

Microsoft

Exam Questions 70-764

Administering a SQL Database Infrastructure (beta)



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 role
- 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 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 3

- (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 4

- (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 5

- (Exam Topic 1)

You have configured Resource Governor with three resource pools.

You have assigned the first resource pool a minimum CPU and memory value of 20%. You have assigned the second resource pool a minimum CPU and memory value of 30%. You want to assign maximum CPU and memory values to the third resource pool.

What is the maximum CPU and memory value you can assign to this resource pool?

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

Answer: B

Explanation:

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

NEW QUESTION 6

- (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 7

- (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>

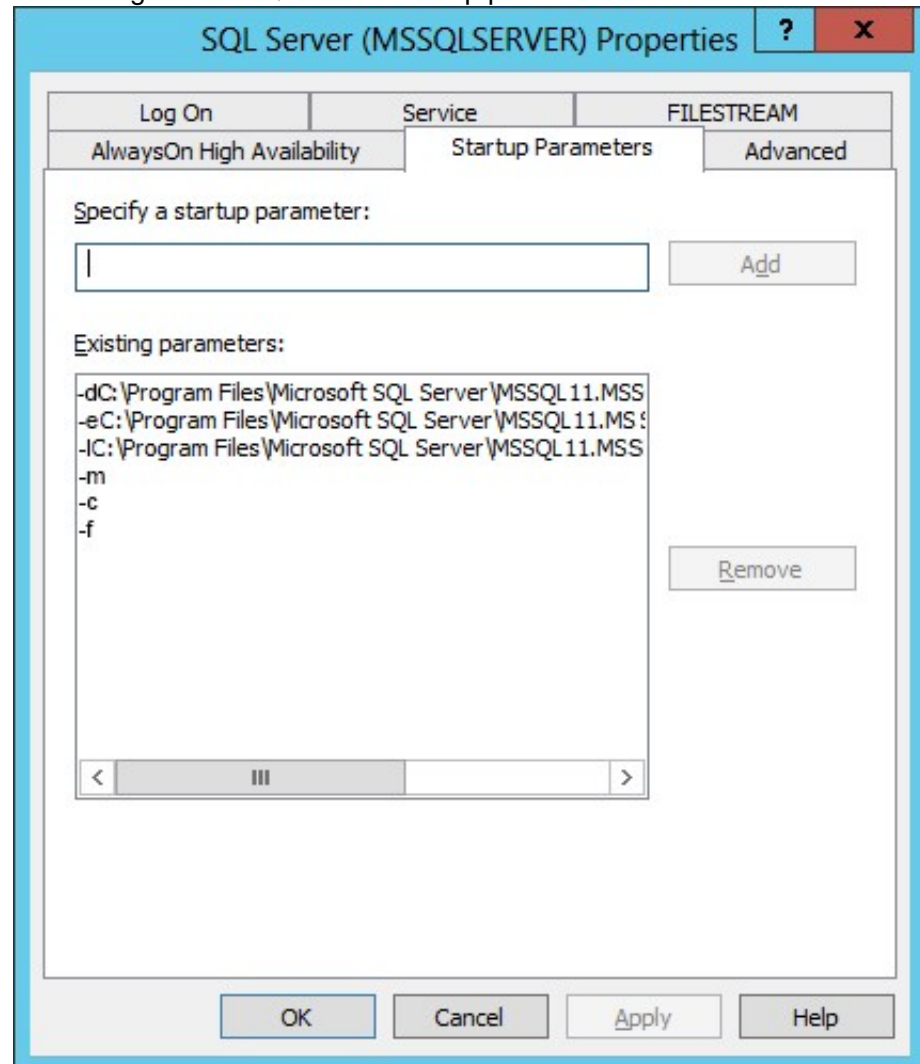
NEW QUESTION 8

- (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 9

- (Exam Topic 1)

You are the database administrator of a Microsoft SQL Server instance. Developers are writing stored procedures to send emails using sp_send_dbmail. Database Mail is enabled.

You need to configure each account's profile security and meet the following requirements:

Account SMTP1_Account must only be usable by logins that have been given explicit permissions to use the SMTP1_profile.

Account SMTP2_Account must only be usable by logins who are a member of the [DatabaseMailUserRole] role in msdb.

In the table below. identify the profile type that must be used for each account. NOTE: Make only one selection in each column.

Answer Area

Profile type	SMTP1_Account	SMTP2_Account
Private Profile	<input type="radio"/>	<input type="radio"/>
Public Profile	<input type="radio"/>	<input type="radio"/>
Default Profile	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

SMTP1_Account1: Private Profile
When no profile_name is specified, sp_send_dbmail uses the default private profile for the current user. I user does not have a default private profile, sp_send_dbmail uses the default public profile for the msdb database.
SMTP1_Account2: Default Profile
Execute permissions forsp_send_dbmail default to all members of the DatabaseMailUser database role in the msdb database.
References:
<https://docs.microsoft.com/en-us/sql/relational-databases/system-stored-procedures/sp-send-dbmail-transact-sql>

NEW QUESTION 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 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

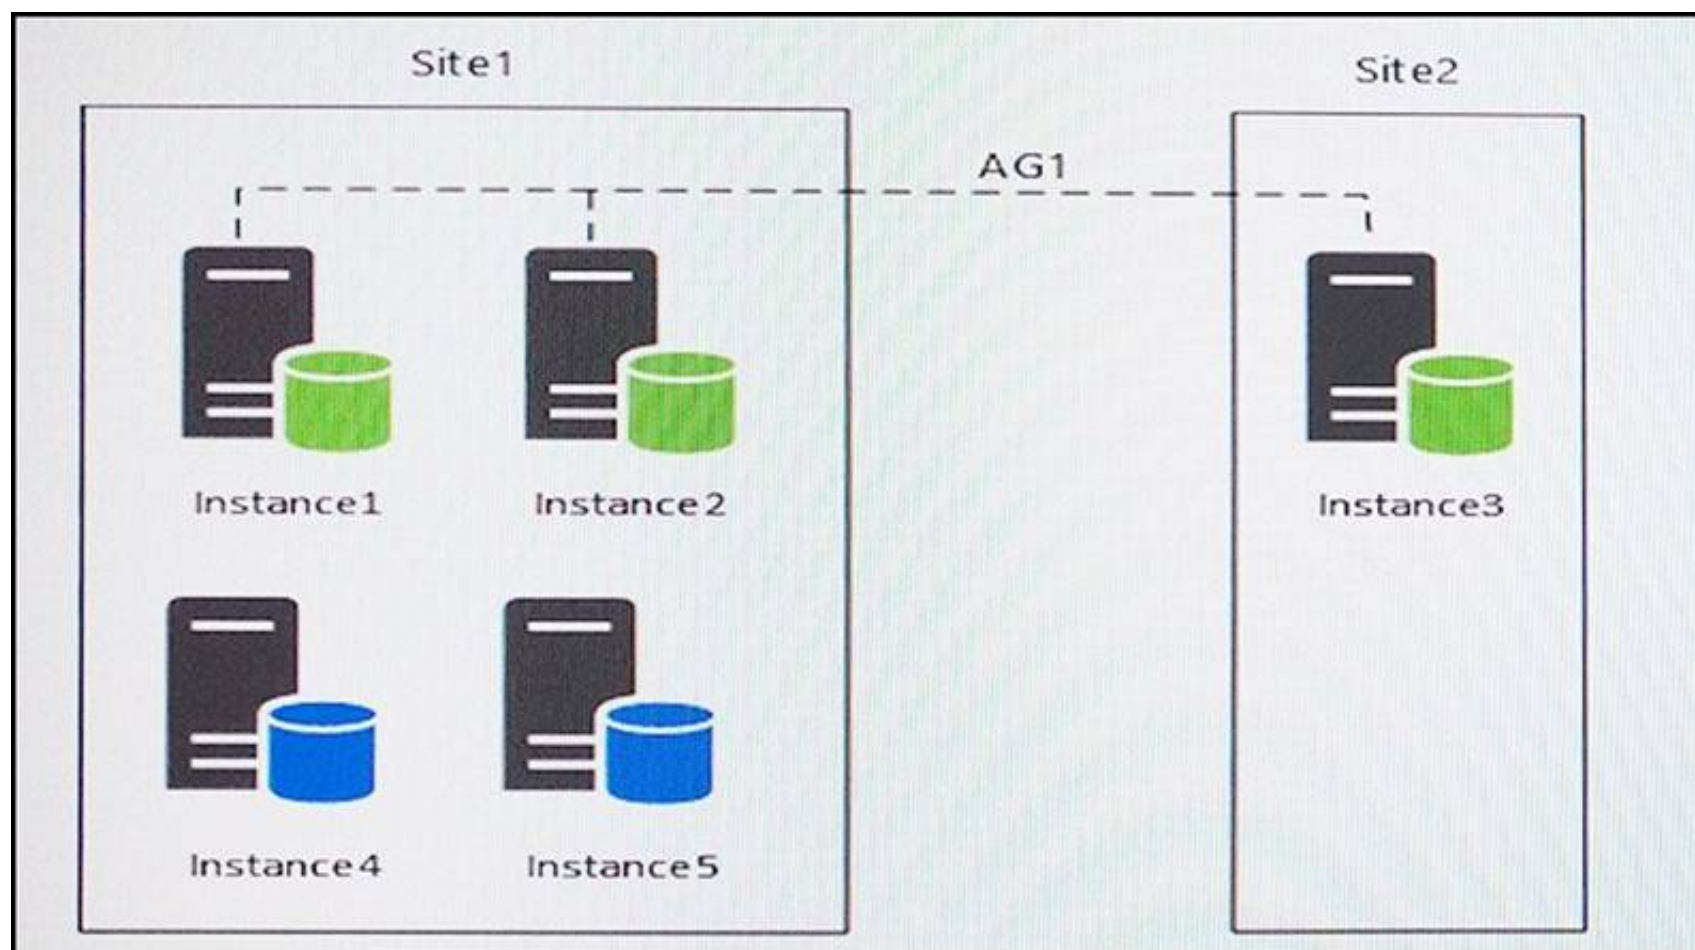
Answer: G

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 10

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



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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Answer: B

Explanation:


From scenario: Instance4 and Instance5 are not part of AG1. Instance4 is engaged in heavy read-write I/O. The recovery point objective of Instance4 is 60

minutes. RecoveryPoint Objectives are commonly described as the amount of data that was lost during the outage and recovery period. You should minimize both the time required to restore the databases and the space required to store backups.
References:
<http://sqlmag.com/blog/sql-server-recovery-time-objectives-and-recovery-point-objectives>

