receive several support calls reporting unexpected behavior in the ERP application. After analyzing the calls, you conclude that users made changes directly to the tables in the database.
Tables
The current database schema contains a table named OrderDetails.
The OrderDetails table contains information about the items sold for each purchase order. OrderDetails stores the product ID, quantities, and discounts applied to each product in a purchase order.
The product price is stored in a table named Products. The Products table was defined by using the SQL_Latin1_General_CP1_CI_AS collation.
A column named ProductName was created by using the varchar data type. The database contains a table named Orders.
Orders contains all of the purchase orders from the last 12 months. Purchase orders that are older than 12 months are stored in a table named OrdersOld.
The previous version of the ERP application relied on table-level security. Stored Procedures
The current version of the database contains stored procedures that change two tables. The following shows the relevant portions of the two stored procedures:

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

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

Customer Problems Installation Issues

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

The current version of the ERP application requires that several SQL Server logins be set up to function correctly. Most customers set up the ERP application in multiple locations and must create logins multiple times.
Index Fragmentation Issues
Customers discover that clustered indexes often are fragmented. To resolve this issue, the customers defragment the indexes more frequently. All of the tables affected by fragmentation have the following columns that are used as the clustered index key:
Backup Issues
Customers who have large amounts of historical purchase order data report that backup time is unacceptable. Search Issues
Users report that when they search product names, the search results exclude product names that contain accents, unless the search string includes the accent.
Missing Data Issues
Customers report that when they make a price change in the Products table, they cannot retrieve the price that the item was sold for in previous orders.
Query Performance Issues
Customers report that query performance degrades very quickly. Additionally, the customers report that users cannot run queries when SQL Server runs maintenance tasks. Import Issues During the monthly import process, database administrators receive many supports call from users who report that they cannot access the supplier data. The database administrators want to reduce the amount of time required to import the data.
Design Requirements
File Storage Requirements
The ERP database stores scanned documents that are larger than 2 MB. These files must only be accessed through the ERP application. File access must have the best possible read and write performance.
Data Recovery Requirements
If the import process fails, the database must be returned to its prior state immediately. Security Requirements
You must provide users with the ability to execute functions within the ERP application, without having direct access to the underlying tables.
Concurrency Requirements
You must reduce the likelihood of deadlocks occurring when Sales.Prod and Sales.Proc2 execute.
You need to recommend a solution that addresses the index fragmentation and index width issue. What should you include in the recommendation? (Each correct answer presents part of the solution. Choose all that apply.)

A. Change the data type of the lastModified column to smalldatetime.
B. Remove the lastModified column from the clustered index.
C. Change the data type of the modifiedBy column to tinyint.
D. Change the data type of the id column to bigint.
E. Remove the modifiedBy column from the clustered index.
F. Remove the id column from the clustered index.

Answer: BE

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

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

NEW QUESTION 179
You administer a single server that contains a Microsoft SQL Server 2016 default instance.
You plan to install a new application that requires the deployment of a database on the server. The application login requires sysadmin permissions.
You need to ensure that the application login is unable to access other production databases. What should you do?

A. Use the SQL Server default instance and configure an affinity mask.
B. Install a new named SQL Server instance on the server.
C. Use the SQL Server default instance and enable Contained Databases.
D. Install a new default SQL Server instance on the server.

Answer: B

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 181
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 file storage requirements.
What should you include in the recommendation?

A. FileStream
B. FileTable
C. The varbinary data type
D. The image data type

Answer: B

Explanation: - Scenario: 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.
- FileTables remove a significant barrier to the use of SQL Server for the storage and management of unstructured data that is currently residing as files on file servers.
Enterprises can move this data from file servers into FileTables to take advantage of integrated administration and services provided by SQL Server. At the same time, they can maintain Windows application compatibility for their existing Windows applications that see this data as files in the file system.

NEW QUESTION 186
You administer a Microsoft SQL Server 2016 database.
You provide temporary securityadmin access to User1 to the database server. You need to know if User1 adds logins to securityadmin.
Which server-level audit action group should you use?

