

70-764 Dumps

Administering a SQL Database Infrastructure (beta)

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



NEW QUESTION 1

- (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 have an on-premises server that runs Microsoft SQL Server 2016 Standard Edition. You need to identify missing indexes.

What should you use?

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

Answer: D

Explanation:

Data Collector can gather performance information from multiple SQL Server instances and store it in a single repository. It has three built-in data collecting specifications (data collectors) designed to collect the most important performance metrics. The information collected by default is about disk usage, query statistics, and server activity.

The Query Statistics data collection set collects information about query statistics, activity, execution plans and text on the SQL Server instance.

Missing indexes can be found with the execution plans.

References: <https://www.sqlshack.com/sql-server-performance-monitoring-data-collector/>

NEW QUESTION 2

- (Exam Topic 1)

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

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

What should you do?

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

Answer: C

Explanation:

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

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

NEW QUESTION 3

- (Exam Topic 1)

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

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

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

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

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

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

Answer: D

NEW QUESTION 4

- (Exam Topic 1)

You plan to install Microsoft SQL Server 2016 for a web hosting company.

The company plans to host multiple web sites, each supported by a SQL Server database.

You need to select an edition of SQL Server that features backup compression of databases, basic data integration features, and low total cost of ownership.

Which edition should you choose?

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

Answer: B

Explanation:

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

NEW QUESTION 5

- (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 6

- (Exam Topic 1)

You administer a Microsoft SQL Server 2016 default instance. The instance is hosted by a server that has a local firewall configured.

The firewall only allows inbound connections on port 1433. The server only hosts a single instance of SQL Server.

You need to ensure that the instance is configured to allow remote connections even if the SQL Server is unresponsive to client connections.

What should you do?

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

Answer: ABC

Explanation:

SQL Server provides a dedicated administrator connection (DAC). The DAC lets an administrator access a running server to execute diagnostic functions or Transact-SQL statements, or to troubleshoot problems on the server, even when the server is locked or running in an abnormal state and not responding to a SQL Server Database Engine connection. By default, the DAC is only available from a client on the server. To enable client applications on remote computers to use the DAC, use the remote admin connections option of `sp_configure`.

By default, the DAC only listens on the loop-back IP address (127.0.0.1), port 1434. The following example enables the DAC from a remote computer.

```
sp_configure 'remote admin connections', 1; GO
```

```
RECONFIGURE; GO
```

References:

<https://docs.microsoft.com/en-us/sql/database-engine/configure-windows/remote-admin-connections-server-con>

NEW QUESTION 7

- (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 8

- (Exam Topic 1)

Note: This question is part of a series of questions that use the same scenario. For your convenience, the

scenario is repeated in each question. Each question presents a different goal and answer choices, but the text of the scenario is exactly the same in each question in this series.

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

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

You need to configure monitoring for Tailspin Toys.
In the table below, identify the monitoring tool that you must use for each activity.
NOTE: Make only one selection in each column.

Answer Area

Monitoring option	Monitoring from application	Trend analysis
Error logs	<input type="radio"/>	<input type="radio"/>
Transact-SQL	<input type="radio"/>	<input type="radio"/>
System Monitor	<input type="radio"/>	<input type="radio"/>
Distributed Replay	<input type="radio"/>	<input type="radio"/>

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Monitoring from application: Transact-SQL

Transact-SQL can be used to monitor a customized application. Trend analysis: System Monitor

System Monitor can provide trend analysis. From question:

Tailspin Toys has a custom application that accesses a hosted database named TSpinDB. The application will monitor TSpinDB and capture information over time about which database objects are accessed and how frequently they are accessed.

Tailspin Toys has a custom application that accesses a hosted database named TSpinDB. The application will monitor TSpinDB and capture information over time about which database objects are accessed and how frequently they are accessed.

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/performance/performance-monitoring-and-tuning-tools>

NEW QUESTION 9

- (Exam Topic 1)

You administer a Microsoft SQL Server 2016 database named Contoso on a server named Server01.

You need to track all SELECT statements issued in the Contoso database only by users in a role named Sales. What should you create?

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

Answer: F

NEW QUESTION 10

- (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 present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

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

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

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

- A. Yes
- B. No

Answer: B

Explanation:

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

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

References:

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

NEW QUESTION 12

- (Exam Topic 1)

You administer a Microsoft SQL Server 2016 database.

