

Microsoft

Exam Questions 70-764

Administering a SQL Database Infrastructure (beta)



NEW QUESTION 1

- (Exam Topic 1)

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

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

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

Which feature should you use?

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

Answer: B

NEW QUESTION 2

- (Exam Topic 1)

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

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

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

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

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

What should you do?

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

Answer: F

Explanation:

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

Note: Requirements for Restoring Pages

A page restore is subject to the following requirements:

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

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

NEW QUESTION 3

- (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 4

- (Exam Topic 1)

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

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

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

Which edition should you choose?

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

Answer: B

Explanation:

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

NEW QUESTION 5

- (Exam Topic 1)

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 6

- (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 7

- (Exam Topic 1)

You administer a Microsoft SQL Server 2016 instance that has multiple databases. You have a two-node SQL Server failover cluster. The cluster uses a storage area network (SAN). You discover I/O issues. The SAN is at capacity and additional disks cannot be added. You need to reduce the I/O workload on the SAN at a minimal cost. What should you do?

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

Answer: D

Explanation:

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

NEW QUESTION 8

- (Exam Topic 1)

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

- A. a certificate
- B. an MD5 hash
- C. a DES key
- D. an AES 256-bit key

Answer: D

Explanation:

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

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

Encryptor: A certificate or asymmetric Key

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

NEW QUESTION 9

- (Exam Topic 1)

You plan to migrate the db to azure.

You verify that all objects are valid for azure sql database. You need to ensure that users and logins are migrated to azure.

What should you do?

- A. Use the Copy Database wizard
- B. Use the Database Transfer wizard
- C. Use the SQL Management Studio to deploy the db to azure
- D. Back up the databases from the local server and restore it to azure

Answer: CD

NEW QUESTION 10

- (Exam Topic 1)

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

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

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

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

Answer: F

Explanation:

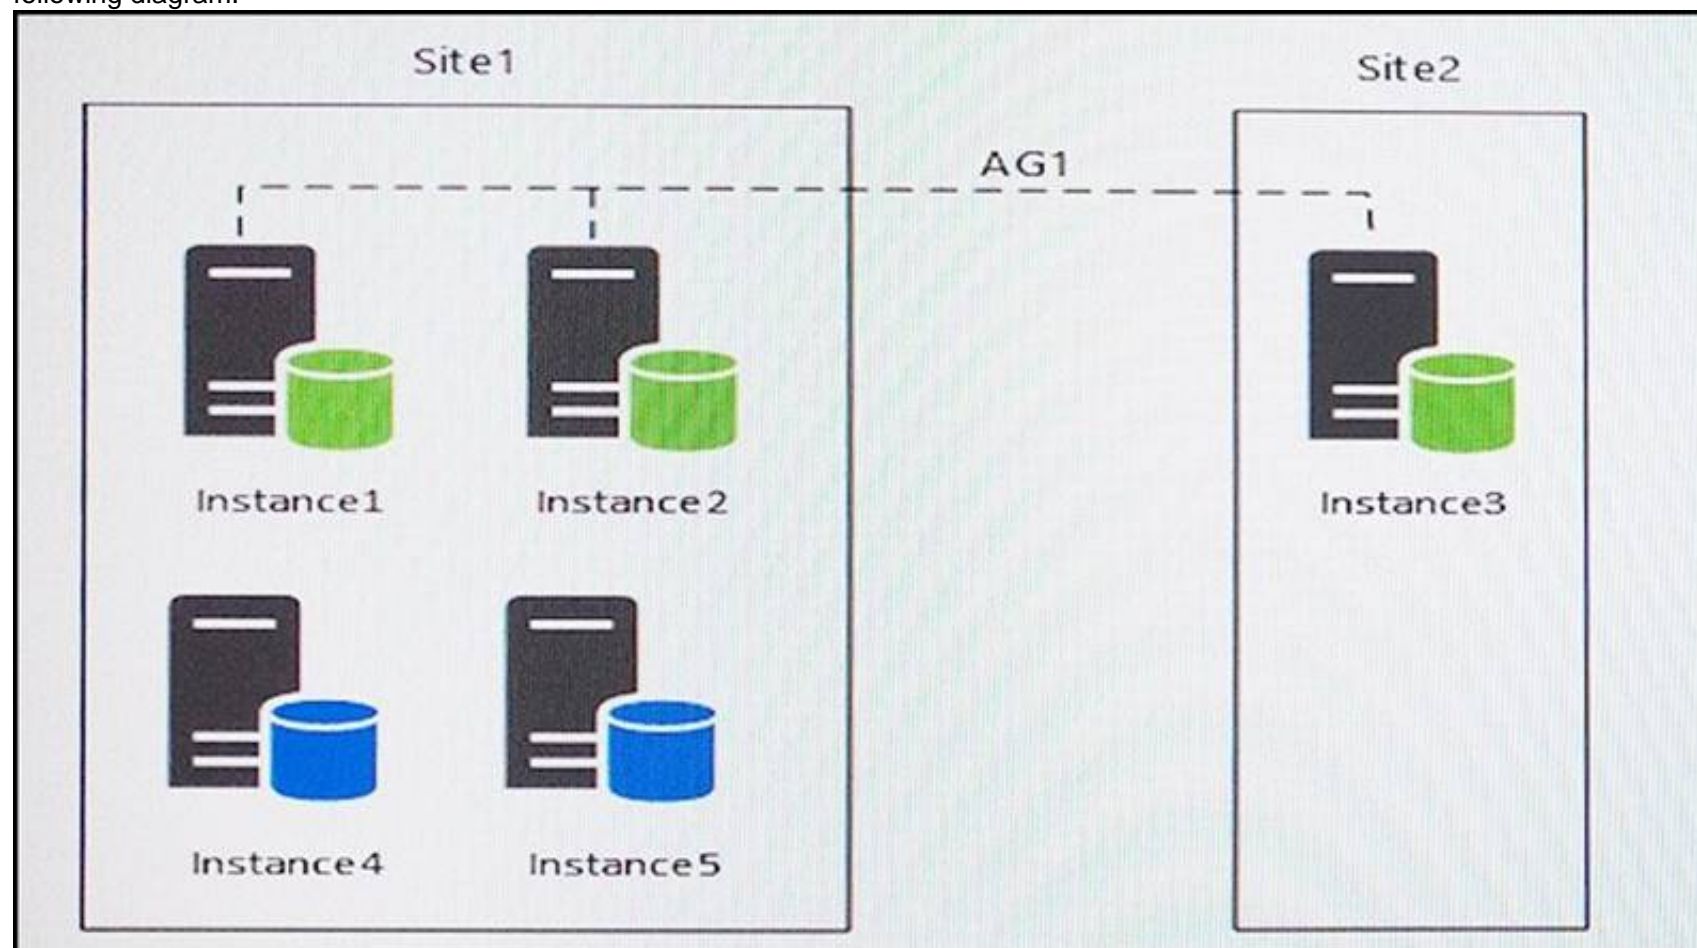
References:

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

NEW QUESTION 10

- (Exam Topic 1)

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



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

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

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

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

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

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

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

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

You should minimize both the time required to restore

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

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

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

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

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

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

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

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

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

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

Answer: D

Explanation:

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

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

References:

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

NEW QUESTION 11

- (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 15

- (Exam Topic 1)

You administer a Microsoft SQL Server 2016 database.

