

Exam Questions 70-764

Administering a SQL Database Infrastructure (beta)

<https://www.2passeasy.com/dumps/70-764/>



NEW QUESTION 1

- (Exam Topic 1)

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

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

What should you use?

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

Answer: D

Explanation:

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

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

Missing indexes can be found with the execution plans.

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

NEW QUESTION 2

- (Exam Topic 1)

You 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 3

- (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 backup
- 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 backup
- G. Then, restore the latest log backup.
- H. Perform a page restore.
- I. Restore the latest full backup
- J. Then, restore each differential backup taken before the time of failure from the most recent full backup.
- K. Restore the latest full backup
- 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 4

- (Exam Topic 1)

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

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

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

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

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

- A. Configure the Subscriber to use the Local System account.
- B. Configure the SQL Server Agent service to run under the Local System account

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

Answer: D

NEW QUESTION 5

- (Exam Topic 1)

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

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

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

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

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

- A. Yes
- B. No

Answer: B

Explanation:

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

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

NEW QUESTION 6

- (Exam Topic 1)

You have configured Resource Governor with three resource pools.

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

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

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

Answer: B

Explanation:

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

NEW QUESTION 7

- (Exam Topic 1)

You administer a Microsoft SQL Server 2016 server.

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

What should you configure?

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

Answer: C

Explanation:

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

References:

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

NEW QUESTION 8

- (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>ErrorDmp[XXXX] .mdmp

Answer: A

NEW QUESTION 9

- (Exam Topic 1)

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

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

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

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

Answer: C

Explanation:

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

It consumes less tempdb space than snapshot isolation. Etc.

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

NEW QUESTION 10

- (Exam Topic 1)

You administer a Microsoft SQL Server 2016 database instance.

You create a new user named UserA. You need to ensure that UserA is able to create SQL Server Agent jobs and to execute SQL Server Agent jobs.

To which role should you add UserA?

- A. Securityadmin
- B. RSExecRole
- C. SQLAgentUserRole
- D. DatabaseMailUserRole

Answer: C

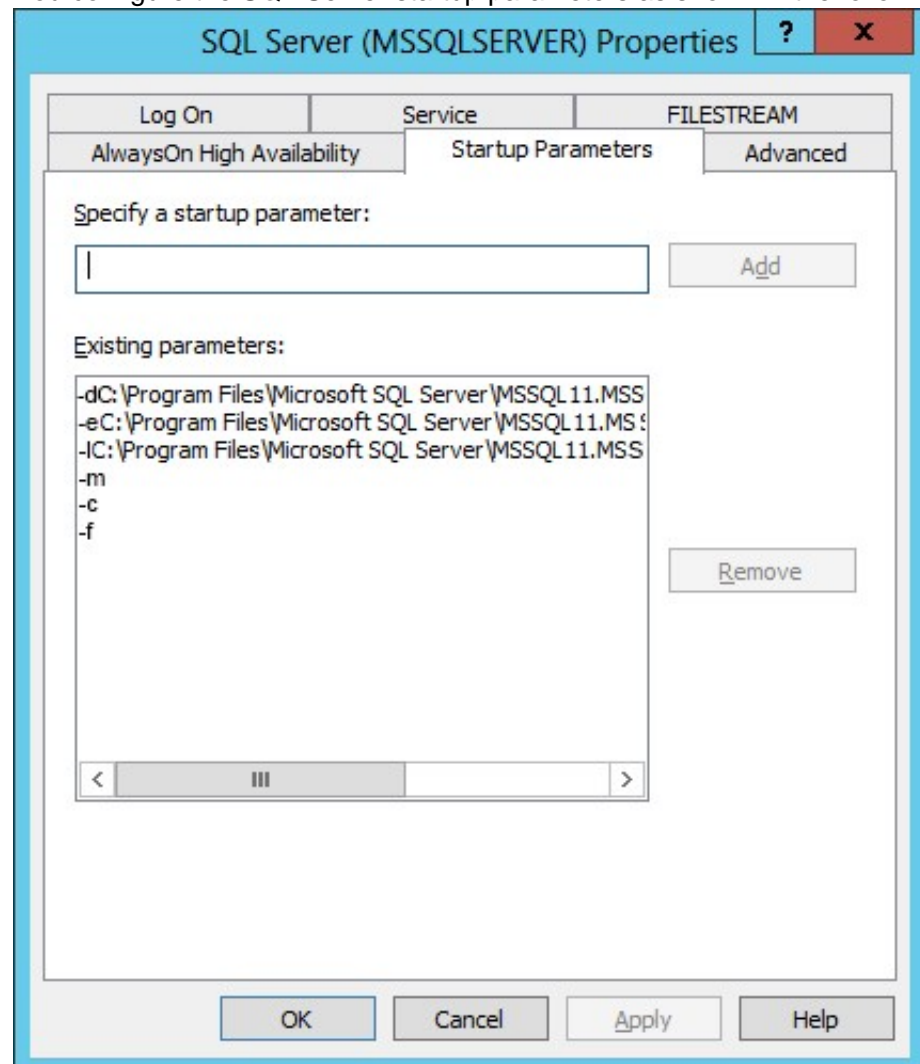
NEW QUESTION 10

- (Exam Topic 1)