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

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

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

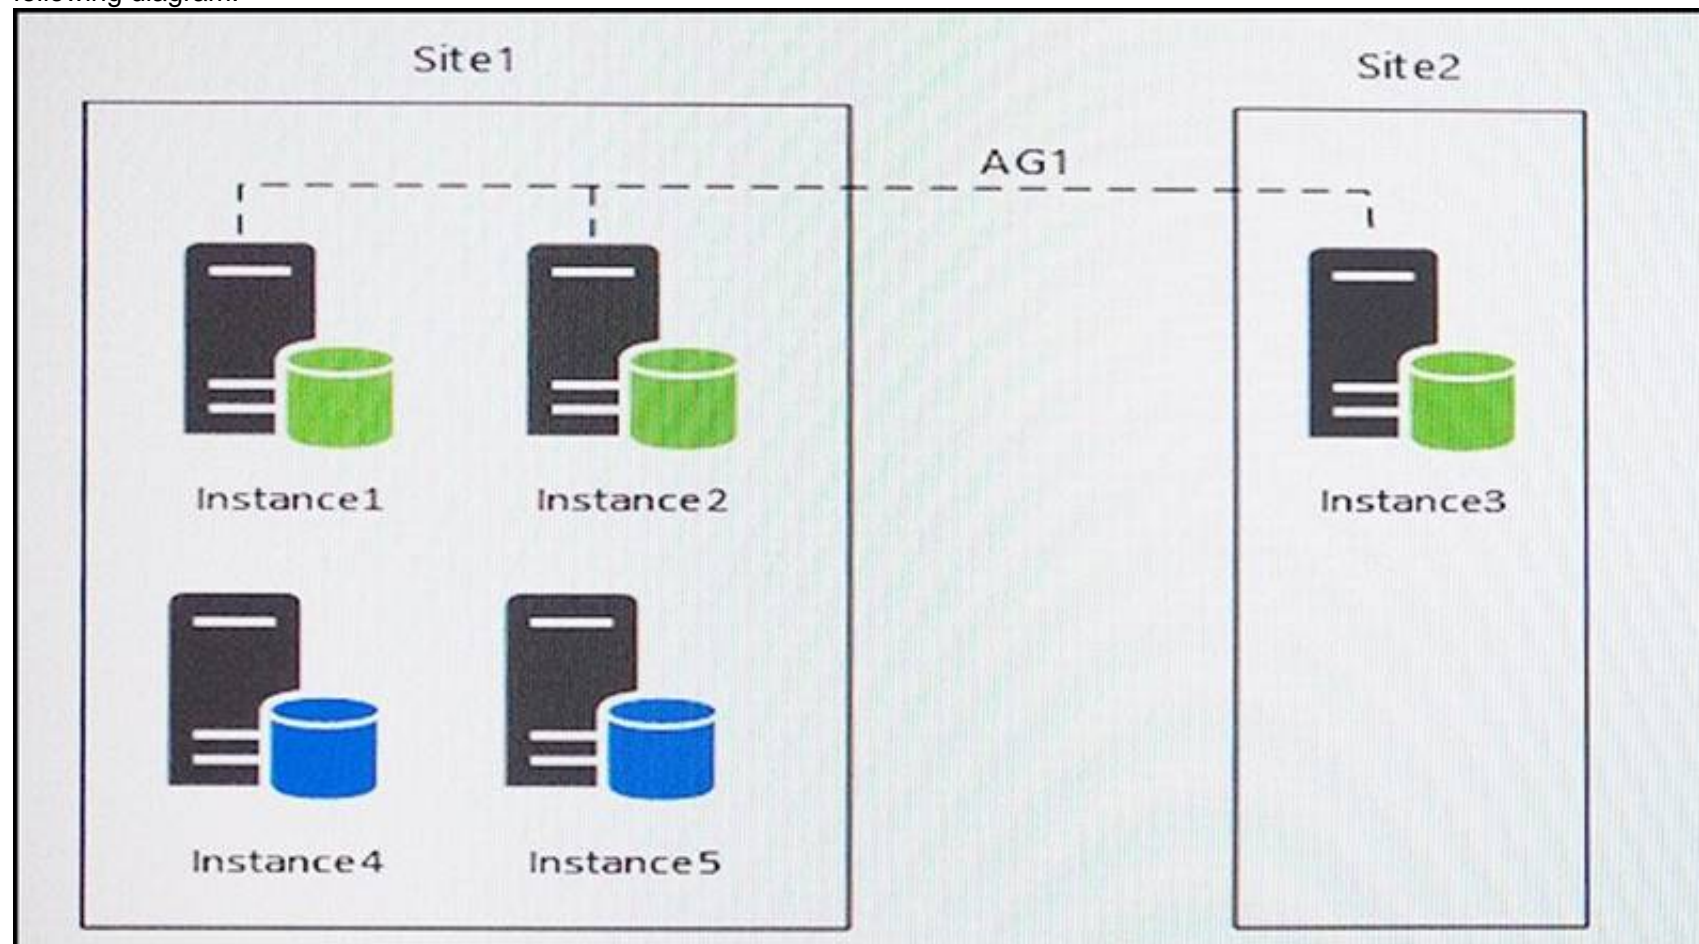
Answer: B

NEW QUESTION 17

- (Exam Topic 1)

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

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



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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Transact-SQL segments