A. SERVER_STATE_CHANGE_GROUP
B. SERVER_PRINCIPAL_IMPERSONATION_GROUP
C. SUCCESSFUL_LOGIN_GROUP
D. SERVER_ROLE_MEMBER_CHANGE_GROUP

Answer: D

Explanation: SERVER_ROLE_MEMBER_CHANGE_GROUP
This event is raised whenever a login is added or removed from a fixed server role. This event is raised for the sp_addsrvrolemember and sp_dropsrvrolemember stored procedures. Equivalent to the Audit Add Login to Server Role Event Class.
References:
http://technet.microsoft.com/en-us/library/cc280663.aspx

NEW QUESTION 187
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 security requirement. What should you recommend?

A. Revoke user permissions on the table
B. Create stored procedures that manipulate dat
C. Grant the users the EXECUTE permission on the stored procedures.
D. Grant the users the SELECT permission on the table
E. Create views that retrieve data from the tables.Grant the users the SELECT permission on the views.
F. Deny the users SELECT permission on the table
G. Create views that retrieve data from the table
H. Grant the users the SELECT permission on the views.
I. Deny the users the SELECT permission on the table
J. Create stored procedures that manipulate data.Grant the users the EXECUTE permission on the stored procedures.

Answer: C

Explanation: - Security Requirements
You must provide users with the ability to execute functions within the ERP application, without having direct access to the underlying tables.

NEW QUESTION 188
Which feature should you enable and configure so session requests addressed to a specific instance can be allocated different processor resources based on session request properties?

A. Resource Governor
B. Windows System Resource Manager
C. Processor affinity
D. I/O affinity

Answer: A

Explanation: Resource Governor enables you to allocate session requests to different resources based on the characteristics of the session request properties.
References: https://docs.microsoft.com/en-us/sql/relational-databases/resource-governor/resource-governor

NEW QUESTION 189
You plan to deploy SQL Server 2012. You must create two tables named Table 1 and Table 2 that will have the following specifications:
Table1 will contain a date column named Column1 that will contain a null value approximately 80 percent of the time.
Table2 will contain a column named Column2 that is the product of two other columns in Table2. Both Table1 and Table2 will contain more than 1 million rows.
You need to recommend which options must be defined for the columns. The solution must minimize the storage requirements for the tables.
Which options should you recommend? To answer, drag the appropriate options to the correct column in the answer area.