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

MSG_AUDIT_FORCED_SHUTDOWN

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



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

Answer Area

In which user mode will the SQL Server instance start?

| | |
|-----------------|---|
| | ▼ |
| single-user | |
| multi-user | |
| restricted-user | |

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

| | |
|-------------|---|
| | ▼ |
| public | |
| serveradmin | |
| sysadmin | |
| setupadmin | |

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: single-user

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

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

References:

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

NEW QUESTION 14

- (Exam Topic 1)

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

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

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

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

- A. Yes
- B. No

Answer: B

Explanation:

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

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

References:

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

NEW QUESTION 15

- (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 18

- (Exam Topic 1)

You administer a Microsoft SQL Server 2016 database.

You have a SQL Server Agent job instance that runs using the service account. You have a job step within the job that requires elevated privileges. You need to ensure that the job step can run using a different user account. What should you use?

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

Answer: D

NEW QUESTION 23

- (Exam Topic 1)

You administer a Microsoft SQL Server 2016 database named Contoso on a server named Server01. You need to collect data for a long period of time to troubleshoot wait statistics when querying Contoso. You also need to ensure minimum impact to the server. What should you create?

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

Answer: C

NEW QUESTION 26

- (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 27

- (Exam Topic 1)

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

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

A company has a server that runs Microsoft SQL Server 2016 Web edition. The server has a default instance that hosts a database named DB1.

You need to ensure that you can perform auditing at the database level for DB1.

Solution: You migrate DB1 to the default instance on a server that runs Microsoft SQL Server 2016 Standard edition.

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

All editions of SQL Server support server level audits. All editions support database level audits beginning with SQL Server 2016 SP1. Prior to that, database level auditing was limited to Enterprise, Developer, and Evaluation editions.

References:

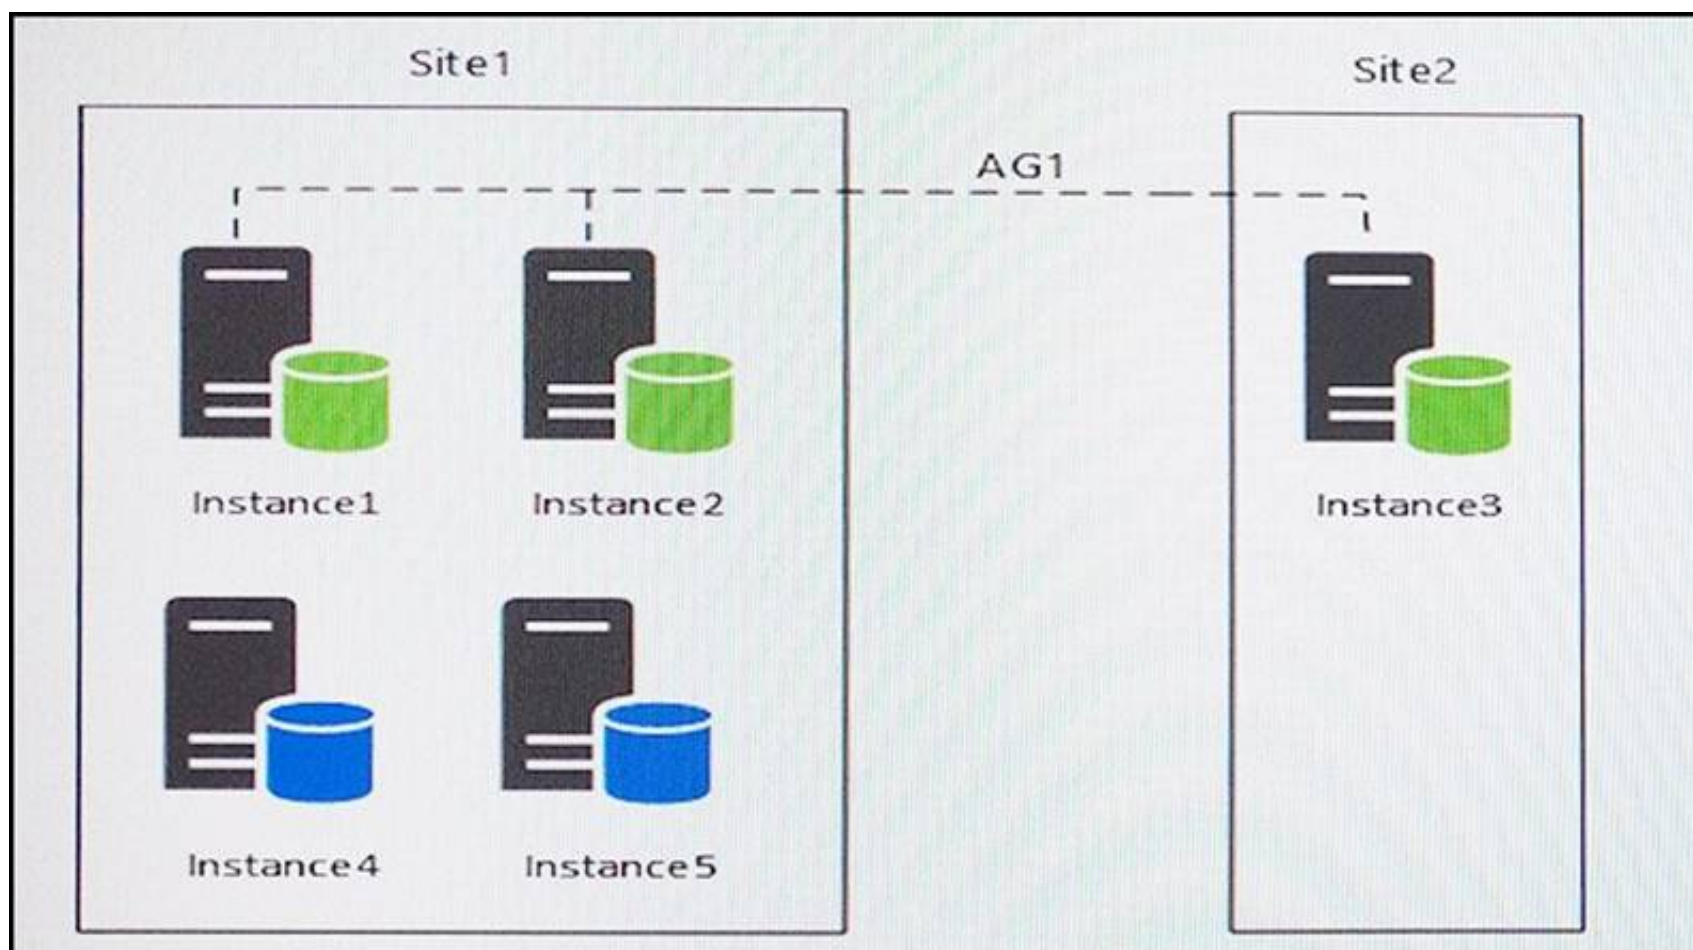
<https://docs.microsoft.com/en-us/sql/relational-databases/security/auditing/sql-server-audit-database-engine>

