



Microsoft

Exam Questions 70-764

Administering a SQL Database Infrastructure (beta)

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)

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 monitor WingDB and gather information for troubleshooting issues. What should you use?

- A. sp_updatestats
- B. sp_lock
- C. sys.dm_os_waiting_tasks
- D. sys.dm_tran_active_snapshot_database_transactions
- E. Activity Monitor

Answer: B

Explanation:

The sp_lock system stored procedure is packaged with SQL Server and will give you insight into the locks that are happening on your system. This procedure returns much of its information from the syslock info in the master database, which is a system table that contains information on all granted, converting, and waiting lock requests.

Note: sp_lock will be removed in a future version of Microsoft SQL Server. Avoid using this feature in new development work, and plan to modify applications that currently use this feature. To obtain information about locks in the SQL Server Database Engine, use the sys.dm_tran_locks dynamic management view. sys.dm_tran_locks returns information about currently active lock manager resources in SQL Server 2008 and later. Each row represents a currently active request to the lock manager for a lock that has been granted or is waiting to be granted.

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/system-stored-procedures/sp-lock-transact-sql>

NEW QUESTION 3

- (Exam Topic 1)

You are a database administrator for a Microsoft SQL Server 2016 environment.

You want to deploy a new application that will scale out the workload to at least five different SQL Server instances.

You need to ensure that for each copy of the database, users are able to read and write data that will then be synchronized between all of the database instances.

Which feature should you use?

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

Answer: B

NEW QUESTION 4

- (Exam Topic 1)

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

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

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

Database	Description
Transactional database	<p>Recovery model:</p> <ul style="list-style-type: none"> • Full <p>Backup schedule:</p> <ul style="list-style-type: none"> • Full database backup: midnight, daily • Differential database backup: on the hour, every two hours starting at 02:00 hours except at 00:00 hours • Log backup: every half hour, except at the times of full and differential backups
Reporting database	<p>Recovery model:</p> <ul style="list-style-type: none"> • Simple <p>Backup schedule:</p> <ul style="list-style-type: none"> • Full database backup: 01:00 hours daily • Differential database backup: 13:00 hours daily <p>Data updates:</p> <ul style="list-style-type: none"> • Changes in data are updated from the transactional database to the reporting database at 00:30 hours and at 12:30 hours • The update takes 15 minutes

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

What should you do?

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

Answer: F

Explanation:

The goal of a page restore is to restore one or more damaged pages without restoring the whole database. Typically, pages that are candidates for restore have been marked as "suspect" because of an error that is encountered when accessing the page.

Note: Requirements for Restoring Pages

A page restore is subject to the following requirements:

The databases must be using the full or bulk-logged recovery model. Etc.

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

NEW QUESTION 5

- (Exam Topic 1)

You administer a Microsoft SQL Server 2016 database.

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

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

Answer: A

NEW QUESTION 6

- (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 7

- (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 8

- (Exam Topic 1)

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

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

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

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

You need to optimize query performance.

Solution: You run the DBCC CHECKDB command. Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

DBCC CHECKDB only checks the logical and physical integrity of all the objects in the specified database. It does not update any indexes, and does not improve query performance.

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

NEW QUESTION 9

- (Exam Topic 1)

You have configured Resource Governor with three resource pools.

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

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

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

Answer: B

Explanation:

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

NEW QUESTION 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.

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 10

- (Exam Topic 1)

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

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

You 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 15

- (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 18

- (Exam Topic 1)

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

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

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

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

You need to optimize query performance. Solution: You rebuild all indexes.

Does the solution meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

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

NEW QUESTION 21

- (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 25

- (Exam Topic 1)

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

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

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

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

Answer: D

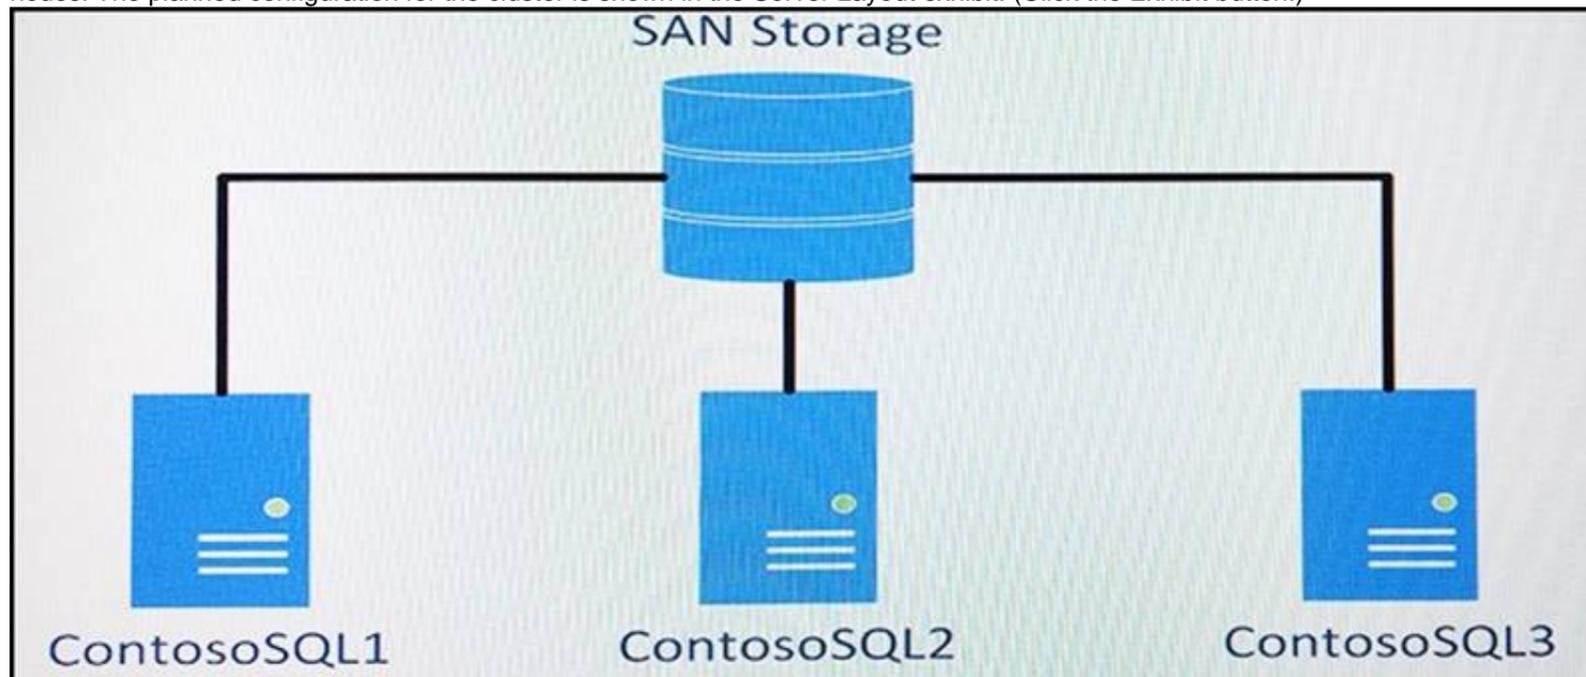
Explanation:

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

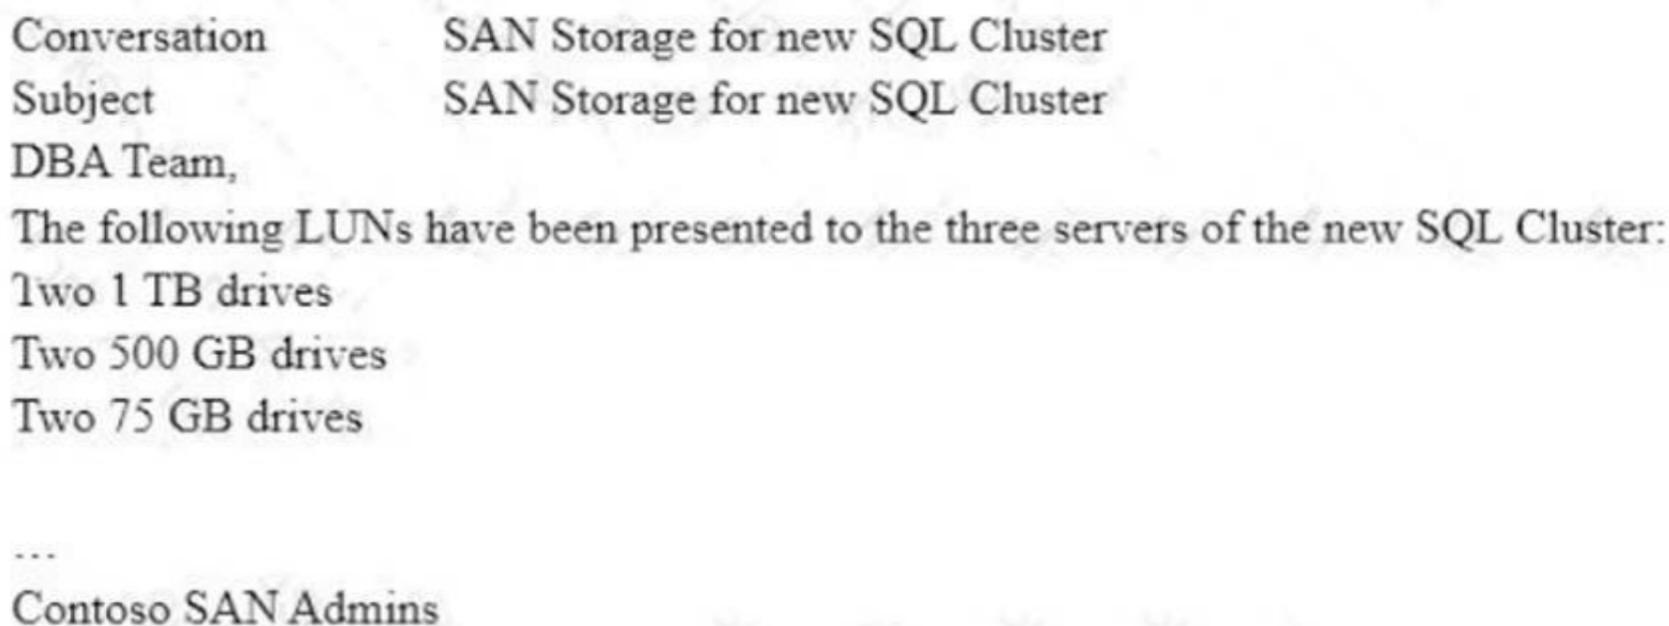
NEW QUESTION 28

- (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.)



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.

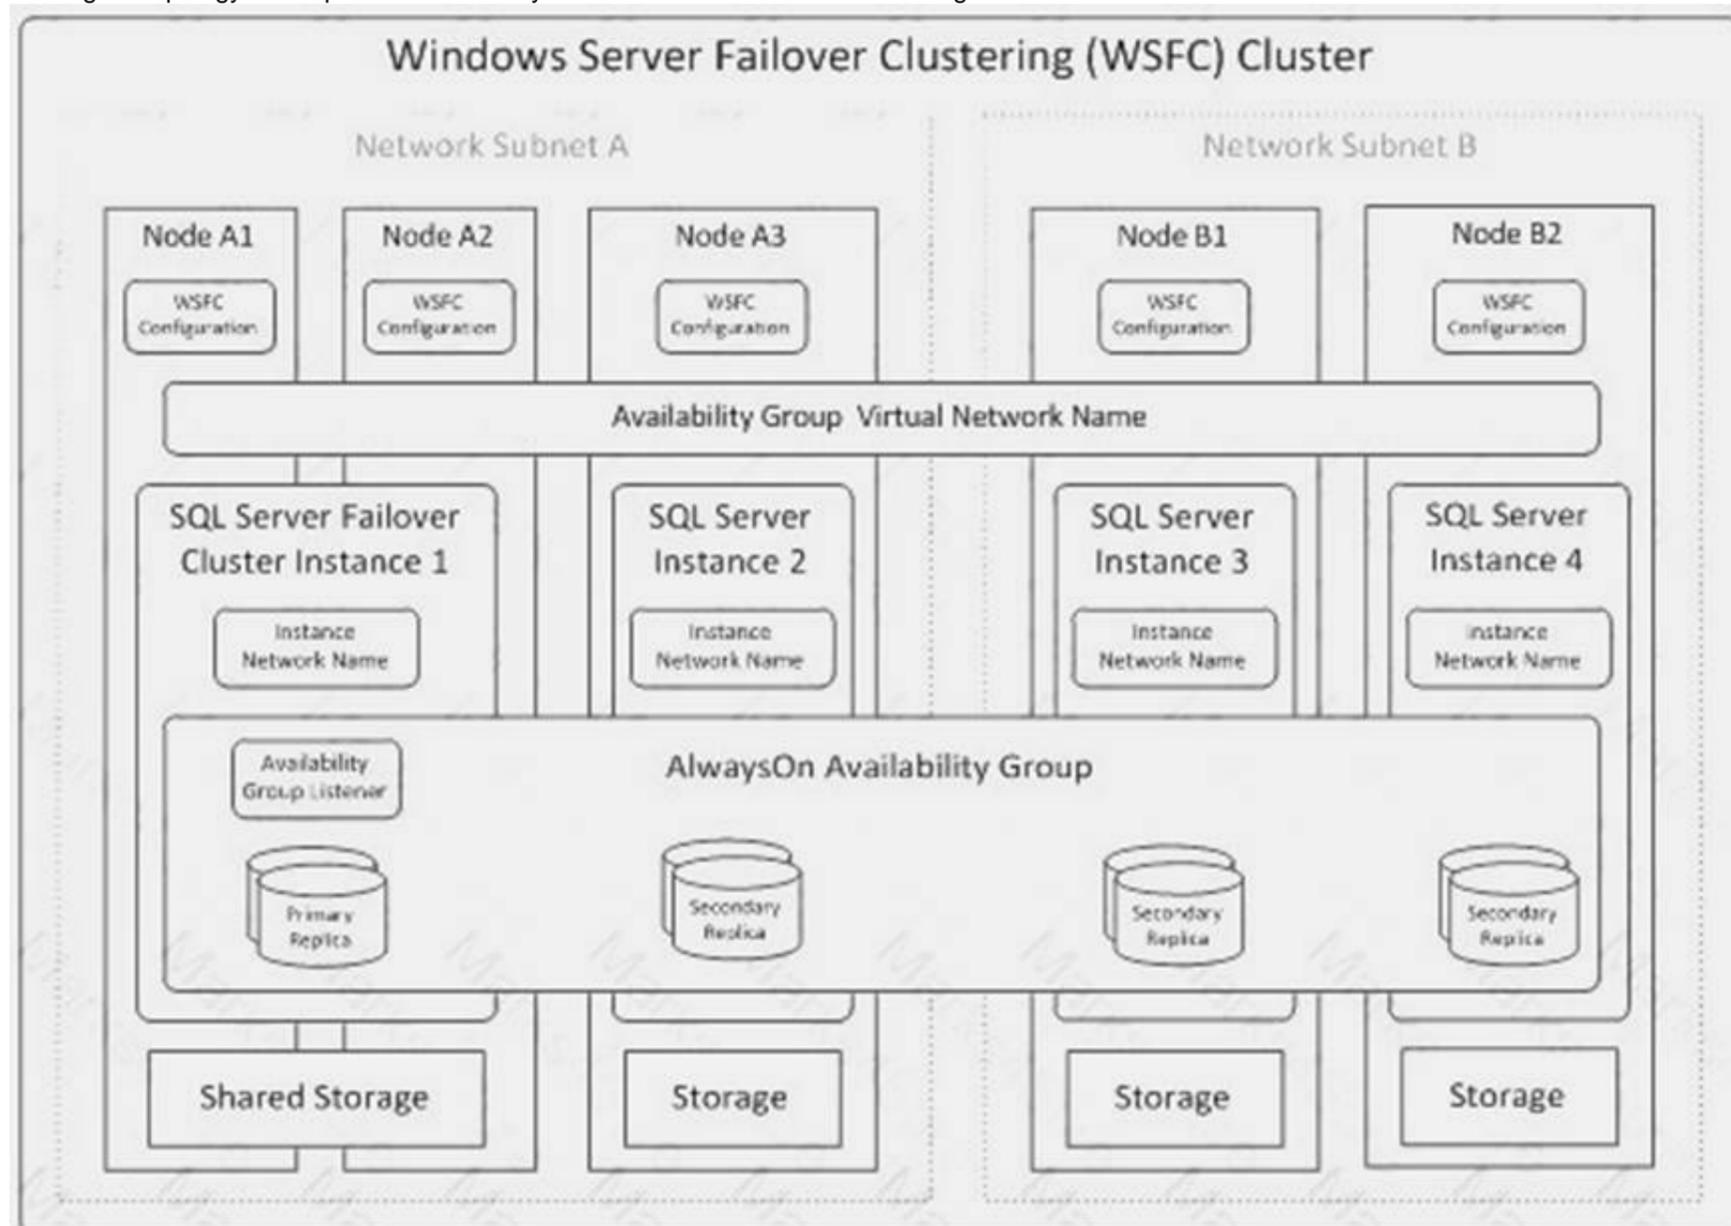
Prior to SQL Server 2012, FCIs required tempdb to be located on a symmetrical shared storage volume that failed over with other system databases.

