

Microsoft

Exam Questions 70-764

Administering a SQL Database Infrastructure (beta)



NEW QUESTION 1

- (Exam Topic 1)

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

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

What should you do?

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

Answer: C

Explanation:

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

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

NEW QUESTION 2

- (Exam Topic 1)

You administer 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 3

- (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.

A company has an on-premises Microsoft SQL Server environment and Microsoft Azure SQL Database instances. The environment hosts a customer database named DB1.

Customers connect to hosted database instances by using line-of-business applications. Developers connect by using SQL Server Management Studio (SSMS).

You need to grant the developers permission to alter views for DB1 while following the principle of least privilege.

Which permission should you grant?

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

Answer: A

Explanation:

To execute ALTER VIEW, at a minimum, ALTER permission on OBJECT is required.

Members of the db_ddladmin fixed database role can run any Data Definition Language (DDL) command in a database.

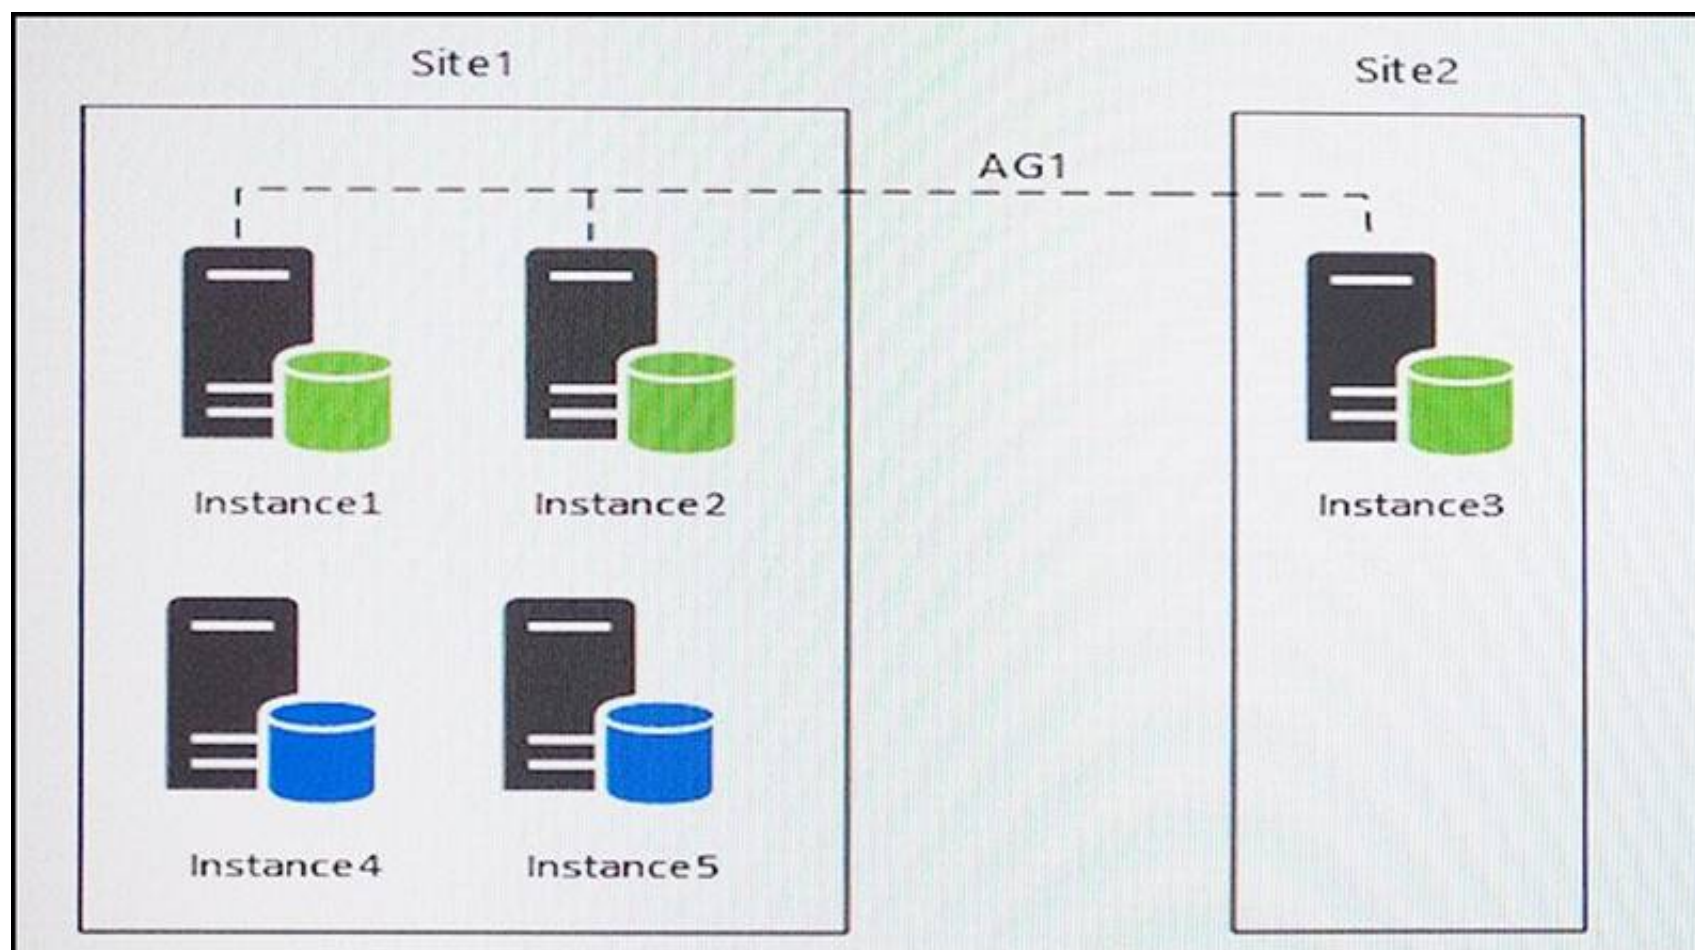
References: [https://technet.microsoft.com/en-us/library/ms190667\(v=sql.90\).aspx](https://technet.microsoft.com/en-us/library/ms190667(v=sql.90).aspx)

NEW QUESTION 4

- (Exam Topic 1)

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

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



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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Actions

Drop all nonclustered indexes on the target table.