NEW QUESTION 32

- (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 OperationsMainthat 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 36

- (Exam Topic 1)

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

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

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

Answer Area

| | | | | |
|----------------|--|-----------------|--|---|
| DBCC | | ('salesOrders', | |) |
| CHECKDB | | CHECKDB | | |
| PHYSICAL_ONLY | | PHYSICAL_ONLY | | |
| REPAIR_FAST | | REPAIR_FAST | | |
| REPAIR_REBUILD | | REPAIR_REBUILD | | |

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: CHECKDB

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

DBCC CHECKDB

[(database_name | database_id | 0 [, NOINDEX
 | , { REPAIR_ALLOW_DATA_LOSS | REPAIR_FAST | REPAIR_REBUILD }]

....

Box 2: REPAIR_REBUILD

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

DBCC CHECKDB repair the found errors.

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

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

NEW QUESTION 38

- (Exam Topic 1)

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

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

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

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

Answer Area

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

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

NEW QUESTION 40

- (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 44

- (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 47

- (Exam Topic 1)

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

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

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

Answer Area

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

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

Box 1: Primary server instance.

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

Box 2: Secondary server instance

Each of the three secondary server instances runs its own copy job to copy the primary log-backup file to its own local destination folder.

copy job: A SQL Server Agent job that copies the backup files from the primary server to a configurable destination on the secondary server and logs history on the secondary server and the monitor server. When log shipping is enabled on a database, the job category "Log Shipping Copy" is created on each secondary server in a log shipping configuration.

Box 3: Secondary server instance.

Each secondary server instance runs its own restore job to restore the log backup from the local destination folder onto the local secondary database.

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

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

NEW QUESTION 48

- (Exam Topic 1)

You administer a SQL Server 2016 database instance.

You need to configure the SQL Server Database Engine service on a failover cluster. Which user account should you use?

- A. A domain user
- B. The BUILTIN\SYSTEM account
- C. A local user with Run as Service permissions
- D. The SQLBrowser account

Answer: A

Explanation:

Account of the person who installs the cluster: The person who installs the cluster must use an account with the following characteristics:

The account must be a domain account. It does not have to be a domain administrator account. It can be a domain user account if it meets the other requirements in this list.