NEW QUESTION 12

- (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,

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

SAN Storage for new SQL Cluster

SAN Storage for new SQL Cluster

...
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

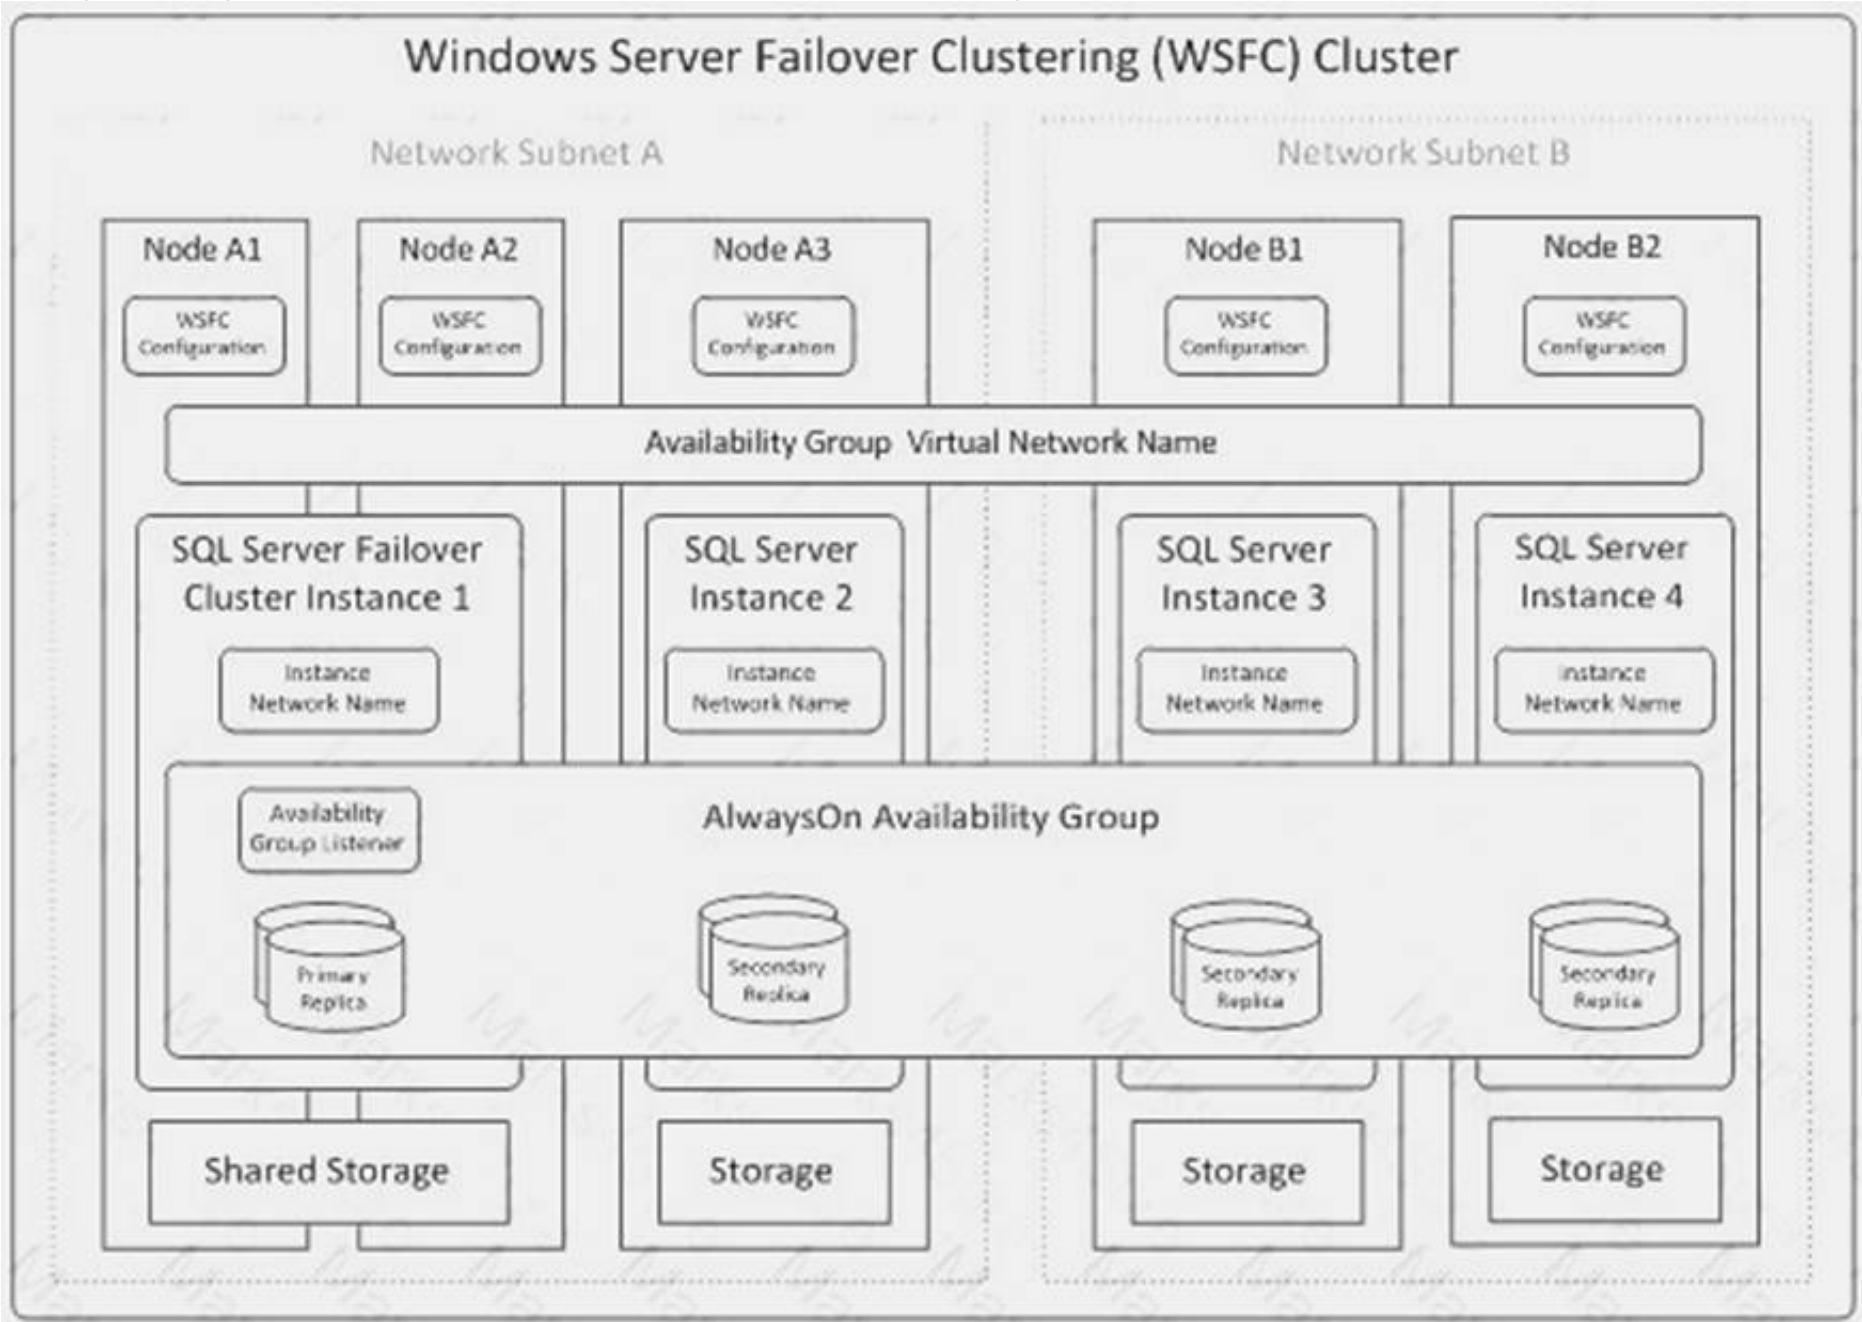
Statements	Yes	No
The Tempdb database for each cluster instance can be placed on the D: drive for the instance.	<input type="radio"/>	<input type="radio"/>
One virtual network name for each SQL Server instance must be configured in the cluster.	<input type="radio"/>	<input type="radio"/>
The shared storage 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 14

- (Exam Topic 1)

You administer all the deployments of Microsoft SQL Server 2016 in your company.

You need to ensure that an OLTP database that includes up-to-the-minute reporting requirements can be off-loaded from the primary database to another server.

You also need to be able to add indexes to the secondary database. Which configuration should you use?

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

Answer: F

Explanation:

References:

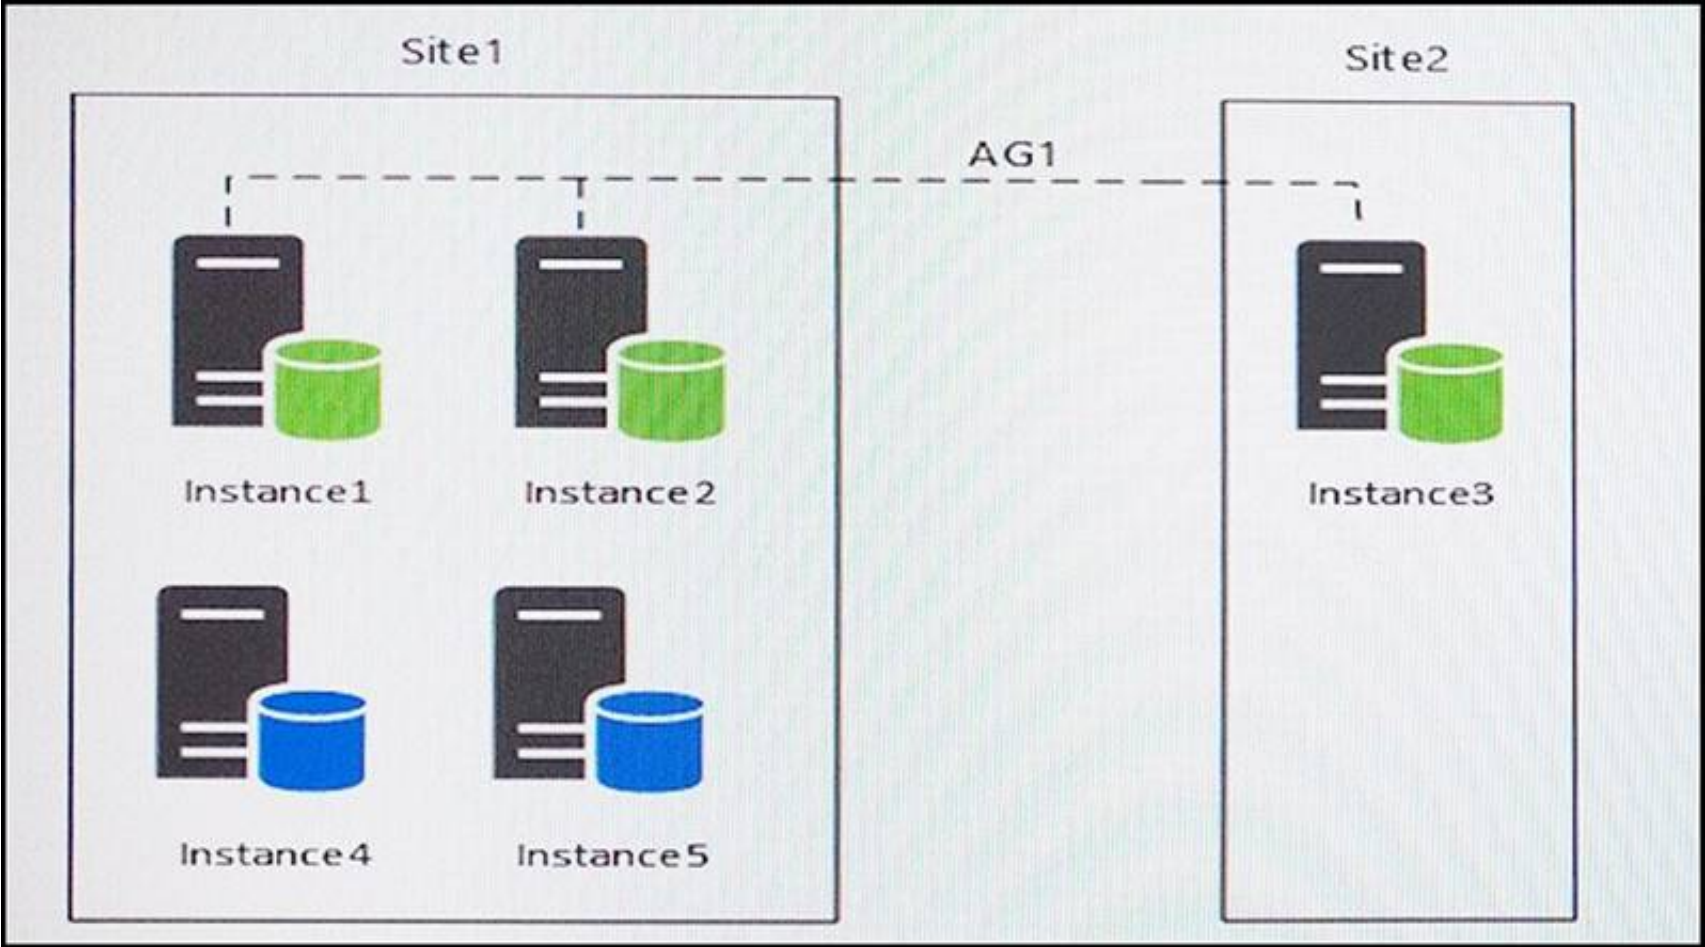
<https://docs.microsoft.com/en-us/sql/relational-databases/replication/transactional/transactional-replication>