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

Run the nightly import process.

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

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

Drop all clustered and nonclustered indexes on the target table.

Recreate any dropped indexes on the target table.

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

Answer Area



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

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

All databases use the full recovery model.

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

NEW QUESTION 5

- (Exam Topic 1)

You administer a Microsoft SQL Server 2016 database named Orders.

Users report that during peak usage periods, certain operations are taking more time than expected. Your initial analysis suggests that blocking is the cause.

You need to gather more data to be able to determine which processes are being blocked and to identify the root cause.

What should you do?

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

Answer: B

NEW QUESTION 6

- (Exam Topic 1)

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

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

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

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

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

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

Answer: A

Explanation:

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

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

NEW QUESTION 7

- (Exam Topic 1)

You plan to install a Microsoft SQL Server 2016 instance.

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

Store Excel workbooks on the file system.

Access the workbooks through Transact-SQL.

Include the workbooks in database backups.

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

Which feature should you use?

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

Answer: B

NEW QUESTION 8

- (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 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)

You administer a Microsoft SQL Server 2016 server.

When transaction logs grow, SQL Server must send an email message to the database administrators. You need to configure SQL Server to send the email messages.

What should you configure?

- A. SQL Mail
- B. An Extended Events session
- C. Alerts and operators in SQL Server Agent
- D. Policies under Policy-Based Management

Answer: C

Explanation:

Operators are aliases for people or groups that can receive electronic notification when jobs have completed or alerts have been raised. The SQL Server Agent service supports the notification of administrators through operators. Operators enable notification and monitoring capabilities of SQL Server Agent.

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/database-mail/configure-sql-server-agent-mail-to-use-d>

NEW QUESTION 10

- (Exam Topic 1)

You install Microsoft SQL Server 2016 on a new server.

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

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

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

Answer: A

NEW QUESTION 15

- (Exam Topic 1)

You use Microsoft SQL Server 2016 to write code for a transaction that contains several statements.

There is high contention between readers and writers on several tables used by your transaction. You need to minimize the use of the tempdb space.

You also need to prevent reading queries from blocking writing queries. Which isolation level should you use?

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

Answer: C

Explanation:

For most applications, read committed isolation using row versioning is recommended over snapshot isolation for the following reasons:

It consumes less tempdb space than snapshot isolation. Etc.

References: <https://msdn.microsoft.com/en-us/library/ms188277.aspx>

NEW QUESTION 19

- (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>

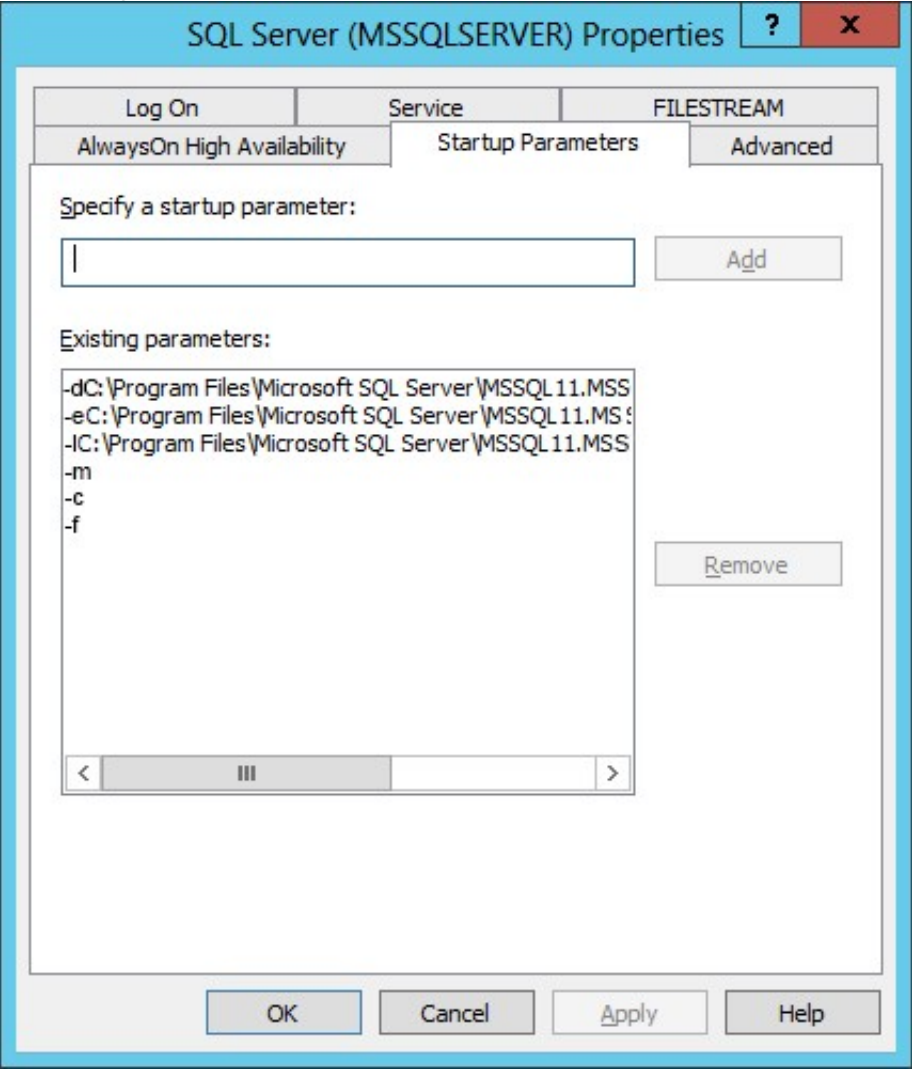
NEW QUESTION 23

- (Exam Topic 1)

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

MSG_AUDIT_FORCED_SHUTDOWN

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



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

Answer Area

In which user mode will the SQL Server instance start?