Etc. References:

<https://docs.microsoft.com/en-us/previous-versions/windows/it-pro/windows-server-2008-R2-and-2008/cc73100>

NEW QUESTION 50

- (Exam Topic 1)

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

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

```
GRANT EXECUTE, SELECT ON Sales.ReviewInvoice 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 52

- (Exam Topic 1)

You create an availability group named HaContoso that has replicas named Server01/HA, Server02/HA, and Server03/HA. Currently, Server01/HA is the primary replica.

You need to ensure that the following requirements are met:

Backup operations occur on Server02/HA.

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

Backup operations do not occur on Server01/HA.

How should you configure HaContoso?

- A. Set the backup preference of HaContoso to Prefer Secondar
 B. Set the backup priority of Server02/HA to20. Set the backup priority of Server03/HA to 10.
 C. Set the backup preference of HaContoso to Secondary onl
 D. Set the backup priority of Server02/HA to20. Set the backup priority of Server03/HA to 10.
 E. Set the backup preference of HaContoso to Secondary onl

- F. Set the backup priority of Server02/HA to10. Set the backup priority of Server03/HA to 20.
- G. set the exclude replica of Server01/HA to tru
- H. Set the backup priority of Server02/HA to 10. Set the backup priority of Server03/HA to 20.

Answer: B

Explanation:

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

Backup Priority (Lowest=1, Highest=100)

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

References:

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

NEW QUESTION 54

- (Exam Topic 1)

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

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

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

Answer options

Add and configure the replica and create an availability group listener named AGL1 on port 1433.

Launch the Failover Cluster Manager and configure AO-AG1 and AO-AG2 as servers in the cluster. Name the cluster WINCL1.

Create the Always On Availability Group and select the user databases for the availability group.

Enable SQL Server 2016 Always On Availability Group feature.

Select the Full data synchronization method and specify the network path: \\DATA-C11\\SQL.

Answer Area

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1: Launch the Failover Cluster Manager and..

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

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

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

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

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

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

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

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

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

NEW QUESTION 57

- (Exam Topic 1)

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

The reporting database is only updated from the transactional database.

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

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

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

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

Answer: H

Explanation:

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

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

NEW QUESTION 58

- (Exam Topic 1)

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

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

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

You enabled log shipping for DB1 and configure backup and restore to occur every 30 minutes.

You discover that the disks on the data server are almost full.

You need to reduce the amount of disk space used by the log shipping process. Solution: You configure log shipping to backup and restore by using shared folder.

Does this meet the goal?

- A. Yes
- B. No

Answer: B

NEW QUESTION 63

- (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.tsq_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 68

- (Exam Topic 1)

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

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

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

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

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

- A. Yes
- B. No

Answer: A

Explanation:

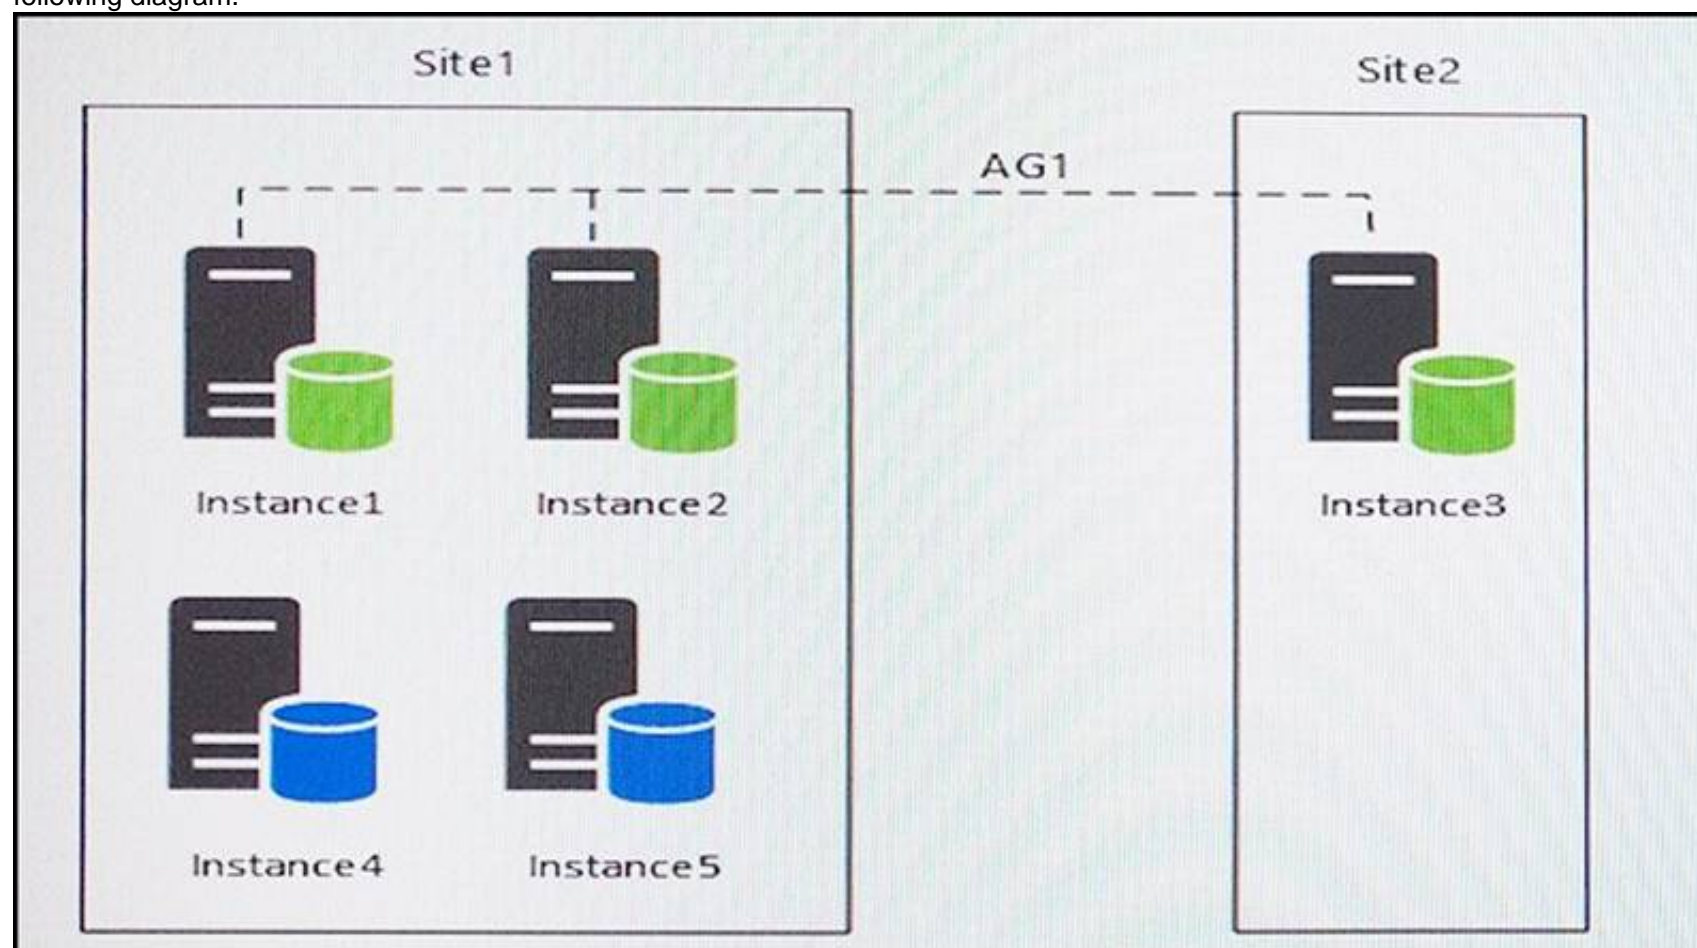
You can remedy index fragmentation by either reorganizing an index or by rebuilding an index. References: [https://msdn.microsoft.com/en-us/library/ms189858\(v=sql.105\).aspx](https://msdn.microsoft.com/en-us/library/ms189858(v=sql.105).aspx)

NEW QUESTION 71

- (Exam Topic 1)

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

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



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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

You need to analyze the wait type and statistics for specific instanced in the environment.

Which object should you use to gather information about each instance? To answer, drag the appropriate

objects to the correct instances. Each object may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Objects