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 29

- (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 33

- (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 35

- (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 37

- (Exam Topic 1)

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

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

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

Answer Area

Actions

User1 server role assignment

Options

▼
diskadmin
serveradmin
securityadmin
setupadmin

Transact-SQL syntax

▼
sys.server_file_audits
sys.server_audit_specifications
sys.server_file_permissions
sys.server_principals

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: securityadmin

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

Box 2: sys.server_audit_specifications

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

NEW QUESTION 40

- (Exam Topic 1)

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

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

You need to determine the changes that the script has on the environment.

How does the script affect the environment? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Answer Area

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

- is
- is not

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

- runs
- does not run

Backup files will be deleted after [answer choice].

- 24 hours
- 48 hours
- 72 hours

The backup job will run every [answer choice].

- 15 minutes
- 60 minutes
- 24 hours

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: is

The dedicated backup file share is \\localhost\Backup Box 2: does not run

The only thing with a name related to ADATM-SQL11 is the schedule name. Box 3: 72 hours

4320 minutes equals 72 hours.

Note: @backup_retention_period=] backup_retention_period

Isthe length of time, in minutes, to retain the log backup file in the backup directory on the primary server. backup_retention_period is int, with no default, and cannot be NULL.

Box 4: 15 minutes.

[@freq_subday_type =] freq_subday_type

Specifies the units for freq_subday_interval. freq_subday_typeis int, with a default of 0, and can be one of these values.

Here it is 4, which means minutes.

[@freq_subday_interval =] freq_subday_interval

The number of freq_subday_type periods to occur between eachexecution of a job. freq_subday_intervalis int, with a default of 0.

Note: Interval should be longer than 10 seconds. freq_subday_interval is ignored in those cases where freq_subday_type is equal to 1.

Here it is 15. References:

<https://docs.microsoft.com/en-us/sql/relational-databases/system-stored-procedures/sp-add-schedule-transact-sql> <https://docs.microsoft.com/en-us/sql/relational-databases/system-stored-procedures/sp-add-log-shipping-primary>

NEW QUESTION 45

- (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 49

- (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 implement a process for importing data into WingDB.

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

Actions

Perform a full backup of the database, and enable the bulk-logged recovery model.

Back up the tail of the transaction log.

Drop any clustered indexes from the tables being imported into.

Perform a full backup of the database and enable the simple recovery model.

Import the data.

Rebuild any indexes on the tables being imported into.

Drop any nonclustered indexes from the tables being imported into.

Answer Area



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1: Perform a full backup of the database and enable the bulk-logged recovery model. Not: Simple recovery model.

With the Simple recovery model we cannot minimize data loss. Step 2: Import the data

Step 3: Backup the tail of the transaction log.

For databases that use full and bulk-logged recovery, database backups are necessary but not sufficient. Transaction log backups are also required.

Note: Three recovery models exist: simple, full, and bulk-logged. Typically, a database uses the full recovery model or simple recovery model. A database can be switched to another recovery model at any time.

References:

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

NEW QUESTION 52

- (Exam Topic 1)

You administer a Microsoft SQL Server 2016 instance that has several SQL Server Agent jobs configured. SQL Server Agent jobs fail, the error messages returned by the job steps are truncated.

The following error message is an example of the truncated error message:

"Executed as user CONTOSO\ServiceAccount. ...0.4035.00 for 64-bit Copyright (C) Microsoft Corp 1984-2011. All rights reserved. Started 63513 PM Error 2012-06-23 183536.87 Code 0XC001000E Source UserImport Description Code 0x00000000 Source Log Import Activity Descript... The package execution fa... The step failed."

You need to ensure that all the details of the job step failures are retained for SQL Server Agent jobs. What should you do?

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

Answer: D

Explanation:

When you have a multiple-step job, then log all steps against a single file. Check the 'Append output to existing file' checkbox for all steps in the job that execute after the initial step. This results in a log file with all of the job steps from the last job execution. Each time the first step executes (each time the job is kicked-off) the file will be overwritten, so we have a record of the last set of output.

References: <https://www.mssqltips.com/sqlservertip/1411/verbose-sql-server-agent-logging/>

NEW QUESTION 57

- (Exam Topic 1)