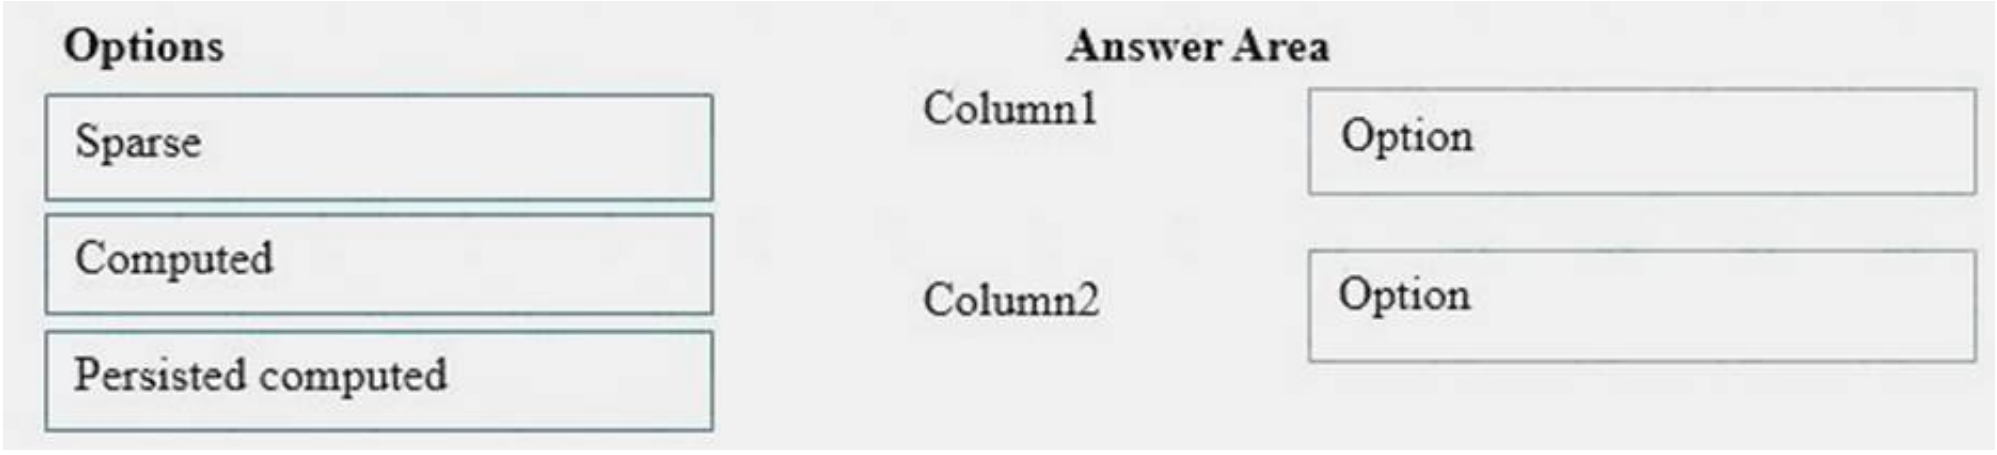

Answer:

Explanation: Column1 – Sparse; Column2 - Computed
- Sparse columns are ordinary columns that have an optimized storage for null values. Sparse columns reduce the space requirements for null values at the cost of more overhead to retrieve nonnull values. Consider using sparse columns when the space saved is at least 20 percent to 40 percent.
- A Persisted column would be faster to retrieve.
- A computed column is computed from an expression that can use other columns in the same table. The expression can be a noncomputed column name, constant, function, and any combination of these connected by one or more operators. Unless otherwise specified, computed columns are virtual columns that are not physically stored in the table. Their values are recalculated every time they are referenced in a query. The Database Engine uses the PERSISTED keyword in the CREATE TABLE and ALTER TABLE statements to physically store computed columns in the table. Their values are updated when any columns that are part of their calculation change.
References:
http://msdn.microsoft.com/en-us/library/cc280604.aspx http://msdn.microsoft.com/en-us/library/ms186241.aspx

NEW QUESTION 190
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 for Application1 that meets the security requirements. What should you include in the recommendation?

A. Encrypted columns
B. Certificate Authentication
C. Secure Socket Layer (SSL)
D. Signed stored procedures

Answer: D

Explanation: - Scenario:
- Data from Database2 will be accessed periodically by an external application named Application1
- Application developers must be denied direct access to the database tables. Applications must be denied direct access to the tables.
Tutorial: Signing Stored Procedures with a Certificate

NEW QUESTION 195
You install a Microsoft SQL Server 2016 instance.
The instance will store data extracted from two databases running on Windows Azure SQL Database. You hire a data steward to perform interactive data cleansing and ad hoc querying and updating of the
database.
You need to ensure that the data steward is given the correct client tools to perform these tasks. Which set of tools should you install?

A. SQL Server Management Studio and Distributed Replay Client
B. Master Data Services and Data Quality Client
C. Data Quality Client and Distributed Replay Client
D. Data Quality Client and SQL Server Management Studio

Answer: B

NEW QUESTION 197
You administer a Microsoft SQL Server 2016 database.
The database contains a Product table created by using the following definition:

```
CREATE TABLE dbo.Product
     (ProductID INT PRIMARY KEY,
     Name VARCHAR(50) NOT NULL,
     Color VARCHAR(15) NOT NULL,
     Size VARCHAR(5) NOT NULL,
     Style CHAR(2) NULL,
     Weight DECIMAL(8,2) NULL);
```

You need to ensure that the minimum amount of disk space is used to store the data in the Product table. What should you do?