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

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

What should you do?

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

Answer: A

Explanation:

Trapping SQL Server Errors with Extended Events

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

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

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

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

References:

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

NEW QUESTION 17

- (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 20

- (Exam Topic 1)

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

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

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

Transact-SQL segments

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

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

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

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

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

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

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

Answer Area



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Transact-SQL segments

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

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

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

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

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

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

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

Answer Area

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

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

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

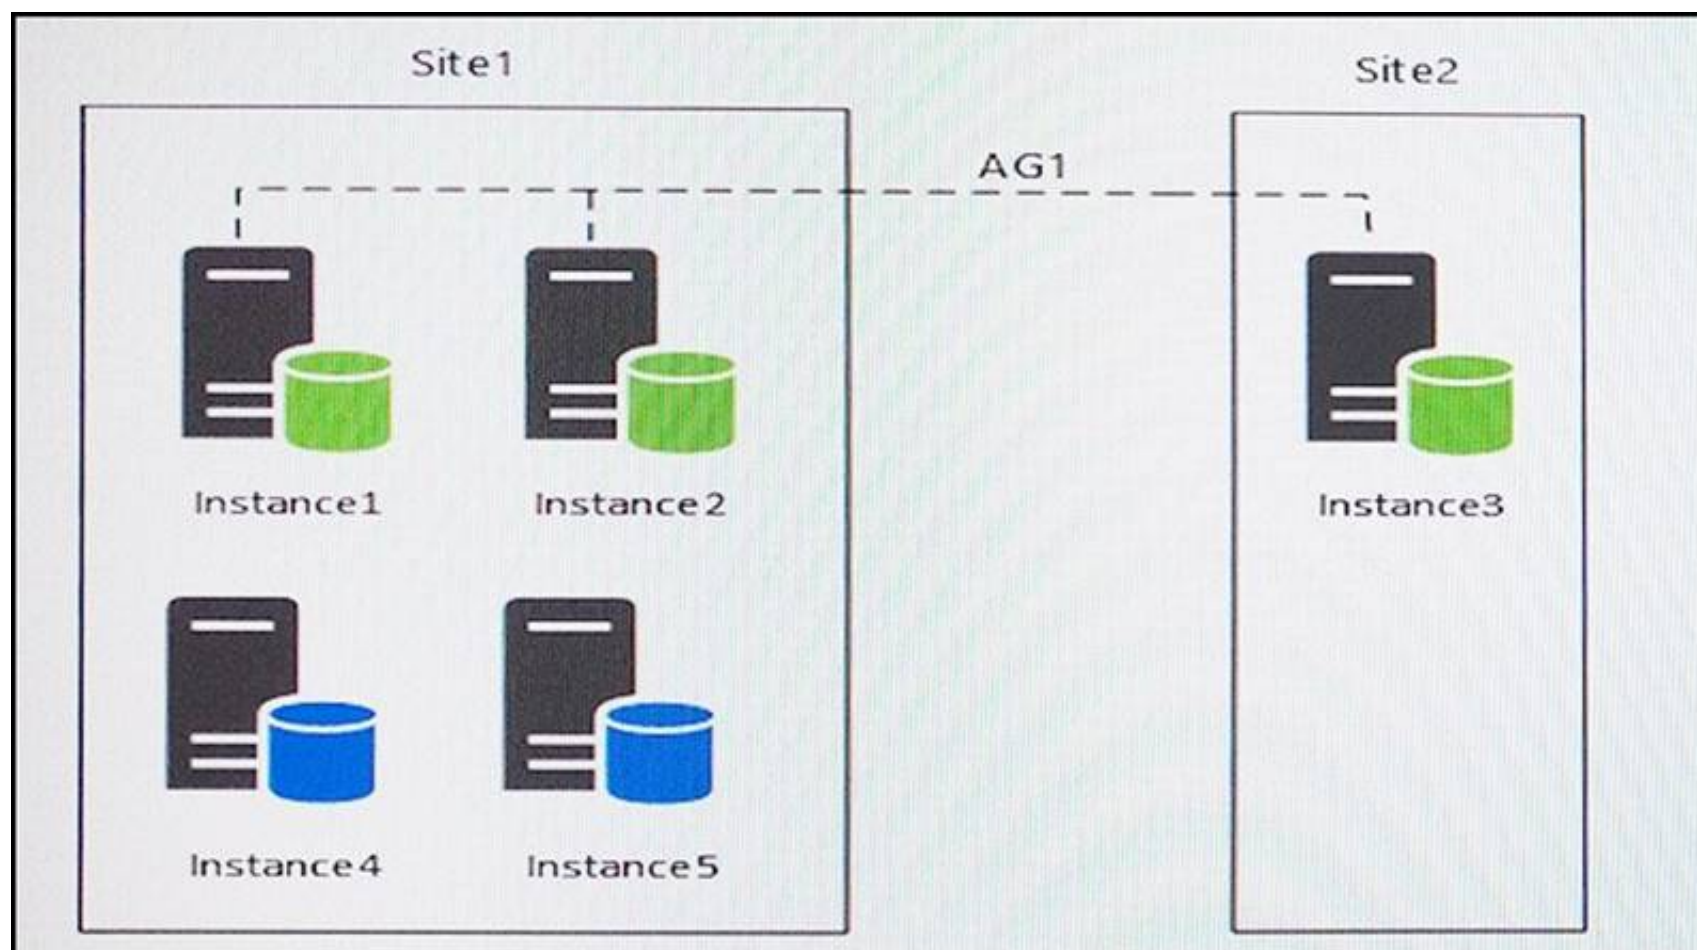


NEW QUESTION 22

- (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 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 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 31

- (Exam Topic 1)

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

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

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

Answer Area

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

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

NEW QUESTION 34

- (Exam Topic 1)
You are the database administrator for a Microsoft SQL Server instance. You develop an Extended Events package to look for events related to application performance.
You need to change the event session to include SQL Server errors that are greater than error severity 15. Which five Transact-SQL segments should you use to develop the solution? To answer, move the appropriate Transact-SQL segments from the list of Transact-SQL segments to the answer area and arrange them in the correct order.