DATABASE

REPLICA

SYNCHRONOUS_COMMIT

ASYNCHRONOUS_COMMIT

PRIMARY

MANUAL

AUTOMATIC

SECONDARY_ONLY

● ● ● ●

Answer Area

```
ALTER AVAILABILITY GROUP AG_1 MODIFY Transact-SQL segment ON 'INSTANCE6'

WITH (AVAILABILITY_MODE = Transact-SQL segment );

ALTER AVAILABILITY GROUP AG_1 MODIFY Transact-SQL segment ON 'INSTANCE6'

WITH (FAILOVER_MODE = Transact-SQL segment );
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

Box 1: REPLICA

MODIFY REPLICA ON modifies any of the replicas of the availability group. Box 2: SYNCHRONOUS_COMMIT

You must minimize latency between the nodes in AG1

AVAILABILITY_MODE = { SYNCHRONOUS_COMMIT | ASYNCHRONOUS_COMMIT }

Specifies whether the primary replica has to wait for the secondary availability group to acknowledge the hardening (writing) of the log records to disk before the primary replica can commit the transaction on a given primary database.

FAILOVER AUTOMATIC (box 4) requires SYNCHRONOUS_COMMIT Box 3: REPLICA

MODIFY REPLICA ON modifies any of the replicas of the availability group. Box 4: AUTOMATIC

You must minimize latency between the nodes in AG1 FAILOVER_MODE = { AUTOMATIC | MANUAL }

Specifies the failover mode of the availability replica that you are defining.

FAILOVER_MODE is required in the ADD REPLICA ON clause and optional in the MODIFY REPLICA ON clause.

AUTOMATIC enables automatic failover. AUTOMATIC is supported only if you also specify

AVAILABILITY_MODE = SYNCHRONOUS_COMMIT.

References: <https://docs.microsoft.com/en-us/sql/t-sql/statements/alter-availability-group-transact-sql>

NEW QUESTION 21

- (Exam Topic 1)

You manage a Microsoft SQL Server environment. You plan to encrypt data when you create backups. You need to configure the encryption options for backups. What should you configure?

- A. a certificate
- B. an MD5 hash

- C. a DES key
- D. an AES 256-bit key

Answer: D

Explanation:

To encrypt during backup, you must specify an encryption algorithm, and an encryptor to secure the encryption key. The following are the supported encryption options:

Encryption Algorithm: The supported encryption algorithms are: AES 128, AES 192, AES 256, and Triple DES


Encryptor: A certificate or asymmetric Key

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

NEW QUESTION 25

- (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 SAN Storage for new SQL Cluster
Subject SAN Storage for new SQL Cluster
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

...
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.

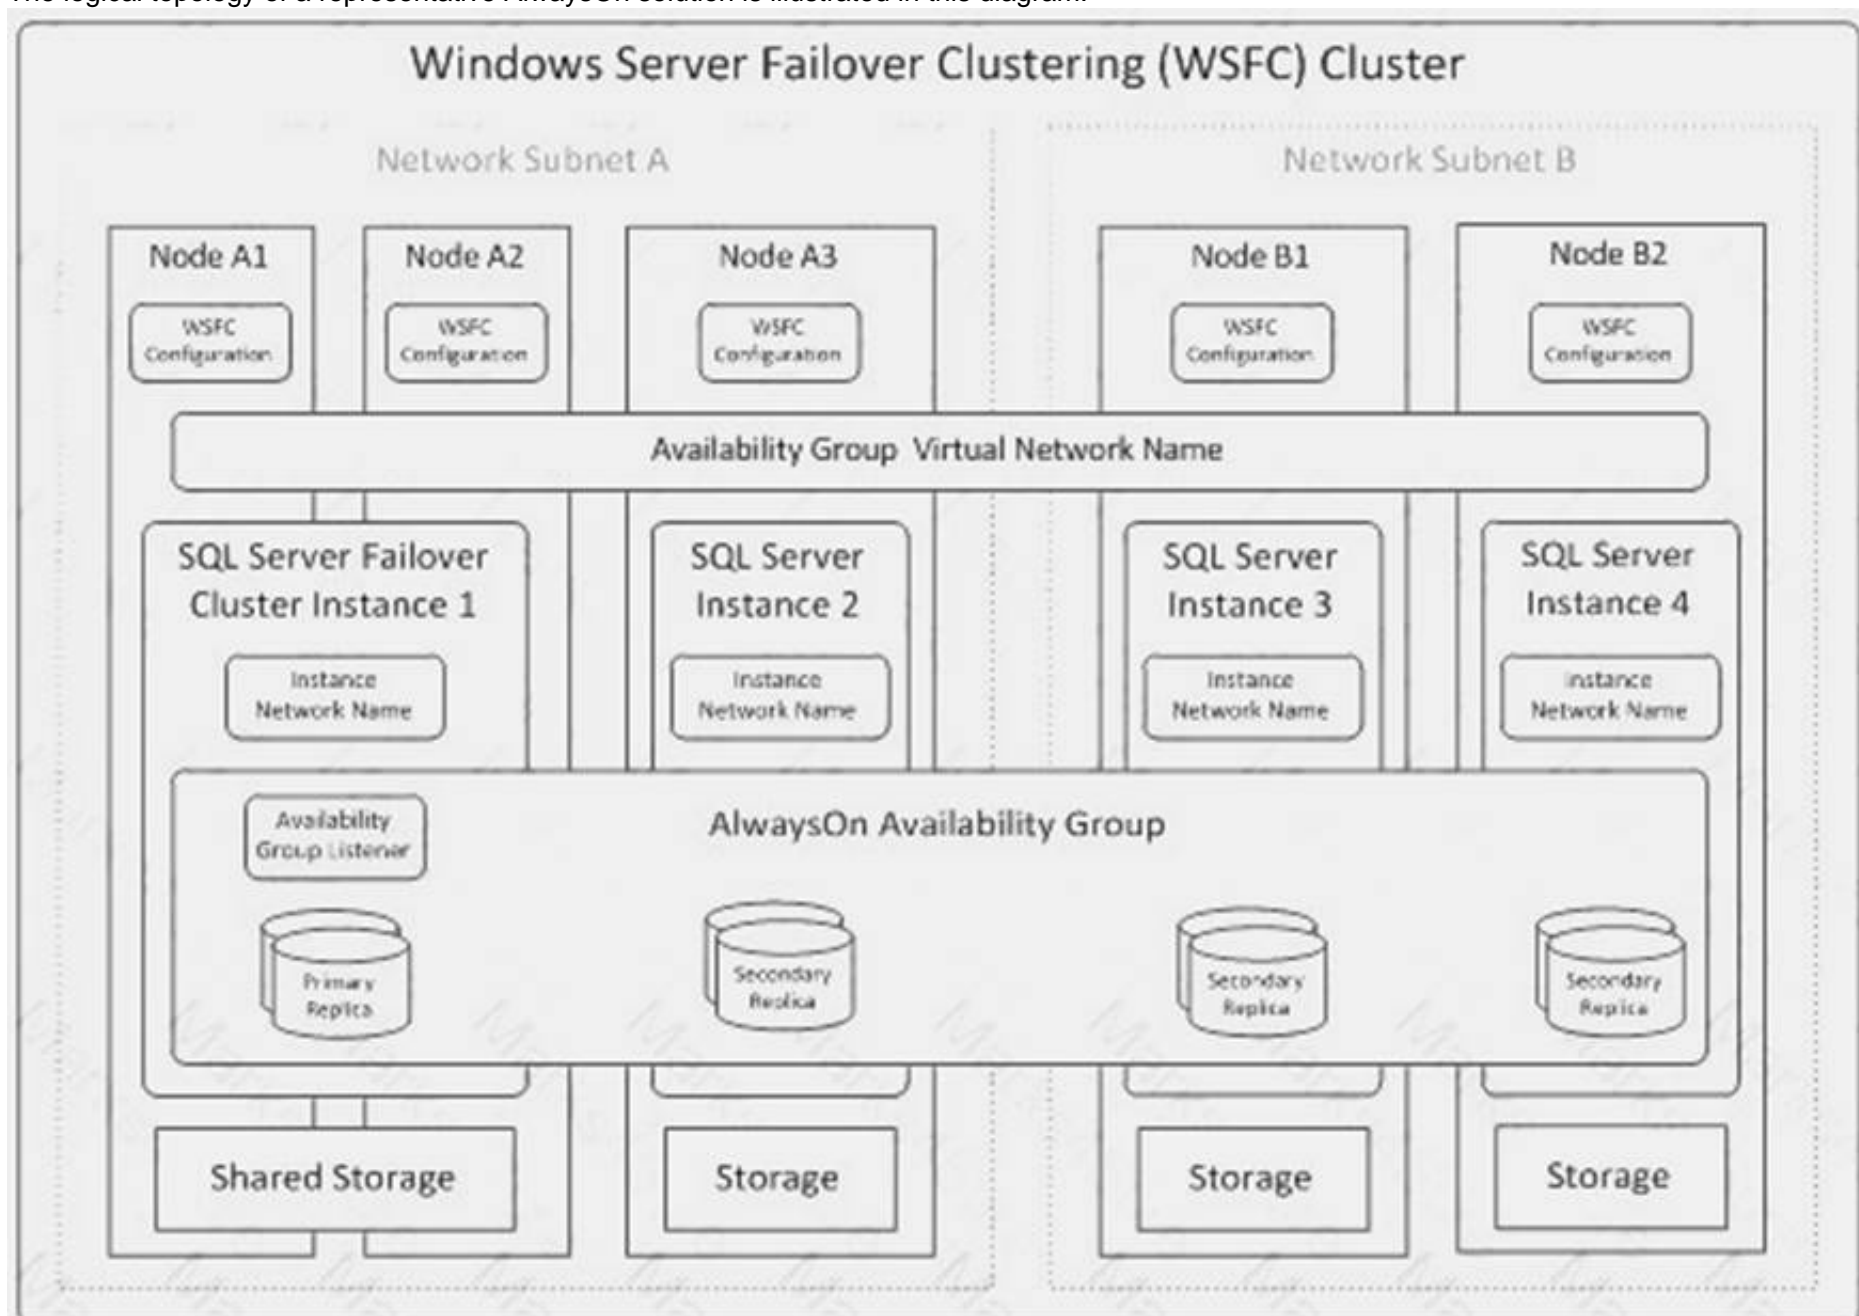
Prior to SQL Server 2012, FCIs required tempdb to be located on a symmetrical shared storage volume that failed over with other system databases.

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 28

- (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 Asynchronous-Commit 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 29

- (Exam Topic 1)

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

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

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

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

Answer Area

Event type	View deletions	Update operations
Data changes	<input type="radio"/>	<input type="radio"/>
Schema changes	<input type="radio"/>	<input type="radio"/>
SQL batch	<input type="radio"/>	<input type="radio"/>
Data access	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

Event type	View deletions	Update operations
Data changes	<input type="radio"/>	<input checked="" type="radio"/>
Schema changes	<input checked="" type="radio"/>	<input type="radio"/>
SQL batch	<input type="radio"/>	<input type="radio"/>
Data access	<input type="radio"/>	<input type="radio"/>

NEW QUESTION 31

- (Exam Topic 1)

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

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

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

You need to configure auditing for the Adventure Works environment.

How should you complete the Transact-SQL statement? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Answer Area

USE master
GO

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

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

GO

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

GO

Use Adventureworks

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

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

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: CREATE SERVER AUDIT

Create the server audit.

You must implement auditing to record access to data that is considered sensitive by the company. Create database audit

Box 2: ALTER SERVER AUDIT

Enable the server audit.

Box 3: CREATE DATABASE AUDIT

Create the database audit specification. Box 4: FOR SERVER AUDIT

