

70-764 Dumps

Administering a SQL Database Infrastructure (beta)

<https://www.certleader.com/70-764-dumps.html>



NEW QUESTION 1

- (Exam Topic 1)

You administer a single server that contains a Microsoft SQL Server 2016 default instance on which several production databases have been deployed.

You plan to install a new ticketing application that requires the deployment of a database on the server. The SQL login for this application requires sysadmin permissions. You need to ensure that the login for the ticketing application cannot access other production databases.

What should you do?

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

Answer: C

Explanation:

SQL Server supports multiple instances of SQL Server on a single server or processor, but only one instance can be the default instance. All others must be named instances. A computer can run multiple instances of SQL Server concurrently, and each instance runs independently of other instances.

References: [https://msdn.microsoft.com/en-us/library/ms143531\(v=SQL.105\).aspx](https://msdn.microsoft.com/en-us/library/ms143531(v=SQL.105).aspx)

NEW QUESTION 2

- (Exam Topic 1)

You administer a Windows Azure SQL Database database named Human_Resources. The database contains 2 tables named Employees and SalaryDetails. You add two Windows groups as logins for the server:

CORP\Employees - All company employees

CORP\HRAdmins - HR administrators only

HR Administrators are also company employees.

You need to grant users access according to the following requirements:

CORP\Employees should have SELECT access to the Employees table.

Only users in CORP\HRAdmins should have SELECT access to the SalaryDetails table.

Logins are based only on Windows security groups.

What should you do?

- A. Create a database role called Employees. Add CORP\Employees to the db_datareader role. Add all company employees except HR administrators to the Employees role. Deny SELECT access to the SalaryDetails table to the Employees role.
- B. Create a database role called HRAdmins. Add all company employees except HR administrators to the db_datareader role. Add all HR administrators to the HRAdmins role. Grant SELECT access to the SalaryDetails table to the HRAdmins role. Deny SELECT access to the SalaryDetails table to the db_datareader role.
- C. Create two database roles: Employees and HRAdmins. Add all company employees to the Employees role. Add HR administrators to the HRAdmins role. Grant SELECT access to all tables except SalaryDetails to the Employees role. Grant SELECT access to the SalaryDetails table to the HRAdmins role. Deny SELECT access to the SalaryDetails table to the Employees role.
- D. Create a database role called Employees. Add all HR administrators to the db_datareader role. Add all company employees to the Employees role. Grant SELECT access to all tables except the SalaryDetails table to the Employees role. Deny SELECT access to the SalaryDetails table to the Employees role.

Answer: D

NEW QUESTION 3

- (Exam Topic 1)

You administer a Microsoft SQL Server 2016 database.

You want to make a full backup of the database to a file on disk. In doing so, you need to output the progress of the backup. Which backup option should you use?

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

Answer: A

NEW QUESTION 4

- (Exam Topic 1)

You administer two Microsoft SQL Server 2016 servers named ProdSrv1 and ProdSrv2. ProdSrv1 is configured as a Distributor.

Both servers are configured to use the Windows NT Service virtual accounts for all SQL Services.

You are configuring snapshot replication from ProdSrv1 to ProdSrv2 by using ProdSrv2 as a pull subscriber.

The distribution agent on ProdSrv2 regularly fails, displaying the following error message: "Cannot access the file. Operating system error code 5 (Access is denied.)."

You need to configure the distribution agent by granting only the minimum required access to all accounts. What should you do?

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

Answer: D

NEW QUESTION 5

- (Exam Topic 1)

You have a database that stores information for a shipping company. You create a table named Customers by running the following Transact-SQL statement. (Line numbers are included for reference only.)

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

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

- A. RETURNS varchar(20)WITH Schemabinding
- B. RETURNS dbo.CustomersORDER BY @salesPerson
- C. RETURNS tableORDER BY @salesPerson
- D. RETURNS tableWITH Schemabinding

Answer: D

Explanation:

The return value can either be a scalar (single) value or a table.

SELECT 1 just selects a 1 for every row, of course. What it's used for in this case is testing whether any rows exist that match the criteria: if a row exists that matches the WHERE clause, then it returns 1, otherwise it returns nothing.

Specify the WITH SCHEMABINDING clause when you are creating the function. This ensures that the objects referenced in the function definition cannot be modified unless the function is also modified.

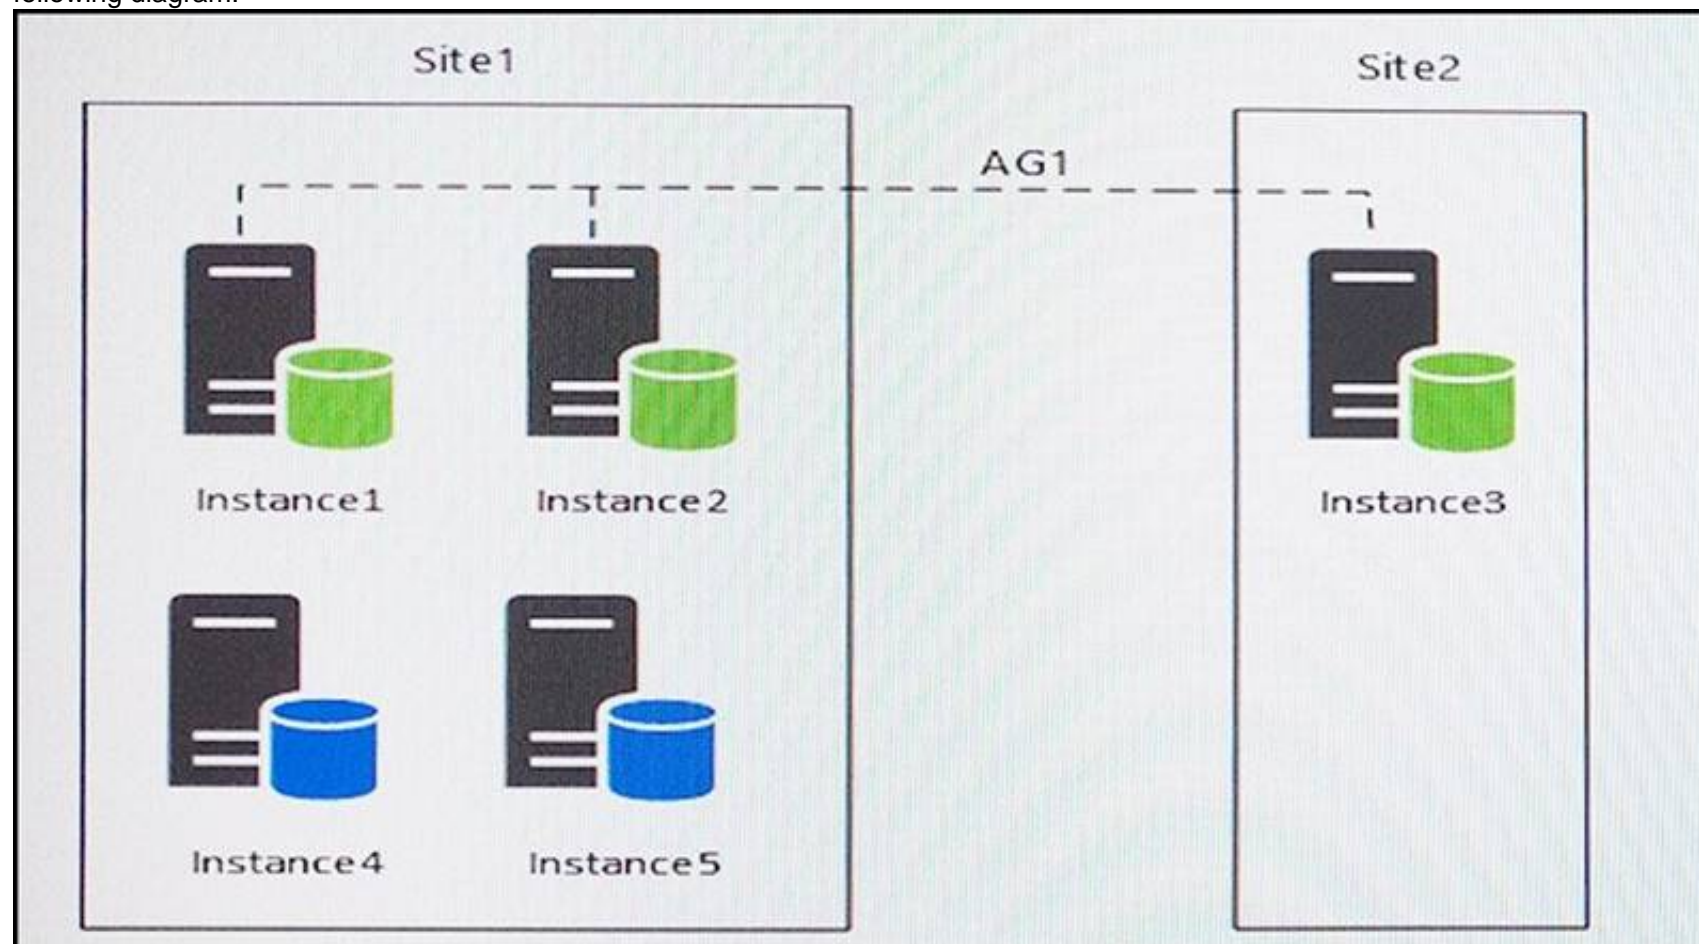
References: <https://docs.microsoft.com/en-us/sql/t-sql/statements/create-function-transact-sql>

NEW QUESTION 6

- (Exam Topic 1)

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

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



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

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

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

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

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

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

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

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

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

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

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

Reporting system: This solution accesses data inDB1with a login that is mapped to a database user that is a member of the db_datareader role. The user has EXECUTE permissions on the database. Queries make no changes to the data. The queries must be load balanced over variable read-only replicas.

Operations system: This solution accesses data inDB1with a login that is mapped to a database user that is a member of the db_datareader and db_datawriter roles. The user has EXECUTE permissions on the database. Queries from the operations system will perform both DDL and DML operations.

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

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

You need to propose a new process for the StagedExternal database.

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

Actions

Drop all nonclustered indexes on the target table.

Create a transaction log backup.
Change the recovery model of **StagedExternal** to **SIMPLE**.

Run the nightly import process.

Change the recovery model of **StagedExternal** to **SIMPLE**.

Change the recovery model of **StagedExternal** to **FULL**. Create a transaction log backup.

Drop all clustered and nonclustered indexes on the target table.

Recreate any dropped indexes on the target table.

Create a transaction log backup.
Change the recovery model of **StagedExternal** to **BULK_LOGGED**.

Answer Area



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

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

All databases use the full recovery model.

References: [https://technet.microsoft.com/en-us/library/ms190421\(v=sql.105\).aspx](https://technet.microsoft.com/en-us/library/ms190421(v=sql.105).aspx)

NEW QUESTION 7

- (Exam Topic 1)

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 8

- (Exam Topic 1)

You are creating an application that will connect to the AgentPortal database by using a SQL login named AgentPortalUser. Stored procedures in the database will use sp_send_dbmail to send email messages.

You create a user account in the msdb database for the AgentPortalUser login.

You use the Database Mail Configuration Wizard to create a Database Mail profile. Security has not been configured for the Database Mail profile.

You need to ensure that AgentPortalUser can send email messages. What should you do?

- A. In the Database Mail Configuration Wizard, configure the Database Mail profile as a private profile for the AgentPortalUser account.
- B. Disable the guest user in the msdb database.
- C. Use the sysmail_help_profileaccount_sp stored procedure to add accounts to the Database Mail profile.

D. In the Database Mail Configuration Wizard, create an email account for each recipient's email address in the Database Mail profile.

Answer: A

Explanation:

You enable and configure Database Mail using the Database Mail Configuration Wizard. Profiles are either public or private. A private profile is accessible only to specific users or roles.

References: <https://docs.microsoft.com/en-us/sql/relational-databases/database-mail/configure-database-mail>

NEW QUESTION 9

- (Exam Topic 1)

You plan to install a Microsoft SQL Server 2016 instance.

The instance will support a database that has the following requirements:

Store Excel workbooks on the file system.

Access the workbooks through Transact-SQL.

Include the workbooks in database backups.

During installation, you need to ensure that the requirements will be met.

Which feature should you use?

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

Answer: B

NEW QUESTION 10

- (Exam Topic 1)

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

You manage a Microsoft SQL Server environment. You implement Transparent Data Encryption (TDE). A user will assist in managing TDE.

You need to ensure that the user can view the TDE metadata while following the principle of 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: G

Explanation:

Viewing the metadata involved with TDE requires the VIEW DEFINITION permission on the certificate. References:

<https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/transparent-data-encryption-tde>

NEW QUESTION 10

- (Exam Topic 1)

You have a database named DB1 that stores more than 700 gigabyte (GB) of data and serves millions of requests per hour.

Queries on DB1 are taking longer than normal to complete. You run the following Transact-SQL statement:

```
SELECT * FROM sys.database_query_store_options
```

You determine that the Query Store is in Read-Only mode.