Transact-SQL segments

WHERE ((sqlserver.data-base_id>(4)) AND (severity>(15)))

(ACTION(sqlserver.client_app_name, sqlserver.data-base_id,sqlserver.session_id)

ALTER EVENT SESSION Contoso1 ON SERVER

)

GO

ADD EVENT sqlserver.error_reported

ADD TARGET sqlserver.error_reported

Answer Area

⏪

⏩

⏴

⏵

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1: ALTER EVENT SESSION Contoso1 ON SERVER

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

Step 5:) GO

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

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

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

```
ACTION
```

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

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

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

```
)
```

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

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

```
) WITH (
```

```
STARTUP_STATE=OFF
```

```
) GO
```

References:

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

NEW QUESTION 37

- (Exam Topic 1)

You administer a Microsoft SQL Server 2016 instance that contains a financial database hosted on a storage area network (SAN). The financial database has the following characteristics:

A data file of 2 terabytes is located on a dedicated LUN (drive D).

A transaction log of 10 GB is located on a dedicated LUN (drive E).

Drive D has 1 terabyte of free disk space.

Drive E has 5 GB of free disk space.

The database is continually modified by users during business hours from Monday through Friday between 09:00 hours and 17:00 hours.

Five percent of the existing data is modified each day.

The Finance department loads large CSV files into a number of tables each business day at 11:15 hours and 15:15 hours by using the BCP or BULK INSERT commands.

Each data load adds 3 GB of data to the database.

These data load operations must occur in the minimum amount of time. A full database backup is performed every Sunday at 10:00 hours.

Backup operations will be performed every two hours (11:00, 13:00, 15:00, and 17:00) during business hours. You implement log shipping of the financial database to another SQL Server 2016 instance. You decide to

failover to this secondary database.

You need to ensure that all transactions will be replicated to the secondary database. Which backup option should you use?

- A. Differential
- B. Transaction Log
- C. FULL

- D. SIMPLE
- E. SKIP
- F. RESTART
- G. STANDBY
- H. CHECKSUM
- I. DBO_ONLY
- J. COPY_ONLY
- K. NORECOVERY
- L. NO_CHECKSUM
- M. CONTINUE_AFTER_ERROR
- N. BULK_LOGGED

Answer: K

Explanation:

Roll back is controlled by the RESTORE statement through the [RECOVERY | NORECOVERY] options: NORECOVERY specifies that roll back not occur. This allows roll forward to continue with the next statement in the sequence.

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

NEW QUESTION 39

- (Exam Topic 1)

You administer several Microsoft SQL Server 2016 database servers.

Merge replication has been configured for an application that is distributed across offices throughout a wide area network (WAN). Many of the tables involved in replication use the XML and varchar (max) data types.

Occasionally, merge replication fails due to timeout errors. You need to reduce the occurrence of these timeout errors. What should you do?

- A. Set the Merge agent on the problem subscribers to use the slow link agent profile.
- B. Create a snapshot publication, and reconfigure the problem subscribers to use the snapshot publication.
- C. Change the Merge agent on the problem subscribers to run continuously.
- D. Set the Remote Connection Timeout on the Publisher to 0.

Answer: A

Explanation:

You might have different profiles for different instances of an agent. For example, a Merge Agent that connects to the Publisher and Distributor over a dialup connection could use a set of parameters that are better suited to the slower communications link by using the slow link profile.

Note: When replication is configured, a set of agent profiles is installed on the Distributor. An agent profile contains a set of parameters that are used each time an agent runs: each agent logs in to the Distributor during its startup process and queries for the parameters in its profile.

References:

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

NEW QUESTION 44

- (Exam Topic 2)

Overview

General Overview

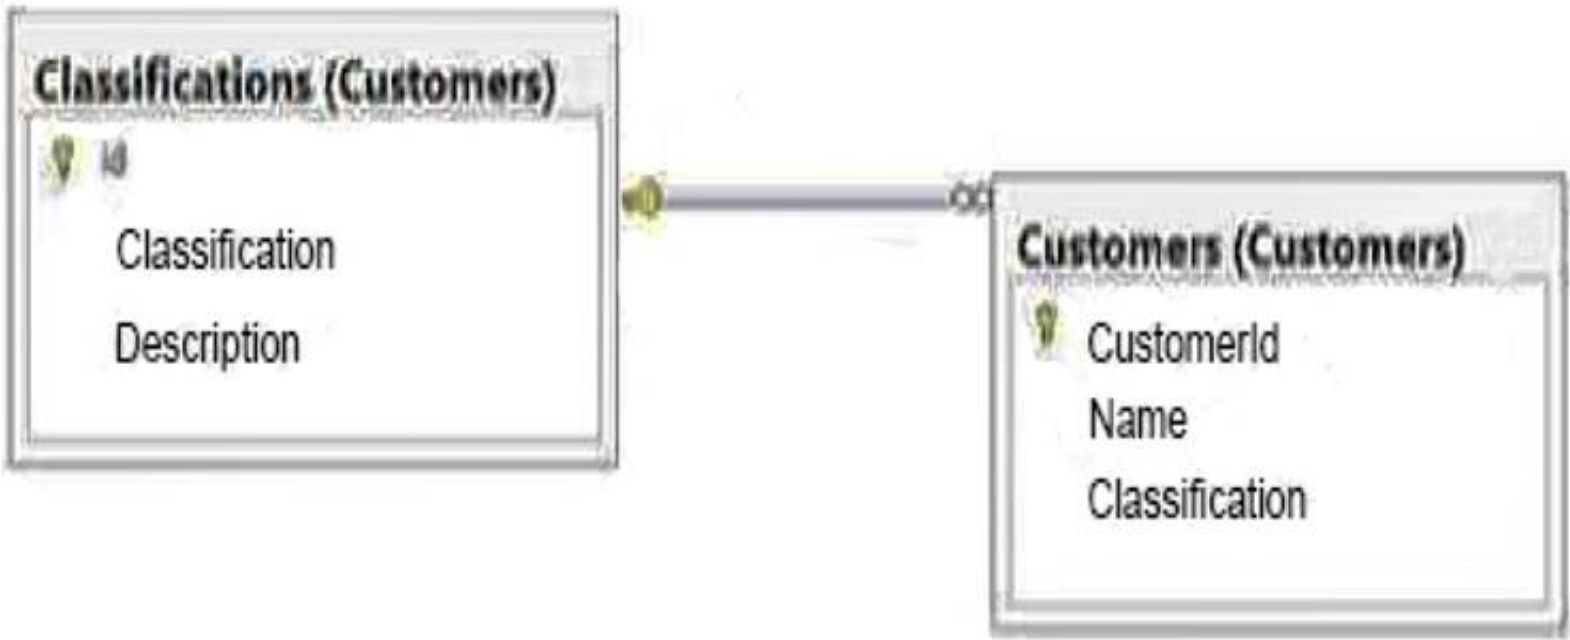
ADatum Corporation has offices in Miami and Montreal.