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

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

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

Transact-SQL segments	Answer Area
<pre>Use Contoso GO CREATE USER Contoso\SalesGroup FOR LOGIN Contoso\SalesGroup</pre>	<div style="display: flex; align-items: center; justify-content: center;"> <div style="margin-right: 20px;">   </div> <div style="margin-left: 20px;">   </div> </div>
<pre>ALTER ROLE db_ddladmin ADD MEMBER Contoso\SalesGroup GRANT VIEW SEFINITION ON Sales.- SalesInvoice TO Contoso\SalesGroup</pre>	
<pre>use master CREATE LOGIN Contoso\SalesGroup FROM WINDOWS GO</pre>	
<pre>GRANT EXECUTE ON Sales.ReviewInvoice TO Contoso\SalesGroup GRANT SELECT ON Sales.SalesInvoice TO Contoso\SalesGroup</pre>	
<pre>use master CREATE LOGIN Contoso\ContosoSearch WITH PASSWORD=N'Pa\$\$w0rd' GO</pre>	
<pre>GRANT EXECUTE ON Sales.ReviewInvoice TO Contoso\SalesGroup GRANT VIEW DEFINITION ON Sales.SalesIn- voice TO Contoso\SalesGroup</pre>	
<pre>GRANT EXECUTE, SELECT ON Sales.Review- Invoice TO Contoso\SalesGroup</pre>	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Transact-SQL segments	Answer Area
<pre>Use Contoso GO CREATE USER Contoso\SalesGroup FOR LOGIN Contoso\SalesGroup</pre>	<pre>use master CREATE LOGIN Contoso\ContosoSearch WITH PASSWORD=N' Pa\$\$w0rd' GO</pre>
<pre>ALTER ROLE db_ddladmin ADD MEMBER Contoso\SalesGroup GRANT VIEW SEFINITION ON Sales.- SalesInvoice TO Contoso\SalesGroup</pre>	<pre>Use Contoso GO CREATE USER Contoso\SalesGroup FOR LOGIN Contoso\SalesGroup</pre>
<pre>use master CREATE LOGIN Contoso\SalesGroup FROM WINDOWS GO</pre>	<pre>GRANT EXECUTE, SELECT ON Sales.Review- Invoice TO Contoso\SalesGroup</pre>
<pre>GRANT EXECUTE ON Sales.ReviewInvoice TO Contoso\SalesGroup GRANT SELECT ON Sales.SalesInvoice TO Contoso\SalesGroup</pre>	<hr style="border-top: 1px dashed red;"/>
<pre>use master CREATE LOGIN Contoso\ContosoSearch WITH PASSWORD=N' Pa\$\$w0rd' GO</pre>	<pre>GRANT EXECUTE, SELECT ON Sales.Review- Invoice TO Contoso\SalesGroup</pre>
<pre>GRANT EXECUTE ON Sales.ReviewInvoice TO Contoso\SalesGroup GRANT VIEW DEFINITION ON Sales.SalesIn- voice TO Contoso\SalesGroup</pre>	<hr style="border-top: 1px dashed red;"/>
<pre>GRANT EXECUTE, SELECT ON Sales.Review- Invoice TO Contoso\SalesGroup</pre>	<hr style="border-top: 1px dashed red;"/>

NEW QUESTION 60

- (Exam Topic 1)

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

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

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 62

- (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 Database Engine instance.

You need to return information about processes that are not idle, that belong to a specific user, or that belong to a specific session.

What should you use?

A. Activity Monitor

B. sp_who3

C. SQL Server Management Studio (SSMS) Object Explorer

D. SQL Server Data Collector

E. SQL Server Data Tools (SSDT)

F. SQL Server Configuration Manager

Answer: B

Explanation:

Use sp_who3 to first view the current system load and to identify a session of interest. You should execute the query several times to identify which session id is most consuming the system resources.

Parameters

sp_who3 null - who is active;

sp_who3 1 or 'memory' - who is consuming the memory;

sp_who3 2 or 'cpu' - who has cached plans that consumed the most cumulative CPU (top 10); sp_who3 3 or 'count' - who is connected and how many sessions it has;

sp_who3 4 or 'idle' - who is idle that has open transactions;

sp_who3 5 or 'tempdb' - who is running tasks that use tempdb (top 5); and, sp_who3 6 or 'block' - who is blocking.

NEW QUESTION 63

- (Exam Topic 1)

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

The reporting database is only updated from the transactional database.

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

Database	Description
Transactional database	<p>Recovery model:</p> <ul style="list-style-type: none"> • Full <p>Backup schedule:</p> <ul style="list-style-type: none"> • Full database backup: midnight, daily • Differential database backup: on the hour, every two hours starting at 02:00 hours except at 00:00 hours • Log backup: every half hour, except at the times of full and differential backups
Reporting database	<p>Recovery model:</p> <ul style="list-style-type: none"> • Simple <p>Backup schedule:</p> <ul style="list-style-type: none"> • Full database backup: 01:00 hours daily • Differential database backup: 13:00 hours daily <p>Data updates:</p> <ul style="list-style-type: none"> • Changes in data are updated from the transactional database to the reporting database at 00:30 hours and at 12:30 hours • The update takes 15 minutes

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

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

Answer: H

Explanation:

At restore time, before you restore a differential backup, you must restore its base. Then, restore only the most recent differential backup to bring the database forward to the time when that differential backup was created. Typically, you would restore the most recent full backup followed by the most recent differential backup that is based on that full backup.

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

NEW QUESTION 65

- (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 67

- (Exam Topic 1)

A company has an on-premises Microsoft SQL Server environment and Microsoft Azure SQL Database instanced. The environments host several customer databases.

You host a local database and a Stretch database that has a table named Members for one specific customer. You need to provide the customer with information about the space used in the databases.

In the table below, identify the query that provides the required information for each database. NOTE: Make only one selection in each column.

Answer Area

Query	Local database	Stretch database
EXEC sp_spaceused @updateusage = N'TRUE'	<input type="radio"/>	<input type="radio"/>
EXEC sp_spaceused N'Company.Members'	<input type="radio"/>	<input type="radio"/>
EXEC sp_spaceused N'Company.Members', @mode = 'REMOTE_ONLY'	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

Query	Local database	Stretch database
EXEC sp_spaceused @updateusage = N'TRUE'	<input type="radio"/>	<input type="radio"/>
EXEC sp_spaceused N'Company.Members'	<input checked="" type="radio"/>	<input type="radio"/>
EXEC sp_spaceused N'Company.Members', @mode = 'REMOTE_ONLY'	<input type="radio"/>	<input checked="" type="radio"/>

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 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 configure log shipping to backup and restore by using shared folder.

Does this meet the goal?

- A. Yes
- B. No

Answer: B

NEW QUESTION 71

- (Exam Topic 1)

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

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

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

Transact-SQL segments	Answer Area
WHERE ((sqlserver.data-base_id>(4)) AND (severity>(15)))	
(ACTION(sqlserver.client_app_name, sqlserver.data-base_id,sqlserver.session_id)	
ALTER EVENT SESSION Contoso1 ON SERVER	
) GO	
ADD EVENT sqlserver.error_reported	
ADD TARGET sqlserver.error_reported	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1: ALTER EVENT SESSION Contoso1 ON SERVER

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

Step 5:) GO

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

```
CREATE EVENT SESSION [error_trap] ON SERVER
ADD EVENT sqlserver.error_reported (
ACTION
(package0.collect_system_time,package0.last_error,sqlserver.client_app_name,sqlserver.client_hostname,sqlserver.plan_handle,sqlserver.query_hash,sqlserver.session_id,sqlserver.sql_text,sqlserver.tsqf_frame,sqlserver)
WHERE ([severity]>10)
)
ADD TARGET package0.event_file (
SET filename=N'D:\Program Files\Microsoft SQL Server\MSSQL11.MSSQLSERVER\MSSQL\XE\error_trap.xel'
) WITH (
STARTUP_STATE=OFF
) GO
```

References:

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

NEW QUESTION 72

- (Exam Topic 1)

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

You need to ensure that an OLTP database that uses a storage area network (SAN) remains available if any of the servers fail.

You also need to minimize the amount of storage used by the database. Which configuration should you use?

- A. Two servers configured in different data centers SQL Server Availability Group configured in Synchronous-Commit Availability Mode One server configured as an Active Secondary
- B. SQL Server that includes an application database configured to perform transactional replication

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

Answer: H

Explanation:

A Windows Server Failover Cluster (WSFC) is a group of independent servers that work together to increase the availability of applications and services. SQL Server takes advantage of WSFC services and capabilities to support Always On availability groups and SQL Server Failover Cluster Instances.