Sys.dm_os_wait_stats

Sys.dm_exec_connections

Sys.dm_exec_requests

Sys.dm_exec_procedure_stats

Sys.dm_exec_sessions

Sys.dm_exec_query_stats

Sys.dm_exec_query_resource_semaphores

Sys.dm_exec_session_wait_stats

Answer Area

Instance Object

Instance1 Object

Instance4 Object

Instance5 Object

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Instance 1: sys.dm_exec_query_stats

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

Instance 4: sys.dm_os_wait_stats

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

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

Instance 5:sys.dm_exec_session_wait_stats

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

NEW QUESTION 73

- (Exam Topic 2)

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

You need to ensure that data can be replicated from a production server to two reporting servers in real time. You also need to ensure that data on the reporting server is always accessible.

Which solution should you use?

- A. Availability Groups
- B. Extended Events
- C. Snapshot Replication
- D. Policy Based Management

Answer: A

NEW QUESTION 77

- (Exam Topic 2)

You use SQL Server 2014 Enterprise Edition.

Your database contains a partitioned table named AuditData. AuditData is partitioned by year. Partition 1 contains data from the year 2010 and prior.

Management has decided to archive all AUDITDATA records from 2010 and prior.

Management wants the records to be removed from the database entirely and provided to the backup team as a zipped text file. The data must no longer reside in the database.

There is very little tolerance for performance degradation in your environment. You need to remove all 2010 and prior data from the AuditData table by using the least amount of system resources possible. Develop the solution by selecting and arranging the required SQL actions in the correct order.

You may not need all of the actions.

| SQL Actions | Answer Area |
|------------------|-------------|
| Drop Table | |
| Select Into | |
| Switch Partition | |
| Move Partition | |
| Merge Range | |
| BCP | |
| Split Range | |
| Create Table | |
| Delete Partition | |
| Drop Partition | |

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Note:

- Create a new partitioned table with the partition function you want, and then insert the data from the old table into the new table by using an INSERT INTO...SELECT FROM statement.

- SPLIT RANGE (boundary_value)

Adds one partition to the partition function. boundary_value determines the range of the new partition, and must differ from the existing boundary ranges of the partition function. Based on boundary_value, the Database Engine splits one of the existing ranges into two.

Of these two, the one where the new boundary_value resides is considered the new partition.

- BCP can be used to produce the zipped text file.

- Example: splitting a partition of a partitioned table or index into two partitions

The following example creates a partition function to partition a table or index into four partitions. ALTER PARTITION FUNCTION splits one of the partitions into two to create a total of five partitions. CREATE PARTITION FUNCTION myRangePF1 (int)

AS RANGE LEFT FOR VALUES (1, 100, 1000); GO

-Split the partition between boundary_values 100 and 1000

-to create two partitions between boundary_values 100 and 500

--and between boundary_values 500 and 1000. ALTER PARTITION FUNCTION myRangePF1 () SPLIT RANGE (500);

NEW QUESTION 82

- (Exam Topic 2)

Overview

General Overview

ADatum Corporation has offices in Miami and Montreal.


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

Databases

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

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

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



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 86

- (Exam Topic 2)

Overview

General Overview

ADatum Corporation has offices in Miami and Montreal.

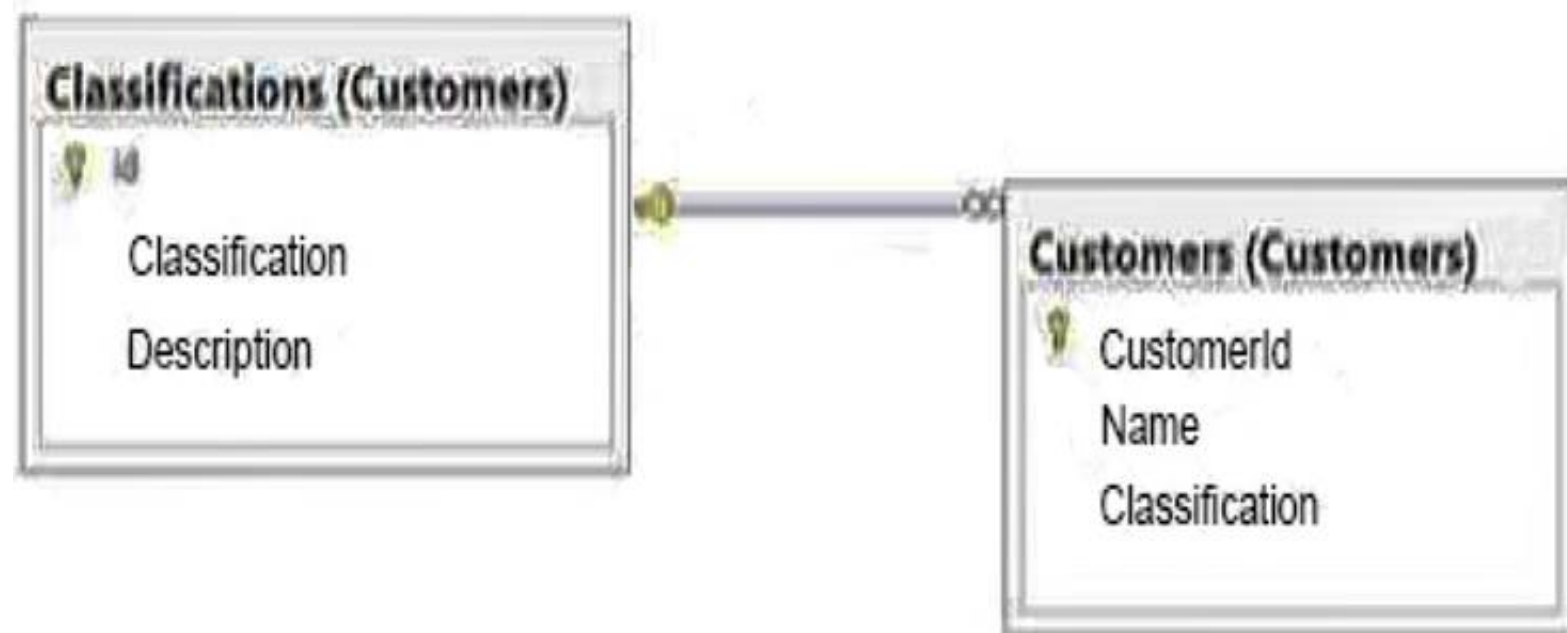
The network contains a single Active Directory forest named adatum.com. The offices connect to each other by using a WAN link that has 5-ms latency. A. Datum standardizes its database platform by using SQL Server 2014 Enterprise edition.

Databases

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

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

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



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

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

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

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

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

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

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

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

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

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

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

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

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

The nested stored procedures are never called directly. Design Requirements

Data Recovery

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

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