NEW QUESTION 18

- (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 in DB1 with 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 in DB1 with 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 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.

References:

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

NEW QUESTION 22

- (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 tempdb database. Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

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

Note: Database Mail was first introduced as a new feature in 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 24

- (Exam Topic 1)

You manage a Microsoft-SQL Server database named sales Orders.

You need to verify the integrity of the database and attempt to repair any errors that are found. Repair must not cause any data to be lost in the database.

How should you complete the DBCC command? To answer, select the appropriate options in the answer area.

Answer Area

DBCC

CHECKDB
PHYSICAL_ONLY
REPAIR_FAST
REPAIR_REBUILD

 ('salesOrders',

CHECKDB
PHYSICAL_ONLY
REPAIR_FAST
REPAIR_REBUILD

)

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: CHECKDB

DBCC CHECKDB checks the logical and physical integrity of all the objects in the specified database. Partial syntax:

DBCC CHECKDB

[(database_name | database_id | 0 [, NOINDEX

| , { REPAIR_ALLOW_DATA_LOSS | REPAIR_FAST | REPAIR_REBUILD }]

....

Box 2: REPAIR_REBUILD

DBCC CHECKDB ...REPAIR_ALLOW_DATA_LOSS | REPAIR_FAST |REPAIR_REBUILD specifies that

DBCC CHECKDB repair the found errors.

REPAIR_REBUILD performs repairs that have no possibility of data loss. This can include quick repairs, such as repairing missing rows in non-clustered indexes, and more time-consuming repairs, such as rebuilding an index.

References: <https://docs.microsoft.com/en-us/sql/t-sql/database-console-commands/dbcc-checkdb-transact-sql>

NEW QUESTION 28

- (Exam Topic 1)

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

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

Clients connect to databases by using line-of-business applications. Developers connect by using SQL Server Management Studio (SSMS).

You need to provide permissions to a service account that will be used to provision a new database for a client. Which permission should you grant?

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

Answer: C

Explanation:

Members of the dbcreator fixed server role can create, alter, drop, and restore any database.

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/security/authentication-access/server-level-roles>

NEW QUESTION 33

- (Exam Topic 1)

You administer a Microsoft SQL Server 2016 database.

Users report that an application that accesses the database displays an error, but the error does not provide meaningful information.

No entries are found in the SQL Server log or Windows event logs related to the error. You need to identify the root cause of the issue by retrieving the error message.

What should you do?

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

Answer: A

Explanation:

Trapping SQL Server Errors with Extended Events

One very useful usage of Extended Events is the ability to trap SQL Server error without the need to have a server trace running (which, btw, is deprecated), with the additional feature of being able to query the data as soon as it comes in. This means that we a solution to monitor and trap errors as soon as they happen can be easily created, in order to help developers to fix problems as soon as they are detected. This is really, really, really helpful especially in very big applications, where the code base is quite old and there is no-one really knowing everything of the solution.

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

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

```
ADD EVENT sqlserver.error_reported Etc.
```

References:

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

NEW QUESTION 36

- (Exam Topic 1)

You are the administrator of a Microsoft SQL Server 2016 server. Some applications consume significant resources.

You need to manage the server workload by restricting resource-intensive applications. You need to dynamically limit resource consumption.

What should you do?

A. Set up Service Broker to ensure that applications are not allowed to consume more than the specified amount of resources.

B. Configure Resource Pools, Workload Groups, and Classifier Function, and then enable the Resource Governor.

C. Configure Extended Events to monitor and restrict resource limits allowed by each application type.

D. Create a new Plan Guide with a Scope Type of sql and define the resource limits for each application.

Answer: B

NEW QUESTION 41

- (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><div>Primary server instance</div><div>Secondary server instance</div><div>Monitor server instance</div><div>Backup share file server</div></div>
Copy the backup job.	<div><div></div><div>Primary server instance</div><div>Secondary server instance</div><div>Monitor server instance</div><div>Backup share file server</div></div>
Restore the backup.	<div><div></div><div>Primary server instance</div><div>Secondary server instance</div><div>Monitor server instance</div><div>Backup share file server</div></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 42

- (Exam Topic 1)

You administer two instances of Microsoft SQL Server 2016.

You deploy an application that uses a database on the named instance.

The application is unable to connect to the database on the named instance. You need to ensure that the application can connect to the named instance. What should you do?

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

Answer: B

Explanation:

The SQL ServerBrowser program runs as a Windows service. SQL Server Browser listens for incoming requests for Microsoft SQL Server resources and provides information about SQL Server instances installed on the computer. SQL Server Browser contributes to the following actions:

Browsing a list of available servers
Connecting to the correct server instance
Etc.

References: <https://docs.microsoft.com/en-us/sql/tools/configuration-manager/sql-server-browser-service>

NEW QUESTION 45

- (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 48

- (Exam Topic 1)

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

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

You plan to delegate encryption operations to a user.

You need to grant the user permission to implement cell-level encryption while following the principle of least privilege.

Which permission should you grant?

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

Answer: G

Explanation:

The following permissions are necessary to perform column-level encryption, or cell-level encryption.

CONTROL permission on the database.

CREATE CERTIFICATE permission on the database. Only Windows logins, SQL Server logins, and application roles can own certificates. Groups and roles cannot own certificates.

ALTER permission on the table.

Some permission on the key and must not have been denied VIEW DEFINITION permission. References:

<https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/encrypt-a-column-of-data>

NEW QUESTION 53

- (Exam Topic 1)

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

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

A company has an on-premises Microsoft SQL Server environment and Microsoft Azure SQL Database instances. The environment hosts several customer databases.

One customer reports that their database is not responding as quickly as the service level agreements dictate. You observe that the database is fragmented.

You need to optimize query performance. Solution: You reorganize all indexes. Does the solution meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

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

NEW QUESTION 57

- (Exam Topic 2)

You are designing a monitoring application for a new SQL Server 2014 instance.

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

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

- A. The SQL Server error log
- B. The sys.dm_os_wait_stats dynamic management view
- C. The DBCC SQLPERF(WAITSTATS) command
- D. SQL Server Profiler

Answer: B

Explanation:

sys.dm_os_wait_stats

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

Columns include: waiting_tasks_count

Number of waits on this wait type.

This counter is incremented at the start of each wait.

NEW QUESTION 58

- (Exam Topic 2)

You plan to deploy SQL Server 2014.

You identify the following security requirements for the deployment:

Users must be prevented from intercepting and reading the T-SQL statements sent from the clients to the database engine.

All database files and log files must be encrypted if the files are moved to another disk on another server.

You need to identify which feature meets each security requirement. The solution must minimize processor overhead.

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

Features	Answer Area
Encrypting File System (EFS)	Users must be prevented from intercepting and reading the T-SQL statements sent from the clients to the database engine. <div>Feature</div>
Policy-Based Management	
Secure Socket Layer (SSL)	
Transparent Data Encryption (TDE)	All database files and log files must be encrypted if the files are moved to another disk on another server. <div>Feature</div>
Windows BitLocker Drive Encryption (BitLocker)	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

NEW QUESTION 60

- (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 63

- (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 66

- (Exam Topic 2)

You administer a Microsoft SQL Server 2016 server. You plan to deploy new features to an application.

You need to evaluate existing and potential clustered and non-clustered indexes that will improve performance.

What should you do?

A. Query the sys.dm_db_index_usage_stats DMV.

B. Query the sys.dm_db_missing_index_details DMV.

C. Use the Database Engine Tuning Advisor.

D. Query the sys.dm_db_missing_index_columns DMV.

Answer: C

Explanation:

The Microsoft Database Engine Tuning Advisor (DTA) analyzes databases and makes recommendations that you can use to optimize query performance. You can use the Database Engine Tuning Advisor to select and create an optimal set of indexes, indexed views, or table partitions without having an expert understanding of the database structure or the internals of SQL Server. Using the DTA, you can perform the following tasks.

Troubleshoot the performance of a specific problem query Tune a large set of queries across one or more databases

Perform an exploratory what-if analysis of potential physical design changes Manage storage space

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/performance/database-engine-tuning-advisor>

NEW QUESTION 71

- (Exam Topic 2)

You administer a Microsoft SQL Server 2016 server.

One of the databases on the server supports a highly active OLTP application. Users report abnormally long wait times when they submit data into the application.

You need to identify which queries are taking longer than 1 second to run over an extended period of time. What should you do?