You must implement auditing for all objects in the ADVSchema.

References:

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

NEW QUESTION 32

- (Exam Topic 1)

You administer all the deployments of Microsoft SQL Server 2016 in your company. A database contains a large product catalog that is updated periodically.

You need to be able to send the entire product catalog to all branch offices on a monthly basis. Which configuration should you use?

- A. 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

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

Answer: E

Explanation:

Snapshot replication distributes data exactly as it appears at a specific moment in time and does not monitor for updates to the data. When synchronization occurs, the entire snapshot is generated and sent to Subscribers.

Using snapshot replication by itself is most appropriate when one or more of the following is true:

- Data changes infrequently.
- It is acceptable to have copies of data that are out of date with respect to the Publisher for a period of time.
- Replicating small volumes of data.
- A large volume of changes occurs over a short period of time.

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

NEW QUESTION 36

- (Exam Topic 1)

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

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

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

Answer Area

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Note: Before you configure log shipping, you must create a share to make the transaction log backups available to the secondary server.

SQL Server Log shipping allows you to automatically send transaction log backups from a primary database on a primary server instance to one or more secondary databases on separate secondary server instances. The transaction log backups are applied to each of the secondary databases individually. An optional third server instance, known as the monitor server, records the history and status of backup and restore operations and, optionally, raises alerts if these operations fail to occur as scheduled.

Box 1: Primary server instance.

The primary server instance runs the backup job to back up the transaction log on the primary database. backup job: A SQL Server Agent job that performs the backup operation, logs history to the local server and the monitor server, and deletes old backup files and history information. When log shipping is enabled, the job category "Log Shipping Backup" is created on the primary server instance.

Box 2: Secondary server instance

Each of the three secondary server instances runs its own copy job to copy the primary log-backup file to its own local destination folder.
copy job: A SQL Server Agent job that copies the backup files from the primary server to a configurable destination on the secondary server and logs history on the secondary server and the monitor server. When log shipping is enabled on a database, the job category "Log Shipping Copy" is created on each secondary server in a log shipping configuration.
Box 3: Secondary server instance.
Each secondary server instance runs its own restore job to restore the log backup from the local destination folder onto the local secondary database.
restore job: A SQL Server Agent job that restores the copied backup files to the secondary databases. It logs history on the local server and the monitor server, and deletes old files and old history information. When log shipping is enabled on a database, the job category "Log Shipping Restore" is created on the secondary server instance.
References: <https://docs.microsoft.com/en-us/sql/database-engine/log-shipping/about-log-shipping-sql-server>

NEW QUESTION 40

- (Exam Topic 1)

You plan to integrate an on-premises Microsoft SQL Server environment with Microsoft Azure. You need to create the authentication object so that you can connect to Azure.

Which Windows PowerShell command or commands should you run?

- A. Invoke-Sqlcmd "CREATE EXTERNAL DATA SOURCE MyAzureStorage WITH (LOCATION = 'wasbs://Azure@myaccount.blob.core.windows.net/', CREDENTIAL = Pa\$\$w0rd)"
- B. New-SqlAzureKeyVaultColumnMasterKeySettings-KeyUri<https://myvault.vault.contoso.net:443/keys/C>
- C. Invoke-Sqlcmd "CREATE CREDENTIAL AzureCred WITH IDENTITY = 'AzureKey', SECRET = 'Pa\$\$w0rd'"
- D. Invoke-Sqlcmd "CREATE LOGIN AzureCred WITH CREDENTIAL = 'AzureKey', PASSWORD = 'Pa\$\$w0rd'"

Answer: C

Explanation:

Invoke-Sqlcmd runs a script containing statements supported by the SQL Server SQLCMD utility.

The following example creates a SQL Server credential for the Database Engine to use when accessing the Azure Key Vault using the SQL Server Connector for Microsoft Azure Key Vault.