	▼
single-user	
multi-user	
restricted-user	

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

	▼
public	
serveradmin	
sysadmin	
setupadmin	

B. Not Mastered

Answer: A

Explanation:

Box 1: single-user

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

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

References:

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

NEW QUESTION 26

- (Exam Topic 1)

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

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

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

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

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

Answer: G

Explanation:

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

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

NEW QUESTION 28

- (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 31

- (Exam Topic 1)

You administer a Microsoft SQL Server 2016 database.

You have a SQL Server Agent job instance that runs using the service account. You have a job step within the job that requires elevated privileges.

You need to ensure that the job step can run using a different user account. What should you use?

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

Answer: D

NEW QUESTION 33

- (Exam Topic 1)

You have a database. The existing backups for the database and their corresponding files are listed in the following table.

Backup type	Backup date/time	File name
Full	05/02/2016 21:00	Full_20160502_2100.bak
Transaction log	05/03/2016 6:00	Log_20160503_0600.trn
Transaction log	05/03/2016 9:00	Log_20160503_0900.trn
Differential	05/03/2016 12:00	Diff_20160503_1200.bak
Transaction log	05/03/2016 15:00	Log_20160503_1500.trn
Differential	05/03/2016 17:00	Diff_20160503_1700.bak
Transaction log	05/03/2016 19:00	Log_20160503_1900.trn

You purchase a new server. You must restore the database to the new server. You need to restore the data to the most recent time possible. Which three files should you restore in sequence? To answer, move the appropriate files from the list of files to the answer area and arrange them in the correct order.

Files

Log_20160503_0600.trn

Log_20160503_1500.trn

Full_20160502_2100.bak

Log_20160503_1900.trn

Log_20160503_0900.trn

Diff_20160503_1200.bak

Diff_20160503_1700.bak

Answer Area



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1: Full.

Start with the full backup.

Step 2: Diff_20160503_1700.bak

Followed by the most recent differential backup. Step 3: Log_20160503_1900.bak

And finally the most recent log backup (the only log backup done after the most recent differential backup).

References:

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

NEW QUESTION 35

- (Exam Topic 1)

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

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

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

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

Answer: F

Explanation:

References:

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

NEW QUESTION 39

- (Exam Topic 1)

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

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

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

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

A. Yes

B. No

Answer: B

Explanation:

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

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

References:

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

NEW QUESTION 43

- (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 47

- (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 48

- (Exam Topic 1)

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

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

- A. Two servers configured in the same data center A primary server configured to perform log-shipping every 10 minutes A backup server configured as a warm standby

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

Answer: E

Explanation:

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

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

Data changes infrequently.

It is acceptable to have copies of data that are out of date with respect to the Publisher for a period of time.

Replicating small volumes of data.

A large volume of changes occurs over a short period of time.

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

NEW QUESTION 53

- (Exam Topic 1)

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

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

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

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

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

Answer: C

Explanation:

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

References:

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

NEW QUESTION 55

- (Exam Topic 1)

You deploy a Microsoft SQL Server instance to support a global sales application. The instance includes the following tables: TableA and TableB.

TableA is a partitioned table that uses an incrementing integer number for partitioning. The table has millions of rows in each partition. Most changes to the data in TableA affect recently added data. The UPDATE STATISTICS for TableA takes longer to complete than the allotted maintenance window.

Thousands of operations are performed against TableB each minute. You observe a large number of Auto Update Statistics events for TableB.

You need to address the performance issues with each table.

In the table below, identify the action that will resolve the issues for each table. NOTE: Make only one selection in each column.

Answer Area

Action

TableA

TableB

Run the following Transact-SQL statement:

`SET AUTO_UPDATE_STATISTICS_ASYNC ON`

☐
☐

Run the following Transact-SQL statement:

`SET AUTO_UPDATE_STATISTICS OFF`

☐
☐

Run the following Transact-SQL statement and then recreate all indexes and statistics using the INCREMENTAL keyword:

`SET AUTO_CREATE_STATISTICS on (INCREMENTAL = ON)`

☐
☐

Run the sp_updatestats procedure instead of the following Transact-SQL statement:

`UPDATE STATISTICS`

☐
☐

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Table A: Auto_update statistics off

Table A does not change much. There is no need to update the statistics on this table. Table B: SET AUTO_UPDATE_STATISTICS_ASYNC ON

You can set the database to update statistics asynchronously: ALTER DATABASE YourDBName

SET AUTO_UPDATE_STATISTICS_ASYNC ON

If you enable this option then the Query Optimizer will run the query first and update the outdated statistics afterwards. When you set this option to OFF, the Query Optimizer will update the outdated statistics before compiling the query. This option can be useful in OLTP environments

References:

<https://www.mssqltips.com/sqlservertip/2766/sql-server-auto-update-and-auto-create-statistics-options/>

NEW QUESTION 56

- (Exam Topic 1)

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

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

What should you do?

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

Answer: B

NEW QUESTION 60

- (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

Is the 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_type is 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 each execution of a job. freq_subday_interval is 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 62

- (Exam Topic 1)

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

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

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

Answer Area

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

Box 1: Primary server instance.

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

Box 2: Secondary server instance

Each of the three secondary server instances runs its own copy job to copy the primary log-backup file to its own local destination folder. copy job: A SQL Server Agent job that copies the backup files from the primary server to a configurable destination on the secondary server and logs history on the secondary server and the monitor server. When log shipping is enabled on a database, the job category "Log Shipping Copy" is created on each secondary server in a log shipping configuration.

Box 3: Secondary server instance.

Each secondary server instance runs its own restore job to restore the log backup from the local destination folder onto the local secondary database. restore job: A SQL Server Agent job that restores the copied backup files to the secondary databases. It logs history on the local server and the monitor server, and deletes old files and old history information. When log shipping is enabled on a database, the job category "Log Shipping Restore" is created on the secondary server instance.