A. use SQL Profiler to trace all queries that are processing on the serve

B. Filter queries that have a Duration value of more than 1,000.

C. Use sp_configure to set a value for blocked process threshol

D. Create an extended event session.

E. Use the Job Activity monitor to review all processes that are actively runnin

F. Review the Job History to find out the duration of each step.

G. Run the sp_who command from a query window.

H. Run the DBCC TRACEON 1222 command from a query window and review the SQL Server event log.

Answer: B

NEW QUESTION 76

- (Exam Topic 2)

Overview

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

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

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

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

Requirements Planned Changes

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

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

Database1 will also contain a stored procedure named usp_UpdateOrderDetails.

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

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

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

Inventory will contain over 100 GB of data.

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

Database2 wilt contains a stored procedure named usp_UpdateInventory. usp_UpdateInventory will manipulate a table that contains a self-join that has an unlimited number of hierarchies. All data in Database2 is recreated each day ad does not change until the next data creation process. Data from Database2 will be

accessed periodically by an external application named Application1. The data from Database2 will be sent to a database named Appl_Dbl as soon as changes occur to the data in Database2. Litware plans to use offsite storage for all SQL Server 2014 backups.

Business Requirements

You have the following requirements:

Costs for new licenses must be minimized.

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

Development effort must be minimized whenever possible.

The storage requirements for databases must be minimized.

System administrators must be able to run real-time reports on disk usage.

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

Database administrators must receive a detailed report that contains allocation errors and data corruption.

Application developers must be denied direct access to the database tables. Applications must be denied direct access to the tables.

You must encrypt the backup files to meet regulatory compliance requirements.

The encryption strategy must minimize changes to the databases and to the applications. You need to recommend a 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 78

- (Exam Topic 2)

You want to reproduce the same SQL Server 2016 installation configuration across five servers. Which of the following files will you generate by using SQL Server Setup to accomplish this goal?

- A. Configuration.xml
- B. Setup.ini
- C. Setup.xml
- D. ConfigurationFile.ini

Answer: D

NEW QUESTION 82

- (Exam Topic 2)

Overview

General Overview

ADatum Corporation has offices in Miami and Montreal.

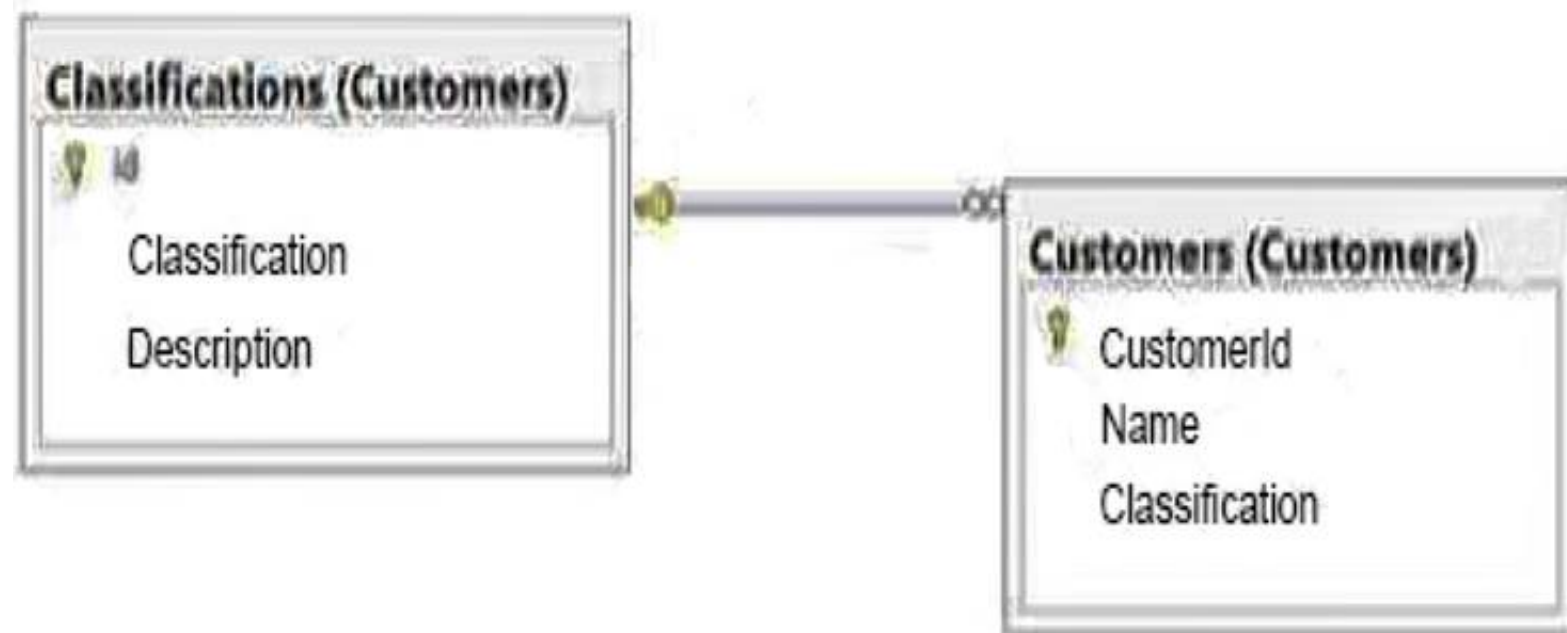
The network contains a single Active Directory forest named adatum.com. The offices connect to each other by using a WAN link that has 5-ms latency. A. Datum standardizes its database platform by using SQL Server 2014 Enterprise edition.

Databases

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

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

Classifications (Customers)



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

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

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

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

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

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

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

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

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

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

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

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

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

The nested stored procedures are never called directly. Design Requirements

Data Recovery

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

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

Classification Changes

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

Storage

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

Error Handling

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

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

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

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

Answer: A

Explanation:

Scenario:

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

NEW QUESTION 87

- (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 feature to support your backup solution. What should you include in the recommendation?

- A. Transparent Data Encryption (TDE)

- B. Column-level encryption
- C. An NTFS file permission
- D. A Secure Sockets Layer (SSL)

Answer: A

Explanation:

- Scenario: You must encrypt the backup files to meet regulatory compliance requirements. The encryption strategy must minimize changes to the databases and to the applications.

- Transparent data encryption (TDE) performs real-time I/O encryption and decryption of the data and log files. The encryption uses a database encryption key (DEK), which is stored in the database boot record for availability during recovery.

Transparent Data Encryption (TDE)

NEW QUESTION 90

- (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 94

- (Exam Topic 2)

Overview

Application Overview

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

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

The new version will use SQL Server 2014.

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

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

Tables

The current database schema contains a table named OrderDetails.

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

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

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

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

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

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

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

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

Customer Problems Installation Issues

The current version of the ERP application requires that several SQL Server logins be set up to function

correctly. Most customers set up the ERP application in multiple locations and must create logins multiple times.

Index Fragmentation Issues

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

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

Backup Issues

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

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

Missing Data Issues

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

Query Performance Issues

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

Design Requirements

File Storage Requirements

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

Data Recovery Requirements

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

Security Requirements

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

Concurrency Requirements

You must reduce the likelihood of deadlocks occurring when Sales.Prod and Sales.Proc2 execute. You need to recommend a solution that addresses the file storage requirements.

What should you include in the recommendation?

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

Answer: B

Explanation:

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

- FileTables remove a significant barrier to the use of SQL Server for the storage and management of unstructured data that is currently residing as files on file servers.

Enterprises can move this data from file servers into FileTables to take advantage of integrated administration and services provided by SQL Server. At the same time, they can maintain Windows application compatibility for their existing Windows applications that see this data as files in the file system.

NEW QUESTION 95

- (Exam Topic 2)

You administer a Microsoft SQL Server 2016 database.

You provide temporary securityadmin access to User1 to the database server. You need to know if User1 adds logins to securityadmin.

Which server-level audit action group should you use?

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

Answer: D

Explanation:

SERVER_ROLE_MEMBER_CHANGE_GROUP

This event is raised whenever a login is added or removed from a fixed server role. This event is raised for the sp_addsrvrolemember and sp_dropsrvrolemember stored procedures. Equivalent to the Audit Add Login to Server Role Event Class.

References:

<http://technet.microsoft.com/en-us/library/cc280663.aspx>

NEW QUESTION 97

- (Exam Topic 2)

You are implementing a SQL Server 2016 five-node failover cluster. You need to choose a quorum configuration.

Which configuration should you use?

- A. Distributed File System (DFS)
- B. Node Majority
- C. Cluster Shared Volume (CSV)
- D. Node and Disk Majority

Answer: D

Explanation:

Node and Disk Majority (recommended for clusters with an even number of nodes)

NEW QUESTION 100

- (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 105

- (Exam Topic 2)

You have two servers named SQL1 and SQL2 that have SQL Server 2012 installed.

SQL1 contains a database that is mirrored asynchronously to SQL2. The database contents are updated once a month.

You need to upgrade the database to SQL Server 2014. The solution must minimize downtime. Which upgrade steps should you recommend? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Fail over
Fail back
Upgrade SQL1
Upgrade SQL2
Establish a mirror
Break the mirror

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

References:<https://docs.microsoft.com/en-us/sql/database-engine/database-mirroring/upgrading-mirrored-instanc>

NEW QUESTION 110

- (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

Index Storage Type

IX_Account

Index Storage Type

IX_Date

Index Storage Type

IX_Amount

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 collocated joins with more tables using different join columns.

NEW QUESTION 113

- (Exam Topic 2)

Overview

Application Overview

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

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

The new version will use SQL Server 2014.

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

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

Tables

The current database schema contains a table named OrderDetails.

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

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

A column named ProductName was created by using the varchar data type. The database contains a table named Orders. Orders contains all of the purchase orders from the last 12 months. Purchase orders that are older than 12 months are stored in a table named OrdersOld. The previous version of the ERP application relied on table-level security. Stored Procedures
 The current version of the database contains stored procedures that change two tables. The following shows the relevant portions of the two stored procedures:

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

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

Customer Problems Installation Issues

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

Index Fragmentation Issues

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

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

Backup Issues

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

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

Missing Data Issues

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

Query Performance Issues

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

Design Requirements

File Storage Requirements

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

Data Recovery Requirements

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

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

Concurrency Requirements

You must reduce the likelihood of deadlocks occurring when Sales.Prod and Sales.Proc2 execute.

You need to recommend a solution that reduces the time it takes to import the supplier data. What should you include in the recommendation?

- A. Enable instant file initialization.
- B. Reorganize the indexes.
- C. Disable Resource Governor.
- D. Enable Auto Update Statistics.

Answer: C

Explanation:

- 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.

NEW QUESTION 117

- (Exam Topic 2)

You plan to deploy three highly available SQL Server environments that will use SQL Server 2014. You identify the following specifications for each environment as shown following table.

Environment	Number of nodes	SQL Server edition	Automatic failover required
ENV1	3	Standard	Yes
ENV2	3	Enterprise	Yes
ENV3	4	Enterprise	Yes

You need to recommend which high-availability feature is required for each environment.
Which features should you identify? To answer, drag the appropriate feature to the correct requirement in the answer area.