You need to maximize the time that the Query Store is in Read-Write mode. Which Transact-SQL statement should you run?

- A. ALTER DATABASE DB1 SET QUERY_STORE (QUERY_CAPTURE_MODE = ALL)
- B. ALTER DATABASE DB1 SET QUERY_STORE (MAX_STORAGE_SIZE_MB = 50)
- C. ALTER DATABASE DB1 SET QUERY_STORE (CLEANUP_POLICY = (STALE_QUERY_THRESHOLD_DAYS = 14));
- D. ALTER DATABASE DB1 SET QUERY_STORE (QUERY_CAPTURE_MODE = NONE)

Answer: C

Explanation:

Stale Query Threshold (Days): Time-based cleanup policy that controls the retention period of persisted runtime statistics and inactive queries.

By default, Query Store is configured to keep the data for 30 days which may be unnecessarily long for your scenario.

Avoid keeping historical data that you do not plan to use. This will reduce changes to read-only status. The size of Query Store data as well as the time to detect and mitigate the issue will be more predictable. Use Management Studio or the following script to configure time-based cleanup policy:

```
ALTER DATABASE [QueryStoreDB]
```

```
SET QUERY_STORE (CLEANUP_POLICY = (STALE_QUERY_THRESHOLD_DAYS = 14));
```

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/performance/best-practice-with-the-query-store>

NEW QUESTION 12

- (Exam Topic 1)

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

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

A company has a server that runs Microsoft SQL Server 2016 Web edition. The server has a default instance that hosts a database named DB1.

You need to ensure that you can perform auditing at the database level for DB1.

Solution: You migrate DB1 to a named instance on a server that runs Microsoft SQL Server 2016 Enterprise edition.

Does the solution meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

All editions of SQL Server support server level audits. All editions support database level audits beginning with SQL Server 2016 SP1. Prior to that, database level auditing was limited to Enterprise, Developer, and Evaluation editions.

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/security/auditing/sql-server-audit-database-engine>

NEW QUESTION 13

- (Exam Topic 1)

You install Microsoft SQL Server 2016 on a new server.

After setup is complete, you attempt to start the SQL Server service.

After being in a starting state for a few moments, the service goes back to a stopped state. You need to determine the cause of the failure. Which file should you use?

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

Answer: A

NEW QUESTION 15

- (Exam Topic 1)

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>

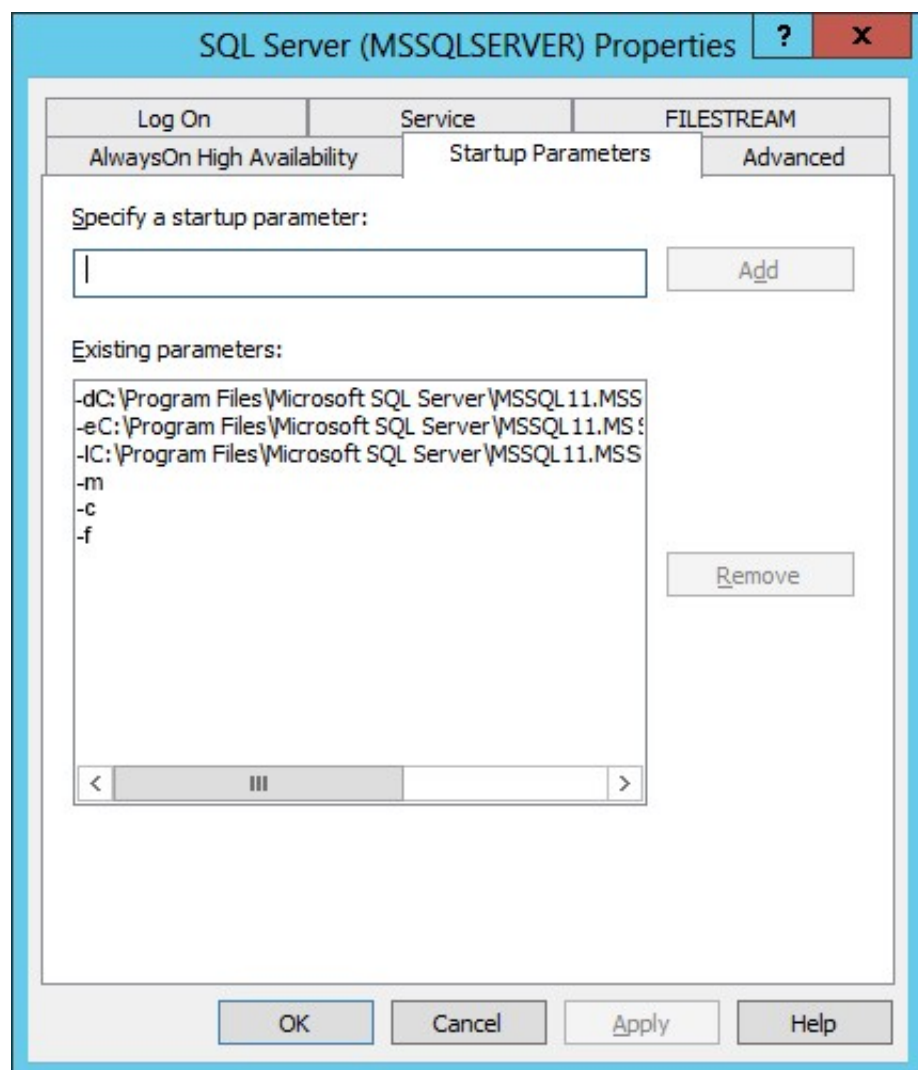
NEW QUESTION 20

- (Exam Topic 1)

You manage a Microsoft SQL Server environment. A server fails and writes the following event to the application event log:

MSG_AUDIT_FORCED_SHUTDOWN

You configure the SQL Server startup parameters as shown in the following graphic:



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

Answer Area

In which user mode will the SQL Server instance start?

	▼
single-user	
multi-user	
restricted-user	

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

	▼
public	
serveradmin	
sysadmin	
setupadmin	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: single-user

The startup option -m starts an instance of SQL Server in single-user mode. Box 2: sysadmin

Starting SQL Server in single-user mode enables any member of the computer's local Administrators group to connect to the instance of SQL Server as a member of the sysadmin fixed server role.

References:

<https://docs.microsoft.com/en-us/sql/database-engine/configure-windows/database-engine-service-startup-option>

NEW QUESTION 24

- (Exam Topic 1)

You administer a Windows Azure SQL Database database named Inventory that contains a stored procedure named p_AddInventory.

Users need to be able to SELECT from all tables in the database and execute the stored procedure. You need to grant only the necessary permissions.

What should you do?

- A. Grant EXECUTE permission on p_AddInventory to all user
- B. Grant VIEW DEFINITION to all users.
- C. Grant EXECUTE permission on p_AddInventory to all user
- D. Add all users to the db_datawriter role.
- E. Add all users to the db_owner role.

F. Grant EXECUTE permission on p_AddInventory to all user
G. Add all users to the db_datareader role.

Answer: D

NEW QUESTION 28

- (Exam Topic 1)

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

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

You need to configure a Microsoft SQL Server instance to ensure that a user named Mail1 can send mail by using Database Mail.

Solution: You add the DatabaseMailUserRole to Mail1 in the master database. Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

Database Mail is guarded by the database role DatabaseMailUserRole in the msdb database, not the master database, in order to prevent anyone from sending arbitrary emails. Database users or roles must be created in the msdb database and must also be a member of DatabaseMailUserRole in order to send emails with the exception of sysadmin who has all privileges.

Note: Database Mail was first introduced as a new feature in SQL Server 2005 and replaces the SQL Mail feature found in previous versions.

References:

http://www.iddevelopment.info/data/SQLServer/DBA_tips/Database_Administration/DBA_20.shtml

NEW QUESTION 30

- (Exam Topic 1)

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

You are the database administrator for a company that hosts Microsoft SQL Server. You manage both on-premises and Microsoft Azure SQL Database environments.

You have a user database named HRDB that contains sensitive human resources data. The HRDB backup files must be encrypted.

You need to grant the correct permission to the service account that backs up the HRDB database. Which permission should you grant?

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

Answer: G

Explanation:

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

References: <https://docs.microsoft.com/en-us/sql/relational-databases/backup-restore/backup-encryption>

NEW QUESTION 31

- (Exam Topic 1)

You administer a Microsoft SQL Server 2016 database.

Users report that a billing application becomes unresponsive during busy times of the day. While investigating, you notice large number of processes taking or waiting for table locks. You suspect that SQL Server is assigning stronger locks to queries.

You start a SQL Profiler trace. Which event should you select?

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

Answer: B

NEW QUESTION 33

- (Exam Topic 1)

You administer a Microsoft SQL Server 2016 instance that has multiple databases. You have a two-node SQL Server failover cluster.

The cluster uses a storage area network (SAN). You discover I/O issues. The SAN is at capacity and additional disks cannot be added.

You need to reduce the I/O workload on the SAN at a minimal cost. What should you do?

- A. Move user databases to a local disk.
- B. Expand the tempdb data and log files.
- C. Modify application code to use table variables.
- D. Move the tempdb files to a local disk.

Answer: D

Explanation:

You can configure TempDB on a local disk when you, for example, installing your SQL Server cluster. References:
<https://www.mssqltips.com/sqlservertip/2817/sql-server-2012-cluster-with-tempdb-on-local-disk/>

NEW QUESTION 38

- (Exam Topic 1)

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

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

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 39

- (Exam Topic 1)

You are planning the deployment of two new Always On Failover Cluster Instances (FCIs) of Microsoft SQL Server to a single Windows Server Cluster with three nodes. The planned configuration for the cluster is shown in the Server Layout exhibit. (Click the Exhibit button.)



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

Conversation

Subject

DBA Team,

SAN Storage for new SQL Cluster

SAN Storage for new SQL Cluster

The following LUNs have been presented to the three servers of the new SQL Cluster:

Two 1 TB drives

Two 500 GB drives

Two 75 GB drives

...