References: <https://docs.microsoft.com/en-us/sql/database-engine/log-shipping/about-log-shipping-sql-server>

NEW QUESTION 65

- (Exam Topic 1)

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

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

H. Two servers configured in the same data center A primary server configured to perform log-shipping every 10 minutes A backup server configured as a warm standby

Answer: A

Explanation:

SQL Server supports the following heterogeneous scenarios for transactional and snapshot replication:

Publishing data from SQL Server to non- SQL Server Subscribers.

Publishing data to and from Oracle has some restrictions.

References:

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

NEW QUESTION 67

- (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:

User 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 72

- (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 73

- (Exam Topic 1)

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

Currently, Server01/HA is the primary replica.

You need to ensure that the following requirements are met:

Backup operations occur on Server02/HA.

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

Backup operations do not occur on Server01/HA.

How should you configure HaContoso?

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

Answer: B

Explanation:

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

Backup Priority (Lowest=1, Highest=100)

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

References:

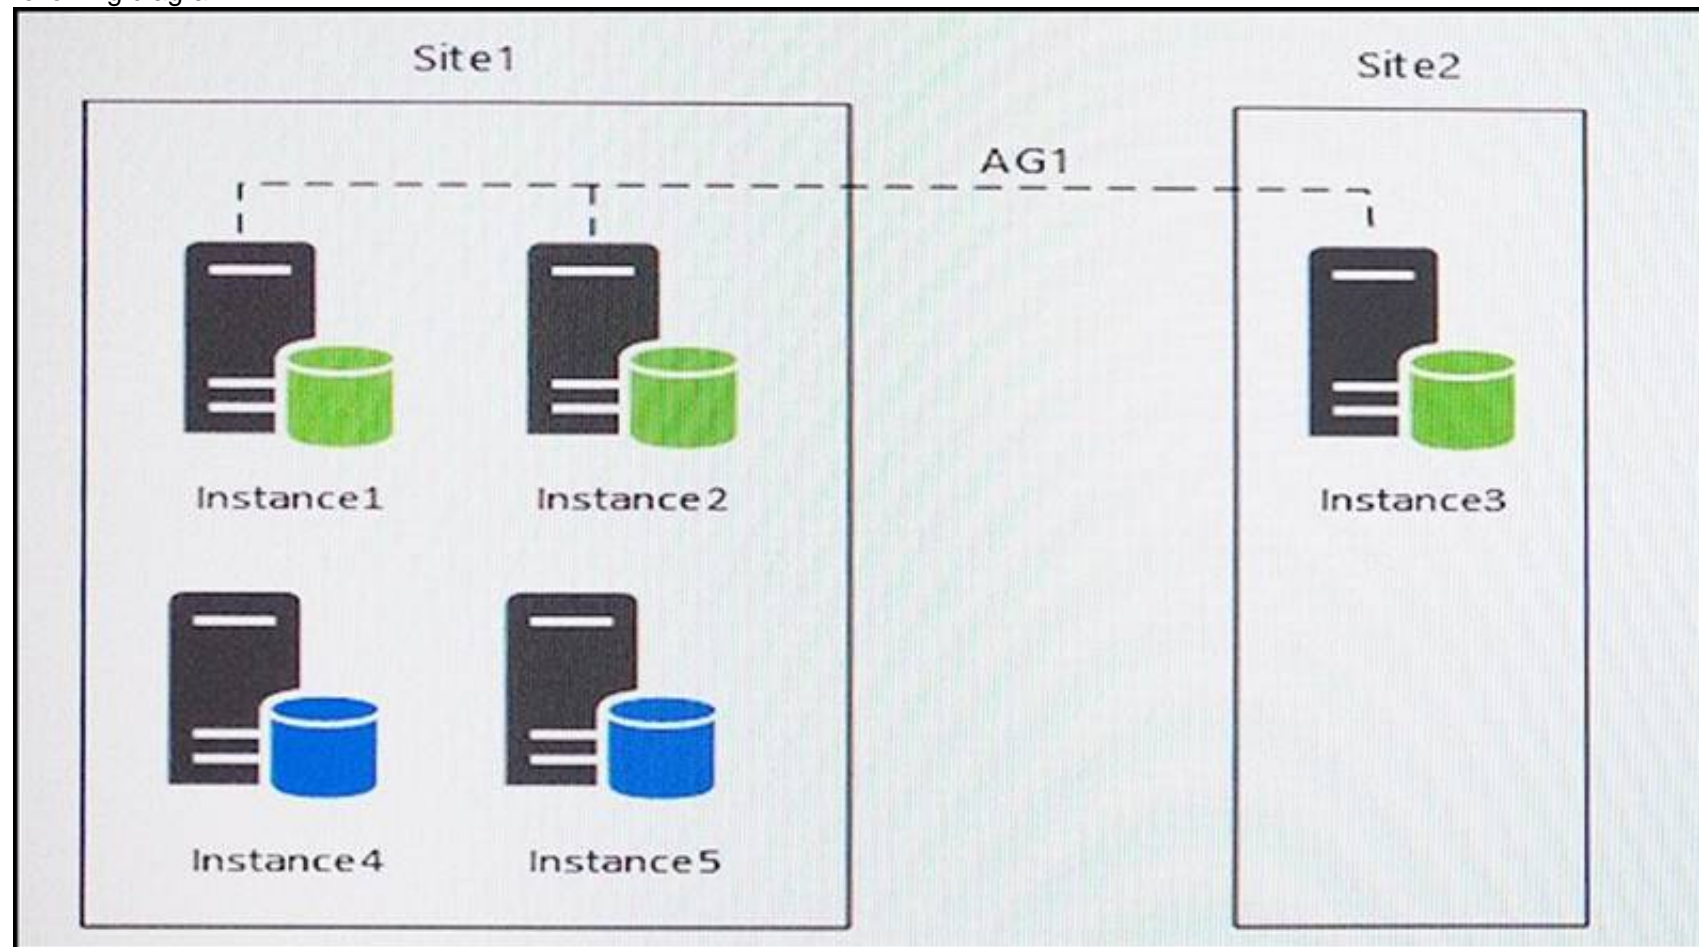
<https://docs.microsoft.com/en-us/sql/database-engine/availability-groups/windows/configure-backup-on-availab>

NEW QUESTION 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 create the connection strings for the operations and reporting systems.
 In the table below, identify the option that must be specified in each connection string. NOTE: Make only one selection in each column.