```
CREATE CREDENTIAL Azure_EKM_TDE_cred WITH IDENTITY = 'ContosoKeyVault',  
SECRET = 'EF5C8E094D2A4A769998D93440D8115DSECRET_DBEngine'  
FOR CRYPTOGRAPHIC PROVIDER AzureKeyVault_EKM_Prov ;
```

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

NEW QUESTION 41

- (Exam Topic 1)

You have a test server that contains a database named DB1. Backups of the database are written to a single backup device. The backup device has a full, differential, and transaction log backup.

You discover that the database is damaged. You restore the database to the point at which the differential backup was taken.

You need to rebuild the database with data stored in the latest transaction logs.

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

Transact-SQL statements

NORECOVERY

RECOVERY

LOG

DBCC CHECKDB

CONTINUE_AFTER_ERROR

RESTORE

RESTORE VERIFYONLY

Answer Area

Transact-SQL segment

DB1 FROM DISK = N'Z:Backups\Backup.bak WITH

Transact-SQL segment

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: RESTORE

Box 2: RECOVERY

The RESTORE ... WITH RECOVERY option puts the database into a useable state, so users can access a restored database.

References:

<https://www.mssqltips.com/sqlservertutorial/112/recovering-a-database-that-is-in-the-restoring-state/>

NEW QUESTION 42

- (Exam Topic 1)

You administer a Microsoft SQL Server 2016 server that has SQL Server Integration Services (SSIS) installed. You plan to deploy new SSIS packages to the server.

The SSIS packages use the Project Deployment Model together with parameters and Integration Services environment variables.

You need to configure the SQL Server environment to support these packages. What should you do?

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

Answer: B

Explanation:

Use can use Project Deployment Model for a project, containing packages and parameters, which is deployed to the SSISDB catalog on an instance of SQL Server.

References:

<https://docs.microsoft.com/en-us/sql/integration-services/packages/deploy-integration-services-ssis-projects-and>

NEW QUESTION 45

- (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.

One instance hosts a user database named HRDB. The database contains sensitive human resources data. You need to grant an auditor permission to view the SQL Server audit logs while following the principle of least privilege.

Which permission should you grant?

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

Answer: F

Explanation:

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

Membership in the sysadmin fixed server role.

The CONTROL SERVER permission.

The VIEW SERVER STATE permission.

The ALTER ANY AUDIT permission.

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

References: [https://technet.microsoft.com/en-us/library/cc280386\(v=sql.110\).aspx](https://technet.microsoft.com/en-us/library/cc280386(v=sql.110).aspx)

NEW QUESTION 50

- (Exam Topic 1)

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

Currently, Server01/HA is the primary replica.

You need to ensure that the following requirements are met:

Backup operations occur on Server02/HA.

If Server02/HA is unavailable, backup operations occur on Server03/HA.

Backup operations do not occur on Server01/HA.

How should you configure HaContoso?

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

Answer: B

Explanation:

Secondary only: Specifies that backups should never be performed on the primary replica. If the primary replica is the only replica online, the backup should not occur.

Backup Priority (Lowest=1, Highest=100)

Specifies your priority for performing backups on this replica relative to the other replicas in the same availability group. The value is an integer in the range of 0..100. 1 indicates the lowest priority, and 100 indicates the highest priority. If Backup Priority = 1, the availability replica would be chosen for performing backups only if no higher priority availability replicas are currently available.

References:

<https://docs.microsoft.com/en-us/sql/database-engine/availability-groups/windows/configure-backup-on-availab>

NEW QUESTION 54

- (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 56

- (Exam Topic 1)

You are configuring a new Microsoft SQL Server Always On Availability Group. You plan to configure a shared network location at \\DATA-C11\\SQL.

You need to create an availability group listener named AGL1 on port 1433.

In which order should you perform the actions? To answer, move all actions from the list of actions to the answer area and arrange them in the correct order.

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1: Launch the Failover Cluster Manager and..

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

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

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

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

Enable the Always On availability groups feature on each server instance that will host an availability replica for any availability group. On a given computer, you can enable as many server instances for Always On availability groups as your SQL Server installation supports.

Step 4: Create the Always On Availability Group and..

Using Transact-SQL to create or configure an availability group listener Step 5: Select the Full data synchronization method and...

References: [https://technet.microsoft.com/en-us/library/jj899851\(v=sc.12\).aspx](https://technet.microsoft.com/en-us/library/jj899851(v=sc.12).aspx)

<https://docs.microsoft.com/en-us/sql/database-engine/availability-groups/windows/create-or-configure-an-availa>

NEW QUESTION 58

- (Exam Topic 1)

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

You are examining information about users, sessions, and processes in an on-premises Microsoft SQL Server 2016 Standard Edition server.

You need to identify waits for resources and return only the following information:

a list of all databases on the SQL Server instance, along with information about the database files, their paths, and names

a list of the queries recently executed that use most of memory, disk, and network resources

What should you use?

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

Answer: E

Explanation:

SQL Server Data Tools (SSDT) is a Microsoft Visual Studio environment for creating business intelligence solutions. SSDT features the Report Designer authoring environment, where you can open, modify, preview, save, and deploy Reporting Services paginated report definitions, shared data sources, shared datasets, and report parts.

References: [https://msdn.microsoft.com/en-us/library/hh272686\(v=vs.103\).aspx](https://msdn.microsoft.com/en-us/library/hh272686(v=vs.103).aspx)

NEW QUESTION 61

- (Exam Topic 1)

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

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

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

Your company has several Microsoft SQL Server instances, Each instance hosts many databases. You observe I/O corruption on some of the instances

You need to perform the following actions:

- Identify databases where the PAGE VERIFY option is not set
- Configure full page protection for the identified databases. Solution: You run the following Transact-SQL statement:

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

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

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

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

NEW QUESTION 64

- (Exam Topic 1)

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

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

What should you use?

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

Answer: C

Explanation:

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

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

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

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

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

To start the rebuild operation, click OK.

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

NEW QUESTION 69

- (Exam Topic 1)

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

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

Your company has several Microsoft SOI Server instances. Each instance hosts many databases. You observe I/O corruption on some of the instances.

You need to perform the following actions:

- Identify databases where the PAGE VERIFY option is not set.
- Configure full page protection for the identified databases. Solution: You run the following Transact-SQL statement:

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

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

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

Does the solution meet the goal?

- A. Yes
- B. NO

Answer: B

NEW QUESTION 70

- (Exam Topic 1)

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

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

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

What should you use?

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

Answer: D

Explanation:

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

NEW QUESTION 75

- (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)

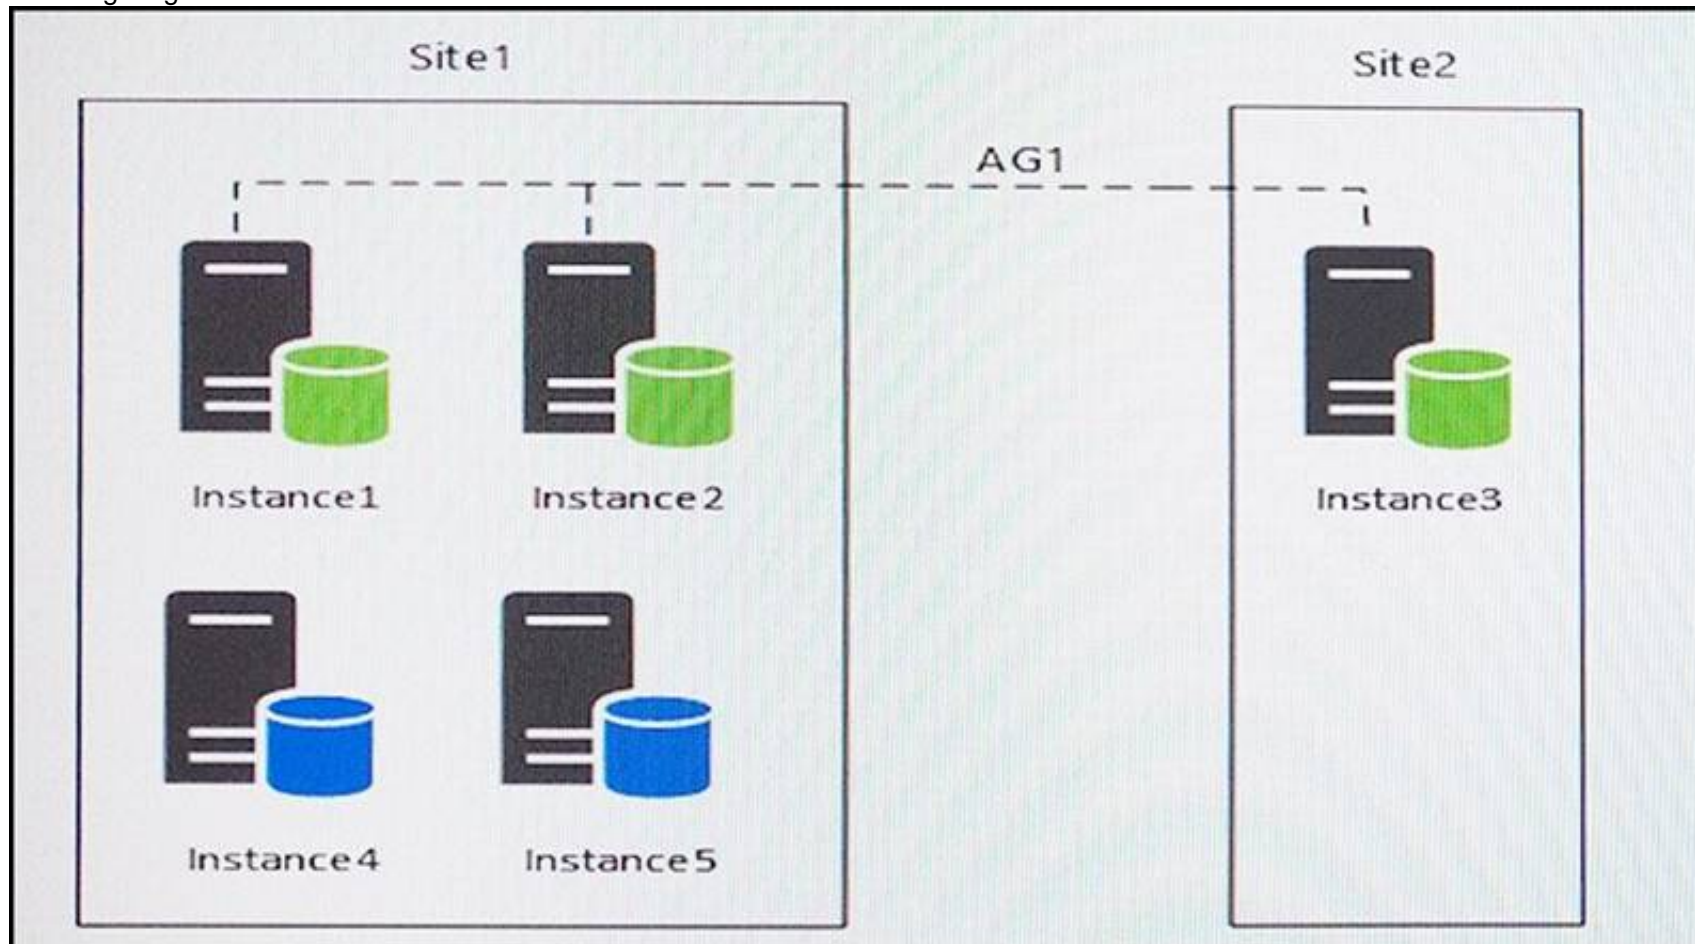
NEW QUESTION 77

- (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 analyze the wait type and statistics for specific instanced in the environment.

Which object should you use to gather information about each instance? To answer, drag the appropriate objects to the correct instances. Each object may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Objects

Sys.dm_os_wait_stats

Sys.dm_exec_connections

Sys.dm_exec_requests

Sys.dm_exec_procedure_stats

Sys.dm_exec_sessions

Sys.dm_exec_query_stats

Sys.dm_exec_query_re-
source_semaphores

Sys.dm_exec_ses-
sion_wait_stats

Answer Area

Instance

Object

Instance1

Object

Instance4

Object

Instance5

Object

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Instance 1: sys.dm_exec_query_stats

From Scenario: Instance1 requirement: Aggregate statistics since last server restart. sys.dm_exec_query_stats returns aggregate performance statistics for cachedquery plans in SQL Server.

Instance 4: sys.dm_os_wait_stats