Contoso SAN Admins

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



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

Answer Area

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

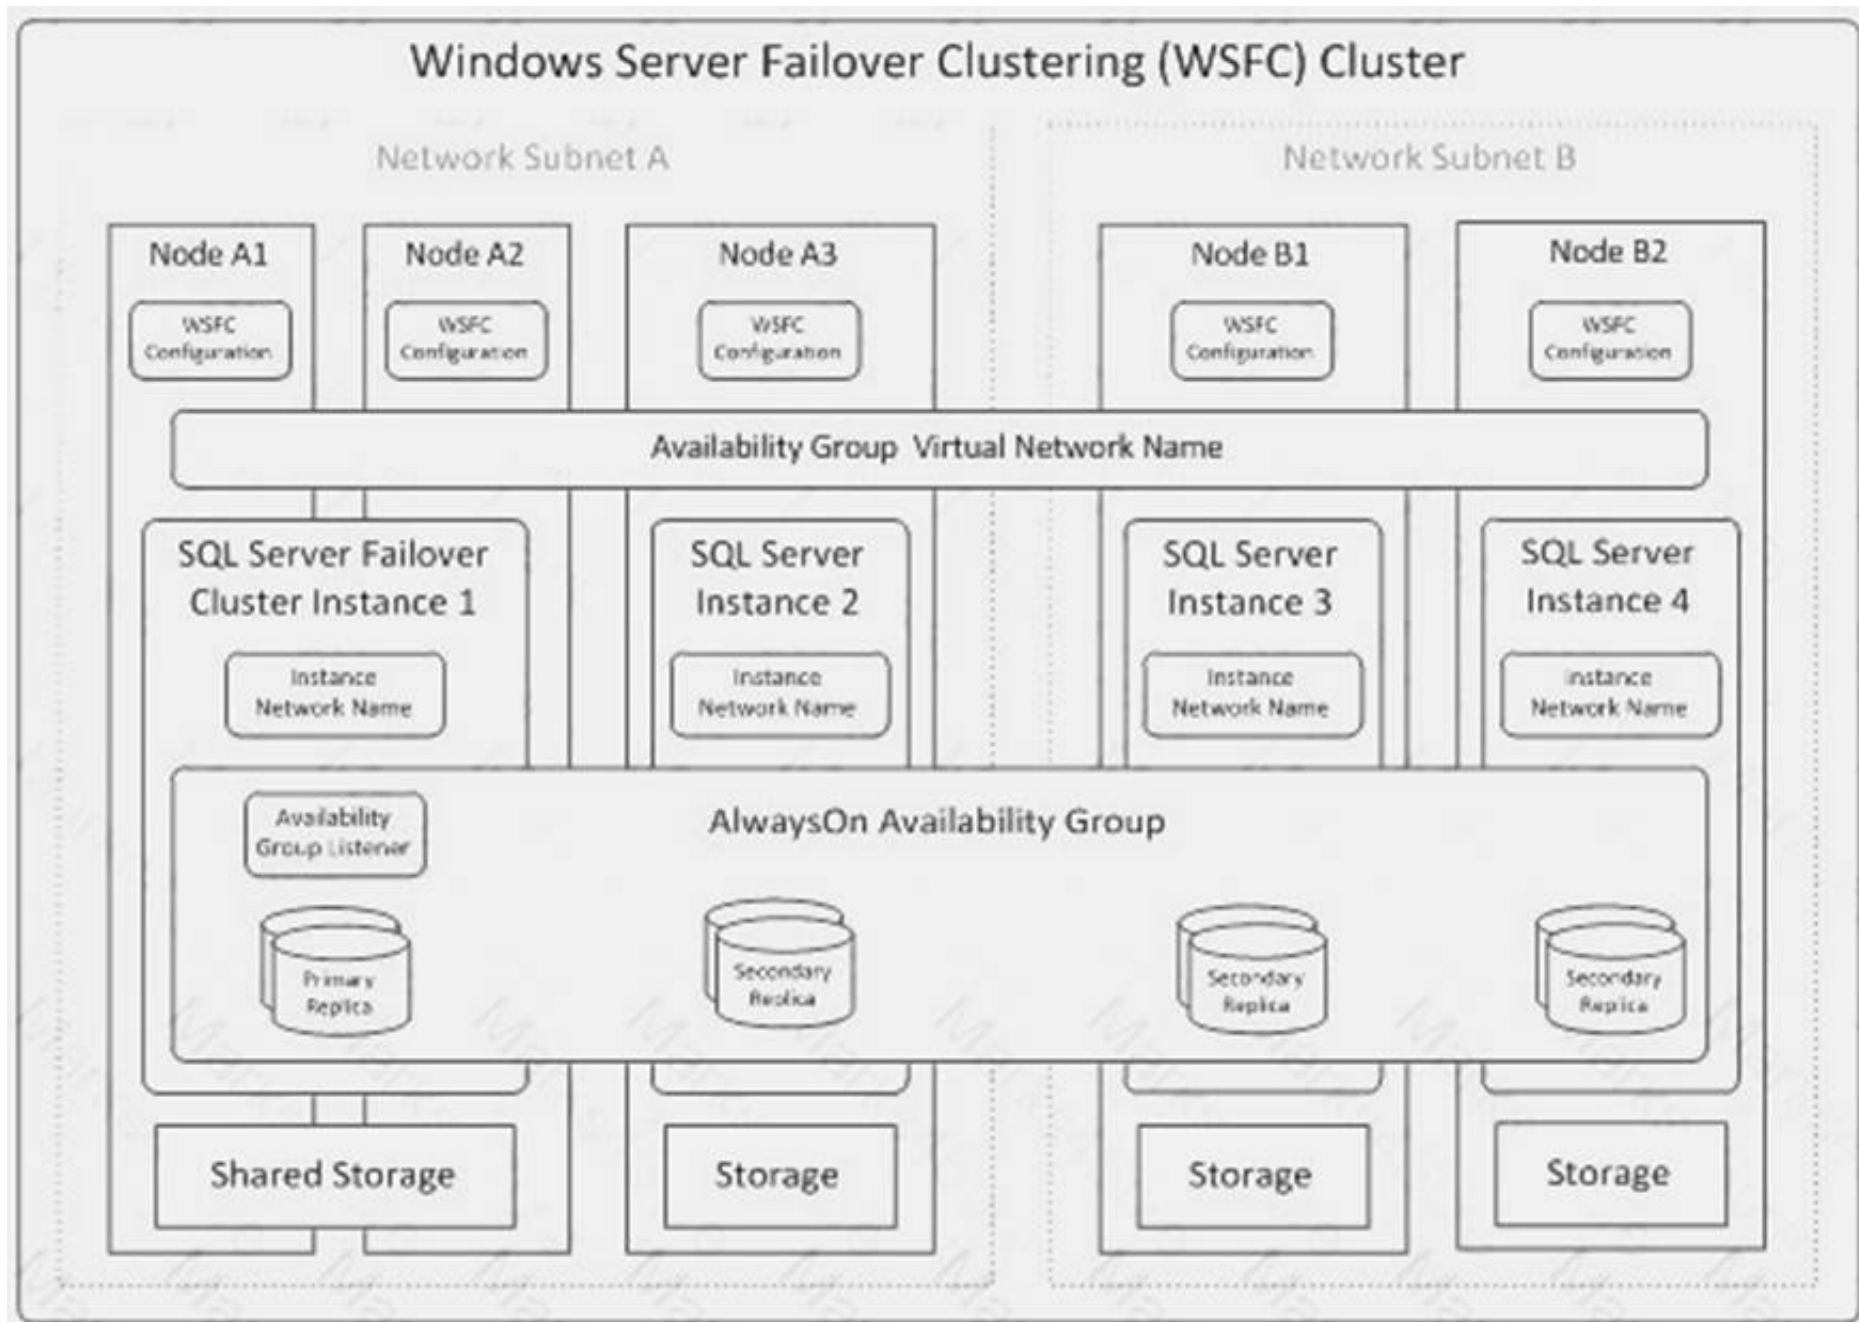
- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Yes
tempdb on local storage. FCIs now support placement of tempdb on local non-shared storage, such as a local solid-state-drive, potentially offloading a significant amount of I/O from a shared SAN.
Prior to SQL Server 2012, FCIs required tempdb to be located on a symmetrical shared storage volume that failed over with other system databases.

Box 2: No
The VNN is set on the group level, not on the instance level.
Database client applications can connect directly to a SQL Server instance network name, or they may connect to a virtual network name (VNN) that is bound to an availability group listener. The VNN abstracts the WSFC cluster and availability group topology, logically redirecting connection requests to the appropriate SQL Server instance and database replica.
The logical topology of a representative AlwaysOn solution is illustrated in this diagram:



Box 3: No

You don't configure the SAN from a SQL Server, instead you can use a Microsoft Server server.

References:

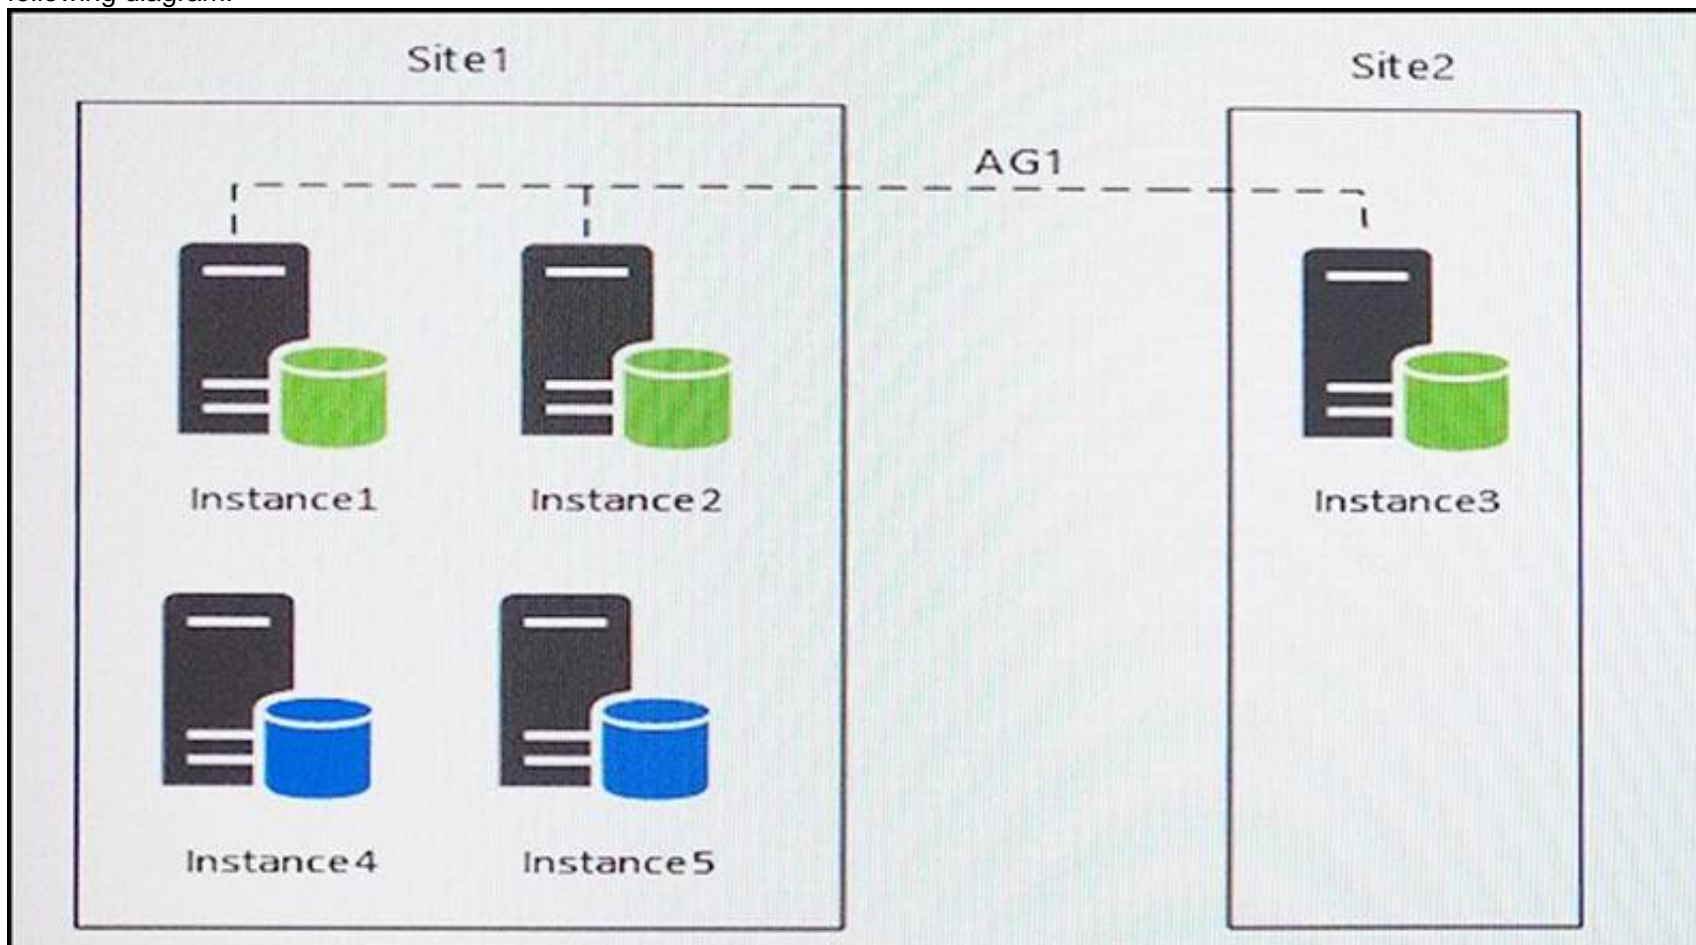
<http://download.microsoft.com/download/d/2/0/d20e1c5f-72ea-4505-9f26-fef9550efd44/microsoft%20sql%20se>

NEW QUESTION 44

- (Exam Topic 1)

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

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



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

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

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

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

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

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

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

All databases use the full recovery model. All backups are written to the network location \\SQLBackup\\. A separate process copies backups to an offsite location.

You should minimize both the time required to restore

the databases and the space required to store backups. The recovery point objective (RPO) for each instance is shown in the following table.

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

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

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

Reporting system: This solution accesses data inDB1with a login that is mapped to a database user that is a member of the db_datareader role. The user has EXECUTE permissions on the database. Queries make no changes to the data. The queries must be load balanced over variable read-only replicas.

Operations system: This solution accesses data inDB1with a login that is mapped to a database user that is a member of the db_datareader and db_datawriter roles. The user has EXECUTE permissions on the database. Queries from the operations system will perform both DDL and DML operations.

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

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

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

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

Answer: D

Explanation:

One of the filegroup is read_only should be as it only need to be backup up once. Partial backups are useful whenever you want to exclude read-only filegroups. A partial backup resembles a full database backup, but a partial backup does not contain all the filegroups. Instead, for a read-write database, a partial backup contains the data in the primary filegroup, every read-write filegroup, and, optionally, one or more read-only files. A partial backup of a read-only database contains only the primary filegroup.

From scenario: Instance1 experiences heavy read-write traffic. The instance hosts a database named OperationsMainthat is four terabytes (TB) in size. The database has multiple data files and filegroups. One of the filegroups is read_only and is half of the total database size.

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/backup-restore/partial-backups-sql-server>

NEW QUESTION 49

- (Exam Topic 1)

You have a database named DB1 that is configured to use the full recovery model. You have a full daily backup job that runs at 02:00. The job backs up data from DB1 to the file B:\DB1.bak.

You need to restore the DB1 database to the point in time of May 25, 2016 at 02:23 and ensure that the database is functional and starts to accept connections.

Which Transact-SQL statement should you run?

A.

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

B.

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

C.

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

D.

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

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: B

NEW QUESTION 54

- (Exam Topic 1)

You manage a Microsoft SQL Server instance. You have a user named User1.

You need to grant the minimum permissions necessary to allow User1 to review audit logs.

For each action, which option should you use? To answer, select the appropriate options in the answer area.