Answer Area

Option	Reporting system	Operations system
Connect to a Listener using ApplicationIntent=ReadOnly.	<input type="radio"/>	<input type="radio"/>
Connect to the current primary replica SQL instance using ApplicationIntent=ReadOnly.	<input type="radio"/>	<input type="radio"/>
Connect to any current read-only replica SQL instance.	<input type="radio"/>	<input type="radio"/>
Connect to a Listener.	<input type="radio"/>	<input type="radio"/>
Connect to the current primary replica SQL instance.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:
 Reporting system: Connect to any current read-only replica instance
 We configure Read-OnlyAccess on an Availability Replica. We select Read-intent only. Only read-only connections are allowed to secondary databases of this replica. The secondary database(s) are all available for read access.
 From Scenario: Reporting system: This solution accesses data inDB1with a login that is mapped to a database user that is a member of the db_datareader role. The user has EXECUTE permissions on the database. Queries make no changes to the data. The queries must be load balanced over variable read-only replicas.
 Operating system: Connect to the current primary replica SQL instance
 By default, both read-write and read-intent access are allowed to the primary replica and no connections are allowed to secondary replicas of an Always On availability group.
 From scenario: Operations system: This solution accesses data inDB1with a login that is mapped to a database user that is a member of the db_datareader and db_datawriter roles. The user has EXECUTE permissions on the database. Queries from the operations system will perform both DDL and DML operations.
 References:
<https://docs.microsoft.com/en-us/sql/database-engine/availability-groups/windows/configure-read-only-access-o>

NEW QUESTION 78

- (Exam Topic 1)
 You are configuring a new Microsoft SQL Server Always On Availability Group. You plan to configure a shared network location at \\DATA-C11\\SQL.
 You need to create an availability group listener named AGL1 on port 1433.
 In which order should you perform the actions? To answer, move all actions from the list of actions to the answer area and arrange them in the correct order.

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1: Launch the Failover Cluster Manager and..

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

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

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

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

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

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

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

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

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

NEW QUESTION 82

- (Exam Topic 1)

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

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

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

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

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

What should you use?

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

Answer: E

Explanation:

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

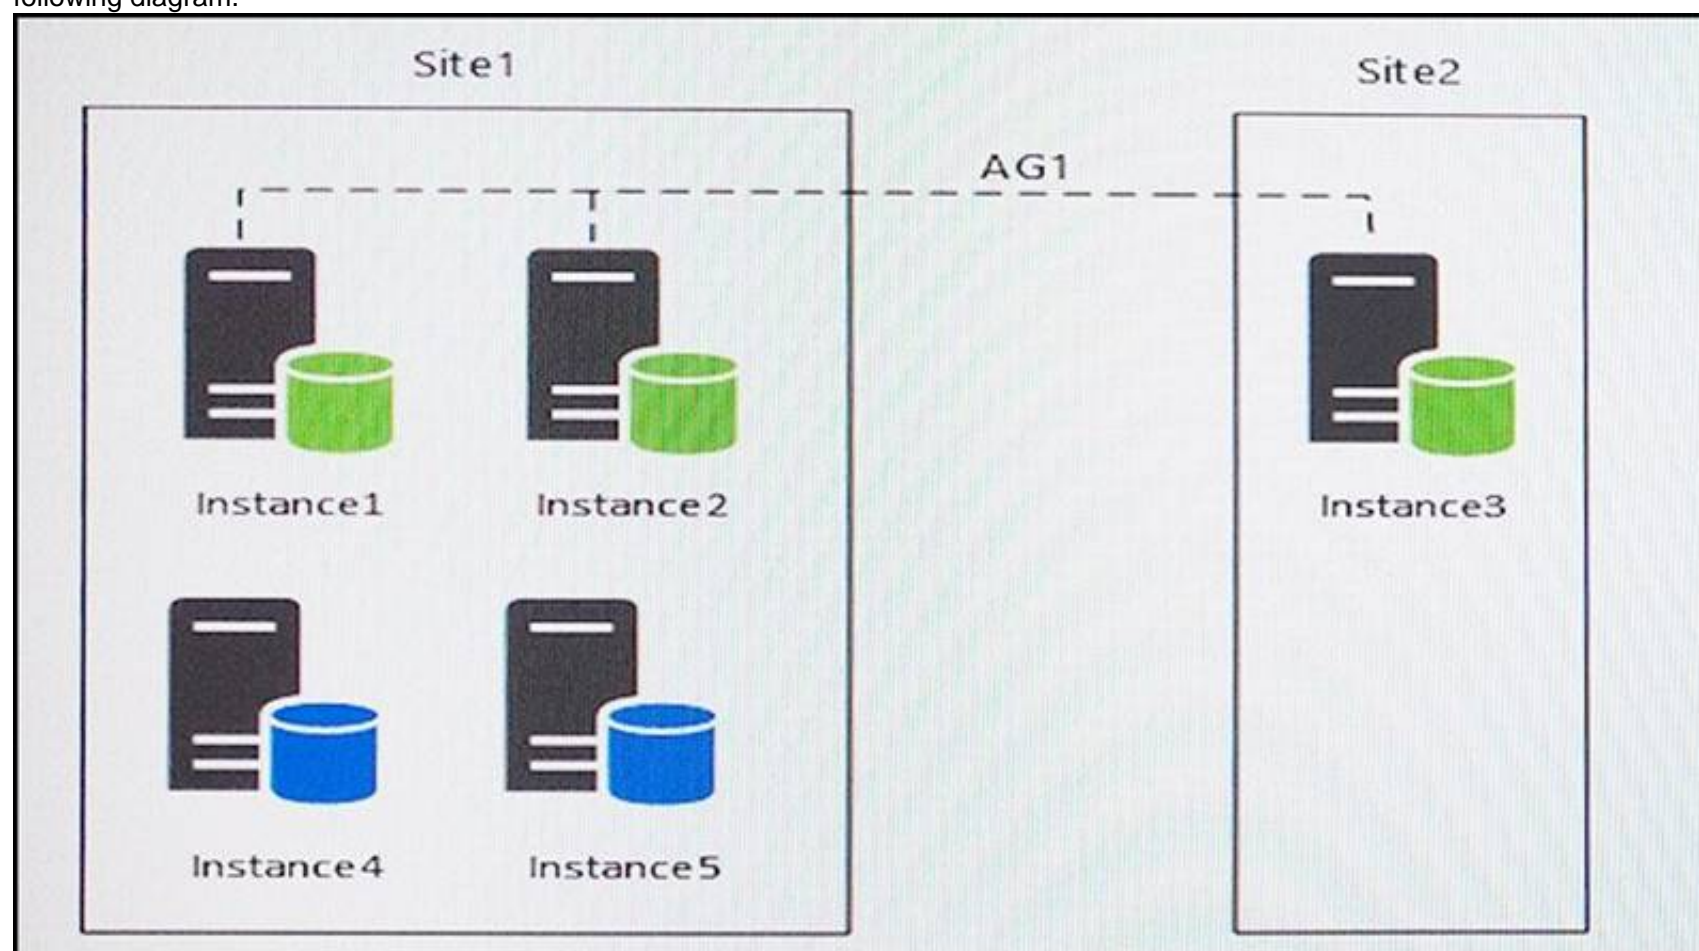
References: [https://msdn.microsoft.com/en-us/library/hh272686\(v=vs.103\).aspx](https://msdn.microsoft.com/en-us/library/hh272686(v=vs.103).aspx)