Features

AlwaysOn availability groups

Database mirroring

Log shipping

Peer-to-peer replication

Answer Area

ENV1

ENV2

ENV3

Feature

Feature

Feature

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:
ENV1 - Database mirroring; ENV2 - AlwaysOn availability groups; ENV3 - AlwaysOn availability groups
- AlwaysOn availability groups
AlwaysOn Availability Groups feature is a high-availability and disaster-recovery solution that provides an enterprise-level alternative to database mirroring. Introduced in SQL Server 2012, AlwaysOn Availability Groups maximizes the availability of a set of user databases for an enterprise. Deploying AlwaysOn Availability Groups requires a Windows Server Failover Clustering (WSFC) cluster.

NEW QUESTION 118

- (Exam Topic 2)
You manage a SQL Server 2014 instance that contains a database named DB1. Users report that some queries to DB1 take longer than expected. Although most queries run in less than one second, some queries take up to 20 seconds to run. You need to view all of the performance statistics for each database file.
Which method should you use?

- A. Query the sys.dm_os_tasks dynamic management view.
- B. Query the sys.dm_os_performance_counters dynamic management view.
- C. Query the sys.dm_io_virtual_file_stats dynamic management function.
- D. Examine the Data File I/O pane in Activity Monitor.

Answer: C

Explanation:
sys.dm_io_virtual_file_stats Returns I/O statistics for data and log files.

NEW QUESTION 122

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

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

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

Customer Problems Installation Issues

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

Index Fragmentation Issues

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

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

Backup Issues

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

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

Missing Data Issues

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

Query Performance Issues

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

Design Requirements

File Storage Requirements

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

Data Recovery Requirements

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

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

Concurrency Requirements

You must reduce the likelihood of deadlocks occurring when Sales.Prod and Sales.Proc2 execute. You need to recommend a solution that addresses the backup issue.

The solution must minimize the amount of development effort. What should you include in the recommendation?

- A. Indexed views
- B. Filegroups
- C. Table partitioning
- D. Indexes

Answer: B

Explanation:

- Backup Issues

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

- For very large databases (and by that, I mean, at least 500gb, but more like 5-10tb or more), it can become too expensive to regularly run a straight full backup. So, where needed, you can choose to backup smaller pieces of the database by choosing to back up one of the files or file groups that make up a database.

NEW QUESTION 124

- (Exam Topic 2)

You administer a Microsoft SQL Server database named Sales. The database is 3 terabytes in size.

The Sales database is configured as shown in the following table. You discover that all files except Sales_2.ndf are corrupt.

You need to recover the corrupted data in the minimum amount of time. What should you do?

Filegroup	File
PRIMARY	<ul style="list-style-type: none"> Sales.mdf
XACTIONS	<ul style="list-style-type: none"> Sales_1.ndf Sales_2.ndf Sales_3.ndf
ARCHIVES	<ul style="list-style-type: none"> SalesArch_1.ndf SalesArch_2.ndf

- A. Perform a restore from a full backup.
- B. Perform a transaction log restore.
- C. Perform a file restore.
- D. Perform a filegroup restore.

Answer: A

Explanation:

References:
<http://technet.microsoft.com/en-us/library/ms187048.aspx>

NEW QUESTION 125

- (Exam Topic 3)

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

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

You attempt to restore a database on a new SQL Server instance and receive the following error message: "Msg 33111, Level 16, State 3, Line 2

Cannot find server certificate with thumbprint '0x7315277C70764B1F252DC7A5101F6F66EFB1069D.'" You need to ensure that you can restore the database successfully.

Solution: You add the backup set password to the restore command. Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

The error is related to the certificate.

References: <https://dba.stackexchange.com/questions/3388/restore-encrypted-database-to-another-server?rq=1>

NEW QUESTION 126

- (Exam Topic 3)

Background Corporate Information

Fabrikam, Inc. is a retailer that sells electronics products on the Internet. The company has a headquarters site and one satellite sales office. You have been hired as the database administrator, and the company wants you to change the architecture of the Fabrikam ecommerce site to optimize performance and reduce downtime while keeping capital expenditures to a minimum. To help with the solution, Fabrikam has decided to use cloud resources as well as on-premise servers.

Physical Locations

All of the corporate executives, product managers, and support staff are stationed at the headquarters office. Half of the sales force works at this location. There is also a satellite sales office. The other half of the sales force works at the satellite office in order to have sales people closer to clients in that area. Only sales people work at the satellite location.

Problem Statement

To be successful, Fabrikam needs a website that is fast and has a high degree of system uptime. The current system operates on a single server and the company is not happy with the single point of failure this presents. The current nightly backups have been failing due to insufficient space on the available drives and manual drive cleanup often needing to happen to get past the errors. Additional space will not be made available for backups on the HQ or satellite servers. During your investigation, you discover that the sales force reports are causing significant contention.

Configuration Windows Logins

The network administrators have set up Windows groups to make it easier to manage security. Users may belong to more than one group depending on their role. The groups have been set up as shown in the following table:

Group	Members
OurDomain\Management	All corporate executives
OurDomain\SalesStaff	All sales people
OurDomain\ProductionStaff	All product managers and support staff
OurDomain\AllUsers	Everyone
OurDomain\CustomerSupport	Customer support representatives

Server Configuration The IT department has configured two physical servers with Microsoft Windows Server 2012 R2 and SQL Server 2014 Enterprise Edition and one Windows Azure Server. There are two tiers of storage available for use by database files only a fast tier and a slower tier. Currently the data and log files are stored on the fast tier of storage only. If a possible use case exists, management would like to utilize the slower tier storage for data files. The servers are configured as shown in the following table:

Location	Server
Company headquarters	HQ_Server
Satellite sales office	Satellite_Server
Microsoft Windows Azure (cloud)	Cloud_File Server

Database

Currently all information is stored in a single database called ProdDB, created with the following script:

```
CREATE DATABASE ProdDB
GO
ALTER DATABASE ProdDB SET RECOVERY SIMPLE
GO
```

The Product table is in the Production schema owned by the ProductionStaff Windows group. It is the main table in the system so access to information in the Product table should be as fast as possible. The columns in the Product table are defined as shown in the following table:

Column	Data type
ProductID	INT
ProductName	VARCHAR(100)
ProductDescription	VARCHAR(MAX)
ProductPrice	SMALLMONEY
QuantityOnHand	INT
ProductCost	SMALLMONEY
ProductSupplierID	INT

The SalesOrderDetail table holds the details about each sale. It is in the Sales schema owned by the SalesStaff Windows group. This table is constantly being updated, inserted into, and read. The columns in the SalesOrderDetail table are defined as shown in the following table:

Column	Data type
SalesOrderDetailID	INT
ProductID	INT
SalePrice	SMALLMONEY
SaleQuantity	INT

Database Issues

The current database does not perform well. Additionally, a recent disk problem caused the system to go down, resulting in lost sales revenue. In reviewing the current system, you found that there are no automated maintenance procedures. The database is severely fragmented, and everyone has read and write access. Requirements Database

The database should be configured to maximize uptime and to ensure that very little data is lost in the event of a server failure. To help with performance, the database needs to be modified so that it can support in-memory data, specifically for the Product table, which the CIO has indicated should be a memoryoptimized table. The auto-update statistics option is set off on this database. Only product managers are allowed to add products or to make changes to the name, description, price, cost, and supplier. The changes are made in an internal database and pushed to the Product table in ProdDB during system maintenance time. Product managers and others working at the headquarters location also should be able to generate reports that include supplier and cost information.

Customer data access

Customers access the company's website to order products, so they must be able to read product information such as name, description, and price from the Product table. When customers place orders, stored procedures called by the website update product quantity-on-hand values. This means the product table is constantly updated at random times.

Customer support data access

Customer support representatives need to be able to view and not update or change product information. Management does not want the customer support representatives to be able to see the product cost or any supplier information.

Sales force data access

Sales people at both the headquarters office and the satellite office must generate reports that read from the Product and SalesOrderDetail tables. No updates or inserts are ever made by sales people. These reports are run at random times and there can be no reporting downtime to refresh the data set except during the monthly maintenance window. The reports that run from the satellite office are process intensive queries with large data sets. Regardless of which office runs a sales force report, the SalesOrderDetail table should only return valid, committed order data; any orders not yet committed should be ignored.

Historical Data

The system should keep historical information about customers who access the site so that sales people can see how frequently customers log in and how long they stay on the site.

The information should be stored in a table called Customer Access. Supporting this requirement should have minimal impact on production website performance.

Backups

The recovery strategy for Fabrikam needs to include the ability to do point in time restores and minimize the risk of data loss by performing transaction log backups every 15 minutes.

Database Maintenance

The company has defined a maintenance window every month when the server can be unavailable. Any maintenance functions that require exclusive access should be accomplished during that window.

Project milestones completed

Revoked all existing read and write access to the database, leaving the schema ownership in place.

Configured an Azure storage container secured with the storage account name MyStorageAccount with the primary access key StorageAccountKey on the cloud file server.

SQL Server 2014 has been configured on the satellite server and is ready for use.

On each database server, the fast storage has been assigned to drive letter F:, and the slow storage has been assigned to drive letter D:.

You need to implement a backup strategy to support the requirements.

Which two actions should you perform? Each correct answer presents part of the solution. (Choose two.)

- A. Create a credential called MyCredential on SQL Server by using a Windows domain account and password.
- B. Schedule a full backup by using the command `BACKUP DATABASE ProdDB TO DISK...`
- C. Create a share on your Windows Azure site by using your Windows Azure storage account information, and grant permission to the SQL Server service login.
- D. Schedule a full backup by using the command `BACKUP DATABASE ProdDB TO URL ... WITH CREDENTIAL=N'MyCredential'`
- E. Create a share on the hot standby site and grant permission to the SQL Server service login.
- F. Create a credential called MyCredential on SQL Server, using MyStorageAccount for the storage account name and StorageAccountKey for the access key.
- G. Schedule a full backup by using the command `BACKUP DATABASE ProdDB TO SHARE ... WITH CREDENTIAL=N' MyCredential'`

Answer: CD

Explanation:

- Scenario: The current nightly backups have been failing due to insufficient space on the available drives and manual drive cleanup often needing to happen to get past the errors. Additional space will not be made available for backups on the HQ or satellite servers.

- Need to store files in the cloud.

- Manage your backups to Windows Azure: Using the same methods used to backup to DISK and TAPE, you can now back up to Windows Azure storage by Specifying URL as the backup destination.

You can use this feature to manually backup or configure your own backup strategy like you would for a local storage or other off-site options.

This feature is also referred to as SQL Server Backup to URL. SQL Server Managed Backup to Windows Azure

NEW QUESTION 129

- (Exam Topic 3)

You have a SQL Server instance on a server named Server1. You need to recommend a solution to perform the following tasks every week:

Rebuild the indexes by using a new fill factor.

Run a custom T-SQL command.

Back up the databases.

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

- A. A trigger
- B. An alert
- C. A maintenance plan
- D. Windows PowerShell
- E. A system policy

Answer: C

Explanation:

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 132

- (Exam Topic 3)

You need to address the requirements for disc usage monitoring for the SQL Servers. What should you do?