Answer Area

Actions	Options
User1 server role assignment	<div>▼</div> <div> diskadmin serveradmin securityadmin setupadmin </div>
Transact-SQL syntax	<div>▼</div> <div> sys.server_file_audits sys.server_audit_specifications sys.server_file_permissions sys.server_principals </div>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: securityadmin

To access log files for instances of SQL Server that are online, this requires membership in the securityadmin fixed server role.

Box 2: sys.server_audit_specifications

sys.server_audit_specifications contains information about the server audit specifications in a SQL Server audit on a server instance.

NEW QUESTION 56

- (Exam Topic 1)

You administer a Microsoft SQL Server 2016 database instance.

You plan to migrate the database to Windows Azure SQL Database.

You verify that all objects contained in the database are compatible with Windows Azure SQL Database. You need to ensure that database users and required

server logins are migrated to Windows Azure SQL Database.
What should you do?

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

Answer: D

NEW QUESTION 60

- (Exam Topic 1)

You administer a Microsoft SQL Server 2016 server that hosts a transactional database and a reporting database.

The transactional database is updated through a web application and is operational throughout the day. The reporting database is only updated from the transactional database.

The recovery model and backup schedule are configured as shown in the following table:

Database	Description
Transactional database	<p>Recovery model:</p> <ul style="list-style-type: none"> Full <p>Backup schedule:</p> <ul style="list-style-type: none"> 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	<p>Recovery model:</p> <ul style="list-style-type: none"> Simple <p>Backup schedule:</p> <ul style="list-style-type: none"> Full database backup: 01:00 hours daily Differential database backup: 13:00 hours daily <p>Data updates:</p> <ul style="list-style-type: none"> 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

- (Exam Topic 1)

You administer a Microsoft SQL Server 2016 database that includes a table named dbo.Log. This table contains millions of records about user activity in an application.

Records in dbo.Log that are more than 90 days old are purged nightly. When records are purged, table locks are causing contention within inserts.

You need to be able to modify dbo.Log without requiring any changes to the applications that utilize dbo.Log. Which type of solution should you use?

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

Answer: C

NEW QUESTION 66

- (Exam Topic 1)

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

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

You have a database named DB1 that is 640 GB and is updated frequently.

You enabled log shipping for DB1 and configure backup and restore to occur every 30 minutes. You discover that the disks on the data server are almost full.

You need to reduce the amount of disk space used by the log shipping process.

Solution: You increase the frequency of the transaction log backups to every 10 minutes. Does this meet the goal?

- A. Yes
- B. No

Answer: B

NEW QUESTION 71

- (Exam Topic 1)

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

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

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

Answer Area

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

- A. Mastered
- B. Not Mastered

Answer: A

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 74

- (Exam Topic 1)

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 79

- (Exam Topic 1)

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 @TSQL
```

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.

Transact-SQL segments

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

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

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

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

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

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

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

Answer Area



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Transact-SQL segments	Answer Area
<pre>Use Contoso GO CREATE USER Contoso\SalesGroup FOR LOGIN Contoso\SalesGroup</pre>	<pre>use master CREATE LOGIN Contoso\ContosoSearch WITH PASSWORD=N'Pa\$\$w0rd' GO</pre>
<pre>ALTER ROLE db_ddladmin ADD MEMBER Contoso\SalesGroup GRANT VIEW SEFINITION ON Sales.- SalesInvoice TO Contoso\SalesGroup</pre>	<pre>Use Contoso GO CREATE USER Contoso\SalesGroup FOR LOGIN Contoso\SalesGroup</pre>
<pre>use master CREATE LOGIN Contoso\SalesGroup FROM WINDOWS GO</pre>	<pre>GRANT EXECUTE, SELECT ON Sales.Review- Invoice TO Contoso\SalesGroup</pre>
<pre>GRANT EXECUTE ON Sales.ReviewInvoice TO Contoso\SalesGroup GRANT SELECT ON Sales.SalesInvoice TO Contoso\SalesGroup</pre>	
<pre>use master CREATE LOGIN Contoso\ContosoSearch WITH PASSWORD=N'Pa\$\$w0rd' GO</pre>	
<pre>GRANT EXECUTE ON Sales.ReviewInvoice TO Contoso\SalesGroup GRANT VIEW DEFINITION ON Sales.SalesIn- voice TO Contoso\SalesGroup</pre>	
<pre>GRANT EXECUTE, SELECT ON Sales.Review- Invoice TO Contoso\SalesGroup</pre>	

NEW QUESTION 84

- (Exam Topic 1)

You create an availability group named HaContoso that has replicas named Server01/HA, Server02/HA, and Server03/HA. Currently, Server01/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 89

- (Exam Topic 1)

You are configuring a new Microsoft SQL Server Always On Availability Group. You plan to configure a shared network location at \\DATA-C11\\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 options	Answer Area
Add and configure the replica and create an availability group listener named AGL1 on port 1433.	
Launch the Failover Cluster Manager and configure AO-AG1 and AO-AG2 as servers in the cluster. Name the cluster WINCL1.	
Create the Always On Availability Group and select the user databases for the availability group.	
Enable SQL Server 2016 Always On Availability Group feature.	
Select the Full data synchronization method and specify the network path: \\DATA-C11\\SQL.	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1: Launch the Failover Cluster Manager and..

To support the Always On availability groups feature, ensure that every computer that is to participate in one or more availability groups meets requirements including:

* Ensure that each computer is a node in a WSFC (Windows Server Failover Clustering). Step 2: Add and configure the replica and...

All the server instances that host availability replicas for an availability group must use the same SQL Server collation.

Step 3: Enable the SQL Server 2016 Always On Availability Group feature.

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://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 93

- (Exam Topic 1)

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

You are examining information about users, sessions, and 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](https://msdn.microsoft.com/en-us/library/hh272686(v=vs.103).aspx)

NEW QUESTION 95

- (Exam Topic 1)

You administer a Microsoft SQL Server 2016 server that hosts a transactional database and a reporting database. The transactional database is updated through a web application and is operational throughout the day.

The reporting database is only updated from the transactional database.

The recovery model and backup schedule are configured as shown in the following table:

Database	Description
Transactional database	<p>Recovery model:</p> <ul style="list-style-type: none"> Full <p>Backup schedule:</p> <ul style="list-style-type: none"> 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	<p>Recovery model:</p> <ul style="list-style-type: none"> Simple <p>Backup schedule:</p> <ul style="list-style-type: none"> Full database backup: 01:00 hours daily Differential database backup: 13:00 hours daily <p>Data updates:</p> <ul style="list-style-type: none"> Changes in data are updated from the transactional database to the reporting database at 00:30 hours and at 12:30 hours The update takes 15 minutes

At 14:00 hours, you discover that pages 71, 520, and 713 on one of the database files are corrupted on the reporting database.

You need to ensure that the databases are restored.

You also need to ensure that data loss is minimal. What should you do?

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

Answer: H

Explanation:

At restore time, before you restore a differential backup, you must restore its base. Then, restore only the most recent differential backup to bring the database

forward to the time when that differential backup was created. Typically, you would restore the most recent full backup followed by the most recent differential backup that is based on that full backup.

References: [https://technet.microsoft.com/en-us/library/ms345448\(v=sql.105\).aspx](https://technet.microsoft.com/en-us/library/ms345448(v=sql.105).aspx)

NEW QUESTION 97

- (Exam Topic 1)

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

You observe that several indexes are fragmented. You need to rebuild the indexes.

What should you use?

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

Answer: C

Explanation:

How to: Rebuild an Index (SQL Server Management Studio) To rebuild an index

In Object Explorer, connect to an instance of the SQL Server Database Engine and then expand that instance.

Expand Databases, expand the database that contains the table with the specified index, and then expand Tables.

Expand the table in which the index belongs and then expand Indexes.

Right-click the index to rebuild and then click Rebuild.

To start the rebuild operation, click OK.

References: [https://technet.microsoft.com/en-us/library/ms187874\(v=sql.105\).aspx](https://technet.microsoft.com/en-us/library/ms187874(v=sql.105).aspx)

NEW QUESTION 98

- (Exam Topic 1)

You administer a Microsoft SQL Server 2016 database.

The database is currently configured to log ship to a secondary server.

You are preparing to cut over to the secondary server by stopping log-shipping and bringing the secondary database online.

You want to perform a tail-log backup. You need to leave the primary database in a restoring state. Which option of the BACKUP LOG command should you use?

- A. NO_TRUNCATE
- B. NORECOVERY
- C. STANDBY
- D. FORMAT

Answer: B

NEW QUESTION 99

- (Exam Topic 1)

You are the database administrator for a Microsoft SQL Server instance. You develop an Extended Events package to look for events related to application performance.

You need to change the event session to include SQL Server errors that are greater than error severity 15. Which five Transact-SQL segments should you use to develop the solution? To answer, move the appropriate

Transact-SQL segments from the list of Transact-SQL segments to the answer area and arrange them in the correct order.

Transact-SQL segments	Answer Area
<pre>WHERE ((sqlserver.data- base_id>(4)) AND (severity> (15)))</pre>	
<pre>(ACTION(sqlserver.client_ap- p_name, sqlserver.data- base_id,sqlserver.session_id)</pre>	
<pre>ALTER EVENT SESSION Contoso1 ON SERVER</pre>	<div>⬅️ ⬆️</div> <div>➡️ ⬆️</div>
<pre>) GO</pre>	
<pre>ADD EVENT sqlserver.error_re- ported</pre>	
<pre>ADD TARGET sqlserver.er- ror_reported</pre>	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1: ALTER EVENT SESSION Contoso1 ON SERVER

Step 2: ADD EVENT ... Step 3: (ACTION ... Step 4: WHERE...

Step 5:) GO

Example: To start an Extended Events sessions in order to trap SQL Server errors with severity greater than 10, just run the following script:

```
CREATE EVENT SESSION [error_trap] ON SERVER
```

```
ADD EVENT sqlserver.error_reported (
```

```
ACTION
```

```
(package0.collect_system_time,package0.last_error,sqlserver.client_app_name,sqlserver.client_hostname,sqlser
```

```
sqlserver.plan_handle,sqlserver.query_hash,sqlserver.session_id,sqlserver.sql_text,sqlserver.tsql_frame,sqlserve
```

```
WHERE ([severity]>10)
```

```
)
```

```
ADD TARGET package0.event_file (
```

```
SET filename=N'D:\Program Files\Microsoft SQL Server\MSSQL11.MSSQLSERVER\MSSQL\XEvents\error_trap.xel'
```

```
) WITH (
```

```
STARTUP_STATE=OFF
```

```
) GO
```

References:

http://sqlblog.com/blogs/davide_mauri/archive/2013/03/17/trapping-sql-server-errors-with-extended-events.aspx

NEW QUESTION 102

- (Exam Topic 1)

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

You collect performance metrics on multiple Microsoft SQL Server instances and store the data in a single repository.

You need to examine disk usage, query statistics, and server activity without building custom counters.

What should you use?

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

Answer: D

Explanation:

The data collector is a core component of the data collection platform for SQL Server 2017 and the tools that are provided by SQL Server. The data collector provides one central point for data collection across your database servers and applications. This collection point can obtain data from a variety of sources and is not limited to performance data

NEW QUESTION 105

- (Exam Topic 1)

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)

- A. Mastered
- B. Not Mastered

Answer: A

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

```
SELECT OBJECT_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 107

- (Exam Topic 1)

You administer a Microsoft SQL Server 2016 failover cluster that contains two nodes named Node A and Node B.

A single instance of SQL Server is installed on the cluster.

An additional node named Node C has been added to the existing cluster.

You need to ensure that the SQL Server instance can use all nodes of the cluster. What should you do?

- A. Create a ConfigurationFile.ini file from Node B, and then run the AddNode command-line tool on Node A.
- B. Use Node A to install SQL Server on Node C.
- C. Run the Add Node to SQL Server Failover Cluster Wizard on Node C.
- D. Use Cluster Administrator to add a new Resource Group to Node B.

Answer: C

Explanation:

To add a node to an existing SQL Server failover cluster

Insert the SQL Server installation media, and from the root folder, double-click Setup.exe. To install from a network share, navigate to the root folder on the share, and then double-click Setup.exe.

The Installation Wizard will launch the SQL Server Installation Center. To add a node to an existing failover cluster instance, click Installation in the left-hand pane. Then, select Add node to a SQL Server failover cluster.

Etc.

References:

<https://docs.microsoft.com/en-us/sql/sql-server/failover-clusters/install/add-or-remove-nodes-in-a-sql-server-fail>

NEW QUESTION 108

- (Exam Topic 1)

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 109

- (Exam Topic 1)

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet goals.

Some question sets might have more than one correct solution, while others might not have a correct solution.

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

You have a database named DB1 that is 640 GB and is updated frequently.

You enabled log shipping for DB1 and configure backup and restore to occur every 30 minutes. You discover that the disks on the data server are almost full.

You need to reduce the amount of disk space used by the log shipping process. Solution: You enable compression for the transaction log backups:

Does this meet the goal?