References:

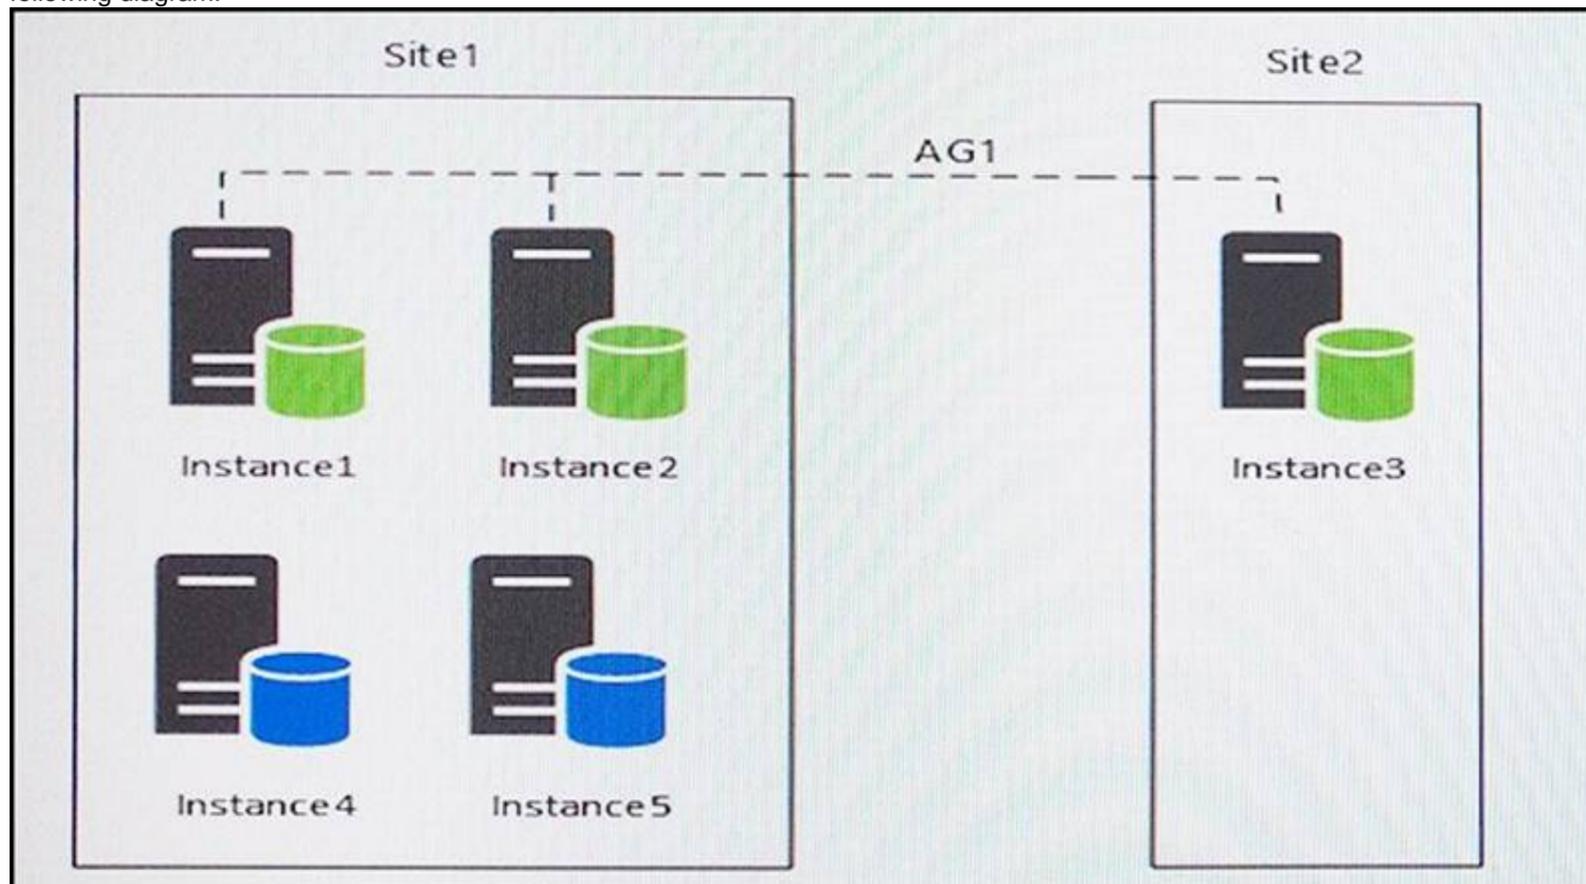
<https://docs.microsoft.com/en-us/sql/sql-server/failover-clusters/windows/windows-server-failover-clustering-ws>

NEW QUESTION 75

- (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_resource_semaphores
- Sys.dm_exec_session_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 80

- (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 83

- (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 85

- (Exam Topic 2)

Overview

General Overview

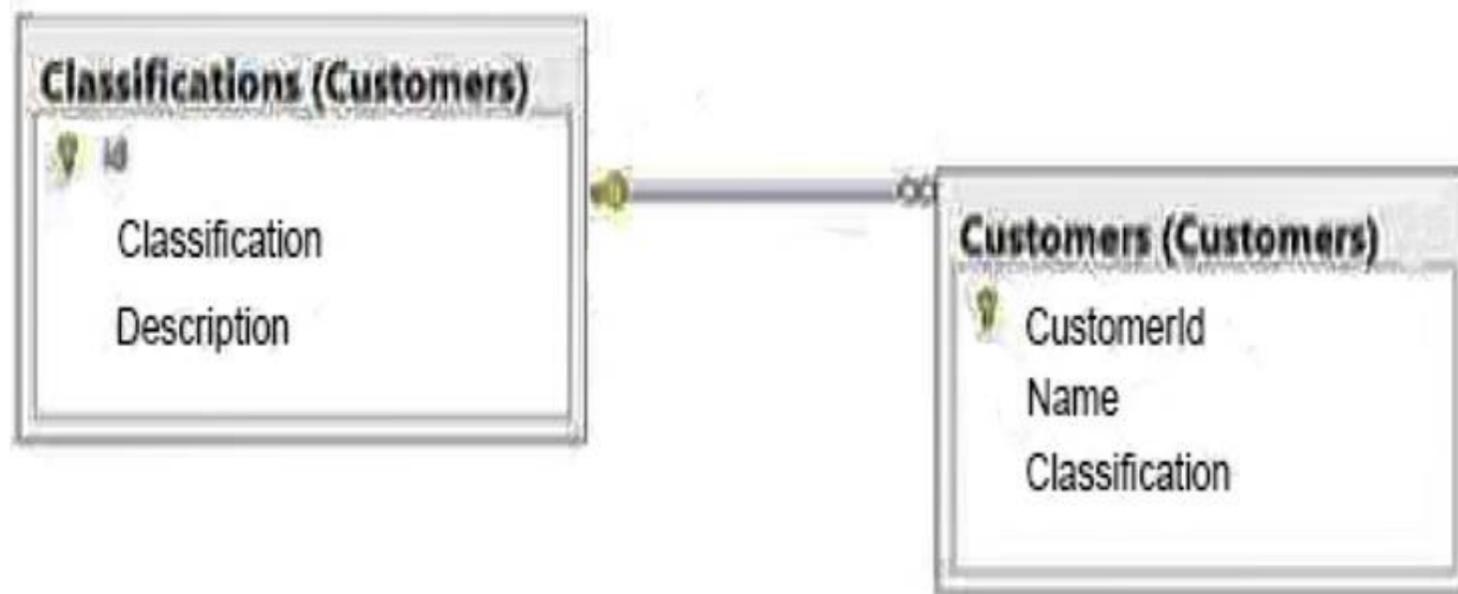
ADatum Corporation has offices in Miami and Montreal.

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

Databases

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

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



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

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

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

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

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

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

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

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

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

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

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

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

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

The nested stored procedures are never called directly. Design Requirements

Data Recovery

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

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

Classification Changes

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

Storage

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

Error Handling

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

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

You need to recommend a solution to minimize the amount of time it takes to execute USP_5. What should you include in the recommendation?

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

Answer: A

Explanation:

Scenario:

A stored procedure named USP_5 changes data in multiple databases. Security checks are performed each time USP_5 accesses a database.

- Cross-database ownership chaining occurs when a procedure in one database depends on objects in another database. A cross-database ownership chain works in the same way as ownership chaining within a single database, except that an unbroken ownership chain requires that all the object owners are mapped to the same login account. If the source object in the source database and the target objects in the target databases are owned by the same login account, SQL Server does not check permissions on the target objects.