NEW QUESTION 84

- (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 89

- (Exam Topic 2)

Overview

General Overview

ADatum Corporation has offices in Miami and Montreal.

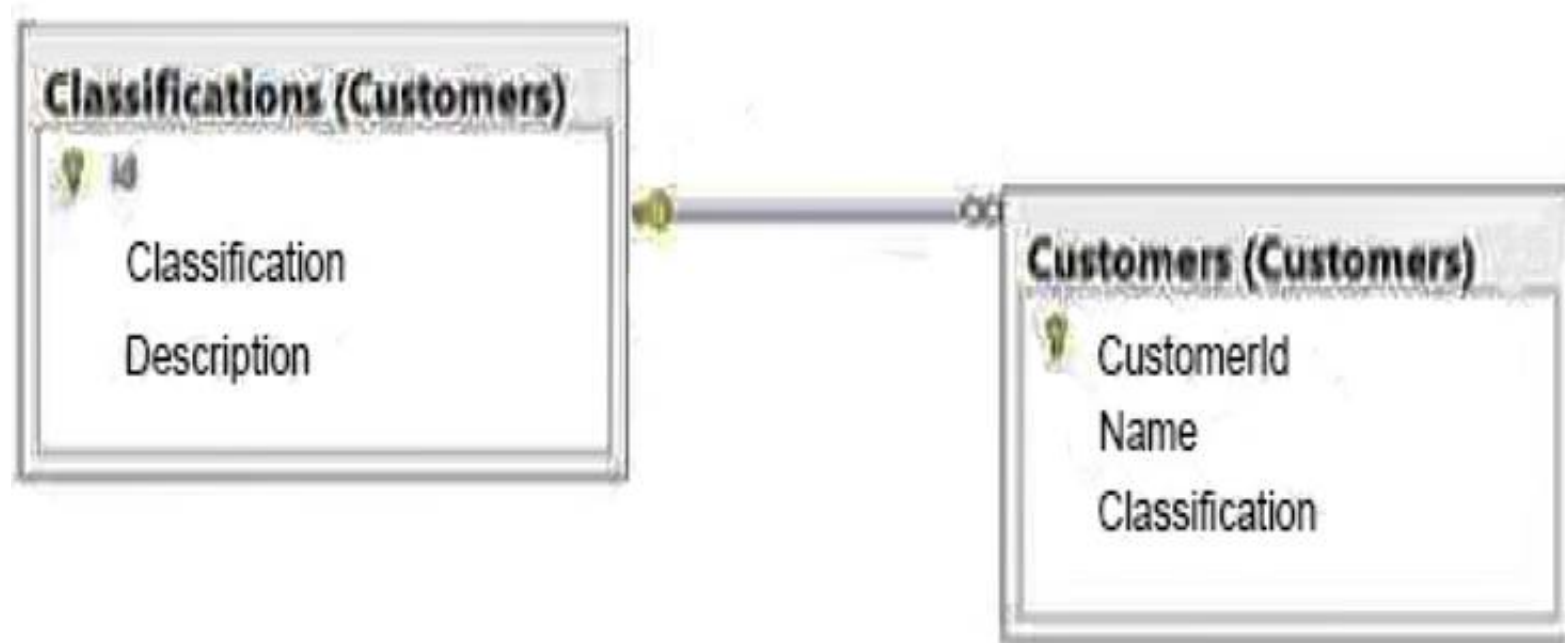
The network contains a single Active Directory forest named adatum.com. The offices connect to each other by using a WAN link that has 5-ms latency. A. Datum standardizes its database platform by using SQL Server 2014 Enterprise edition.

Databases

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

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

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



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

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

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

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

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

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

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

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

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

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

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

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

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

The nested stored procedures are never called directly. Design Requirements

Data Recovery

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

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

Classification Changes

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

Storage

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

Error Handling

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

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

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

Which change should you recommend?

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

Answer: A

Explanation:

- Scenario: A stored procedure named USP_3 is used to update prices. USP_3 is composed of several UPDATE statements called in sequence from within a transaction. Currently, if one of the UPDATE statements fails, the stored procedure continues to execute.

- When SET XACT_ABORT is OFF, in some cases only the Transact-SQL statement that raised the error is rolled back and the transaction continues processing.