A. Convert all indexes to Column Store indexes.
B. Implement Unicode Compression.
C. Implement row-level compression.
D. Implement page-level compression.

Answer: D

NEW QUESTION 200
You want to simulate read, write, checkpoint, backup, sort, and read-ahead activities for your organization's SQL Server 2016 deployment.
Which of the following tools would you use to accomplish this goal?

A. SQLIO
B. SQLIOSim
C. SQLIOStress
D. chkdsk

Answer: B

Explanation: The SQLIOSim utility has been upgraded from the SQLIOStress utility. The SQLIOSim utility more accurately simulates the I/O patterns of Microsoft SQL Server.
References:

https://support.microsoft.com/en-us/help/231619/how-to-use-the-sqliosim-utility-to-simulate-sql-server-activity-

NEW QUESTION 201
You are creating a database that will store usernames and passwords for an application. You need to recommend a solution to store the passwords in the database.
What should you recommend? More than one answer choice may achieve the goal. Select the BEST answer.

A. One-way encryption
B. Transparent Data Encryption (TDE)
C. Encrypting File System (EFS)
D. Reversible encryption

Answer: B

Explanation: Transparent Data Encryption (TDE) is a special case of encryption using a symmetric key. TDE encrypts an entire database using that symmetric key called the database encryption key. The database encryption key is protected by other keys or certificates which are protected either by the database master key or by an asymmetric key stored in an EKM module.
SQL Server provides the following mechanisms for encryption:
Transact-SQL functions
Asymmetric keys
Symmetric keys
Certificates
Transparent Data Encryption

NEW QUESTION 205
You have a server named SQL1 that has SQL Server 2012 installed. SQL1 hosts a database named Database1.
Database1 contains a table named Table1. Table1 is partitioned across five filegroups based on the Date field. The schema of Table1 is configured as shown in the following table.

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

Table1 contains the indexes shown in the following table.

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

You need to recommend an index strategy to maximize performance for the queries that consume the indexes available to Table1.
Which type of index storage should you recommend? To answer, drag the appropriate index storage type to the correct index in the answer area.