sys.dm_os_wait_statsreturns information about all the waits encountered by threads that executed. From Scenario: Instance4 requirement: Identify the most prominent wait types.

Identify the most prominent wait types for all the commands originating from a session, between session connections, or between application pool resets.

Instance 5:sys.dm_exec_session_wait_stats

From Scenario: Instance5 requirement: Identify all wait types for queries currently running on the server. sys.dm_exec_session_wait_stats returns information about all the waits encountered by threads that executed for each session.

NEW QUESTION 82

- (Exam Topic 1)

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

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

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

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

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

Answer: C

Explanation:

To add a node to an existing SQL Server failover cluster

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

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

Etc.

References:

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

NEW QUESTION 87

- (Exam Topic 1)

A Microsoft SQL Server database named DB1 has two filegroups named FG1 and FG2. You implement a backup strategy that creates backups for the filegroups. DB1 experiences a failure. You must restore FG1 and then FG2.

You need to ensure that the database remains in the RECOVERING state until the restoration of FG2 completes. After the restoration of FG2 completes, the database must be online.

What should you specify when you run the recovery command?

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

Answer: A

NEW QUESTION 90

- (Exam Topic 1)

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

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

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

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

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

Does this meet the goal?

- A. Yes
- B. No

Answer: A

NEW QUESTION 94

- (Exam Topic 1)

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

You need to examine information about logins, CPU times, and Disk I/O on a particular database in Microsoft Azure.

What should you use?

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

Answer: A

Explanation:

Activity Monitor displays information about SQL Server processes and how these processes affect the current instance of SQL Server.

Activity Monitor is a tabbed document window with the following expandable and collapsible panes: Overview, Active User Tasks, Resource Waits, Data File I/O, and Recent Expensive Queries.

The Activity User Tasks Pane shows information for active user connections to the instance, including the following column:

* Login: The SQL Server login name under which the session is currently executing.

The Recent Expensive Queries Pane shows information about the most expensive queries that have been run on the instance over the last 30 seconds, including the following column:

* CPU (ms/sec): The rate of CPU use by the query

References: [https://technet.microsoft.com/en-us/library/cc879320\(v=sql.105\).aspx](https://technet.microsoft.com/en-us/library/cc879320(v=sql.105).aspx)

NEW QUESTION 97

- (Exam Topic 1)

You administer a database that is used for reporting purposes. The database has a large fact table that contains three hundred million rows. The table includes a clustered columnstore index and a nonclustered index on the ProductID column. New rows are inserted into the table every day.

Performance of queries that filter the Product ID column have degraded significantly. You need to improve the performance of the queries.

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

Actions	Answer Area
Drop the clustered columnstore index.	
Create a nonclustered index on ProductID.	
Drop and recreate the clustered columnstore index.	
Create a nonclustered columnstore index on ProductID.	
Recreate the clustered columnstore index using DROP EXISTING.	
Create a clustered rowstore index on ProductID.	
Rebuild the clustered columnstore index.	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1: Drop the clustered columnstore index

Step 2: Create a clustered rowstore index on ProductID.

Rowstore indexes perform best on queries that seek into the data, searching for a particular value, or for queries on a small range of values. Use rowstore indexes with transactional workloads since they tend to require mostly table seeks instead of table scans.

Step 3: Create a nonclustered index on ProductID

NEW QUESTION 100

- (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 resolves the missing data issue.

The solution must minimize the amount of development effort. What should you recommend?

- A. Denormalize the Products table.
- B. Denormalize the OrderDetails table.
- C. Normalize the OrderDetails table.
- D. Normalize the Products table.

Answer: D

Explanation:

- Scenario:

- 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.

- 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.

NEW QUESTION 102

- (Exam Topic 2)

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

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

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

Fulfilled
Shipped
Ordered
Received

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

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

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

Answer: A

NEW QUESTION 104

- (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 App1_Db1 as soon as changes occur to the data in Database2. Litware plans to use offsite storage for all SQL Server 2014 backups.

Business Requirements

You have the following requirements:

Costs for new licenses must be minimized.

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

Development effort must be minimized whenever possible.

The storage requirements for databases must be minimized.

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

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

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

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

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

The encryption strategy must minimize changes to the databases and to the applications. You need to recommend a solution to improve the performance of usp_UpdateInventory.

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

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

Answer: A

Explanation:

- Scenario: Database2 will contain a stored procedure named usp_UpdateInventory. Usp_UpdateInventory will manipulate a table that contains a self-join that has an unlimited number of hierarchies.

- A table variable can be very useful to store temporary data and return the data in the table format.

- Example: The following example uses a self-join to find the products that are supplied by more than one vendor. Because this query involves a join of the

ProductVendor table with itself, the ProductVendor table appears in two roles. To distinguish these roles, you must give the ProductVendor table two different aliases (pv1 and pv2) in the FROM clause. These aliases are used to qualify the column names in the rest of the query. This is an example of the self-join Transact-SQL statement:

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

NEW QUESTION 105

- (Exam Topic 2)

You plan to create a database.

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

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

Answer: D

NEW QUESTION 109

- (Exam Topic 2)

Overview

General Overview

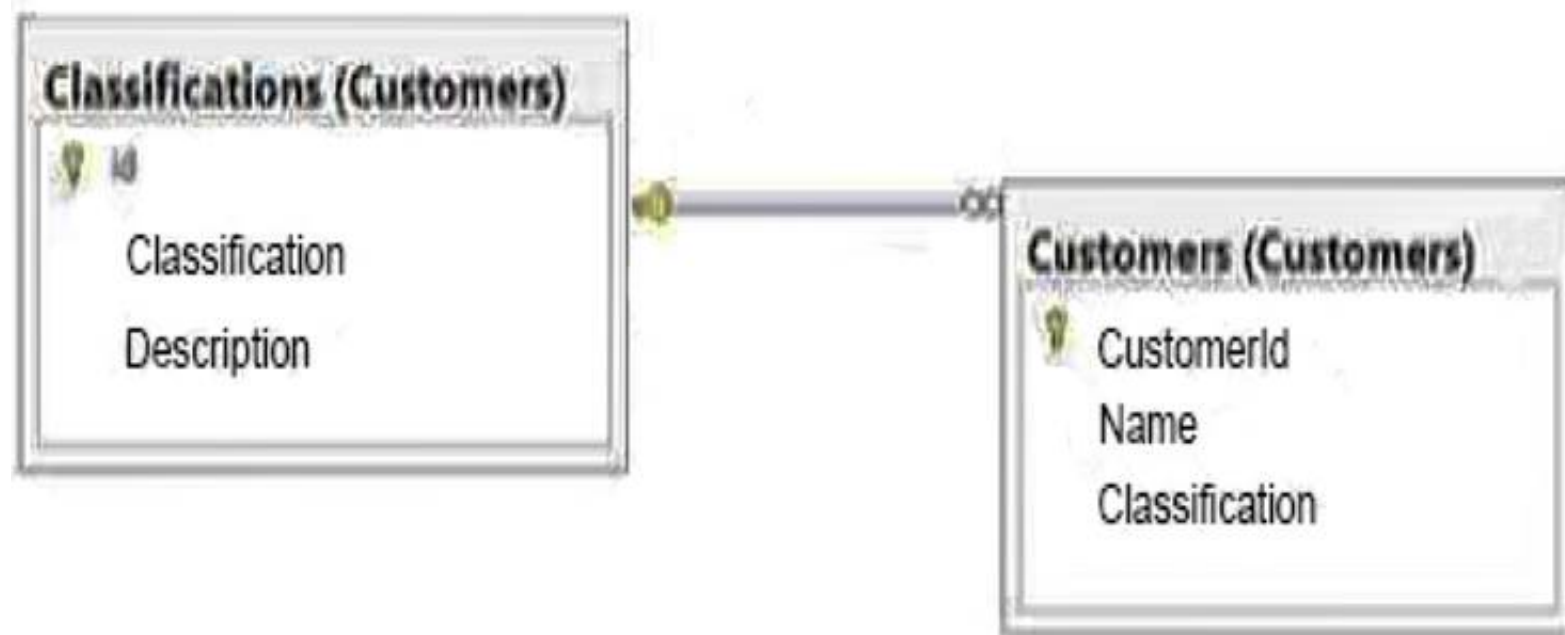
ADatum Corporation has offices in Miami and Montreal.

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

Databases

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

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



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

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

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

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

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

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

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

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

A stored procedure named USP_3 is used to update prices. USP_3 is composed of several UPDATE statements called in sequence from within a transaction. Currently, if one of the UPDATE statements fails, the stored procedure fails. A stored procedure named USP_4 calls stored procedures in the Sales, Customers, and Inventory databases.

The nested stored procedures read tables from the Sales, Customers, and Inventory databases. USP_4 uses an EXECUTE AS clause. All nested stored procedures handle errors by using structured exception handling. A stored procedure named USP_5 calls several stored procedures in the same database. Security checks are performed each time USP_5 calls a stored procedure. You suspect that the security checks are slowing down the performance of USP_5. All stored procedures accessed by user applications call nested stored procedures. The nested stored procedures are never called directly. Design Requirements

Data Recovery
You must be able to recover data from the Inventory database if a storage failure occurs. You have a Recovery Time Objective (RTO) of 5 minutes. You must be able to recover data from the Dev database if data is lost accidentally. You have a Recovery Point Objective (RPO) of one day.

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

Storage
ADatum has limited storage. Whenever possible, all storage space should be minimized for all databases and all backups. Error Handling
There is currently no error handling code in any stored procedure. You plan to log errors in called stored procedures and nested stored procedures. Nested stored procedures are never called directly. You need to recommend a disaster recovery solution for the Dev database. What should you include in the recommendation?

- A. The simple recovery model and full backups
- B. The full recovery model, full backups, and transaction log backups
- C. The full recovery model, full backups, and differential backups
- D. The bulk-logged recovery model and full backups

Answer: A

Explanation:

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