NEW QUESTION 90

- (Exam Topic 2)

Overview

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

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

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

will use the application.

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

Requirements Planned Changes

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

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

Database1 will also contain a stored procedure named usp_UpdateOrderDetails.

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

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

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

Inventory will contain over 100 GB of data.

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

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

Business Requirements

You have the following requirements:

Costs for new licenses must be minimized.

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

Development effort must be minimized whenever possible.

The storage requirements for databases must be minimized.

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

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

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

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

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

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

You need to recommend a solution for 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 91

- (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 96

- (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 97

- (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 101

- (Exam Topic 2)

You administer a Microsoft SQL Server 2016 server.

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

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

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

Answer: CD

Explanation:

How To Enable Instant File Initialization

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

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

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

References:

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

NEW QUESTION 103

- (Exam Topic 2)

You administer a Microsoft SQL Server 2016 instance.

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

What should you do?

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

Answer: B

NEW QUESTION 108

- (Exam Topic 2)

Overview

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

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

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

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

Requirements Planned Changes

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

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

Database1 will also contain a stored procedure named usp_UpdateOrderDetails.

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

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

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

Inventory will contain over 100 GB of data.

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

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

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

Business Requirements

You have the following requirements:

Costs for new licenses must be minimized.

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

Development effort must be minimized whenever possible.

The storage requirements for databases must be minimized.

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

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

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

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

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

The encryption strategy must minimize changes to the databases and to the applications. You need to recommend an isolation level for usp_UpdateOrderDetails. Which isolation level should you recommend?

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

Answer: B

Explanation:

- Scenario: Database1 will also contain a stored procedure named usp_UpdateOrderDetails. The stored procedure is used to update order information. The stored procedure queries the Orders table twice each time the procedure executes. The rows returned from the first query must be returned on the second query unchanged along with any rows added to the table between the two read operations.

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

NEW QUESTION 109

- (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 113

- (Exam Topic 2)

Overview

Application Overview

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

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

The new version will use SQL Server 2014.

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

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

Tables

The current database schema contains a table named OrderDetails.

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

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

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

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

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

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

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

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

Customer Problems Installation Issues

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

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

Index Fragmentation Issues

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

Backup Issues

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

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

Missing Data Issues

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

Query Performance Issues

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

Design Requirements

File Storage Requirements

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

Data Recovery Requirements

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

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

Concurrency Requirements

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

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

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

Answer: BE

Explanation:

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

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

NEW QUESTION 114

- (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 115

- (Exam Topic 2)

You administer a Microsoft SQL Server 2016 database.

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

Which server-level audit action group should you use?

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

Answer: D

Explanation:

SERVER_ROLE_MEMBER_CHANGE_GROUP

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

References:

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

NEW QUESTION 116

- (Exam Topic 2)

You plan to deploy a database to SQL Azure. You are designing two stored procedures named USP_1 and USP_2 that have the following requirements:

Prevent data read by USP_1 from being modified by other active processes.

Allow USP_2 to perform dirty reads.

You need to recommend the isolation level for the stored procedures. The solution must maximize concurrency.

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

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

- A. Mastered
 B. Not Mastered

Answer: A

Explanation:

SP1 – repeatable read; SP2 – read uncommitted Note:

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

NEW QUESTION 117

- (Exam Topic 2)

You are implementing a SQL Server 2016 five-node failover cluster. You need to choose a quorum configuration. Which configuration should you use?

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

Answer: D

Explanation:

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

NEW QUESTION 118

- (Exam Topic 2)

You have a database named DB1.

You plan to create a stored procedure that will insert rows into three different tables. Each insert must use the same identifying value for each table, but the value must increase from one invocation of the stored procedure to the next.

Occasionally, the identifying value must be reset to its initial value. You need to design a mechanism to hold the identifying values for the stored procedure to use. What should you do? More than one answer choice may achieve the goal. Select the BEST answer.

- A. Create a sequence object that holds the next value in the sequence. Retrieve the next value by using the stored procedure. Reset the value by using an ALTER SEQUENCE statement as needed.
- B. Create a sequence object that holds the next value in the sequence. Retrieve the next value by using the stored procedure. Increment the sequence object to the next value by using an ALTER SEQUENCE statement. Reset the value as needed by using a different ALTER SEQUENCE statement.
- C. Create a fourth table that holds the next value in the sequence. At the end each transaction, update the value by using the stored procedure. Reset the value as needed by using an UPDATE statement.
- D. Create an identity column in each of the three tables. Use the same seed and the same increment for each table. Insert new rows into the tables by using the stored procedure. Use the DBCC CHECKIDENT command to reset the columns as needed.

Answer: A

Explanation:

- an application can obtain the next sequence number without inserting the row by calling the NEXT VALUE FOR function.
- ALTER SEQUENCE Includes argument: RESTART [WITH <constant>]

The next value that will be returned by the sequence object. If provided, the RESTART WITH value must be an integer that is less than or equal to the maximum and greater than or equal to the minimum value of the sequence object. If the WITH value is omitted, the sequence numbering restarts based on the original CREATE SEQUENCE options.

- CREATE SEQUENCE

Creates a sequence object and specifies its properties. A sequence is a user-defined schema bound object that generates a sequence of numeric values according to the specification with which the sequence was created. The sequence of numeric values is generated in an ascending or descending order at a defined interval and can be configured to restart (cycle) when exhausted.

NEW QUESTION 122

- (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 126

- (Exam Topic 2)

You administer a Microsoft SQL Server 2016 database.

The database contains a Product table created by using the following definition:

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

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

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

Answer: D

NEW QUESTION 131

- (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 133

- (Exam Topic 2)

Overview

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

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

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

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

Requirements Planned Changes

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

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

Database1 will also contain a stored procedure named usp_UpdateOrderDetails.

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

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

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

Inventory will contain over 100 GB of data.

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

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

Business Requirements

You have the following requirements:

Costs for new licenses must be minimized.

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

Development effort must be minimized whenever possible.

The storage requirements for databases must be minimized.

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

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

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

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

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

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

You need to recommend a solution to allow application users to perform UPDATE operations on the database tables. The solution must meet the business requirements.