- A. Yes
- B. No

Answer: A

NEW QUESTION 111

- (Exam Topic 2)

You are designing a SQL Server database for an order fulfillment system. You create a table named Sales.Orders by using the following script:

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

Each order is tracked by using one of the following statuses:

- Fulfilled
- Shipped
- Ordered
- Received

You need to design the database to ensure that you can retrieve the status of an order on a given date. The solution must ensure that new statuses can be added in the future.

What should you do? More than one answer choice may achieve the goal. Select the BEST answer.

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

Answer: A

NEW QUESTION 114

- (Exam Topic 2)

Overview

You are a database administrator for a company named Litware, Inc.

Litware is a book publishing house. Litware has a main office and a branch office.

You are designing the database infrastructure to support a new web-based application that is being developed. The web application will be accessed at www.litwareinc.com. Both internal employees and external partners will use the application.

You have an existing desktop application that uses a SQL Server 2008 database named App1_DB. App1_DB will remain in production.

Requirements Planned Changes

You plan to deploy a SQL Server 2014 instance that will contain two databases named Database1 and Database2.

All database files will be stored in a highly available SAN. Database1 will contain two tables named Orders and OrderDetails.

Database1 will also contain a stored procedure named usp_UpdateOrderDetails.

The stored procedure is used to update order information. The stored procedure queries the Orders table twice each time the procedure executes.

The rows returned from the first query must be returned on the second query unchanged along with any rows added to the table between the two read operations.

Database1 will contain several queries that access data in the Database2 tables. Database2 will contain a table named Inventory.

Inventory will contain over 100 GB of data.

The Inventory table will have two indexes: a clustered index on the primary key and a nonclustered index. The column that is used as the primary key will use the identity property.

Database2 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. All data in Database2 is recreated each day and 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 119

- (Exam Topic 2)

You plan to create a database.

The database will be used by a Microsoft .NET application for a special event that will last for two days. During the event, data must be highly available. After the event, the database will be deleted. You need to

recommend a solution to implement the database while minimizing costs. The solution must not affect any existing applications.

What should you recommend? More than one answer choice may achieve the goal. Select the BEST answer.

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

Answer: D

NEW QUESTION 123

- (Exam Topic 2)

Overview

You are a database administrator for a company named Litware, Inc.

Litware is a book publishing house. Litware has a main office and a branch office.

You are designing the database infrastructure to support a new web-based application that is being developed. The web application will be accessed at www.litwareinc.com. Both internal employees and external partners will use the application.

You have an existing desktop application that uses a SQL Server 2008 database named App1_DB. App1_DB will remain in production.

Requirements Planned Changes

You plan to deploy a SQL Server 2014 instance that will contain two databases named Database1 and Database2.

All database files will be stored in a highly available SAN. Database1 will contain two tables named Orders and OrderDetails.

Database1 will also contain a stored procedure named usp_UpdateOrderDetails.

The stored procedure is used to update order information. The stored procedure queries the Orders table twice each time the procedure executes.

The rows returned from the first query must be returned on the second query unchanged along with any rows added to the table between the two read operations.

Database1 will contain several queries that access data in the Database2 tables. Database2 will contain a table named Inventory.

Inventory will contain over 100 GB of data.

The Inventory table will have two indexes: a clustered index on the primary key and a nonclustered index. The column that is used as the primary key will use the identity property.

Database2 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. All data in Database2 is recreated each day and 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_Db1 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 the deployment of SQL Server 2014. The solution must meet the business requirements. What should you include in the recommendation?

A. Create a new instance of SQL Server 2014 on the server that hosts the SQL Server 2008 instance.

B. Upgrade the existing SQL Server 2008 instance to SQL Server 2014.

C. Deploy two servers that have SQL Server 2014 installed and implement Failover Clustering.

D. Deploy two servers that have SQL Server 2014 installed and implement database mirroring.

Answer: C

Explanation:

Scenario: The databases must be available if the SQL Server service fails.

NEW QUESTION 126

- (Exam Topic 2)

Overview

Application Overview

Contoso, Ltd., is the developer of an enterprise resource planning (ERP) application.

Contoso is designing a new version of the ERP application. The previous version of the ERP application used SQL Server 2008 R2.

The new version will use SQL Server 2014.

The ERP application relies on an import process to load supplier data. The import process updates thousands of rows simultaneously, requires exclusive access to the database, and runs daily.

You receive several support calls reporting unexpected behavior in the ERP application. After analyzing the calls, you conclude that users made changes directly to the tables in the database.

Tables

The current database schema contains a table named OrderDetails.

The OrderDetails table contains information about the items sold for each purchase order. OrderDetails stores the product ID, quantities, and discounts applied to each product in a purchase order.

The product price is stored in a table named Products. The Products table was defined by using the SQL_Latin1_General_CP1_CI_AS collation.

A column named ProductName was created by using the varchar data type. The database contains a table named Orders.

Orders contains all of the purchase orders from the last 12 months. Purchase orders that are older than 12 months are stored in a table named OrdersOld.

The previous version of the ERP application relied on table-level security. Stored Procedures

The current version of the database contains stored procedures that change two tables. The following shows the relevant portions of the two stored procedures:

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

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

Customer Problems Installation Issues

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

Index Fragmentation Issues

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

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

Backup Issues

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

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

Missing Data Issues

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

Query Performance Issues

Customers report that query performance degrades very quickly. Additionally, the customers report that users cannot run queries when SQL Server runs maintenance tasks. Import Issues During the monthly import process, database administrators receive many supports call from users who report that they cannot access the supplier data. The database administrators want to reduce the amount of time required to import the data.

Design Requirements

File Storage Requirements

The ERP database stores scanned documents that are larger than 2 MB. These files must only be accessed through the ERP application. File access must have the best possible read and write performance.

Data Recovery Requirements

If the import process fails, the database must be returned to its prior state immediately. Security Requirements

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

Concurrency Requirements

You must reduce the likelihood of deadlocks occurring when Sales.Prod and Sales.Proc2 execute. What should you recommend for the updates to Sales.TransactionHistory?

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

Answer: A

NEW QUESTION 131

- (Exam Topic 2)

You are designing a database named DB1.

Changes will be deployed to DB1 every Wednesday night.

You need to recommend a strategy to deploy the changes to DB1. The strategy must meet the following requirements:

The strategy must not disrupt backup operations.

DB1 must be unavailable to users while the changes are deployed.

You must be able to undo quickly the entire operation.

What should you recommend? More than one answer choice may achieve the goal. Select the BEST answer.

- A. Perform a copy-only database backup before the changes are deployed.If the deployment fails, restore the database to another server and recover the original.Objects from the restored database.
- B. Create a database snapshot.If the deployment fails, recover the objects from the database snapshot.
- C. Create a database snapshot.If the deployment fails, revert the database to the database snapshot.
- D. Perform a full database backup before the changes are deployed.If the deployment fails, restore the database to another server and recover the original objects from the restored database.

Answer: C

NEW QUESTION 132

- (Exam Topic 2)

You administer a Microsoft SQL Server 2016 database that has Trustworthy set to On.

You create a stored procedure that returns database-level information from Dynamic Management Views. You grant User1 access to execute the stored procedure.

You need to ensure that the stored procedure returns the required information when User1 executes the stored procedure.

You need to achieve this goal by granting the minimum permissions required.

What should you do? (Each correct answer presents a complete solution. Choose all that apply.)

- A. Create a SQL Server login that has VIEW SERVER STATE permission
- B. Create an application role and a secured password for the role.
- C. Modify the stored procedure to include the EXECUTE AS OWNER statemen
- D. Grant VIEW SERVER STATE permissions to the owner of the stored procedure.
- E. Create a SQL Server login that has VIEW SERVER STATE permission
- F. Modify the stored procedure to include the EXECUTE AS {newlogin} statement.
- G. Grant the db_owner role on the database to User1.
- H. Grant the sysadmin role on the database to User1.

Answer: BC

Explanation:

References:

<http://msdn.microsoft.com/en-us/library/ms187861.aspx> <http://msdn.microsoft.com/en-us/library/ms191291.aspx>

NEW QUESTION 134

- (Exam Topic 2)

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 137

- (Exam Topic 2)

You plan to deploy a database by using SQL Server 2014. Your company identifies the following requirements for the database:

The name of all stored procedures must start with "usp_"always.

All distribution statistics must be updated daily

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

Which features should you identify? To answer, drag the appropriate feature to the correct database requirement in the answer area.

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

- Policy-Based Management Each Stored Procedure that are created and that will be created has to have prefix "USP_".
- Maintenance plans create a workflow of the tasks required to make sure that your database is optimized, regularly backed up, and free of inconsistencies.

NEW QUESTION 138

- (Exam Topic 2)

Your company has offices in Seattle and Montreal.

The network contains two servers named Server1 and Server2 that have SQL Server 2012 installed. The servers are located in separate building within your campus.

The latency of the WAN link between the buildings is less than 10 ms.

You plan to implement an AlwaysOn availability group on both servers. You need to recommend a failover type for the availability group.

What should you recommend?

- A. Asynchronous automatic failover
- B. Synchronous manual failover
- C. Asynchronous manual failover
- D. Synchronous automatic failover

Answer: D

NEW QUESTION 139

- (Exam Topic 2)

You administer a Microsoft SQL Server 2016 database named Contoso on a server named Server01. You need to prevent users from disabling server audits in Server01.

What should you create?

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

Answer: D

NEW QUESTION 140

- (Exam Topic 2)

You are troubleshooting an application that runs a query. The application frequently causes deadlocks. You need to identify which transaction causes the deadlock.

What should you do? More than one answer choice may achieve the goal. Select the BEST answer.

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

Answer: D

Explanation:

Troubleshooting deadlocks

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

Configures deadlock event tracking for the session.

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

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

NEW QUESTION 145

- (Exam Topic 2)

Overview

You are a database administrator for a company named Litware, Inc.

Litware is a book publishing house. Litware has a main office and a branch office.

You are designing the database infrastructure to support a new web-based application that is being developed. The web application will be accessed at www.litwareinc.com. Both internal employees and external partners will use the application.

You have an existing desktop application that uses a SQL Server 2008 database named App1_DB. App1_DB will remain in production.

Requirements Planned Changes

You plan to deploy a SQL Server 2014 instance that will contain two databases named Database1 and Database2.

All database files will be stored in a highly available SAN. Database1 will contain two tables named Orders and OrderDetails.

Database1 will also contain a stored procedure named usp_UpdateOrderDetails.

The stored procedure is used to update order information. The stored procedure queries the Orders table twice each time the procedure executes.

The rows returned from the first query must be returned on the second query unchanged along with any rows added to the table between the two read operations.

Database1 will contain several queries that access data in the Database2 tables. Database2 will contain a table named Inventory.

Inventory will contain over 100 GB of data.

The Inventory table will have two indexes: a clustered index on the primary key and a nonclustered index. The column that is used as the primary key will use the identity property.

Database2 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. All data in Database2 is recreated each day and 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 disk monitoring solution that meets the business requirements. What should you include in the recommendation?

- A. a SQL Server Agent alert
- B. a dynamic management view
- C. a maintenance plan
- D. an audit

Answer: B

Explanation:

Dynamic Management Views and Functions (Transact-SQL)

NEW QUESTION 147

- (Exam Topic 2)

You administer a Microsoft SQL Server 2016 instance.

After a routine shutdown, the drive that contains tempdb fails. You need to be able to start the SQL Server.

What should you do?

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

Answer: B

NEW QUESTION 149

- (Exam Topic 2)

You have two SQL Server instances named SQLDev and SQLProd that have access to various storage media. You plan to synchronize SQLDev and SQLProd.

You need to recommend a solution that meets the following requirements:

The database schemas must be synchronized from SQLDev to SQLProd.

The database on SQLDev must be deployed to SQLProd by using a package.

The package must support being deployed to SQL Azure.

What should you recommend? More than one answer choice may achieve the goal. Select the BEST answer.

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

Answer: B

Explanation:

*SIS supports connections to SQL Database by using the ADO.NET provider. OLEDB is not supported at this time. You can build the SSIS package connecting to SQL Database and create the data flow tasks the same way as you would against a typical on premise SQL Server.

<http://technet.microsoft.com/en-us/library/ee210546.aspx>

NEW QUESTION 153

- (Exam Topic 2)

Overview

You are a database administrator for a company named Litware, Inc.

Litware is a book publishing house. Litware has a main office and a branch office.

You are designing the database infrastructure to support a new web-based application that is being developed. The web application will be accessed at

www.litwareinc.com. Both internal employees and external partners

will use the application.

You have an existing desktop application that uses a SQL Server 2008 database named App1_DB. App1_DB will remain in production.

Requirements Planned Changes

You plan to deploy a SQL Server 2014 instance that will contain two databases named Database1 and Database2.

All database files will be stored in a highly available SAN. Database1 will contain two tables named Orders and OrderDetails.

Database1 will also contain a stored procedure named usp_UpdateOrderDetails.

The stored procedure is used to update order information. The stored procedure queries the Orders table twice each time the procedure executes.

The rows returned from the first query must be returned on the second query unchanged along with any rows added to the table between the two read operations.