Classification Changes

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

Storage

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

Error Handling

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

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

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

Which change should you recommend?

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

Answer: A

Explanation:

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

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

NEW QUESTION 90

- (Exam Topic 2)

Overview

Application Overview

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

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

The new version will use SQL Server 2014.

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

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

Tables

The current database schema contains a table named OrderDetails.

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

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

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

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

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

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

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

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

Customer Problems Installation Issues

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

Index Fragmentation Issues

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

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

Backup Issues

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

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

Missing Data Issues

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

Query Performance Issues

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

Design Requirements

File Storage Requirements

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

Data Recovery Requirements

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

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

Concurrency Requirements

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

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

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

Answer: B

Explanation:

- Concurrency Requirements

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

NEW QUESTION 92

- (Exam Topic 2)

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

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

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

Answer: D

Explanation:

Troubleshooting deadlocks

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

Configures deadlock event tracking for the session.

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

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

NEW QUESTION 94

- (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 an isolation level for usp_UpdateOrderDetails.

Which isolation level should 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 95

- (Exam Topic 2)

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

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

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

Answer: A

Explanation:

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

NEW QUESTION 96

- (Exam Topic 2)

You administer a Microsoft SQL Server 2016 instance.

You need to stop a blocking process that has an SPID of 64 without stopping other processes. What should you do?

- A. Execute the following Transact-SQL statement: EXECUTE sp_KillSPID 64
- B. Restart the SQL Server service.
- C. Execute the following Transact-SQL statement: KILL 64
- D. Execute the following Transact-SQL statement: ALTER SESSION KILL '64'

Answer: C

Explanation:

KILL can be used to terminate a normal connection, which internally terminates the transactions that are associated with the specified session ID.

References:

<http://msdn.microsoft.com/en-us/library/ms173730.aspx>

NEW QUESTION 100

- (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 101

- (Exam Topic 2)

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

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

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

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

- A. Mastered
- B. Not Mastered

Answer: A

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

NEW QUESTION 104

- (Exam Topic 2)
You are designing a database for a university.
The database will contain two tables named Classes and Enrollment that have the following specifications:
Classes will store brochures in the XPS format. The brochures must be structured in folders and must be accessible by using UNC paths.
Enrollment will store information about students and their classes. Performance must be enhanced for queries of the current enrollments.
You need to identify which SQL Server technology meets the specifications of each table. Which technologies should you identify? To answer, drag the appropriate technology to the correct table in the answer area.

| Technologies | Answer area |
|--------------------|-----------------------|
| FileStream | Classes Technology |
| FileTable | Enrollment Technology |
| Partitioned tables | |
| Partitioned views | |