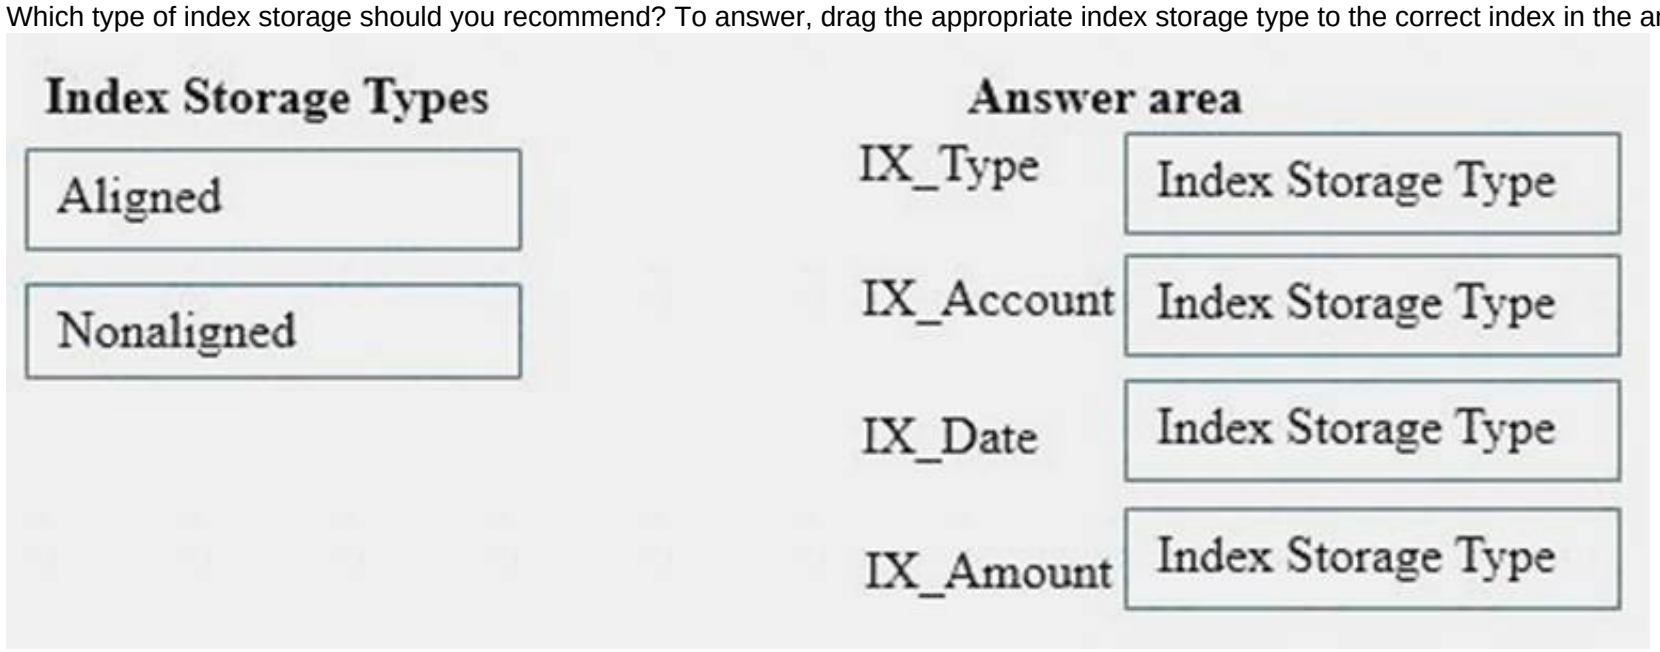

Answer:

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

NEW QUESTION 206
You administer a Microsoft SQL Server 2016 instance. The instance contains a database that supports a retail sales application.
The application generates hundreds of transactions per second and is online 24 hours per day and 7 days per week. You plan to define a backup strategy for the database.
You need to ensure that the following requirements are met:
No more than 5 minutes worth of transactions are lost.
Data can be recovered by using the minimum amount of administrative effort.
What should you do? Choose all that apply.

A. Configure the database to use the SIMPLE recovery model.
B. Create a DIFFERENTIAL database backup every 4 hours.
C. Create a LOG backup every 5 minutes.
D. Configure the database to use the FULL recovery model.
E. Create a FULL database backup every 24 hours.
F. Create a DIFFERENTIAL database backup every 24 hours.

Answer: BCDE
Explanation: If there are only three options, the CDE (exclude differential backup), is the best answer.

NEW QUESTION 207
You administer a Microsoft SQL Server 2016 database named Contoso that contains a single user-defined database role namedBillingUsers.
All objects in Contoso are in the dbo schema.
You need to grant EXECUTE permissions for all stored procedures in Contoso to BillingUsers. Which Transact-SQL statement should you use?

A. EXEC sp_addrolemember'db_procexecutor', 'BillingUsers'
B. CREATE ROLE proc_caller GRANT EXECUTE ON ALL PROCEDURES TO proc_caller ALTER MEMBER BillingUsers ADD TO ROLE proc_caller
C. GRANT EXECUTE ON Schema::dbo TO BillingUsers
D. GRANT EXECUTE ON Contoso::dbo TOBillingUsers

Answer: D

NEW QUESTION 210
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 changes to the ERP application to resolve the search issue.
The solution must minimize the impact on other queries generated from the ERP application. What should you recommend changing?

A. The collation of the Products table
B. The index on the ProductName column
C. The collation of the ProductName column
D. The data type of the ProductName column

Answer: C

Explanation: References:
http://technet.microsoft.com/en-us/library/aa214408(v=sql.80).aspx

NEW QUESTION 213
.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questons and Answers in PDF Format

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