Database1 will contain several queries that access data in the Database2 tables. Database2 will contain a table named Inventory. Inventory will contain over 100 GB of data.

The Inventory table will have two indexes: a clustered index on the primary key and a nonclustered index.

The column that is used as the primary key will use the identity property.

Database2 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. All data in Database2 is recreated each day and 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 an isolation level for usp_UpdateOrderDetails. Which isolation level should you recommend?

- A. Read committed
- B. Repeatable read
- C. Read uncommitted
- D. Serializable

Answer: B

Explanation:

- Scenario: 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.

- REPEATABLE READ Specifies that statements cannot read data that has been modified but not yet committed by other transactions and that no other transactions can modify data that has been read by the current transaction until the current transaction completes.

NEW QUESTION 155

- (Exam Topic 2)

You are creating a database that will store usernames and credit card numbers for an application. You need to recommend a solution to store and reuse the credit card numbers in the database.

What should you recommend? More than one answer choice may achieve the goal. Select the BEST answer.

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

Answer: A

Explanation:

If we are going to encrypt credit card number for storage, then we should have Data Encryption Key(DEK) for encrypting the credit card number.

NEW QUESTION 158

- (Exam Topic 2)

Overview

Application Overview

Contoso, Ltd., is the developer of an enterprise resource planning (ERP) application.

Contoso is designing a new version of the ERP application. The previous version of the ERP application used SQL Server 2008 R2.

The new version will use SQL Server 2014.

The ERP application relies on an import process to load supplier data. The import process updates thousands of rows simultaneously, requires exclusive access to the database, and runs daily.

You receive several support calls reporting unexpected behavior in the ERP application. After analyzing the calls, you conclude that users made changes directly to the tables in the database.

Tables

The current database schema contains a table named OrderDetails.

The OrderDetails table contains information about the items sold for each purchase order. OrderDetails stores the product ID, quantities, and discounts applied to each product in a purchase order.

The product price is stored in a table named Products. The Products table was defined by using the SQL_Latin1_General_CP1_CI_AS collation.

A column named ProductName was created by using the varchar data type. The database contains a table named Orders.

Orders contains all of the purchase orders from the last 12 months. Purchase orders that are older than 12 months are stored in a table named OrdersOld.

The previous version of the ERP application relied on table-level security. Stored Procedures

The current version of the database contains stored procedures that change two tables. The following shows the relevant portions of the two stored procedures:


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

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

Customer Problems Installation Issues

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

Index Fragmentation Issues

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

Column	Data type
id	uniqueidentifier
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 installation issues.

What should you include in the recommendation?

- A. Windows logins
- B. Server roles
- C. Contained users
- D. Database roles

Answer: C

Explanation:

- Scenario: 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.
- Creating contained users enables the user to connect directly to the contained database. This is a very significant feature in high availability and disaster recovery scenarios such as in an AlwaysOn solution. If the users are contained users, in case of failover, people would be able to connect to the secondary without creating logins on the instance hosting the secondary. This provides an immediate benefit.

NEW QUESTION 163

- (Exam Topic 2)

Overview

General Overview

ADatum Corporation has offices in Miami and Montreal.

The network contains a single Active Directory forest named adatum.com. The offices connect to each other by using a WAN link that has 5-ms latency. A. Datum standardizes its database platform by using SQL Server 2014 Enterprise edition.

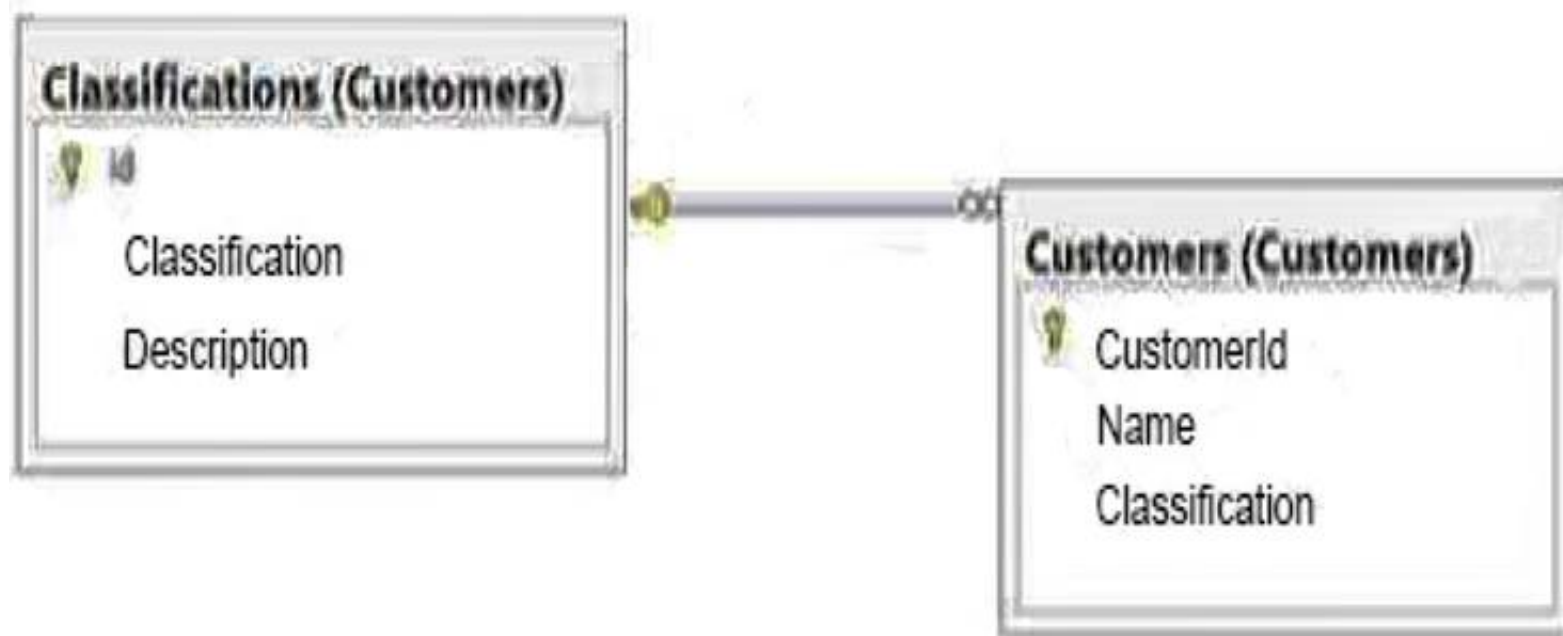
Databases

Each office contains databases named Sales, Inventory, Customers, Products, Personnel, and Dev. Servers and databases are managed by a team of database administrators. Currently, all of the database

administrators have the same level of permissions on all of the servers and all of the databases.

The Customers database contains two tables named Customers and Classifications. The following graphic shows the relevant portions of the tables:

Classifications (Customers)



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

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

The Inventory database is updated frequently. The database is often used for reporting.

A full backup of the database currently takes three hours to complete. Stored Procedures

A stored procedure named USP_1 generates millions of rows of data for multiple reports. USP_1 combines data from five different tables from the Sales and Customers databases in a table named Table1.

After Table1 is created, the reporting process reads data from Table1 sequentially several times. After the process is complete, Table1 is deleted.

A stored procedure named USP_2 is used to generate a product list. The product list contains the names of products grouped by category.

USP_2 takes several minutes to run due to locks on the tables the procedure accesses. The locks are caused by USP_1 and USP_3.

A stored procedure named USP_3 is used to update prices. USP_3 is composed of several UPDATE statements called in sequence from within a transaction.

Currently, if one of the UPDATE statements fails, the stored procedure fails. A stored procedure named USP_4 calls stored procedures in the Sales, Customers, and Inventory databases.

The nested stored procedures read tables from the Sales, Customers, and Inventory databases. USP_4 uses an EXECUTE AS clause.

All nested stored procedures handle errors by using structured exception handling. A stored procedure named USP_5 calls several stored procedures in the same database. Security checks are performed each time USP_5 calls a stored procedure.

You suspect that the security checks are slowing down the performance of USP_5. All stored procedures accessed by user applications call nested stored procedures.

The nested stored procedures are never called directly. Design Requirements

Data Recovery

You must be able to recover data from the Inventory database if a storage failure occurs. You have a Recovery Time Objective (RTO) of 5 minutes.

You must be able to recover data from the Dev database if data is lost accidentally. You have a Recovery Point Objective (RPO) of one day.

Classification Changes

You plan to change the way customers are classified. The new classifications will have four levels based on the number of orders. Classifications may be removed or added in the future. Management requests that historical data be maintained for the previous classifications. Security A group of junior database administrators must be able to manage security for the Sales database. The junior database administrators will not have any other administrative rights. A. Datum wants to track which users run each stored procedure.

Storage

A Datum has limited storage. Whenever possible, all storage space should be minimized for all databases and all backups.

Error Handling

There is currently no error handling code in any stored procedure.

You plan to log errors in called stored procedures and nested stored procedures. Nested stored procedures are never called directly.

You need to recommend a disaster recovery strategy for the Inventory database. What should you include in the recommendation?

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

Answer: A

Explanation:

Scenario:

- You must be able to recover data from the Inventory database if a storage failure occurs. You have a Recovery Point Objective (RPO) of one hour.
- A. Datum Corporation has offices in Miami and Montreal.
- SQL Server Log shipping allows you to automatically send transaction log backups from a primary database on a primary server instance to one or more secondary databases on separate secondary server instances. The transaction log backups are applied to each of the secondary databases individually.

NEW QUESTION 164

- (Exam Topic 2)

You administer a Microsoft SQL Server 2016 instance.

You need to stop a blocking process that has an SPID of 64 without stopping other processes. What should you do?

- A. Execute the following Transact-SQL statement: EXECUTE sp_KillSPID 64
- B. Restart the SQL Server service.
- C. Execute the following Transact-SQL statement: KILL 64
- D. Execute the following Transact-SQL statement: ALTER SESSION KILL '64'

Answer: C

Explanation:

KILL can be used to terminate a normal connection, which internally terminates the transactions that are associated with the specified session ID.

References:

<http://msdn.microsoft.com/en-us/library/ms173730.aspx>

NEW QUESTION 169

- (Exam Topic 2)

You have a SQL Server 2014 environment That contains 20 servers.

The corporate security policy states that all SQL Server 2014 instances must meet specific security standards. You need to recommend a management strategy for the SQL Server 2014 servers.

What should you include in the recommendation? More than one answer choice may achieve the goal. Select the BEST answer.

- A. Multi server jobs
- B. Policy-Based Management
- C. Common criteria compliance
- D. Maintenance plans

Answer: B

Explanation:

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

NEW QUESTION 172

- (Exam Topic 2)

You administer a Microsoft SQL Server 2016 failover cluster.

You need to ensure that a failover occurs when the server diagnostics returns query_processing error. Which server configuration property should you set?

- A. SqlOumperDumpFlags
- B. FailureConditionLevel
- C. HealthCheckTimeout
- D. SqlDumperDumpPath

Answer: B

Explanation:

The SQL Server Database Engine resource DLL determines whether the detected health status is a condition for failure using the FailureConditionLevel property. The FailureConditionLevel property defines which detected health statuses cause restarts or failovers. Multiple levels of options are available, ranging from no automatic restart or failover to all possible failure conditions resulting in an automatic restart or failover.