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:
The FileTable feature brings support for the Windows file namespace and compatibility with Windows applications to the file data stored in SQL Server. FileTable lets an application integrate its storage and data management components, and provides integrated SQL Server services including full-text search and semantic search-over unstructured data and metadata. In other words, you can store files and documents in special tables in SQL Server called FileTables, but access them from Windows applications as if they were stored in the file system, without making any changes to your client applications.

NEW QUESTION 108

- (Exam Topic 2)
You plan to deploy SQL Server 2014. Your company identifies the following monitoring requirements:
Tempdb must be monitored for insufficient free space.
Deadlocks must be analyzed by using Deadlock graphs.
You need to identify which feature meets each monitoring requirement.
Which features should you identify? To answer, drag the appropriate feature to the correct monitoring requirement in the answer area.

| Features | Answer area |
|-------------------------|---|
| Dynamic management view | Tempdb must be monitored for insufficient free space. |
| Activity Monitor | Deadlocks must be analyzed by using Deadlock graphs. |
| Resource Governor | |
| SQL Trace | |

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

You can use the sys.dm_db_file_space_usage dynamic management view to monitor the disk space used by the user objects, internal objects, and version stores in the tempdb files. Additionally, to monitor the page allocation or deallocation activity in tempdb at the session or task level, you can use the sys.dm_db_session_space_usage and sys.dm_db_task_space_usage dynamic management views. These views can be used to identify large queries, temporary tables, or table variables that are using a large amount of tempdb disk space.

Use SQL Server Profiler to identify the cause of a deadlock. A deadlock occurs when there is a cyclic dependency between two or more threads, or processes, for some set of resources within SQL Server. Using SQL Server Profiler, you can create a trace that records, replays, and displays deadlock events for analysis.

SQL Server Profiler and SQL Server Management Studio use a deadlock wait-for graph to describe a deadlock. The deadlock wait-for graph contains process nodes, resource nodes, and edges representing the relationships between the processes and the resources. References: Troubleshooting Insufficient Disk Space in tempdb

References: Analyze Deadlocks with SQL Server Profiler

NEW QUESTION 113

- (Exam Topic 2)

Overview

Application Overview

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

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

The new version will use SQL Server 2014.

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

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

Tables

The current database schema contains a table named OrderDetails.

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

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

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

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

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

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


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

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

Customer Problems Installation Issues

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

Index Fragmentation Issues

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

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

Backup Issues

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

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

Missing Data Issues

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

Query Performance Issues

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

Design Requirements

File Storage Requirements

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

Data Recovery Requirements

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

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

Concurrency Requirements

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

You need to recommend a solution that meets the data recovery requirement. What should you include in the recommendation?

- A. A differential backup
- B. A transaction log backup
- C. Snapshot isolation
- D. A database snapshot

Answer: D

NEW QUESTION 114

- (Exam Topic 2)

You plan to deploy SQL Server 2014.

Your company identifies the following monitoring requirements for the database:

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

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

You need to identify which feature meets each monitoring requirement.

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

| | | |
|--|---|---------|
| Policy-Based Management | An e-mail message must be sent when a user logs in. | Feature |
| A SQL Server Agent alert | | |
| SQL Server Integration Services (SSIS) | An e-mail message must be sent if CPU utilization exceeds 90 percent. | Feature |
| trace flags | | |
| triggers | | |

- A. Mastered
 B. Not Mastered

Answer: A

Explanation:

| | | |
|--|---|--------------------------|