- A. You should configure disc quotas.
- B. You should configure a Dynamic Management View.
- C. You should configure alerts sent by the SQL Server Agent.
- D. You should configure a SQL Server Maintenance Plan.

Answer: B

NEW QUESTION 133

- (Exam Topic 3)

General Overview

You are the Senior Database Administrator (DBA) for a software development company named Leaffield Solutions. The company develops software applications custom designed to meet customer requirements.

Requirements Leaffield Solutions has been asked by a customer to develop a web-based Enterprise Resource Planning and Management application. The new application will eventually replace a desktop application that the customer is currently using. The current application will remain in use while the users are trained to use the new webbased application.

You need to design the SQL Server and database infrastructure for the web-based application. Databases

You plan to implement databases named Customers, Sales, Products, Current_Inventory, and TempReporting. The Sales database contains a table named OrderTotals and a table named SalesInfo.

A stored procedure named SPUpdateSalesInfo reads data in the OrderTotals table and modifies data in the SalesInfo table.

The stored procedure then reads data in the OrderTotals table a second time and makes further changes to the information in the SalesInfo table.

The Current_Inventory database contains a large table named Inv_Current. The Inv_Current table has a clustered index for the primary key and a nonclustered index. The primary key column uses the identity property.

The data in the Inv_Current table is over 120GB in size. The tables in the Current_Inventory database are accessed by multiple queries in the Sales database.

Another table in the Current_Inventory database contains a self-join with an unlimited number of hierarchies. This table is modified by a stored procedure named SPUpdate2.

An external application named ExternalApp1 will periodically query the Current_Inventory database to generate statistical information. The TempReporting database contains a single table named GenInfo.

A stored procedure named SPUPdateGenInfo combines data from multiple databases and generates millions of rows of data in the GenInfo table.

The GenInfo table is used for reports.

When the information in GenInfo is generated, a reporting process reads data from the Inv_Current table and queries information in the GenInfo table based on that data.

The GenInfo table is deleted after the reporting process completes. The Products database contains tables named ProductNames and ProductTypes.

Current System

The current desktop application uses data stored in a SQL Server 2005 database named DesABCOppAppDB. This database will remain online and data from the Current_Inventory database will be copied to it as soon as data is changed in the Current_Inventory database.

SQL Servers

A new SQL Server 2012 instance will be deployed to host the databases for the new system. The databases will be hosted on a Storage Area Network (SAN) that provides highly available storage.

Design Requirements

Your SQL Server infrastructure and database design must meet the following requirements:

Confidential information in the Current_Inventory database that is accessed by ExternalApp1 must be securely stored.

Direct access to database tables by developers or applications must be denied.

The account used to generate reports must have restrictions on the hours when it is allowed to make a connection.

Deadlocks must be analyzed with the use of Deadlock Graphs.

In the event of a SQL Server failure, the databases must remain available.

Software licensing and database storage costs must be minimized.

Development effort must be minimized.

The Tempdb databases must be monitored for insufficient free space.

Failed authentication requests must be logged.

Every time a new row is added to the ProductTypes table in the Products database, a user defined function that validates the row must be called before the row is added to the table.

When SPUpdateSalesInfo queries data in the OrderTotals table the first time, the same rows must be returned along with any newly added rows when

SPUpdateSalesInfo queries data in the OrderTotals table the second time.

You need to recommend a solution to meet the design requirement of logging all failed authentication requests.

What should you recommend?

- A. Object Access Auditing
- B. C2 Audit Mode
- C. Logon Triggers.
- D. Login Auditing.

Answer: D

NEW QUESTION 136

- (Exam Topic 3)

You work as a Database Administrator (DBA) for a company named ABC.com. The company uses a Microsoft SQL Server 2012 infrastructure.

You have a database named CorpDB.

You plan to create a stored procedure to access data in CorpDB.

You need to ensure that the stored procedure supports dirty reads. What should you do?

- A. You should configure the stored procedure to use the READ UNCOMMITTED isolation level.
- B. You should configure the stored procedure to use the READ COMMITTED isolation level.
- C. You should configure the stored procedure to use the READ SERIALIZABLE isolation level.
- D. You should configure the stored procedure to use the REPEATABLE READ isolation level.
- E. You should configure the stored procedure to use the SNAPSHOT isolation level.

Answer: A

NEW QUESTION 138

- (Exam Topic 3)

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

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

You have a server named Server1 that has Microsoft SQL Server installed. Server1 has SQL Server Audit configured to send audit even records to a file. You need to ensure that a database user named User1 can review the audit data. Solution: You grant the VIEW SERVER STATE permission to User1.

Does this meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

Each feature and command for SQL Server Audit has individual permission requirements.

Unless otherwise specified, viewing catalog views requires a principal to have one of the following:

The VIEW SERVER STATE permission.

The VIEW AUDIT STATE permission (gives only the principal access to the sys.server_audits catalog view).

Membership in the sysadmin fixed server role.

The CONTROL SERVER permission.

The ALTER ANY AUDIT permission.

A principal must have the VIEW SERVER STATE or ALTER ANY AUDIT permission to use the Dynamic Management Views.

References: [https://technet.microsoft.com/en-us/library/cc280665\(v=sql.105\).aspx](https://technet.microsoft.com/en-us/library/cc280665(v=sql.105).aspx)

NEW QUESTION 143

- (Exam Topic 3)

You work as a Database Administrator (DBA) for a company named ABC.com. The company uses a Microsoft SQL Server 2012 infrastructure.

You have a database named CorpDB. CorpDB contains a table named SalesInfo.

You discover that some table has been deleted from the SalesInfo table. You are unable to find out who deleted the information.

You need to implement a solution to monitor the deletion of any further information from the SalesInfo table. You want to minimize the development effort required for the solution.

What should you configure?

- A. You should configure table permissions.
- B. You should configure a user role.
- C. You should configure change data capture.
- D. You should configure a trigger.

Answer: C

NEW QUESTION 144

- (Exam Topic 3)

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

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

You have a data warehouse that stored sales data. One fact table has 100 million rows. You must reduce storage needs for the data warehouse.

You need to implement a solution that uses column-based storage and provides real-time analytics for the operational workload.

Solution: You remove all clustered indexes, sort the transactions in the table, and create a clustered index on the table, so that the table is not a heap.

Does the solution meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

Columnstore indexes are the standard for storing and querying large data warehousing fact tables. It uses column-based data storage and query processing to achieve up to 10x query performance gains in your data warehouse over traditional row-oriented storage, and up to 10x data compression over the uncompressed data size.

In SQL Server, rowstore refers to table where the underlying data storage format is a heap, a clustered index, or a memory-optimized table.

References: <https://docs.microsoft.com/en-us/sql/relational-databases/indexes/columnstore-indexes-overview>

NEW QUESTION 145

- (Exam Topic 3)

General Overview

You are the Senior Database Administrator (DBA) for a software development company named Leaffield Solutions. The company develops software applications custom designed to meet customer requirements.

Requirements Leaffield Solutions has been asked by a customer to develop a web-based Enterprise Resource Planning and Management application. The new application will eventually replace a desktop application that the customer is currently using. The current application will remain in use while the users are trained to use the new webbased application.

You need to design the SQL Server and database infrastructure for the web-based application. Databases

You plan to implement databases named Customers, Sales, Products, Current_Inventory, and TempReporting. The Sales database contains a table named OrderTotals and a table named SalesInfo.

A stored procedure named SPUpdateSalesInfo reads data in the OrderTotals table and modifies data in the SalesInfo table.

The stored procedure then reads data in the OrderTotals table a second time and makes further changes to the information in the SalesInfo table.

The Current_Inventory database contains a large table named Inv_Current. The Inv_Current table has a clustered index for the primary key and a nonclustered index. The primary key column uses the identity property.

The data in the Inv_Current table is over 120GB in size. The tables in the Current_Inventory database are accessed by multiple queries in the Sales database.

Another table in the Current_Inventory database contains a self-join with an unlimited number of hierarchies. This table is modified by a stored procedure named SPUpdate2.

An external application named ExternalApp1 will periodically query the Current_Inventory database to generate statistical information. The TempReporting database contains a single table named GenInfo.

A stored procedure named SPUPdateGenInfo combines data from multiple databases and generates millions of rows of data in the GenInfo table.

The GenInfo table is used for reports.

When the information in GenInfo is generated, a reporting process reads data from the Inv_Current table and queries information in the GenInfo table based on that data.

The GenInfo table is deleted after the reporting process completes. The Products database contains tables named ProductNames and ProductTypes.

Current System

The current desktop application uses data stored in a SQL Server 2005 database named DesABCOppAppDB. This database will remain online and data from the Current_Inventory database will be copied to it as soon as data is changed in the Current_Inventory database.

SQL Servers

A new SQL Server 2012 instance will be deployed to host the databases for the new system. The databases will be hosted on a Storage Area Network (SAN) that provides highly available storage.

Design Requirements

Your SQL Server infrastructure and database design must meet the following requirements:

Confidential information in the Current_Inventory database that is accessed by ExternalApp1 must be securely stored.

Direct access to database tables by developers or applications must be denied.

The account used to generate reports must have restrictions on the hours when it is allowed to make a connection.

Deadlocks must be analyzed with the use of Deadlock Graphs.

In the event of a SQL Server failure, the databases must remain available.

Software licensing and database storage costs must be minimized.

Development effort must be minimized.

The Tempdb databases must be monitored for insufficient free space.

Failed authentication requests must be logged.

Every time a new row is added to the ProductTypes table in the Products database, a user defined

function that validates the row must be called before the row is added to the table.

When SPUpdateSalesInfo queries data in the OrderTotals table the first time, the same rows must be returned along with any newly added rows when SPUpdateSalesInfo queries data in the OrderTotals table the second time.

The performance of the SPUpdate2 stored procedure needs to be improved. Your solution must meet the design requirements.

What should your solution include?

- A. A common table expression.
- B. A derived table.
- C. A Cursor.
- D. A table variable.

Answer: A

NEW QUESTION 147

- (Exam Topic 3)

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

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

You attempt to restore a database on a new SQL Server instance and receive the following error message: "Msg 33111, Level 16, State 3, Line 2

Cannot find server certificate with thumbprint '0x7315277C70764B1F252DC7A5101F6F66EFB1069D.'" You need to ensure that you can restore the database successfully.

Solution: You restore the certificate on the new instance. Does this meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

The error is related to the certificate.

References: <https://dba.stackexchange.com/questions/3388/restore-encrypted-database-to-another-server?rq=1>

NEW QUESTION 149

- (Exam Topic 3)

You have two databases named DB1 and DB2 that are located on the same server.

You plan to create a stored procedure named SProc1 in DB1. SProc1 will query a table named Table2 in DB2. You need to recommend a solution to ensure that SProc1 can access Table2 without granting users direct access to Table2.