References:

<https://docs.microsoft.com/en-us/sql/sql-server/failover-clusters/windows/failover-policy-for-failover-cluster-ins>

NEW QUESTION 177

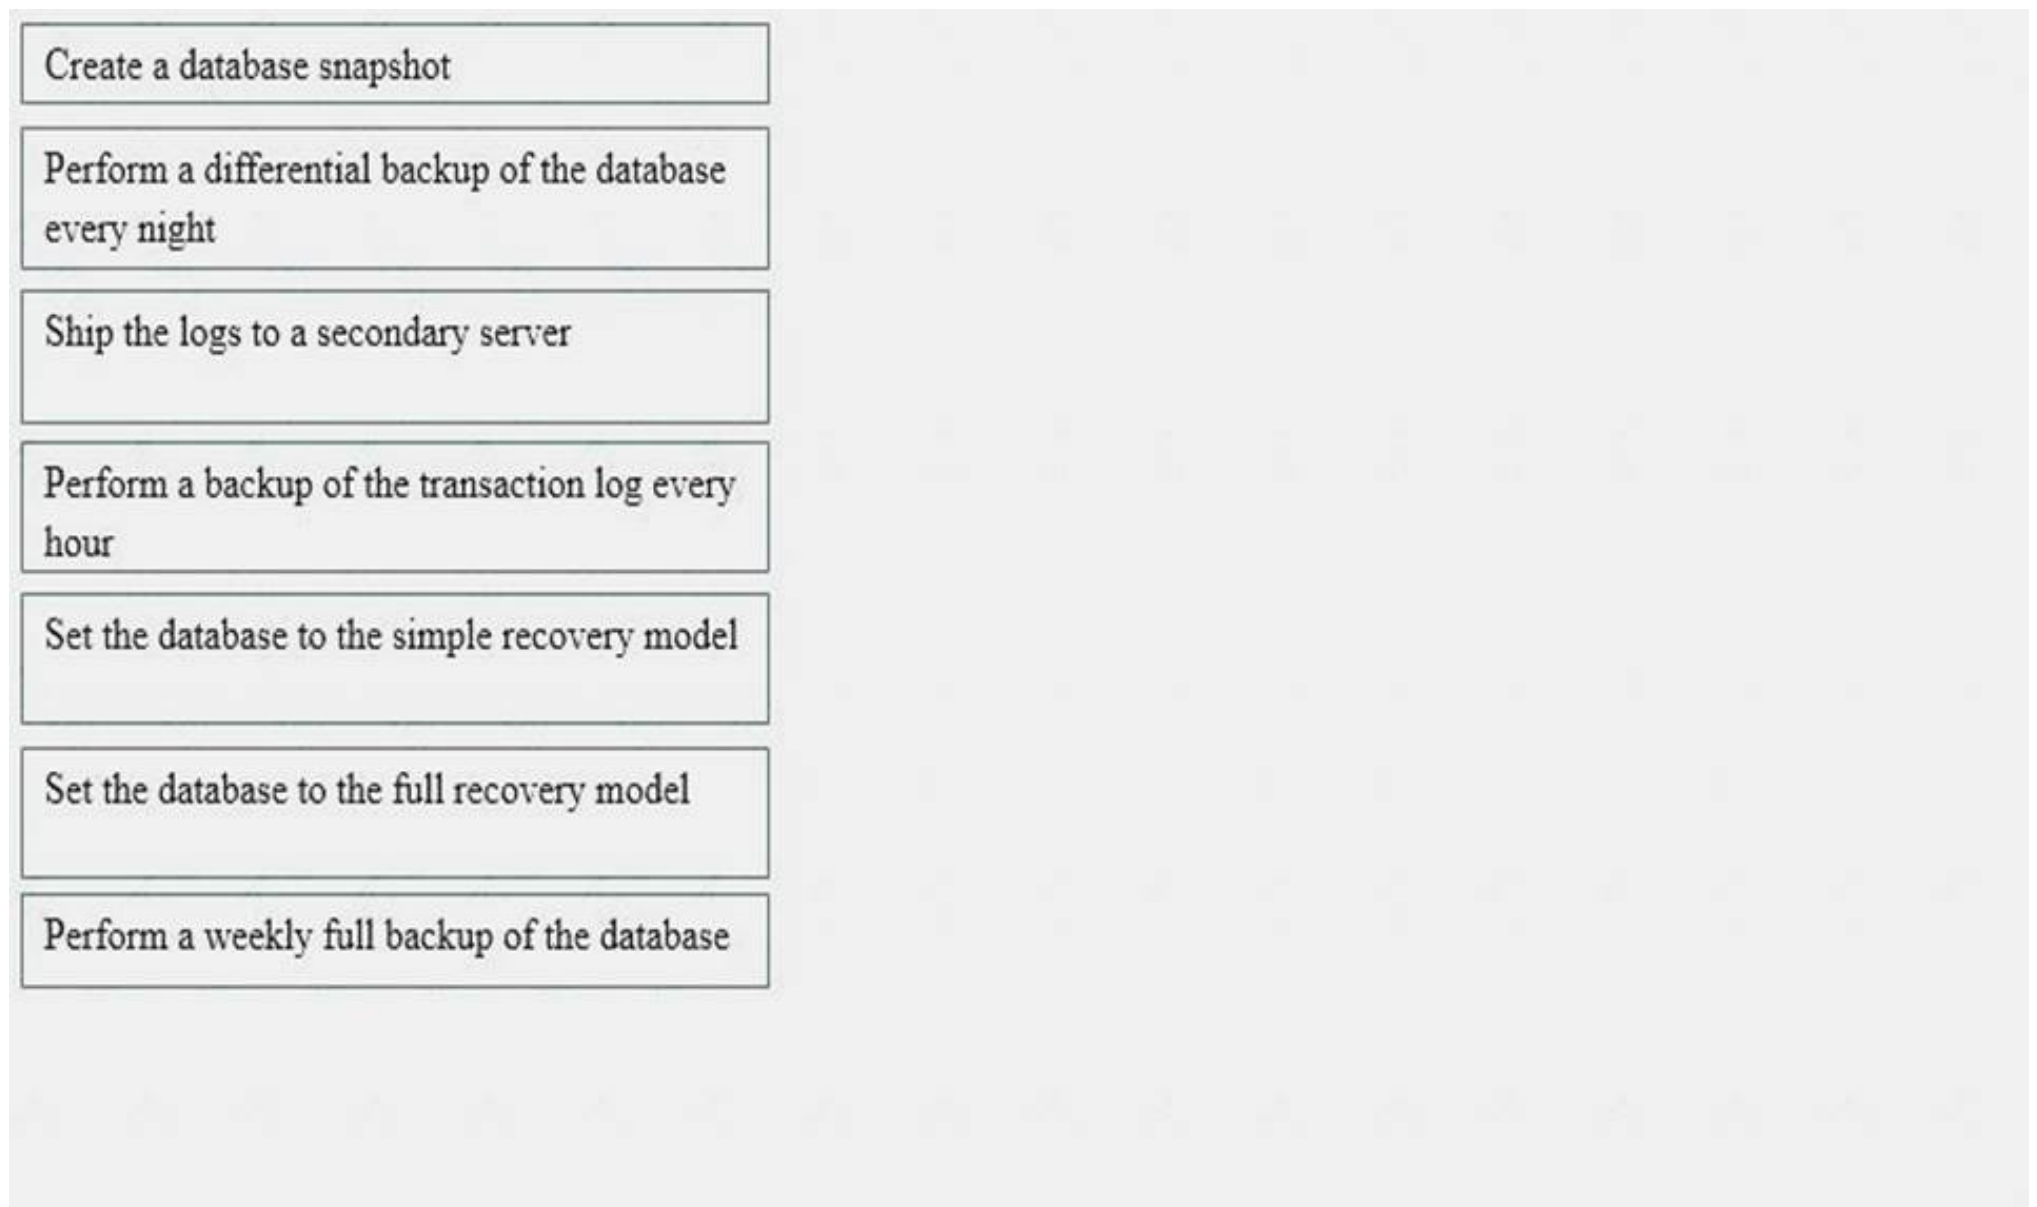
- (Exam Topic 2)

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.

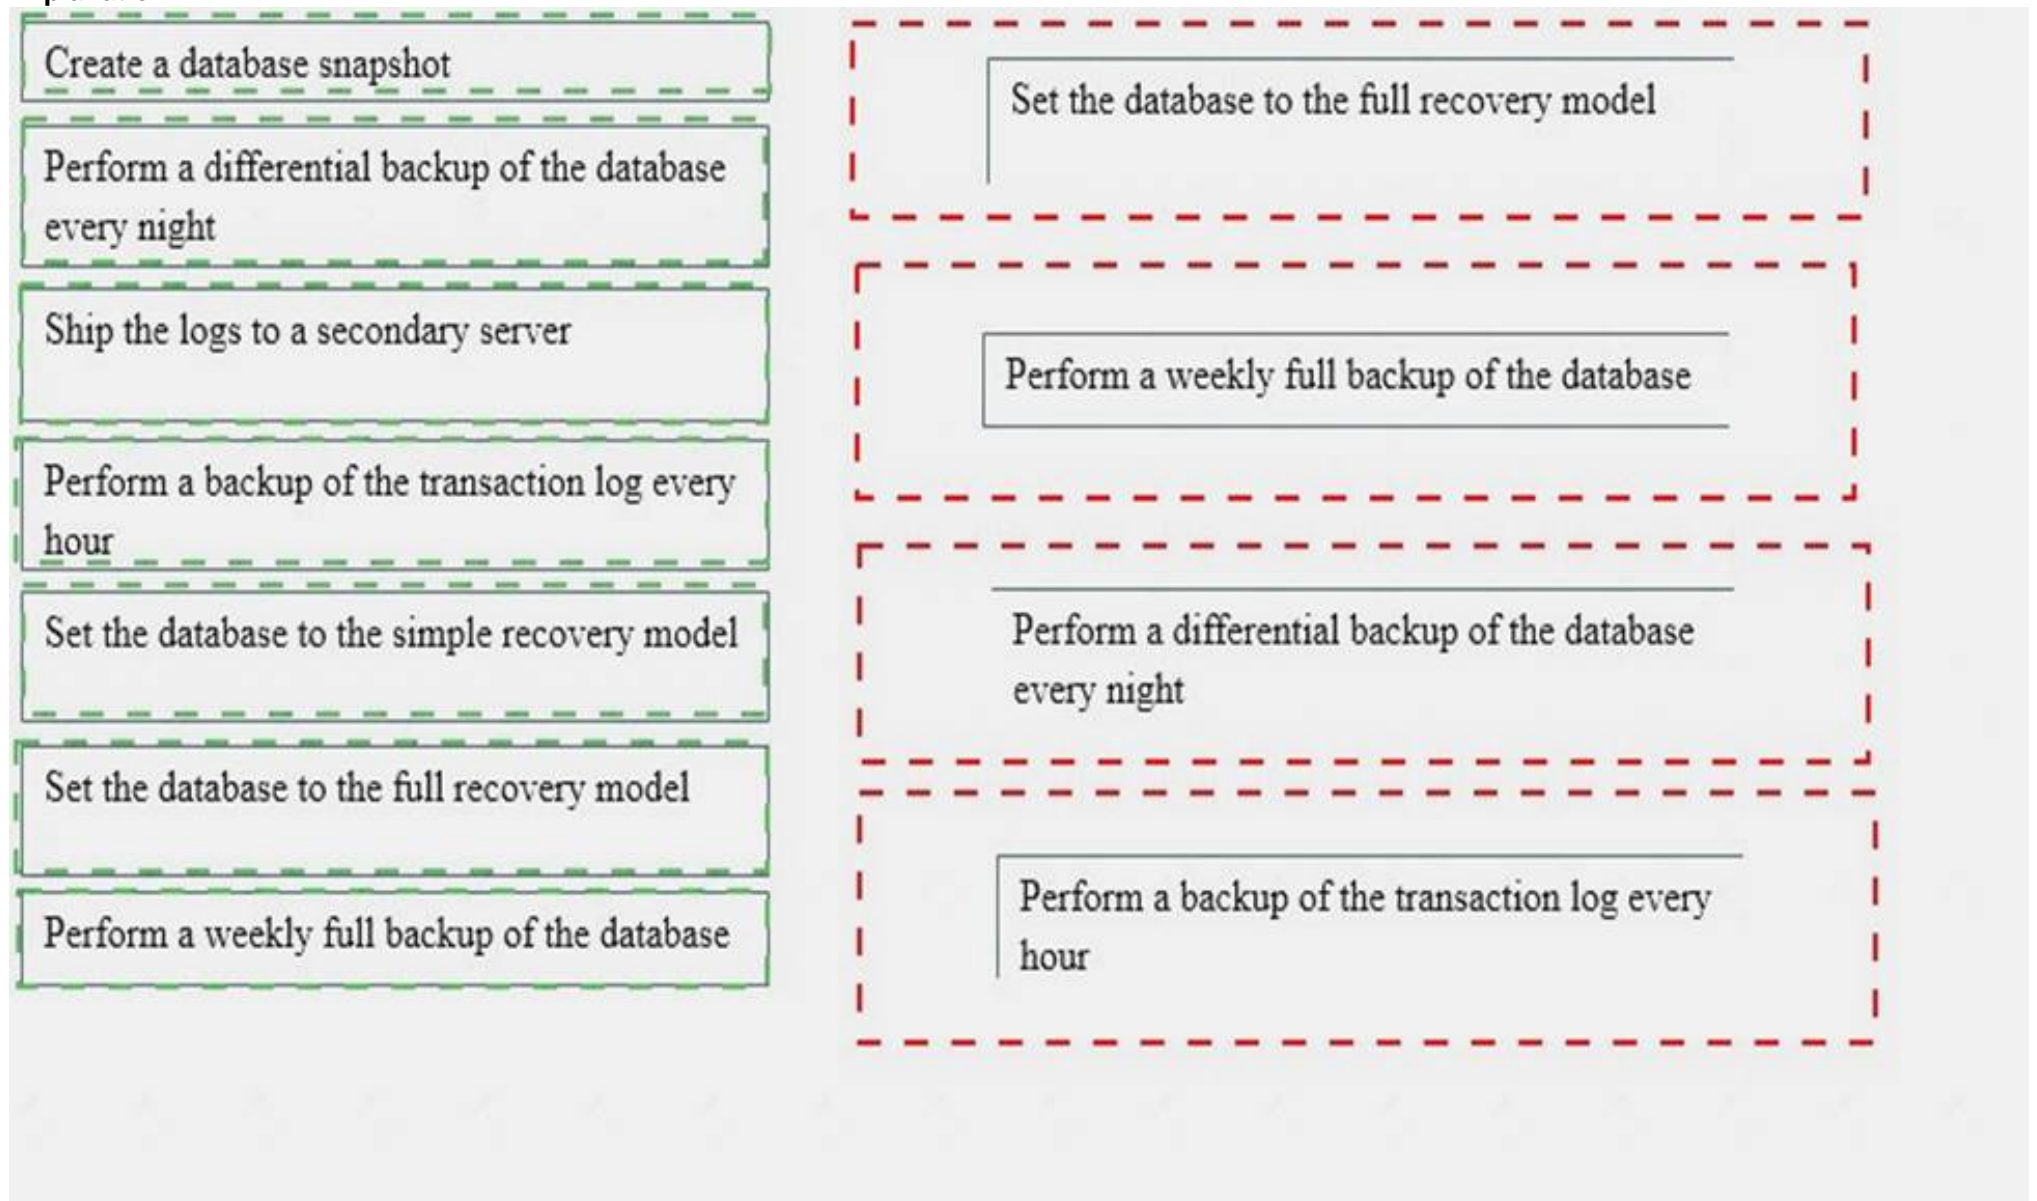
What should you include in the recommendation? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:



NEW QUESTION 179

- (Exam Topic 2)

You administer a Microsoft SQL Server 2016 database that contains a table named OrderDetail. You discover that the NCI_OrderDetail_CustomerID non-clustered

index is fragmented.
You need to reduce fragmentation.
You need to achieve this goal without taking the index offline. Which Transact-SQL batch should you use?