What should you recommend?

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

Answer: A

Explanation:

- EXECUTE AS Clause (Transact-SQL)

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

NEW QUESTION 137

- (Exam Topic 2)

You administer a Microsoft SQL Server 2016 instance. The instance contains a database that supports a retail sales application.

The application generates hundreds of transactions per second and is online 24 hours per day and 7 days per week. You plan to define a backup strategy for the database.

You need to ensure that the following requirements are met:

No more than 5 minutes worth of transactions are lost.

Data can be recovered by using the minimum amount of administrative effort.

What should you do? Choose all that apply.

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

Answer: BCDE

Explanation:

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

NEW QUESTION 141

- (Exam Topic 2)

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

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

Answer: D

NEW QUESTION 145

- (Exam Topic 2)

You deploy a database by using SQL Server 2012. The database contains a table named Table1.

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

Stores the most recent data from Table1 by using the fastest storage solution possible.

Stores the historical data from Table1 by using a slower storage solution.

What should you recommend?

- A. partitioned views
- B. a database snapshot
- C. change data capture
- D. table partitioning

Answer: D

NEW QUESTION 147

- (Exam Topic 2)

You manage a SQL Server 2014 instance that contains a database named DB1. Users report that some queries to DB1 take longer than expected.

Although most queries run in less than one second, some queries take up to 20 seconds to run. You need to view all of the performance statistics for each database file.

Which method should you use?

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

Answer: C

Explanation:

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

NEW QUESTION 148

- (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 151

- (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 156

- (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 160

- (Exam Topic 2)

You plan to deploy SQL Server 2014.

Your company identifies the following monitoring requirements for the database:

An e-mail message must be sent when a user logs in.

An e-mail message must be sent if CPU utilization exceeds 90 percent

You need to identify which feature meets each monitoring requirement.

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

Policy-Based Management

A SQL Server Agent alert

SQL Server Integration Services (SSIS)

trace flags

triggers

An e-mail message must be sent when a user logs in.

Feature

An e-mail message must be sent if CPU utilization exceeds 90 percent.

Feature

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Policy-Based Management

A SQL Server Agent alert

SQL Server Integration Services (SSIS)

trace flags

triggers

An e-mail message must be sent when a user logs in.

Policy-Based Management

An e-mail message must be sent if CPU utilization exceeds 90 percent.

A SQL Server Agent alert

NEW QUESTION 161

- (Exam Topic 2)

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

You need to write messages to the Application Log when users are added to or removed from a fixed server role in Server01.

What should you create?

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

Answer: G

Explanation:

The SQL Server Audit feature enables you to audit server-level and database-level groups of events and individual events. Audits can have the following categories of actions:

Server-level. These actions include server operations, such as management changes and logon and logoff operations.

Database-level. These actions encompass data manipulation languages (DML) and data definition language (DDL) operations.

Audit-level. These actions include actions in the auditing process.

References:

[http://technet.microsoft.com/en-us/library/cc280663\(v=sql.105\).aspx](http://technet.microsoft.com/en-us/library/cc280663(v=sql.105).aspx)

NEW QUESTION 165

- (Exam Topic 2)

You administer a Microsoft SQL Server database named Sales. The database is 3 terabytes in size. The Sales database is configured as shown in the following table. You discover that all files except Sales_2.ndf are corrupt. You need to recover the corrupted data in the minimum amount of time. What should you do?

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

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

Answer: A

Explanation:

References:

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

NEW QUESTION 170

- (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 173

- (Exam Topic 3)

You work as a Database Administrator (DBA) for a company named ABC.com. The company uses a Microsoft SQL Server 2012 infrastructure. You have a database named CorpDB. CorpDB contains 2TB of data. You plan to import a large amount of data into tables in CorpDB. You want to minimize the size of the transaction log while the data is imported. What should you do?

- A. You should configure the recovery model of the database to Full.
- B. You should configure the recovery model of the database to Bulk-Logged.
- C. You should start a new transaction log file.
- D. You should configure a new filegroup for the existing log file.

Answer: B

NEW QUESTION 174

- (Exam Topic 3)

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

- A. You should configure disc quotas.
- B. You should configure a Dynamic Management View.

- C. You should configure alerts sent by the SQL Server Agent.
D. You should configure a SQL Server Maintenance Plan.

Answer: B

NEW QUESTION 175

- (Exam Topic 3)

General Overview

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

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

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

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

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

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

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

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

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

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

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

The GenInfo table is used for reports.

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

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

Current System

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

SQL Servers

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

Design Requirements

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

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

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

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

Deadlocks must be analyzed with the use of Deadlock Graphs.

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

Software licensing and database storage costs must be minimized.

Development effort must be minimized.

The Tempdb databases must be monitored for insufficient free space.

Failed authentication requests must be logged.

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

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

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

What should you recommend?

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

Answer: D

NEW QUESTION 180

- (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 enable users to modify data in the database tables using UPDATE operations. You need to implement a solution that meets the design requirements.

What should you configure?

A. You should configure a server role.

B. You should configure a database role.

C. You should configure functions that use the EXECUTE AS statement.

D. You should configure stored procedures that use the EXECUTE AS statement.

Answer: D

NEW QUESTION 184

- (Exam Topic 3)

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

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

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

Does this meet the goal?

A. Yes

B. No

Answer: B

Explanation:

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

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

The VIEW SERVER STATE permission.

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

Membership in the sysadmin fixed server role.

The CONTROL SERVER permission.

The ALTER ANY AUDIT permission.

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

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

NEW QUESTION 187

- (Exam Topic 3)

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

A. Resource Governor

B. Windows System Resource Manager

C. Processor affinity

D. I/O affinity

Answer: C

NEW QUESTION 192

- (Exam Topic 3)

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

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

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

Does this meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

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

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

The VIEW SERVER STATE permission.

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

Membership in the sysadmin fixed server role.

The CONTROL SERVER permission.

The ALTER ANY AUDIT permission.

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

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

NEW QUESTION 197

.....

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](#)