NEW QUESTION 88

- (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 93

- (Exam Topic 2)

Overview

General Overview

ADatum Corporation has offices in Miami and Montreal.

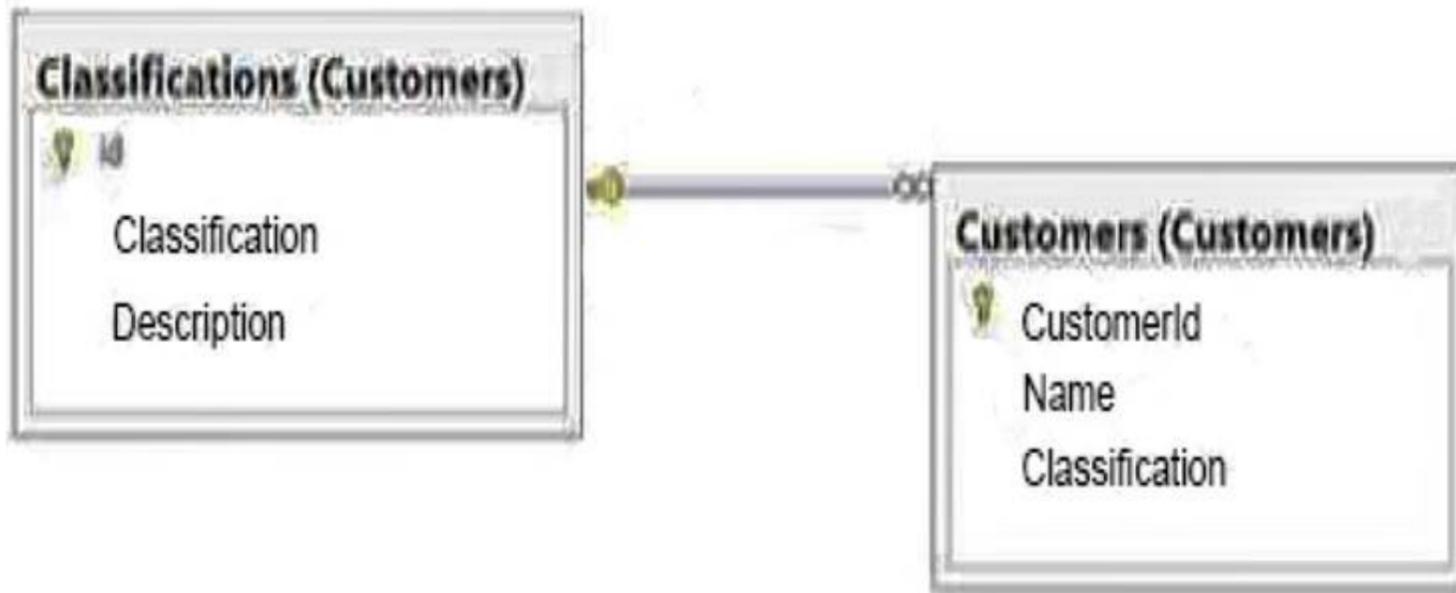
The network contains a single Active Directory forest named adatum.com. The offices connect to each other by using a WAN link that has 5-ms latency. A. Datum standardizes its database platform by using SQL Server 2014 Enterprise edition.

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 97

- (Exam Topic 2)

Overview

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

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

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

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

Requirements Planned Changes

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

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

Database1 will also contain a stored procedure named usp_UpdateOrderDetails.

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

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

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

Inventory will contain over 100 GB of data.

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

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

Business Requirements

You have the following requirements:

Costs for new licenses must be minimized.

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

Development effort must be minimized whenever possible.

The storage requirements for databases must be minimized.

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

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

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

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

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

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

You need to recommend a solution for the deployment of SQL Server 2014. The solution must meet the business requirements. What should you include in the recommendation?

- A. Create a new instance of SQL Server 2014 on the server that hosts the SQL Server 2008 instance.
- B. Upgrade the existing SQL Server 2008 instance to SQL Server 2014.
- C. Deploy two servers that have SQL Server 2014 installed and implement Failover Clustering.
- D. Deploy two servers that have SQL Server 2014 installed and implement database mirroring.

Answer: C

Explanation:

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

NEW QUESTION 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 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 101

- (Exam Topic 2)

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

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

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

You need to achieve this goal by granting the minimum permissions required.
What should you do? (Each correct answer presents a complete solution. Choose all that apply.)

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

Answer: BC

Explanation:

References:

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

NEW QUESTION 104

- (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 107

- (Exam Topic 2)

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

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

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

Answer: D

Explanation:

Troubleshooting deadlocks

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

Configures deadlock event tracking for the session.

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

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

NEW QUESTION 110

- (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 server
- B. Filter queries that have a Duration value of more than 1,000.
- C. Use sp_configure to set a value for blocked process threshold
- 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 111

- (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 115

- (Exam Topic 2)

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

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

Answer: D

NEW QUESTION 117

- (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 database reporting solution that meets the business requirements. What should you include in the recommendation?

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

Answer: A

Explanation:

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

NEW QUESTION 120

- (Exam Topic 2)

You are building a stored procedure for a SQL Azure database. The procedure will add multiple rows to a table. You need to design the stored procedure to meet the following requirements:

If any of the new rows violates a table constraint, then no further additions must be attempted and all changes made by the stored procedure must be discarded.
If any errors occur, a row must be added to an audit table, and the original error must be returned to the caller of the stored procedure.
What should you include in the design?

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

Answer: D

Explanation:

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

NEW QUESTION 125

- (Exam Topic 2)

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

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

Answer: A

Explanation:

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

NEW QUESTION 129

- (Exam Topic 2)

Overview

General Overview

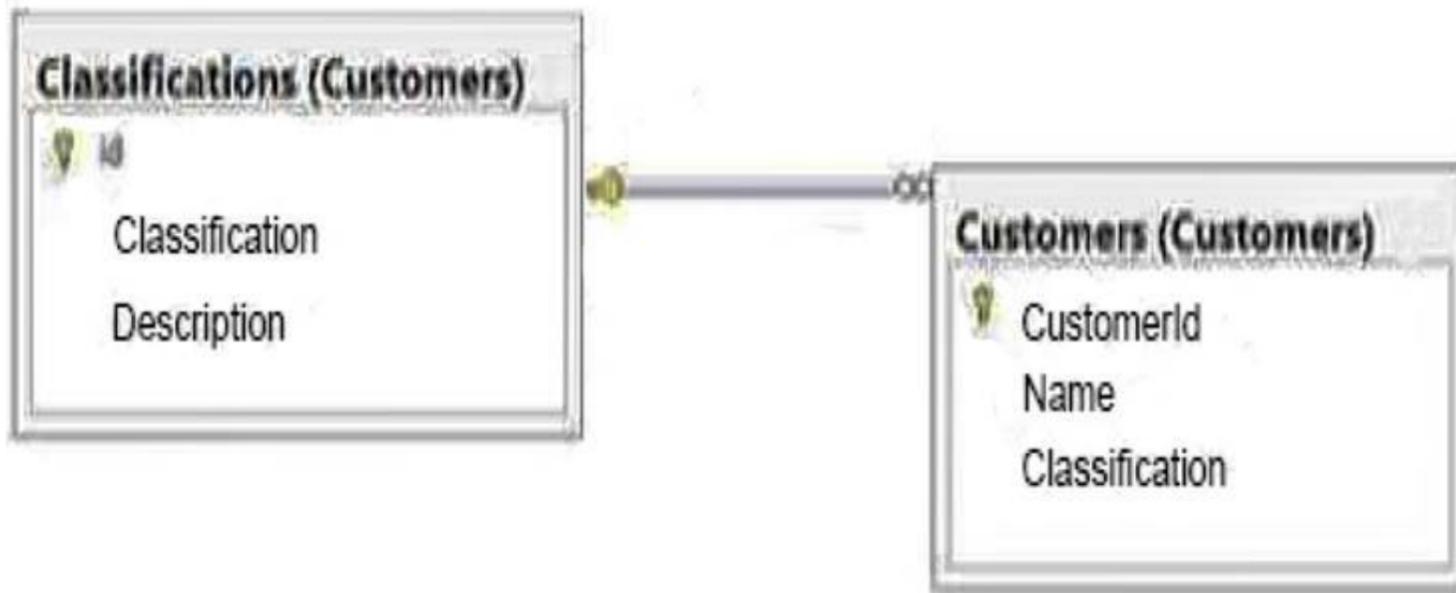
ADatum Corporation has offices in Miami and Montreal.