- A. CREATE INDEX NCI_OrderDetail_CustomerID ON OrderDetail.CustomerID WITH DROP EXISTING
- B. ALTER INDEX NCI_OrderDetail_CustomerID ON OrderDetail.CustomerID REORGANIZE
- C. ALTER INDEX ALL ON OrderDetail REBUILD
- D. ALTER INDEX NCI_OrderDetail_CustomerID ON OrderDetail.CustomerID REBUILD

Answer: B

Explanation:

References:
<http://msdn.microsoft.com/en-us/library/ms188388.aspx>

NEW QUESTION 180

- (Exam Topic 2)
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](https://msdn.microsoft.com/en-us/library/ms143531(v=SQL.105).aspx)

NEW QUESTION 182

- (Exam Topic 2)
You plan to deploy a database to SQL Azure. You are designing two stored procedures named USP_1 and USP_2 that have the following requirements:
Prevent data read by USP_1 from being modified by other active processes.
Allow USP_2 to perform dirty reads.
You need to recommend the isolation level for the stored procedures. The solution must maximize concurrency.
Which isolation levels should you recommend? To answer, drag the appropriate isolation level to the correct stored procedure in the answer area.

Isolation Levels		Answer area
Read committed	SP1	Isolation level
Read uncommitted	SP2	Isolation level
Repeatable read		
Serializable		

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

SP1 – repeatable read; SP2 – read uncommitted Note:
- SP1: repeatable read a repeatable read scan retains locks on every row it touches until the end of the transaction. Even rows that do not qualify for the query result remain locked. These locks ensure that the rows touched by the query cannot be updated or deleted by a concurrent session until the current transaction completes (whether it is committed or rolled back).
- SP2: read uncommitted permits repeatable reads

NEW QUESTION 185

- (Exam Topic 2)
You plan to deploy SQL Server 2014. You are designing two stored procedures named SP1 and SP2 that have the following requirements:
- Prevent data read by SP1 from being modified by other active processes.
- Prevent SP2 from performing dirty reads.
You need to recommend the isolation level for each stored procedure.
The solution must maximize concurrency. Which isolation levels should you recommend? To answer, drag the appropriate isolation level to the correct stored procedure in the answer area.

Isolation Levels		Answer area
Read committed	SP1	Isolation level
Read uncommitted	SP2	Isolation level
Repeatable read		
Serializable		

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

SP1 – repeatable read; SP2 – read committed

- REPEATABLE READ

This isolation level includes the guarantees given by SNAPSHOT isolation level. In addition, REPEATABLE READ guarantees that for any row that is read by the transaction, at the time the transaction commits the row has not been changed by any other transaction. Every read operation in the transaction is repeatable up to the end of the transaction.

- Committed Read is SQL Server's default isolation level. It ensures that an operation will never read data another application has changed but not yet committed.

NEW QUESTION 187

- (Exam Topic 2)

Overview

Application Overview

Contoso, Ltd., is the developer of an enterprise resource planning (ERP) application.

Contoso is designing a new version of the ERP application. The previous version of the ERP application used SQL Server 2008 R2.

The new version will use SQL Server 2014.

The ERP application relies on an import process to load supplier data. The import process updates thousands of rows simultaneously, requires exclusive access to the database, and runs daily.

You receive several support calls reporting unexpected behavior in the ERP application. After analyzing the calls, you conclude that users made changes directly to the tables in the database.

Tables

The current database schema contains a table named OrderDetails.

The OrderDetails table contains information about the items sold for each purchase order. OrderDetails stores the product ID, quantities, and discounts applied to each product in a purchase order.

The product price is stored in a table named Products. The Products table was defined by using the SQL_Latin1_General_CP1_CI_AS collation.

A column named ProductName was created by using the varchar data type. The database contains a table named Orders.

Orders contains all of the purchase orders from the last 12 months. Purchase orders that are older than 12 months are stored in a table named OrdersOld.

The previous version of the ERP application relied on table-level security. Stored Procedures

The current version of the database contains stored procedures that change two tables. The following shows the relevant portions of the two stored procedures:

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

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

Customer Problems Installation Issues

The current version of the ERP application requires that several SQL Server logins be set up to function correctly. Most customers set up the ERP application in

multiple locations and must create logins multiple times.

Index Fragmentation Issues

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

Column	Data type
id	uniqueidentifier
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 191

- (Exam Topic 2)

You are designing an authentication strategy for a new server that has SQL Server 2014 installed. The strategy must meet the following business requirements:

The account used to generate reports must be allowed to make a connection during certain hours only.

Failed authentication requests must be logged.

You need to recommend a technology that meets each business requirement. The solution must minimize the amount of events that are logged.

Which technologies should you recommend? To answer, drag the appropriate solution to the correct business requirement in the answer area.

Isolation Levels	Answer area
Login auditing	The account used to generate reports must be allowed to make a connection during certain hours only. Technology
Logon triggers	
C2 audit tracing	Failed authentication requests must be logged. Technology
Policy-Based Management	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

1. Logon triggers fire stored procedures in response to a LOGON event. This event is raised when a user session is established with an instance of SQL Server.

Logon triggers fire after the authentication phase of logging in finishes, but before the user session is actually established.

You can use logon triggers to audit and control server sessions, such as by tracking login activity, restricting logins to SQL Server, or limiting the number of sessions for a specific login.

2. Login auditing can be configured to write to the error log on the following events.

- Failed logins
- Successful logins
- Both failed and successful logins

NEW QUESTION 195

- (Exam Topic 2)

You plan to create a database.

The database will be used by a Microsoft .NET application for a special event that will last for two days. During the event, data must be highly available.

After the event, the database will be deleted.

You need to recommend a solution to implement the database while minimizing costs. The solution must not affect any existing applications.

What should you recommend? More than one answer choice may achieve the goal. Select the BEST answer.

- A. SQL Server 2014 Enterprise
- B. SQL Server 2014 Standard
- C. SQL Azure
- D. SQL Server 2014 Express with Advanced Services

Answer: B

Explanation:

Programmability (AMO, ADOMD.Net, OLEDB, XML/A, ASSL) supported by Standard and Enterprise editions only. References: Features Supported by the Editions of SQL Server 2014.

NEW QUESTION 196

- (Exam Topic 2)

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

- (Exam Topic 2)

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 198

- (Exam Topic 2)

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 202

- (Exam Topic 2)

You have a SQL Server 2012 database named DB1.

You plan to import a large number of records from a SQL Azure database to DB1.

You need to recommend a solution to minimize the amount of space used in the transaction log during the import operation.

What should you include in the recommendation?

- A. a new log file
- B. a new filegroup
- C. the full recovery model
- D. a new partitioned table
- E. the bulk-logged recovery model

Answer: E

Explanation:

Compared to the full recovery model, which fully logs all transactions, the bulk-logged recovery model minimally logs bulk operations, although fully logging other transactions. The bulk-logged recovery model protects against media failure and, for bulk operations, provides the best performance and least log space usage.

Note:

The bulk-logged recovery model is a special-purpose recovery model that should be used only intermittently to improve the performance of certain large-scale bulk operations, such as bulk imports of large amounts of data. Recovery Models (SQL Server)

NEW QUESTION 205

- (Exam Topic 2)

You are the lead database administrator (DBA) of a Microsoft SQL Server 2016 environment. All DBAs are members of the DOMAIN\JrDBAs Active Directory group.

You grant DOMAIN\JrDBAs access to the SQL Server.

You need to create a server role named SpecialDBARole that can perform the following functions:

View all databases.

View the server state.

Assign GRANT, DENY, and REVOKE permissions on logins.

You need to add DOMAIN\JrDBAs to the server role.

You also need to provide the least level of privileges necessary.

Which SQL statement or statements should you use? Choose all that apply.

- A. CREATE SERVER ROLE [SpecialDBARole] AUTHORIZATION setupadmin;
- B. ALTER SERVER ROLE [SpecialDBARole] ADD MEMBER [DOMAIN\JrDBAs];
- C. CREATE SERVER ROLE [SpecialDBARole] AUTHORIZATION securityadmin;
- D. GRANT VIEW DEFINITION TO [SpecialDBARole];
- E. CREATE SERVER ROLE [SpecialDBARole] AUTHORIZATION serveradmin;
- F. GRANT VIEW SERVER STATE, VIEW ANY DATABASE TO [SpecialDBARole];

Answer: BCF

NEW QUESTION 209

- (Exam Topic 2)

Overview

You are a database administrator for a company named Litware, Inc.

Litware is a book publishing house. Litware has a main office and a branch office.

You are designing the database infrastructure to support a new web-based application that is being developed. The web application will be accessed at www.litwareinc.com. Both internal employees and external partners will use the application.

You have an existing desktop application that uses a SQL Server 2008 database named App1_DB. App1_DB will remain in production.

Requirements Planned Changes

You plan to deploy a SQL Server 2014 instance that will contain two databases named Database1 and Database2.

All database files will be stored in a highly available SAN. Database1 will contain two tables named Orders and OrderDetails.

Database1 will also contain a stored procedure named usp_UpdateOrderDetails.

The stored procedure is used to update order information. The stored procedure queries the Orders table twice each time the procedure executes.

The rows returned from the first query must be returned on the second query unchanged along with any rows added to the table between the two read operations.

Database1 will contain several queries that access data in the Database2 tables. Database2 will contain a table named Inventory.

Inventory will contain over 100 GB of data.

The Inventory table will have two indexes: a clustered index on the primary key and a nonclustered index. The column that is used as the primary key will use the identity property.

Database2 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. All data in Database2 is recreated each day and does not change until the next data creation process. Data from Database2 will be accessed periodically by an external application named Application1. The data from Database2 will be sent to a database named Appl_Dbl as soon as changes occur to the data in Database2. Litware plans to use offsite storage for all SQL Server 2014 backups.

Business Requirements

You have the following requirements:

Costs for new licenses must be minimized.

Private information that is accessed by Application must be stored in a secure format.

Development effort must be minimized whenever possible.

The storage requirements for databases must be minimized.

System administrators must be able to run real-time reports on disk usage.
The databases must be available if the SQL Server service fails.
Database administrators must receive a detailed report that contains allocation errors and data corruption.
Application developers must be denied direct access to the database tables. Applications must be denied direct access to the tables.
You must encrypt the backup files to meet regulatory compliance requirements.
The encryption strategy must minimize changes to the databases and to the applications.
You need to recommend a solution to allow application users to perform UPDATE operations on the database tables. The solution must meet the business requirements.
What should you recommend?

- A. Create stored procedures that use EXECUTE AS clauses.
- B. Create a user-defined database role and add users to the role.
- C. Create functions that use EXECUTE AS clauses.
- D. Create a Policy-Based Management Policy.

Answer: A

Explanation:

- EXECUTE AS Clause (Transact-SQL)

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

NEW QUESTION 214

- (Exam Topic 2)

You need to recommend the actions that are required to partition a table.

In which order should the four actions be performed? To answer, move the actions from the list of actions to the answer area and arrange them in the correct order.

Actions	Answer Area
Create a partition scheme	
Create a partition function	
Create filegroups	
Create the table	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

References:

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

NEW QUESTION 219

- (Exam Topic 2)

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.

Index Storage Types

Aligned

Nonaligned

Answer area

IX_Type

IX_Account

IX_Date

IX_Amount

Index Storage Type

Index Storage Type

Index Storage Type

Index Storage Type

- A. Mastered
- B. Not Mastered

Answer: A

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 colocated joins with more tables using different join columns.

NEW QUESTION 224

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your 70-764 Exam with Our Prep Materials Via below:

<https://www.certleader.com/70-764-dumps.html>