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

Databases

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

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



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

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

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

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

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

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

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

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

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

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

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

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

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

The nested stored procedures are never called directly. Design Requirements

Data Recovery

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

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

Classification Changes

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

Storage

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

Error Handling

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

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

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

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

Answer: A

Explanation:

Scenario:

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

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

NEW QUESTION 46

- (Exam Topic 2)

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

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

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

- Fulfilled
- Shipped
- Ordered
- Received

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

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

- A. To the Sales.Orders table, add a column named Status that will store the order statu
- B. Update the Status column as the order status changes.
- C. Create a new table named Sales.OrderStatus that contains three columns named OrderID, StatusDate, and Statu
- D. Insert new rows into the table as the order status changes.

E. Implement change data capture on the Sales.Orders table.

F. To the Sales.Orders table, add three columns named FulfilledDate, ShippedDate, and ReceivedDate. Update the value of each column from null to the appropriate date as the order status changes.

Answer: A

NEW QUESTION 48

- (Exam Topic 2)

Overview

General Overview

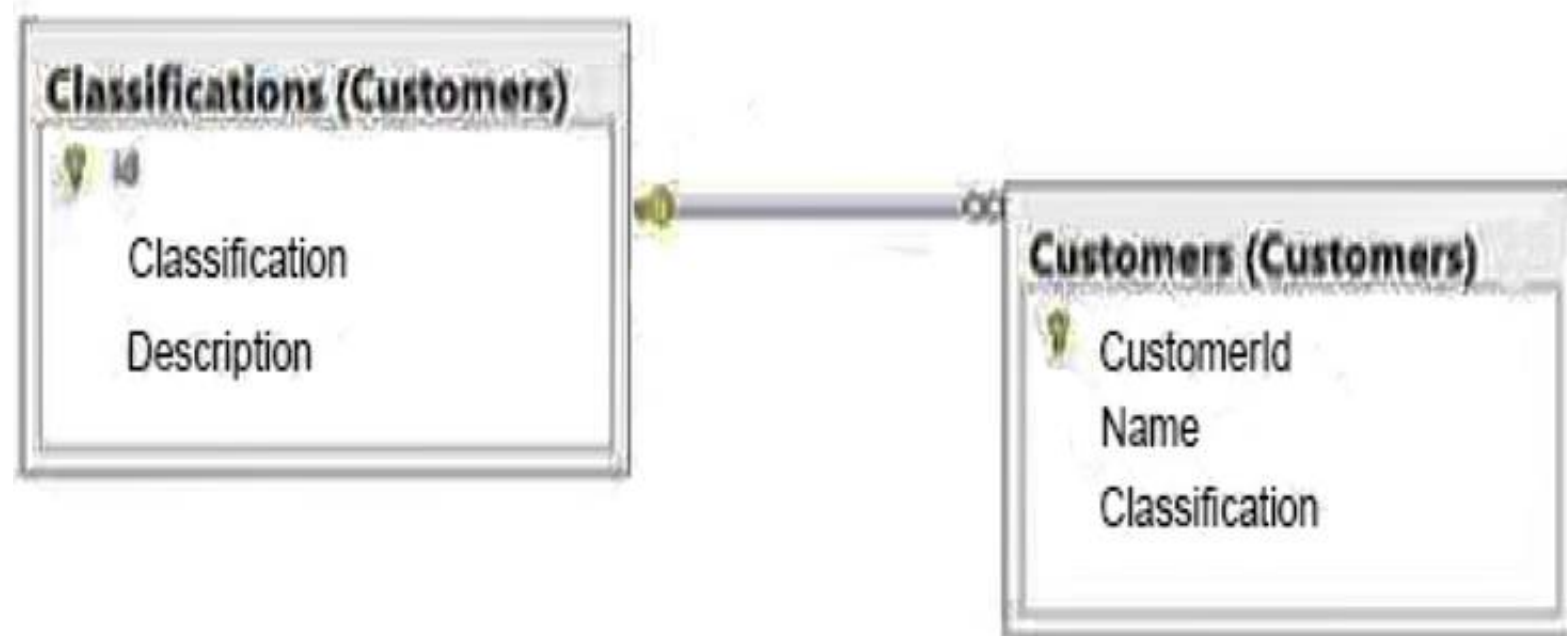
ADatum Corporation has offices in Miami and Montreal.

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

Databases

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

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



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

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

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

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

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

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

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

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

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

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

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

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

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

The nested stored procedures are never called directly. Design Requirements

Data Recovery

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

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

Classification Changes

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

Storage

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

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

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

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

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

Answer: A

Explanation:

Scenario:

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

- The simple recovery model provides the simplest form of backup and restore. This recovery model supports both database backups and file backups, but does not support log backups. Transaction log data is backed up only with the associated user data.

The absence of log backups simplifies managing backup and restore. However, a database can be restored only to the end of the most recent backup.

NEW QUESTION 52

- (Exam Topic 2)

Overview

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

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

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

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

Requirements Planned Changes

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

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

Database1 will also contain a stored procedure named usp_UpdateOrderDetails.

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

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

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

Inventory will contain over 100 GB of data.

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

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

Business Requirements

You have the following requirements:

Costs for new licenses must be minimized.

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

Development effort must be minimized whenever possible.

The storage requirements for databases must be minimized.

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

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

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

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

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

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

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

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

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

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

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

Answer: C

Explanation:

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

NEW QUESTION 57

- (Exam Topic 2)

Overview

General Overview

ADatum Corporation has offices in Miami and Montreal.