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

- A. Contained databases
- B. Application roles
- C. Cross-database ownership chaining
- D. Digital certificates

Answer: B

Explanation:

An application role is a database principal that enables an application to run with its own, userlike permissions. You can use application roles to enable access to specific data to only those users who connect through a particular application. Unlike database roles, application roles contain no members and are inactive by default.

NEW QUESTION 151

- (Exam Topic 3)

You administer a SQL Server 2014 instance.

Users report that the SQL Server has seemed slow today.

A large database was being restored for much of the day, which could be causing issues. You want to write a query of the system views that will report the following:

Number of users that have a connection to the server

Whether a user's connection is active

Whether any connections are blocked

What queries are being executed

Whether the database restore is still executing and, if it is, what percentage of the restore is complete. Which system objects should you use in your query to best achieve this task?

- A. sys.dm_exec_requests, sys.dm_exec_sessions, sys.objects
- B. sys.dm_exec_sessions, sys.dm_exec_query_stats, sys.dm_exec_query_text, sys.objects
- C. sys.sysprocesses, sys.dm_exec_query_text, sys.objects
- D. sys.dm_exec_requests, sys.dm_exec_sessions, sys.dm_exec_query_text

Answer: D

Explanation:

- sys.dm_exec_requests

Returns information about each request that is executing within SQL Server.

- sys.dm_exec_sessions

Returns one row per authenticated session on SQL Server. sys.dm_exec_sessions is a serverscope view that shows information about all active user connections and internal tasks. This information includes client version, client program name, client login time, login user, current session setting, and more.

- sys.dm_exec_query_text

Returns the text of the SQL batch that is identified by the specified sql_handle.

NEW QUESTION 155

- (Exam Topic 3)

You use SQL Server 2014 to maintain the data used by applications at your company. You want to execute two statements.

You need to guarantee that either both statements succeed, or both statements fail together as a batch. Which code should you use?

Option A.

```
BEGIN TRY
INSERT TABLE1 (FIELD1) VALUES ('ONE')
INSERT TABLE2 (FIELD1) VALUES ('TWO')
END TRY
BEGIN CATCH
ROLLBACK TRANSACTION
THROW
END CATCH
```

Option B.

```
BEGIN TRY
INSERT TABLE1 (FIELD1) VALUES ('ONE')
INSERT TABLE2 (FIELD1) VALUES ('TWO')
END TRY
BEGIN CATCH
THROW
ROLLBACK TRANSACTION
END CATCH
```

Option C.

```
BEGIN TRANSACTION
INSERT TABLE1 (FIELD1) VALUES ('ONE')
INSERT TABLE2 (FIELD1) VALUES ('TWO')
IF @@ERROR = 0
COMMIT TRANSACTION
ELSE
ROLLBACK TRANSACTION
```

Option D.

```
BEGIN TRY
BEGIN TRANSACTION
INSERT TABLE1 (FIELD1) VALUES ('ONE')
INSERT TABLE2 (FIELD1) VALUES ('TWO')
COMMIT TRANSACTION
END TRY
BEGIN CATCH
    ROLLBACK TRANSACTION
END CATCH
```

Option E.

```
BEGIN TRY
INSERT TABLE1 (FIELD1) VALUES ('ONE')
INSERT TABLE2 (FIELD1) VALUES ('TWO')
END TRY
BEGIN CATCH
THROW
END CATCH
```

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

Answer: D

Explanation:

xplanation: Structure should be:

```
BEGIN TRY
BEGIN TRANSACTION
..
COMMIT TRANSACTION
END TRY
BEGIN CATCH
ROLLBACK TRANSACTION
END CATCH.
```

NEW QUESTION 159

- (Exam Topic 3)

You work as a Database Administrator (DBA) for a company named ABC.com.

The company uses a Microsoft SQL Server 2012 infrastructure. Computer users use custom applications that store data in SQL Server databases hosted on a server named ABC-SQL1.

You discover that the Windows Application event log on ABC-SQL1 contains error events.

You need to configure a solution that will run an application whenever an event with a specific event ID number is generated in the Windows Application event log. What should you configure?

- A. You should configure an alert and a job.
- B. You should configure an alert and a maintenance plan.
- C. You should configure a trigger and a maintenance plan.
- D. You should configure a trigger and a job.
- E. You should configure an alert and a trigger.

Answer: A

NEW QUESTION 161

- (Exam Topic 3)

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. A service account has the required permissions to backup all databases.

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 service account that backs up the HRDB database the permission necessary to access the encryption key. Your solution must use the principle of least privilege.

Which permission should you grant?

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

Answer: G

Explanation:

The 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 162

- (Exam Topic 3)

You have a SQL Server 2014 environment That includes four servers. The servers are configured as shown in the following table.

Server name	SQL Server 2014 edition	SQL Server version	Details
Server1	Enterprise	SQL Server 2014	A production Online Transaction Processing (OLTP) server
Server2	Web	SQL Server 2014	A test server
Server3	Standard	SQL Server 2012	A production report server
Server4	Express	SQL Server 2008 R2	A witness server

You plan to configure Policy-Based Management to enforce the following rules:

On Server1, enable SQL Server password policies and enable the default trace.

On Server3, ensure that the names of user-defined stored procedures begin with the prefix "usp_" and ensure that all databases use a casesensitive collation.

You need to recommend which server you must configure as a Central Management Server.

Which server should you recommend? (Each correct answer presents a complete solution. Choose all that apply.)

- A. Server1
- B. Server2

C. Server3
D. Seiver4

Answer: AC

Explanation:
Need Standard or Enterprise edition of SQL Server.

NEW QUESTION 163

- (Exam Topic 3)

Overview

General Overview

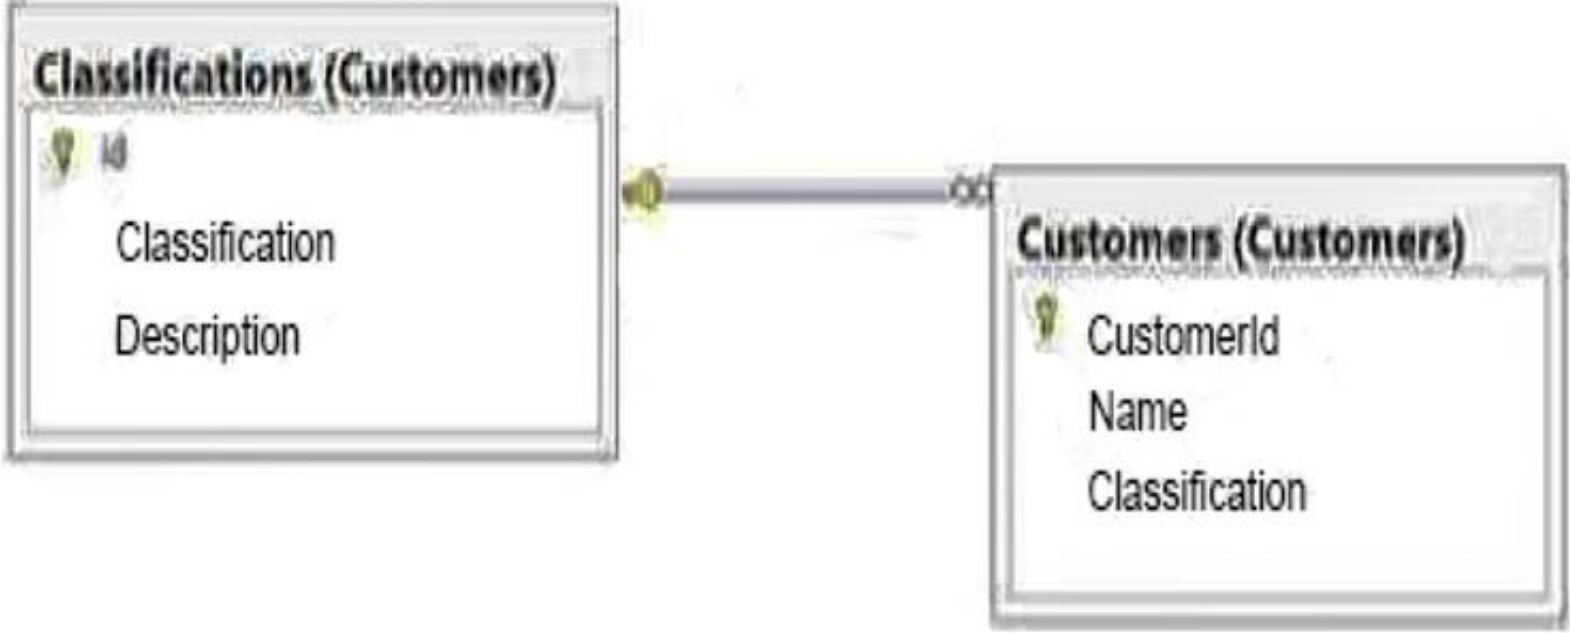
ADatum Corporation has offices in Miami and Montreal.

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

Databases

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

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



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

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

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

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

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

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

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

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

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

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

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

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

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

The nested stored procedures are never called directly. Design Requirements

Data Recovery

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

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

Classification Changes

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

Storage

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

Error Handling

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

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

You need to recommend a solution to minimize the amount of time it takes to execute USP_1. With what should you recommend replacing Table1?

A. An indexed view
B. A function

- C. A table variable
- D. A temporary table

Answer: D

Explanation:

Scenario:

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 a table in the Products database and searches for information in Table1 based on input from the Products table. After the process is complete, Table1 is deleted.

NEW QUESTION 164

- (Exam Topic 3)

General Overview

You are the Senior Database Administrator (DBA) for a software development company named Leaffield Solutions. The company develops software applications custom designed to meet customer requirements.

Requirements Leaffield Solutions has been asked by a customer to develop a web-based Enterprise Resource Planning and Management application. The new application will eventually replace a desktop application that the customer is currently using. The current application will remain in use while the users are trained to use the new webbased application.

You need to design the SQL Server and database infrastructure for the web-based application. Databases

You plan to implement databases named Customers, Sales, Products, Current_Inventory, and TempReporting. The Sales database contains a table named OrderTotals and a table named SalesInfo.

A stored procedure named SPUUpdateSalesInfo reads data in the OrderTotals table and modifies data in the SalesInfo table.

The stored procedure then reads data in the OrderTotals table a second time and makes further changes to the information in the SalesInfo table.

The Current_Inventory database contains a large table named Inv_Current. The Inv_Current table has a clustered index for the primary key and a nonclustered index. The primary key column uses the identity property.

The data in the Inv_Current table is over 120GB in size. The tables in the Current_Inventory database are accessed by multiple queries in the Sales database.

Another table in the Current_Inventory database contains a self-join with an unlimited number of hierarchies. This table is modified by a stored procedure named SPUUpdate2.

An external application named ExternalApp1 will periodically query the Current_Inventory database to generate statistical information. The TempReporting database contains a single table named GenInfo.

A stored procedure named SPUPdateGenInfo combines data from multiple databases and generates millions of rows of data in the GenInfo table.

The GenInfo table is used for reports.

When the information in GenInfo is generated, a reporting process reads data from the Inv_Current table and queries information in the GenInfo table based on that data.