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

Databases

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

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



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

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

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

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

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

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

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

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

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

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

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

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

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

The nested stored procedures are never called directly. Design Requirements

Data Recovery

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

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

Classification Changes

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

Storage

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

Error Handling

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

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

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

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

Answer: A

Explanation:

Scenario:

- You must be able to recover data from the Inventory database if a storage failure occurs. You have a Recovery Point Objective (RPO) of one hour.

- A. Datum Corporation has offices in Miami and Montreal.

- SQL Server Log shipping allows you to automatically send transaction log backups from a primary database on a primary server instance to one or more secondary databases on separate secondary server instances. The transaction log backups are applied to each of the secondary databases individually.

NEW QUESTION 133

- (Exam Topic 2)

You are planning to deploy a database to Windows Azure SQL Database.

You need to design a stored procedure to update rows. The stored procedure must meet the following requirements:

If more than one row is updated, an error must be raised to the application and the update must be discarded.

The stored procedure must be designed to maximize concurrency.

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

Actions	Work Area
Raise an error in a catch block	
Commit the transaction in a finally block	
Read the @@ROWCOUNT system variable	
Perform the update in a try block	
Raise an error and roll back the transaction if the row count is less than 1	
Issue a SELECT statement to count the number of rows	
Set the isolation level to serializable	
Begin an explicit transaction	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Note:

- Read Committed is SQL Server's default isolation level.
- @@ROWCOUNT returns the number of rows affected by the last statement.
- Using TRY...CATCH in a transaction

The following example shows how a TRY...CATCH block works inside a transaction. The statement inside the TRY block generates a constraint violation error.

```

- BEGIN TRANSACTION;
- BEGIN TRY

DELETE FROM Production.Product
WHERE ProductID = 980;
END TRY
BEGIN CATCH
SELECT
ERROR_NUMBER() AS ErrorNumber
,ERROR_SEVERITY() AS ErrorSeverity
,ERROR_STATE() AS ErrorState
,ERROR_PROCEDURE() AS ErrorProcedure
,ERROR_LINE() AS ErrorLine
,ERROR_MESSAGE() AS ErrorMessage;
IF @@TRANCOUNT > 0
ROLLBACK TRANSACTION;
END CATCH;
IF @@TRANCOUNT > 0
COMMIT TRANSACTION;
GO
    
```

NEW QUESTION 138

- (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 141

- (Exam Topic 2)

You administer a single server that contains a Microsoft SQL Server 2016 default instance.

You plan to install a new application that requires the deployment of a database on the server. The application login requires sysadmin permissions.

You need to ensure that the application login is unable to access other production databases. What should you do?

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

Answer: B

Explanation:

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

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

NEW QUESTION 146

- (Exam Topic 2)

You plan to deploy SQL Server 2014. You are designing two stored procedures named SP1 and SP2 that have the following requirements:

- Prevent data read by SP1 from being modified by other active processes.

- Prevent SP2 from performing dirty reads.

You need to recommend the isolation level for each stored procedure.

The solution must maximize concurrency. Which isolation levels should you recommend? To answer, drag the appropriate isolation level to the correct stored procedure in the answer area.

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

SP1 – repeatable read; SP2 – read committed

- REPEATABLE READ

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

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

NEW QUESTION 149

- (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 151

- (Exam Topic 2)

You install a Microsoft SQL Server 2016 instance.

The instance will store data extracted from two databases running on Windows Azure SQL Database. You hire a data steward to perform interactive data cleansing and ad hoc querying and updating of the database.

You need to ensure that the data steward is given the correct client tools to perform these tasks. Which set of tools should you install?

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

Answer: B

NEW QUESTION 153

- (Exam Topic 2)

You want to simulate read, write, checkpoint, backup, sort, and read-ahead activities for your organization's SQL Server 2016 deployment.

Which of the following tools would you use to accomplish this goal?

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

Answer: B

Explanation:

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

References:

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

NEW QUESTION 156

- (Exam Topic 2)

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

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

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

Answer: B

Explanation:

Transparent Data Encryption (TDE) is a special case of encryption using a symmetric key. TDE encrypts an entire database using that symmetric key called the

database encryption key. The database encryption key is protected by other keys or certificates which are protected either by the database master key or by an asymmetric key stored in an EKM module.

SQL Server provides the following mechanisms for encryption:

- Transact-SQL functions
- Asymmetric keys
- Symmetric keys
- Certificates
- Transparent Data Encryption

NEW QUESTION 157

- (Exam Topic 2)

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

You grant DOMAIN\JrDBAs access to the SQL Server.

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

- View all databases.
 - View the server state.
 - Assign GRANT, DENY, and REVOKE permissions on logins.
- You need to add DOMAIN\JrDBAs to the server role.
 You also need to provide the least level of privileges necessary.
 Which SQL statement or statements should you use? Choose all that apply.

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

Answer: BCF

NEW QUESTION 160

- (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 reduces the time it takes to import the supplier data. What should you include in the recommendation?

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

Answer: C

Explanation:

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

NEW QUESTION 165

- (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. Server1 is located in the Seattle office. Server2 is located in the Montreal office. The latency of the WAN link between the Montreal office and the Seattle office is more than 200 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. Synchronous manual failover
- B. Synchronous automatic failover
- C. Asynchronous automatic failover
- D. Asynchronous manual failover

Answer: D

NEW QUESTION 166

- (Exam Topic 2)

Overview

Application Overview

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

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

The new version will use SQL Server 2014.

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

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

Tables

The current database schema contains a table named OrderDetails.

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

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

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

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

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

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

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

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

Customer Problems Installation Issues

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

Index Fragmentation Issues

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

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

Backup Issues

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

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

Missing Data Issues

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

Query Performance Issues

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

Design Requirements

File Storage Requirements

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

Data Recovery Requirements

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

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

Concurrency Requirements

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

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

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

Answer: B

Explanation:

- Backup Issues

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

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

NEW QUESTION 171

- (Exam Topic 2)

You have a server that has SQL Server 2014 installed. The server contains 100 user databases.

You need to recommend a backup solution for the user databases. The solution must meet the following requirements:

Perform a transaction log backup every hour.

Perform a full backup of each database every week.

Perform a differential backup of each database every day.

Ensure that new user databases are added automatically to the backup solution.

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

- A. A maintenance plan
- B. SQL Server Agent jobs
- C. Policy-Based Management

D. A Data Definition Language (DDL) trigger

Answer: A

Explanation:

Maintenance plans create a workflow of the tasks required to make sure that your database is optimized, regularly backed up, and free of inconsistencies. Maintenance plans can be created to perform the following task (among others): Back up the database and transaction log files. Database and log backups can be retained for a specified period. This lets you create a history of backups to be used if you have to restore the database to a time earlier than the last database backup. You can also perform differential backups.