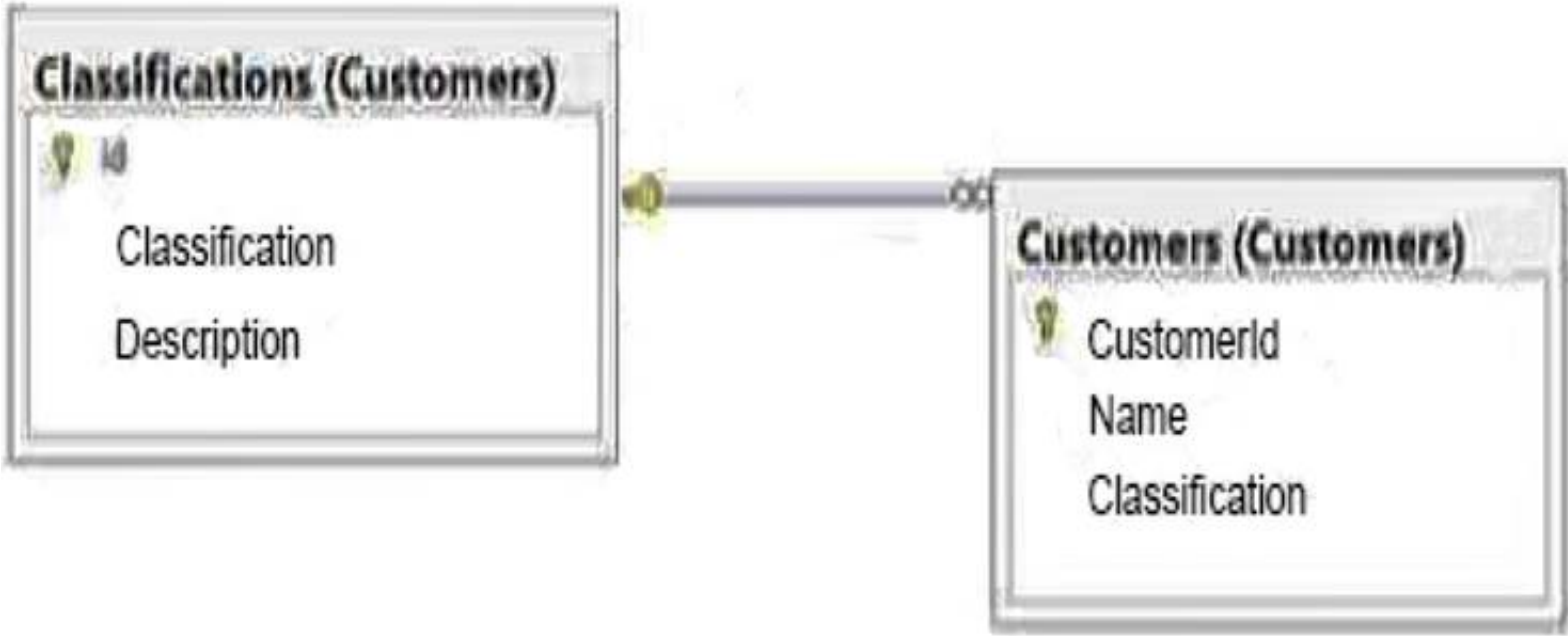
The network contains a single Active Directory forest named adatum.com. The offices connect to each other by using a WAN link that has 5-ms latency. A. Datum standardizes its database platform by using SQL Server 2014 Enterprise edition.

Databases

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

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

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



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

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

The Inventory database is updated frequently. The database is often used for reporting. A full backup of the database currently takes three hours to complete. Stored Procedures

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

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

A stored procedure named USP_2 is used to generate a product list. The product list contains the names of products grouped by category. USP_2 takes several minutes to run due to locks on the tables the procedure accesses. The locks are caused by USP_1 and USP_3.

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

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

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

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

The nested stored procedures are never called directly. Design Requirements

Data Recovery

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

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

Classification Changes

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

Storage

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

Error Handling

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

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

You need to recommend a solution for the planned changes to the customer classifications. What should you recommend? (Each correct answer presents part of the solution. Choose all that apply.)

A. Add a row to the Customers table each time a classification changes.

B. Add columns for each classification to the Customers table.

C. Add a table to track any changes made to the classification of each customer.

D. Add a column to the Classifications table to track the status of each classification.

E. Implement change data capture.

Answer: CD

Explanation:

Scenario:

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.

NEW QUESTION 59

- (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 60

- (Exam Topic 2)

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

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

What should you recommend?

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

Answer: A

Explanation:

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

NEW QUESTION 65

- (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 69

- (Exam Topic 2)

Overview

General Overview

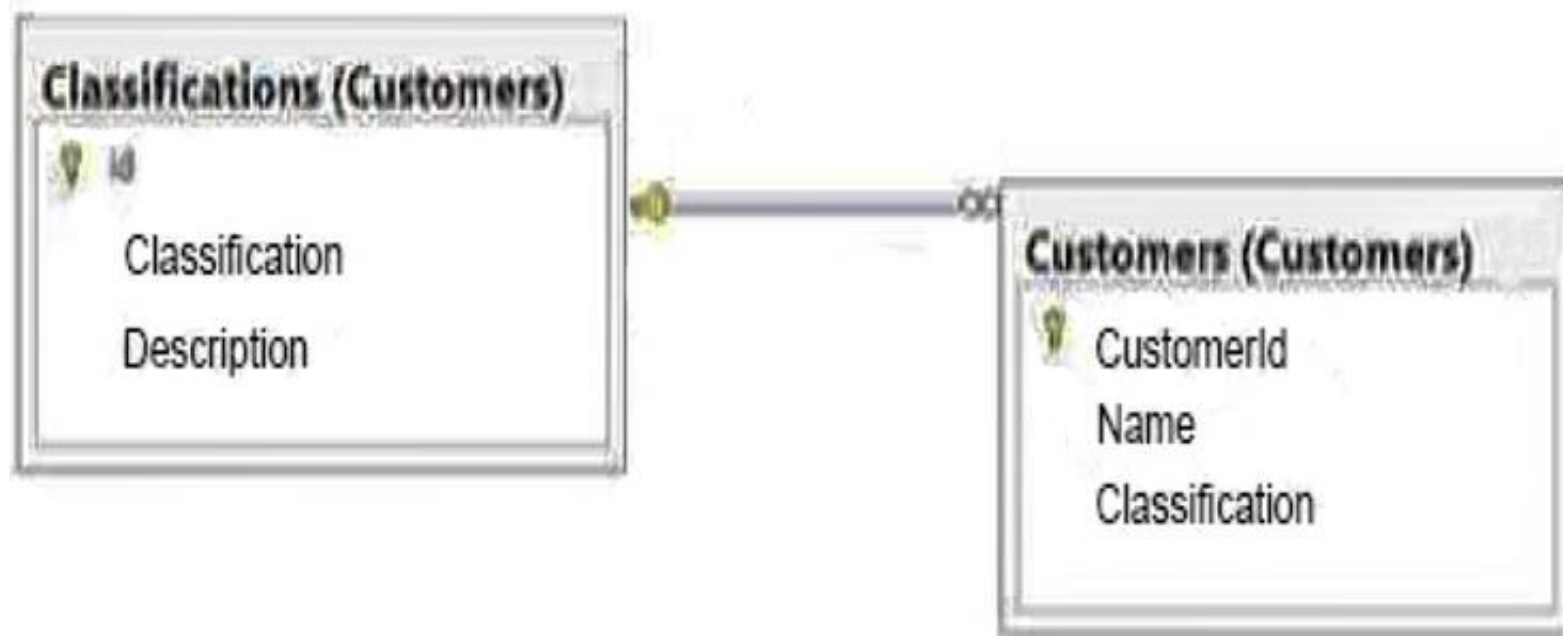
ADatum Corporation has offices in Miami and Montreal.