The GenInfo table is deleted after the reporting process completes. The Products database contains tables named ProductNames and ProductTypes.

Current System

The current desktop application uses data stored in a SQL Server 2005 database named DesABCOppAppDB. This database will remain online and data from the Current_Inventory database will be copied to it as soon as data is changed in the Current_Inventory database.

SQL Servers

A new SQL Server 2012 instance will be deployed to host the databases for the new system. The databases will be hosted on a Storage Area Network (SAN) that provides highly available storage.

Design Requirements

Your SQL Server infrastructure and database design must meet the following requirements:

Confidential information in the Current_Inventory database that is accessed by ExternalApp1 must be securely stored.

Direct access to database tables by developers or applications must be denied.

The account used to generate reports must have restrictions on the hours when it is allowed to make a connection.

Deadlocks must be analyzed with the use of Deadlock Graphs.

In the event of a SQL Server failure, the databases must remain available.

Software licensing and database storage costs must be minimized.

Development effort must be minimized.

The Tempdb databases must be monitored for insufficient free space.

Failed authentication requests must be logged.

Every time a new row is added to the ProductTypes table in the Products database, a user defined function that validates the row must be called before the row is added to the table.

When SPUUpdateSalesInfo queries data in the OrderTotals table the first time, the same rows must be returned along with any newly added rows when SPUUpdateSalesInfo queries data in the OrderTotals table the second time.

You need to ensure that the account used to generate reports can only connect during certain hours. What should you configure?

- A. A CHECK constraint.
- B. Windows Server Resource Manager (WSRM).
- C. Logon Triggers.
- D. Login Auditing.

Answer: C

NEW QUESTION 165

- (Exam Topic 3)

You use SQL Server 2014.

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

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

Code Blocks

```
END
```

```
CREATE FUNCTION dbo.Func_Table1  
@InsertWord Varchar (10), @Var1  
int OUTPUT AS BEGIN
```

```
CREATE FUNCTION dbo.Func_Table1  
(@InsertWord Varchar (10)),  
Returns INT  
AS BEGIN
```

```
CREATE PROCEDURE dbo.Spr_Table1  
@InsertWord Varchar (10), @Var1  
INT  
AS BEGIN
```

```
Select @Var1 = count(*) from  
TABLE1
```

```
Declare @Var1 INT
```

```
CREATE PROCEDURE dbo.Spr_Table1  
@InsertWord Varchar (10), @Var1  
int OUTPUT AS BEGIN
```

```
INSERT INTO TABLE1 (FIELD1)  
values (@InsertWord)
```

Answer Area

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

NEW QUESTION 166

- (Exam Topic 3)

You have a database hosted on SQL Server 2012 R2. The database contains 5 million rows.

You need to recommend a repeatable method to migrate the database to SQL Azure.

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

- A. Create a SQL Server Integration Services (SSIS) package, and then run the package.
- B. Back up the database, and then restore the database.
- C. Extract a data-tier application, and then import the application.
- D. Generate scripts to create all of the all database objects and all of the data, and then execute the scripts by using SQL Azure.

Answer: A

Explanation:

SQL Server Integration Services Most flexibility

Data Transfer Efficiency: Good

- SSIS can be used to perform a broad range of data migration tasks. SSIS provides support for complex workflow and data transformation between the source and destination. It is a good choice to transfer of data for databases that require many changes to work on Microsoft Azure SQL Database.

You can use SSIS data transfer packages with another mechanism for transferring the database schema, such as a Data-tier Application package.

NEW QUESTION 168

- (Exam Topic 3)

Background Corporate Information

Fabrikam, Inc. is a retailer that sells electronics products on the Internet. The company has a headquarters site and one satellite sales office. You have been hired as the database administrator, and the company wants you to change the architecture of the Fabrikam ecommerce site to optimize performance and reduce downtime while keeping capital expenditures to a minimum. To help with the solution, Fabrikam has decided to use cloud resources as well as on-premise servers.

Physical Locations

All of the corporate executives, product managers, and support staff are stationed at the headquarters office. Half of the sales force works at this location. There is also a satellite sales office. The other half of the sales force works at the satellite office in order to have sales people closer to clients in that area. Only sales people work at the satellite location.

Problem Statement

To be successful, Fabrikam needs a website that is fast and has a high degree of system uptime. The current system operates on a single server and the company is not happy with the single point of failure this presents. The current nightly backups have been failing due to insufficient space on the available drives and manual drive cleanup often needing to happen to get past the errors. Additional space will not be made available for backups on the HQ or satellite servers. During your investigation, you discover that the sales force reports are causing significant contention.

Configuration Windows Logins

The network administrators have set up Windows groups to make it easier to manage security. Users may belong to more than one group depending on their role. The groups have been set up as shown in the following table:

Group	Members
OurDomain\Management	All corporate executives
OurDomain\SalesStaff	All sales people
OurDomain\ProductionStaff	All product managers and support staff
OurDomain\AllUsers	Everyone
OurDomain\CustomerSupport	Customer support representatives

Server Configuration The IT department has configured two physical servers with Microsoft Windows Server

2012 R2 and SQL Server 2014 Enterprise Edition and one Windows Azure Server. There are two tiers of storage available for use by database files only a fast tier and a slower tier. Currently the data and log files are stored on the fast tier of storage only. If a possible use case exists, management would like to utilize the slower tier storage for data files. The servers are configured as shown in the following table:

Location	Server
Company headquarters	HQ_Server
Satellite sales office	Satellite_Server
Microsoft Windows Azure (cloud)	Cloud_File Server

Database

Currently all information is stored in a single database called ProdDB, created with the following script:

```
CREATE DATABASE ProdDB
```

```
GO
```

```
ALTER DATABASE ProdDB SET RECOVERY SIMPLE
```

```
GO
```

The Product table is in the Production schema owned by the ProductionStaff Windows group. It is the main table in the system so access to information in the Product table should be as fast as possible. The columns in the Product table are defined as shown in the following table:

Column	Data type
ProductID	INT
ProductName	VARCHAR(100)
ProductDescription	VARCHAR(MAX)
ProductPrice	SMALLMONEY
QuantityOnHand	INT
ProductCost	SMALLMONEY
ProductSupplierID	INT

The SalesOrderDetail table holds the details about each sale. It is in the Sales schema owned by the SalesStaff Windows group. This table is constantly being updated, inserted into, and read. The columns in the SalesOrderDetail table are defined as shown in the following table:

Column	Data type
SalesOrderDetailID	INT
ProductID	INT
SalePrice	SMALLMONEY
SaleQuantity	INT

Database Issues

The current database does not perform well. Additionally, a recent disk problem caused the system to go down, resulting in lost sales revenue. In reviewing the current system, you found that there are no automated maintenance procedures. The database is severely fragmented, and everyone has read and write access.

Requirements Database

The database should be configured to maximize uptime and to ensure that very little data is lost in the event of a server failure. To help with performance, the database needs to be modified so that it can support in-memory data, specifically for the Product table, which the CIO has indicated should be a memoryoptimized table. The auto-update statistics option is set off on this database. Only product managers are allowed to add products or to make changes to the name,

description, price, cost, and supplier. The changes are made in an internal database and pushed to the Product table in ProdDB during system maintenance time. Product managers and others working at the headquarters location also should be able to generate reports that include supplier and cost information.

Customer data access

Customers access the company's website to order products, so they must be able to read product information such as name, description, and price from the Product table. When customers place orders, stored procedures called by the website update product quantity-on-hand values. This means the product table is constantly updated at random times.

Customer support data access

Customer support representatives need to be able to view and not update or change product information. Management does not want the customer support representatives to be able to see the product cost or any supplier information.

Sales force data access

Sales people at both the headquarters office and the satellite office must generate reports that read from the Product and SalesOrderDetail tables. No updates or inserts are ever made by sales people. These reports are run at random times and there can be no reporting downtime to refresh the data set except during the monthly maintenance window. The reports that run from the satellite office are process intensive queries with large data sets. Regardless of which office runs a sales force report, the SalesOrderDetail table should only return valid, committed order data; any orders not yet committed should be ignored.

Historical Data

The system should keep historical information about customers who access the site so that sales people can see how frequently customers log in and how long they stay on the site.

The information should be stored in a table called Customer Access. Supporting this requirement should have minimal impact on production website performance.

Backups

The recovery strategy for Fabrikam needs to include the ability to do point in time restores and minimize the risk of data loss by performing transaction log backups every 15 minutes.

Database Maintenance

The company has defined a maintenance window every month when the server can be unavailable. Any maintenance functions that require exclusive access should be accomplished during that window.

Project milestones completed

Revoked all existing read and write access to the database, leaving the schema ownership in place.

Configured an Azure storage container secured with the storage account name MyStorageAccount with the primary access key StorageAccountKey on the cloud file server.

SQL Server 2014 has been configured on the satellite server and is ready for use.

On each database server, the fast storage has been assigned to drive letter F:, and the slow storage has been assigned to drive letter D:.

You need to implement changes to the system to reduce contention and improve performance of the SalesOrderDetail table.

Which three actions should you perform? Each correct answer presents part of the solution. (Choose three.)

- A. Use (SNAPSHOT) hints in the report queries
- B. ALTER DATABASE [ProdDB] SET READ_COMMITTED_SNAPSHOT ON
- C. ALTER DATABASE [ProdDB] SET READ_COMMITTED_SNAPSHOT OFF
- D. SET TRANSACTION ISOLATION LEVEL SNAPSHOT
- E. Use (TABLOCK) hints in the report queries
- F. SET TRANSACTION ISOLATION LEVEL SERIALIZABLE
- G. ALTER DATABASE [ProdDB] SET ALLOW_SNAPSHOT_ISOLATION ON
- H. Use (SNAPSHOT) hints in the update statements

Answer: ABF

Explanation:

- Scenario: The SalesOrderDetail table holds the details about each sale. It is in the Sales schema owned by the SalesStaff Windows group. This table is constantly being updated, inserted into, and read.
 - Regardless of which office runs a sales force report, the SalesOrderDetail table should only return valid, committed order data; any orders not yet committed should be ignored.
 - READ_COMMITTED_SNAPSHOT { ON | OFF } ON Enables Read-Committed Snapshot option at the database level. When it is enabled, DML statements start generating row versions even when no transaction uses Snapshot Isolation. Once this option is enabled, the transactions specifying the read committed isolation level use row versioning instead of locking.
- When a transaction runs at the read committed isolation level, all statements see a snapshot of data as it exists at the start of the statement. OFF Turns off Read-Committed Snapshot option at the database level. Transactions specifying the READ COMMITTED isolation level use locking. ALTER DATABASE SET Options (Transact-SQL) SET Statements (Transact-SQL)

NEW QUESTION 170

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

70-764 Practice Exam Features:

- * 70-764 Questions and Answers Updated Frequently
- * 70-764 Practice Questions Verified by Expert Senior Certified Staff
- * 70-764 Most Realistic Questions that Guarantee you a Pass on Your First Try
- * 70-764 Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The 70-764 Practice Test Here](#)