NEW QUESTION 174

- (Exam Topic 2)

You administer a Microsoft SQL Server 2016 database that contains a table named AccountTransaction. You discover that query performance on the table is poor due to fragmentation on the

IDX_AccountTransaction_AccountCode non-clustered index.

You need to defragment the index.

You also need to ensure that user queries are able to use the index during the defragmenting process. Which Transact-SQL batch should you use?

- A. ALTER INDEX IDX_AccountTransaction_AccountCode ON AccountTransaction.AccountCode REORGANIZE
- B. ALTER INDEX ALL ON AccountTransaction REBUILD
- C. ALTER INDEX IDX_AccountTransaction_AccountCode ON AccountTransaction.AccountCode REBUILD
- D. CREATE INDEX IDXAccountTransactionAccountCode ON AccountTransaction.AccountCode WITH DROP EXISTING

Answer: A

NEW QUESTION 179

- (Exam Topic 3)

Overview

General Overview

ADatum Corporation has offices in Miami and Montreal.

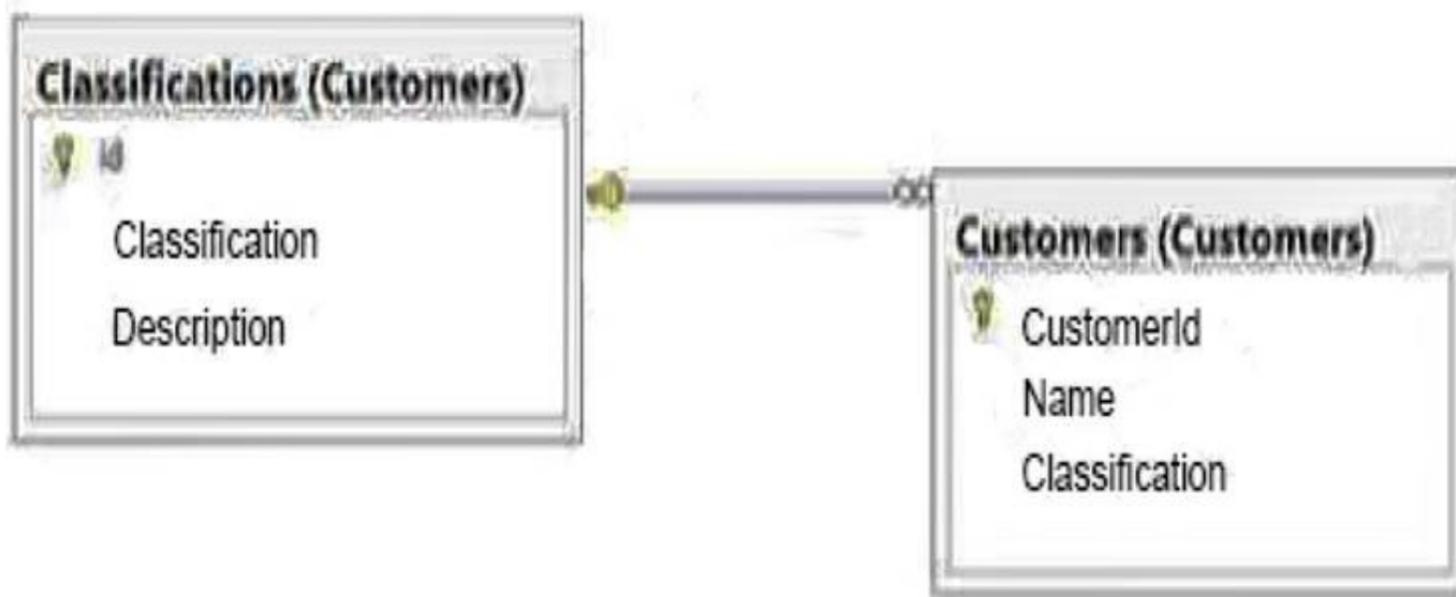
The network contains a single Active Directory forest named adatum.com. The offices connect to each other by using a WAN link that has 5-ms latency. A. Datum standardizes its database platform by using SQL Server 2014 Enterprise edition.

Databases

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

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

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



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

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

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

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

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

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

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

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

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

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

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

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

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

The nested stored procedures are never called directly. Design Requirements

Data Recovery

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

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

Classification Changes

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

Storage

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

Error Handling

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

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

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

- A. A view
- B. A temporary table
- C. A table variable
- D. A function

Answer: A

Explanation:

- 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.

NEW QUESTION 181

- (Exam Topic 3)

You need to recommend a backup process for data warehouse database. The solution must meet the following requirements:

Ensure that if a hardware failure occurs, you can bring the database online without losing more than 24 hours of transactions.

Minimize the amount of administrative effort required to restore any lost data.

Minimize the space used by transaction logs.

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

Actions	Work Area
Perform a weekly full backup of the database	
Ship the logs to a secondary server	
Set the database to the simple recovery model	
Set the database to the full recovery model	
Create a database snapshot	
Perform a backup of the transaction log every hour	
Perform a differential backup of the database every night	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Set the database to the simple recovery model; Perform a weekly full backup of the database;

Perform a differential backup of the database every night Note:

- Simple recovery model No log backups.
- Full recovery model Requires log backups

NEW QUESTION 182

- (Exam Topic 3)

General Overview

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

Requirements Leaffield Solutions has been asked by a customer to develop a web-based Enterprise Resource Planning and Management application. The new

application will eventually replace a desktop application that the customer is currently using. The current application will remain in use while the users are trained to use the new webbased application.

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

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

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

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

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

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

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

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

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

The GenInfo table is used for reports.

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

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

Current System

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

SQL Servers

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

Design Requirements

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

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

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

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

Deadlocks must be analyzed with the use of Deadlock Graphs.

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

Software licensing and database storage costs must be minimized.

Development effort must be minimized.

The Tempdb databases must be monitored for insufficient free space.

Failed authentication requests must be logged.

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

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

SPUpdateSalesInfo queries data in the OrderTotals table the second time.

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

What should your solution include?

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

Answer: A

NEW QUESTION 183

- (Exam Topic 3)

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

Rebuild the indexes by using a new fill factor.

Run a custom T-SQL command.

Back up the databases.

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

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

Answer: C

Explanation:

Maintenance plans create a workflow of the tasks required to make sure that your database is optimized, regularly backed up, and free of inconsistencies.

NEW QUESTION 187

- (Exam Topic 3)

You manage database servers in a high security environment. Your company has the following auditing requirements:

SQL Server auditing must be enabled on all server instances.

Auditing results must be logged in the Windows Security event log.

A routine review shows that a SQL Server is writing auditing entries to Windows Application event log. You change the SQL Server audit target to Windows Security event log. SQL Server auditing stops working on the server.

You need to ensure that the server meets the auditing requirements.

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

- A. Grant the manage auditing and security log permission to the SQL Server service account.
- B. Grant the generate security audits permission on the SQL Server service account.

- C. Update Windows security policy to audit object access.
- D. Restart the SQL Server Agent service.

Answer: BC

Explanation:

There are two key requirements for writing SQL Server server audits to the Windows Security log:

The audit object access setting must be configured to capture the events.

The account that the SQL Server service is running under must have the generate security audits permission to write to the Windows Security log.

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/security/auditing/write-sql-server-audit-events-to-the-s>

NEW QUESTION 188

- (Exam Topic 3)

General Overview

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

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

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

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

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

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

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

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

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

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

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

The GenInfo table is used for reports.

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

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

Current System

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

SQL Servers

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

Design Requirements

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

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

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

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

Deadlocks must be analyzed with the use of Deadlock Graphs.

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

Software licensing and database storage costs must be minimized.

Development effort must be minimized.

The Tempdb databases must be monitored for insufficient free space.

Failed authentication requests must be logged.

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

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

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

What should you recommend?

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

Answer: D

NEW QUESTION 192

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

70-764 Practice Exam Features:

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

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