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

Databases

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

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



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

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

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

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

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

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

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

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

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

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

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

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

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

The nested stored procedures are never called directly. Design Requirements

Data Recovery

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

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

Classification Changes

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

Storage

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

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

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

You need to recommend a solution to meet the security requirements of the junior database administrators. What should you include in the recommendation?

- A. A server role
- B. A database role
- C. A credential
- D. A shared login

Answer: C

Explanation:

- Scenario: A group of junior database administrators must be able to view the server state of the SQL Server instance that hosts the Sales database. The junior database administrators will not have any other administrative rights.

- Credentials provide a way to allow SQL Server Authentication users to have an identity outside of SQL Server. Credentials can also be used when a SQL Server Authentication user needs access to a domain resource, such as a file location to store a backup.

NEW QUESTION 74

- (Exam Topic 2)

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

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

Answer: D

NEW QUESTION 79

- (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 82

- (Exam Topic 2)

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

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

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

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

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

Answer: E

Explanation:

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

NEW QUESTION 84

- (Exam Topic 2)

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

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

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

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

Answer: B

Explanation:

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

NEW QUESTION 86

- (Exam Topic 2)

You administer a Microsoft SQL Server 2016 failover cluster.

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

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

Answer: B

Explanation:

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

References:

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

NEW QUESTION 90

- (Exam Topic 2)

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

You need to diagnose deadlocks that happen when executing a specific set of stored procedures by recording events and playing them back on a different test server.

What should you create?

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

Answer: F

Explanation:

References:

<http://msdn.microsoft.com/en-us/library/ms188246.aspx>

NEW QUESTION 91

- (Exam Topic 2)

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

- Prevent data read by SP1 from being modified by other active processes.
- Prevent SP2 from performing dirty reads.

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

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

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

A. Mastered

B. Not Mastered

Answer: A

Explanation:

SP1 – repeatable read; SP2 – read committed

- REPEATABLE READ

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

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

NEW QUESTION 92

- (Exam Topic 2)

Overview

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

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

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

will use the application.

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

Requirements Planned Changes

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

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

Database1 will contain two tables named Orders and OrderDetails.

Database1 will also contain a stored procedure named usp_UpdateOrderDetails.

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

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

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

Inventory will contain over 100 GB of data.

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

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

Business Requirements

You have the following requirements:

Costs for new licenses must be minimized.

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

Development effort must be minimized whenever possible.

The storage requirements for databases must be minimized.

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

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

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

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

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

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

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

A. Encrypted columns

B. Certificate Authentication

C. Secure Socket Layer (SSL)

D. Signed stored procedures

Answer: D

Explanation:

- Scenario:

- Data from Database2 will be accessed periodically by an external application named Application1

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

Tutorial: Signing Stored Procedures with a Certificate

NEW QUESTION 95

- (Exam Topic 2)

You plan to create a database.

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

After the event, the database will be deleted.

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

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

A. SQL Server 2014 Enterprise

B. SQL Server 2014 Standard

C. SQL Azure

D. SQL Server 2014 Express with Advanced Services

Answer: B

Explanation:

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

NEW QUESTION 100

- (Exam Topic 2)

You install a Microsoft SQL Server 2016 instance.

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

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

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

Answer: B

NEW QUESTION 103

- (Exam Topic 2)

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

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

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

Answer: B

Explanation:

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

References:

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

NEW QUESTION 108

- (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 110

- (Exam Topic 2)

Overview

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

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

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

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

Requirements Planned Changes

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

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

Database1 will also contain a stored procedure named usp_UpdateOrderDetails.

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

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

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

Inventory will contain over 100 GB of data.

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

Database2 will contain a stored procedure named usp_UpdateInventory. usp_UpdateInventory will manipulate a table that contains a self-join that has an unlimited number of hierarchies. All data in Database2 is recreated each day and does not change until the next data creation process. Data from Database2 will be accessed periodically by an external application named Application1. The data from Database2 will be sent to a database named App1_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 113

- (Exam Topic 2)

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

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

You need to recommend which high-availability feature is required for each environment.

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

Features		Answer Area
AlwaysOn availability groups	ENV1	Feature
Database mirroring	ENV2	Feature
Log shipping	ENV3	Feature
Peer-to-peer replication		

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

ENV1 - Database mirroring; ENV2 - AlwaysOn availability groups; ENV3 - AlwaysOn availability groups

- AlwaysOn availability groups

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

NEW QUESTION 115

- (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. During performance testing, you discover that database INSERT operations against the Inventory table are slow. You need to recommend a solution to reduce the amount of time it takes to complete the INSERT operations. What should you recommend?

- A. Partition the nonclustered index.
- B. Partition the Inventory table.snapshot replication
- C. Create a column store index.Master Data Services
- D. Drop the clustered index.change data capture

Answer: A

Explanation:

Scenario:

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.