NEW QUESTION 110

- (Exam Topic 2)

Overview

General Overview

ADatum Corporation has offices in Miami and Montreal.

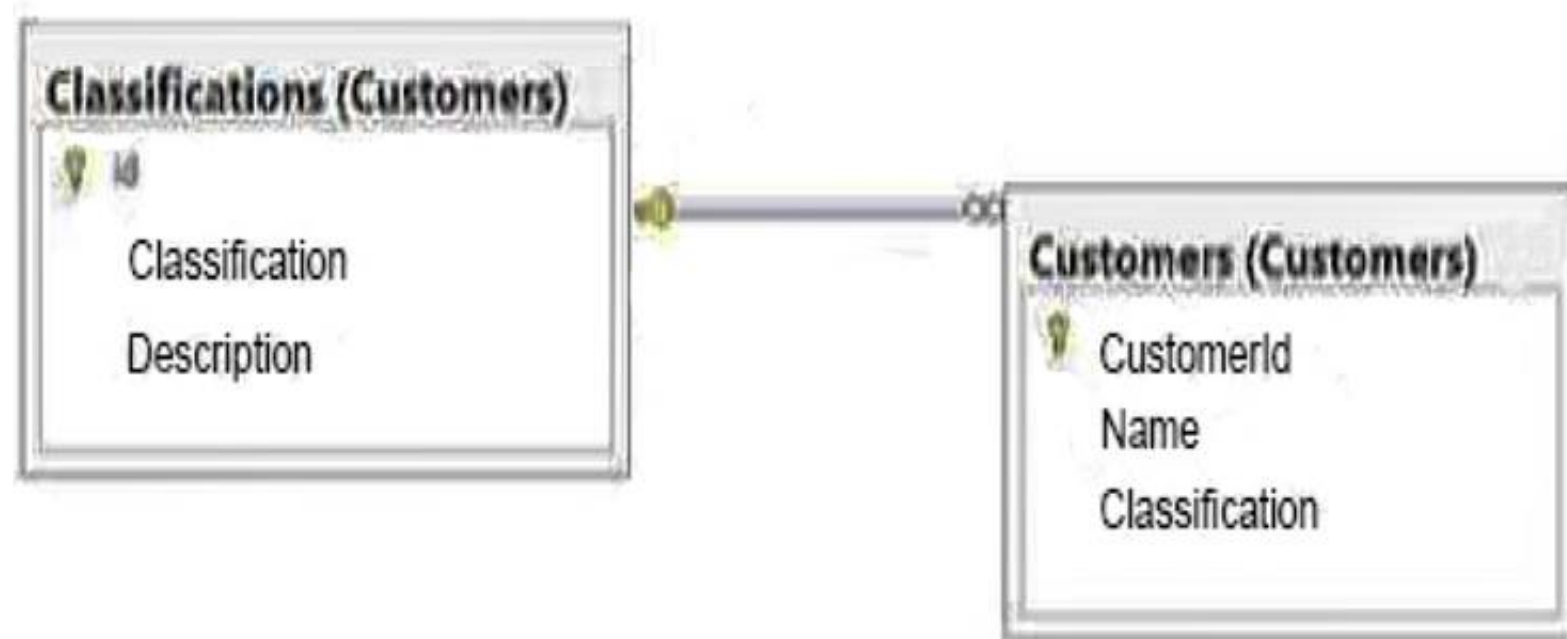
The network contains a single Active Directory forest named adatum.com. The offices connect to each other by using a WAN link that has 5-ms latency. A. Datum standardizes its database platform by using SQL Server 2014 Enterprise edition.

Databases

Each office contains databases named Sales, Inventory, Customers, Products, Personnel, and Dev.

Servers and databases are managed by a team of database administrators. Currently, all of the database administrators have the same level of permissions on all of the servers and all of the databases.

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



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

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

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

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

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

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

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

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

A stored procedure named USP_3 is used to update prices. USP_3 is composed of several UPDATE statements called in sequence from within a transaction. Currently, if one of the UPDATE statements fails, the stored procedure fails. A stored procedure named USP_4 calls stored procedures in the Sales, Customers, and Inventory databases.

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

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

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

The nested stored procedures are never called directly. Design Requirements

Data Recovery

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

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

Classification Changes

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

Storage

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

Error Handling

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

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

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

Which change should you recommend?

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

Answer: A

Explanation:

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

NEW QUESTION 115

- (Exam Topic 2)

Overview

Application Overview

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

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

The new version will use SQL Server 2014.

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

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

Tables

The current database schema contains a table named OrderDetails.

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

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

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

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

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

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

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

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

Customer Problems Installation Issues

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

Index Fragmentation Issues

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

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

Backup Issues

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

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

Missing Data Issues

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

Query Performance Issues

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

Design Requirements

File Storage Requirements

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

Data Recovery Requirements

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

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

Concurrency Requirements

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

You need to recommend a solution that addresses the concurrency requirement. What should you recommend?

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

Answer: B

Explanation:

- Concurrency Requirements

You must reduce the likelihood of deadlocks occurring when Sales.Proc1 and Sales.Proc2 execute.

NEW QUESTION 116

- (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 119

- (Exam Topic 2)

You deploy a database by using SQL Server 2014. The database contains a table named Table1.

You need to recommend a solution to track all of the deletions executed on Table1. The solution must minimize the amount of custom code required.

What should you recommend?

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

Answer: A

Explanation:

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

NEW QUESTION 122

- (Exam Topic 2)

Your company has offices in Seattle and Montreal.

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

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

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

What should you recommend?

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

Answer: D

NEW QUESTION 126

- (Exam Topic 2)

You administer a Microsoft SQL Server 2016 server.

The MSSQLSERVER service uses a domain account named CONTOSO\SQLService. You plan to configure Instant File Initialization.

You need to ensure that Data File Autogrow operations use Instant File Initialization. What should you do? Choose all that apply.

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

Answer: CD

Explanation:

How To Enable Instant File Initialization

Open Local Security Policy and go to Local Policies → User Rights Assignment.

Double click Perform Volume Maintenance Tasks and add your SQL Server database engine service account.

Restart the SQL Server service using SQL Server Configuration Manager and this setting should now be enabled.

References:

<http://msdn.microsoft.com/en-us/library/ms175935.aspx>

NEW QUESTION 131

- (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 running.
- 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 134

- (Exam Topic 2)

Overview

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

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

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

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

Requirements Planned Changes

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

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

Database1 will also contain a stored procedure named usp_UpdateOrderDetails.

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

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

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

Inventory will contain over 100 GB of data.

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

Database2 will contain a stored procedure named usp_UpdateInventory. usp_UpdateInventory will manipulate a table that contains a self-join that has an unlimited number of hierarchies. All data in Database2 is recreated each day and does not change until the next data creation process. Data from Database2 will be accessed periodically by an external application named Application1. The data from Database2 will be sent to a database named Appl_Dbl as soon as changes occur to the data in Database2. Litware plans to use offsite storage for all SQL Server 2014 backups.

Business Requirements

You have the following requirements:

Costs for new licenses must be minimized.

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

Development effort must be minimized whenever possible.

The storage requirements for databases must be minimized.

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

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

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

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

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

The encryption strategy must minimize changes to the databases and to the applications. You need to recommend a disk monitoring solution that meets the business requirements. What should you include in the recommendation?

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

Answer: B

Explanation:

Dynamic Management Views and Functions (Transact-SQL)

NEW QUESTION 136

- (Exam Topic 2)

Overview

Application Overview

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

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

The new version will use SQL Server 2014.

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

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

Tables

The current database schema contains a table named OrderDetails.

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

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

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

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

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

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

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

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

Customer Problems Installation Issues

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

Index Fragmentation Issues

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

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

Backup Issues

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

Users report that when they search product names, the search results exclude product names that contain

accents, unless the search string includes the accent. Missing Data Issues

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

Query Performance Issues

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

Design Requirements

File Storage Requirements

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

Data Recovery Requirements

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

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

Concurrency Requirements

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

What should you include in the recommendation?

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

Answer: C

Explanation:

- Scenario: Installation Issues The current version of the ERP application requires that several SQL Server logins be set up to function correctly. Most customers set up the ERP application in multiple locations and must create logins multiple times.
- Creating contained users enables the user to connect directly to the contained database. This is a very significant feature in high availability and disaster recovery scenarios such as in an AlwaysOn solution. If the users are contained users, in case of failover, people would be able to connect to the secondary without creating logins on the instance hosting the secondary. This provides an immediate benefit.

NEW QUESTION 138

- (Exam Topic 2)

Overview

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

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

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

will use the application.

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

Requirements Planned Changes

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

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

Database1 will also contain a stored procedure named usp_UpdateOrderDetails.

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

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

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

Inventory will contain over 100 GB of data.

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

Database2 will contain a stored procedure named usp_UpdateInventory. usp_UpdateInventory will manipulate a table that contains a self-join that has an unlimited number of hierarchies. All data in Database2 is recreated each day and does not change until the next data creation process. Data from Database2 will be accessed periodically by an external application named Application1. The data from Database2 will be sent to a database named Appl_Db1 as soon as changes occur to the data in Database2. Litware plans to use offsite storage for all SQL Server 2014 backups.

Business Requirements

You have the following requirements:

Costs for new licenses must be minimized.

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

Development effort must be minimized whenever possible.

The storage requirements for databases must be minimized.

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

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

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

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

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

The encryption strategy must minimize changes to the databases and to the applications.

You need to recommend a 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 139

- (Exam Topic 2)

You have two SQL Server 2012 instances named SQLDev and SQLProd.

You plan to create a new database on SQLProd that will use SQL Server Authentication.

You need to ensure that when the new database is copied from SQLProd to SQLDev, users can connect to the database on SQLDev if they do not have a login on the SQLDev instance.

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

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

Answer: E

Explanation:

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

NEW QUESTION 140

- (Exam Topic 2)

You administer a Microsoft SQL Server 2016 failover cluster.

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

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

Answer: B

Explanation:

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

References:

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

NEW QUESTION 145

- (Exam Topic 2)

You administer a Microsoft SQL Server 2016 database named Contoso on a server named Server01. You need to be notified immediately when fatal errors occur on Server01.

What should you create?

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

Answer: A

Explanation:

References:

<http://www.sqlskills.com/blogs/glenn/creating-sql-server-agent-alerts-for-critical-errors/>

NEW QUESTION 147

- (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

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

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

Index Fragmentation Issues

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

Backup Issues

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

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

Missing Data Issues

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

Query Performance Issues

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

Design Requirements**File Storage Requirements**

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

Data Recovery Requirements

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

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

Concurrency Requirements

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

You need to recommend a solution that addresses the index fragmentation and index width issue. What should you include in the recommendation? (Each correct answer presents part of the solution. Choose all that apply.)

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

Answer: BE

Explanation:

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

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

NEW QUESTION 149

- (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 153

- (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 security requirement. What should you recommend?

A. Revoke user permissions on the table

B. Create stored procedures that manipulate dat

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

Answer: C

Explanation:

- Security Requirements

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

NEW QUESTION 158

- (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 160

- (Exam Topic 2)

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

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

Answer: A

Explanation:

Resource Governor enables you to allocate session requests to different resources based on the characteristics of the session request properties.

References: <https://docs.microsoft.com/en-us/sql/relational-databases/resource-governor/resource-governor>

NEW QUESTION 162

- (Exam Topic 2)

You plan to deploy SQL Server 2012. You must create two tables named Table 1 and Table 2 that will have the following specifications:

Table1 will contain a date column named Column1 that will contain a null value approximately 80 percent of the time.

Table2 will contain a column named Column2 that is the product of two other columns in Table2. Both Table1 and Table2 will contain more than 1 million rows.

You need to recommend which options must be defined for the columns. The solution must minimize the storage requirements for the tables.

Which options should you recommend? To answer, drag the appropriate options to the correct column in the answer area.

Options	Answer Area	
Sparse	Column1	Option
Computed	Column2	Option
Persisted computed		

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Column1 – Sparse; Column2 - Computed

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

References:

<http://msdn.microsoft.com/en-us/library/cc280604.aspx> <http://msdn.microsoft.com/en-us/library/ms186241.aspx>

NEW QUESTION 165

- (Exam Topic 2)

Overview

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

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

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

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

Requirements Planned Changes

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

All database files will be stored in a highly available SAN.

Database1 will contain two tables named Orders and OrderDetails.

Database1 will also contain a stored procedure named usp_UpdateOrderDetails.

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

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

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

Inventory will contain over 100 GB of data.

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

Database2 will contain a stored procedure named usp_UpdateInventory. usp_UpdateInventory will manipulate a table that contains a self-join that has an unlimited number of hierarchies. All data in Database2 is recreated each day and does not change until the next data creation process. Data from Database2 will be accessed periodically by an external application named Application1. The data from Database2 will be sent to a database named Appl_Dbl as soon as changes occur to the data in Database2. Litware plans to use offsite storage for all SQL Server 2014 backups.

Business Requirements

You have the following requirements:

- Costs for new licenses must be minimized.
- Private information that is accessed by Application must be stored in a secure format.
- Development effort must be minimized whenever possible.
- The storage requirements for databases must be minimized.
- System administrators must be able to run real-time reports on disk usage.
- The databases must be available if the SQL Server service fails.
- Database administrators must receive a detailed report that contains allocation errors and data corruption.
- Application developers must be denied direct access to the database tables. Applications must be denied direct access to the tables.
- You must encrypt the backup files to meet regulatory compliance requirements.
- The encryption strategy must minimize changes to the databases and to the applications.

You need to recommend a solution for Application1 that meets the security requirements. What should you include in the recommendation?

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

Answer: D

Explanation:

- Scenario:
- Data from Database2 will be accessed periodically by an external application named Application1
- Application developers must be denied direct access to the database tables. Applications must be denied direct access to the tables.

Tutorial: Signing Stored Procedures with a Certificate

NEW QUESTION 170

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

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

* One year free update

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

* Trusted by Millions

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

* Shop Securely

All transactions are protected by VeriSign!

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

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