NEW QUESTION 120

- (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 123

- (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 126

- (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 130

- (Exam Topic 2)

You have a SQL Server 2014 instance named SQL1. SQL1 creates error events in the Windows Application event log. You need to recommend a solution that will run an application when SQL1 logs a specific error in the Application log. Which SQL elements should you include in the recommendation? (Each correct answer presents part of the solution. Choose all that apply.)

- A. A policy
- B. A maintenance plan
- C. An alert
- D. A job
- E. A trigger

Answer: DE

Explanation:

Use a trigger that starts a job which executes the application.
References:
<http://technet.microsoft.com/en-us/library/hh849759.aspx>

NEW QUESTION 132

- (Exam Topic 3)

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

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

You attempt to restore a database on a new SQL Server instance and receive the following error message: "Msg 33111, Level 16, State 3, Line 2 Cannot find server certificate with thumbprint '0x7315277C70764B1F252DC7A5101F6F66EFB1069D.'" You need to ensure that you can restore the database successfully.

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

- A. Yes
- B. No

Answer: B

Explanation:

The error is related to the certificate.
References: <https://dba.stackexchange.com/questions/3388/restore-encrypted-database-to-another-server?rq=1>

NEW QUESTION 134

- (Exam Topic 3)

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

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

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

Minimize the space used by transaction logs.

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

Actions

Perform a weekly full backup of the database

Ship the logs to a secondary server

Set the database to the simple recovery model

Set the database to the full recovery model

Create a database snapshot

Perform a backup of the transaction log every hour

Perform a differential backup of the database every night

Work Area

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Set the database to the simple recovery model; Perform a weekly full backup of the database;
Perform a differential backup of the database every night Note:
- Simple recovery model No log backups.
- Full recovery model Requires log backups

NEW QUESTION 137

- (Exam Topic 3)

Background Corporate Information

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

Physical Locations

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

Problem Statement

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

Configuration Windows Logins

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

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

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

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

Database

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

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

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

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

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

Column	Data type
SalesOrderDetailID	INT
ProductID	INT
SalePrice	SMALLMONEY
SaleQuantity	INT

Database Issues

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

Requirements Database

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

Customer data access

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

Customer support data access

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

Sales force data access

Sales people at both the headquarters office and the satellite office must generate reports that read from the Product and SalesOrderDetail tables. No updates or inserts are ever made by sales people. These reports are run at random times and there can be no reporting downtime to refresh the data set except during the monthly

maintenance window. The reports that run from the satellite office are process intensive queries with large data sets. Regardless of which office runs a sales force report, the SalesOrderDetail table should only return valid, committed order data; any orders not yet committed should be ignored.

Historical Data

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

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

Backups

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

Database Maintenance

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

Project milestones completed

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

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

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

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

You need to distribute functionality across the three servers. Which function should you assign to each server? To answer, drag the appropriate functions to the correct servers. Each function 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.

Functions	Answer Area
AlwaysOn primary replica	HQ office server. Function
AlwaysOn secondary replica	Satellite office server. Function
file backup server	Cloud server. Function
witness server	
mirroring primary	
mirroring secondary	
log shipping primary	
log shipping secondary	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:
 HQ office server - AlwaysOn primary replica;
 Satellite office server - AlwaysOn secondary replica; Cloud server - file backup server

NEW QUESTION 142

- (Exam Topic 3)
 Background Corporate Information
 Fabrikam, Inc. is a retailer that sells electronics products on the Internet. The company has a headquarters site and one satellite sales office. You have been hired as the database administrator, and the company wants you to change the architecture of the Fabrikam ecommerce site to optimize performance and reduce downtime while keeping capital expenditures to a minimum. To help with the solution, Fabrikam has decided to use cloud resources as well as on-premise servers.
 Physical Locations
 All of the corporate executives, product managers, and support staff are stationed at the headquarters office. Half of the sales force works at this location. There is also a satellite sales office. The other half of the sales force works at the satellite office in order to have sales people closer to clients in that area. Only sales people work at the satellite location.
 Problem Statement
 To be successful, Fabrikam needs a website that is fast and has a high degree of system uptime. The current system operates on a single server and the company is not happy with the single point of failure this presents. The current nightly backups have been failing due to insufficient space on the available drives and manual drive cleanup often needing to happen to get past the errors. Additional space will not be made available for backups on the HQ or satellite servers. During your investigation, you discover that the sales force reports are causing significant contention.
 Configuration Windows Logins
 The network administrators have set up Windows groups to make it easier to manage security. Users may belong to more than one group depending on their role. The groups have been set up as shown in the following table:

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

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

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

Database

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

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

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

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

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

Column	Data type
SalesOrderDetailID	INT
ProductID	INT
SalePrice	SMALLMONEY
SaleQuantity	INT

Database Issues

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

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

Customer data access

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

Customer support data access

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

Sales force data access

Sales people at both the headquarters office and the satellite office must generate reports that read from the Product and SalesOrderDetail tables. No updates or inserts are ever made by sales people. These reports are run at random times and there can be no reporting downtime to refresh the data set except during the monthly

maintenance window. The reports that run from the satellite office are process intensive queries with large data sets. Regardless of which office runs a sales force report, the SalesOrderDetail table should only return valid, committed order data; any orders not yet committed should be ignored.

Historical Data

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

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

Backups

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

Database Maintenance

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

Project milestones completed

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

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

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

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

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

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

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

Answer: CD

Explanation:

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

- Need to store files in the cloud.

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

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

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

NEW QUESTION 144

- (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 147

- (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 148

- (Exam Topic 3)

Overview

General Overview

ADatum Corporation has offices in Miami and Montreal.

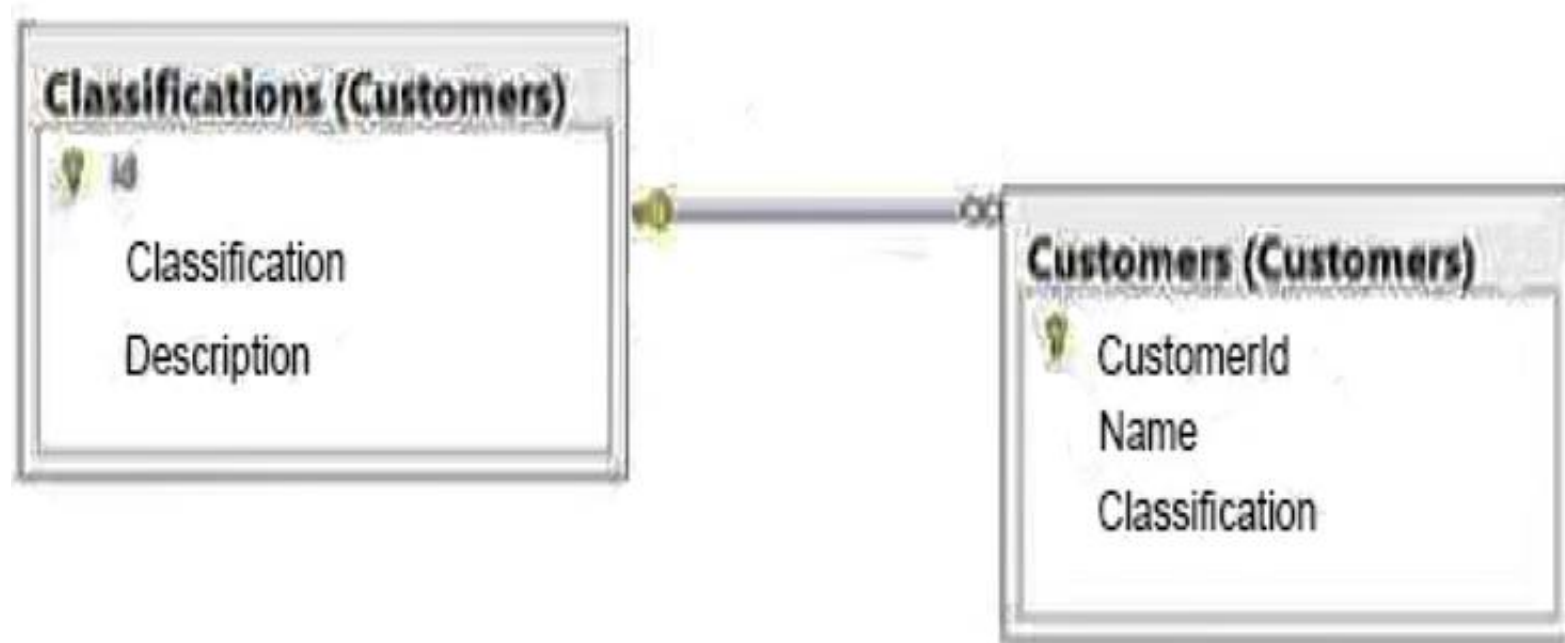
The network contains a single Active Directory forest named adatum.com. The offices connect to each other by using a WAN link that has 5-ms latency. A. Datum standardizes its database platform by using SQL Server 2014 Enterprise edition.

Databases

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

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

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



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

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

The Inventory database is updated frequently.

The database is often used for reporting.

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

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

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

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

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

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

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

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

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

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

The nested stored procedures are never called directly. Design Requirements

Data Recovery

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

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

Classification Changes

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

Storage

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

Error Handling

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

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

You need to recommend a solution for the error handling of USP_3. The solution must minimize the amount of custom code required. What should you recommend?

- A. Use the @@ERROR variable in the nested stored procedures.
- B. Use a TRY CATCH block in the called stored procedures.
- C. Use the @@ERROR variable in the called stored procedures.
- D. Use the RAISERROR command in the nested stored procedures.

Answer: B

Explanation:

- Must catch and handle the error. 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.

NEW QUESTION 150

- (Exam Topic 3)

General Overview

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

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

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

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

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

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

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

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

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

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

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

The GenInfo table is used for reports.

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

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

Current System

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

SQL Servers

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

Design Requirements

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

- Confidential information in the Current_Inventory database that is accessed by ExternalApp1 must be securely stored.
- Direct access to database tables by developers or applications must be denied.
- The account used to generate reports must have restrictions on the hours when it is allowed to make a connection.
- Deadlocks must be analyzed with the use of Deadlock Graphs.
- In the event of a SQL Server failure, the databases must remain available.
- Software licensing and database storage costs must be minimized.
- Development effort must be minimized.
- The Tempdb databases must be monitored for insufficient free space.
- Failed authentication requests must be logged.

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

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

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

What should you recommend?

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

Answer: D

NEW QUESTION 154

- (Exam Topic 3)

You are designing two stored procedures named Procedure1 and Procedure2. You identify the following requirements:

Procedure1 must take a parameter that ensures that multiple rows of data can pass into the stored procedure.

Procedure2 must use business logic that resides in a Microsoft .NET Framework assembly. You need to identify the appropriate technology for each stored procedure.

Which technologies should you identify? To answer, drag the appropriate technology to the correct stored procedure in the answer area. (Answer choices may be used once, more than once, or not at all.)

Technologies	Answer Area
Common language runtime (CLR)	Procedure 1 Technology
Extensible Markup Language (XML)	Procedure 2 Technology
A table-valued parameter (TVP)	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Procedure 1 - A table-valued parameter (TVP); Procedure 2 - Common language runtime (CLR) References:

<http://msdn.microsoft.com/en-us/library/ms131102.aspx> <http://msdn.microsoft.com/en-us/library/bb522446.aspx> <http://msdn.microsoft.com/en-us/library/bb510489.aspx>

NEW QUESTION 158

- (Exam Topic 3)

Background

Corporate Information

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

Physical Locations

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

Problem Statement

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

Configuration Windows Logins

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

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

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

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

Database

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

```
CREATE DATABASE ProdDB
```

```
GO
```

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

```
GO
```

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

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

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

Column	Data type
SalesOrderDetailID	INT
ProductID	INT
SalePrice	SMALLMONEY
SaleQuantity	INT

Database Issues

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

Requirements

Database

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

Customer data access

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

Customer support data access

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

Sales force data access

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

Historical Data

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

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

Backups

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

Database Maintenance

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

Project milestones completed

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

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

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

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

You need to configure security on the Product table for customer support representatives.

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

- A. Create a view called CustProduct that includes columns ProductID, ProductName, Product Description, QuantityOnHand, ProductPrice, ProductCost, and ProductSupplierID.
- B. GRANT ALL on CustProduct TO OurDomain\CustomerSupport
- C. Create a user-defined data type called CustProduct that includes columns ProductID, ProductName, Product Description, and ProductPrice.
- D. Create a view called CustProduct that includes columns ProductID, ProductName, Product Description, QuantityOnHand, and ProductPrice.
- E. GRANT SELECT on CustProduct TO OurDomain\CustomerSupport.
- F. GRANT SELECT on CustProduct TO public.

Answer: AE

Explanation:

Give access to CustomerSupport through a view. The view must include all these columns (refer to scenario). GRANT Object Permissions (Transact-SQL)

NEW QUESTION 160

.....

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](#)