| Policy-Based Management | An e-mail message must be sent when a user logs in. | Policy-Based Management |
| A SQL Server Agent alert | | |
| SQL Server Integration Services (SSIS) | An e-mail message must be sent if CPU utilization exceeds 90 percent. | A SQL Server Agent alert |
| trace flags | | |
| triggers | | |

NEW QUESTION 115

- (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 120

- (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 124

- (Exam Topic 3)

You need to recommend a solution to meet the recovery requirements for the manufacturing database. Your solution must minimize costs. What should you recommend?

- A. Database snapshots
- B. Transaction log backups
- C. Differential backups
- D. SQL Server Failover Clustering
- E. Peer-to-peer replication

Answer: A

NEW QUESTION 125

- (Exam Topic 3)

You need to move a database in between servers.

You need to guarantee database users will be able to login in the database with a minimum of management effort.

- A. Application role
- B. Database user
- C. Server login
- D. Server role

Answer: B

NEW QUESTION 129

- (Exam Topic 3)

You administer a Microsoft SQL Server 2012 instance.

You need to configure a new database to support FILETABLES. What should you do? Choose all that apply.

- A. Disable FILESTREAM on the Database.
- B. Enable FILESTREAM on the Server Instance.
- C. Configure the Database for Partial Containment.
- D. Create a non-empty FILESTREAM file group.
- E. Enable Contained Databases on the Server Instance.
- F. Set the FILESTREAM directory name on the Database.

Answer: BDF

Explanation:

References:

<http://msdn.microsoft.com/en-us/library/gg509097.aspx>

NEW QUESTION 131

- (Exam Topic 3)

You need to address the Sales Director's requirements regarding the customer classification. You need to recommend a solution for changing the classifications. What should you recommend?

- A. Add each classification change to a new row in the Customers table.
- B. Record each change to the classification of each customer in a new row in the Customers table.
- C. Add a new row to the Customers table for each new classification.
- D. Record each change to the classification of each customer in a new table in the Customers database.

Answer: D

NEW QUESTION 136

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

You need to configure a synchronization solution to copy data from the Current_Inventory database the DesABCopAppDB database.

What should you configure?

- A. Transactional Replication.
- B. Database Mirroring.
- C. Snapshot Replication.
- D. Incremental Backups

Answer: A

NEW QUESTION 139

- (Exam Topic 3)

You create an availability group that has replicas named HA/Server01 and HA/Server02. Currently, HA/Server01 is the primary replica.

You have multiple queries that read data and produce reports from the database.

You need to offload the reporting workload to the secondary replica when HA/Server01 is the primary replica. What should you do?

- A. Set the Availability Mode property of HA/Server02 to Asynchronous commit.
- B. Set the Readable Secondary property of HA/Server02 to Read-intent only.
- C. Set the Connections in Primary Role property of HA/Server01 to Allow read/write connections.
- D. Set the Availability Mode property of HA/Server01 to Asynchronous commit.

Answer: B

Explanation:

References: <https://msdn.microsoft.com/en-us/library/jj542414.aspx>

NEW QUESTION 141

- (Exam Topic 3)

You are using dynamic management views to monitor an SQL Server server named SQL1. A database administrator named Dba1 must monitor the health of SQL1.

You need to ensure that Dba1 can access dynamic management views for SQL1.

The solution must use the principle of least privilege. Which permissions should you assign to Dba1?

- A. VIEW ANY DEFINITION
- B. VIEW SERVER STATE
- C. VIEW DEFINITION
- D. CONTROL SERVER

Answer: B

Explanation:

To query a dynamic management view or function requires SELECT permission on object and VIEW SERVER STATE or VIEW DATABASE STATE permission.

References: <https://msdn.microsoft.com/en-us/library/ms188754.aspx>

NEW QUESTION 145

- (Exam Topic 3)

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

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

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

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

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

- A. Yes
- B. No

Answer: A

Explanation:

The error is related to the certificate.

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

NEW QUESTION 149

- (Exam Topic 3)

You administer a SQL Server 2014 instance.

Users report that the SQL Server has seemed slow today.

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

Number of users that have a connection to the server

Whether a user's connection is active

Whether any connections are blocked

What queries are being executed

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

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

Answer: D

Explanation:

- sys.dm_exec_requests

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

- sys.dm_exec_sessions

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

- sys.dm_exec_query_text

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

NEW QUESTION 152

- (Exam Topic 3)

Overview

General Overview

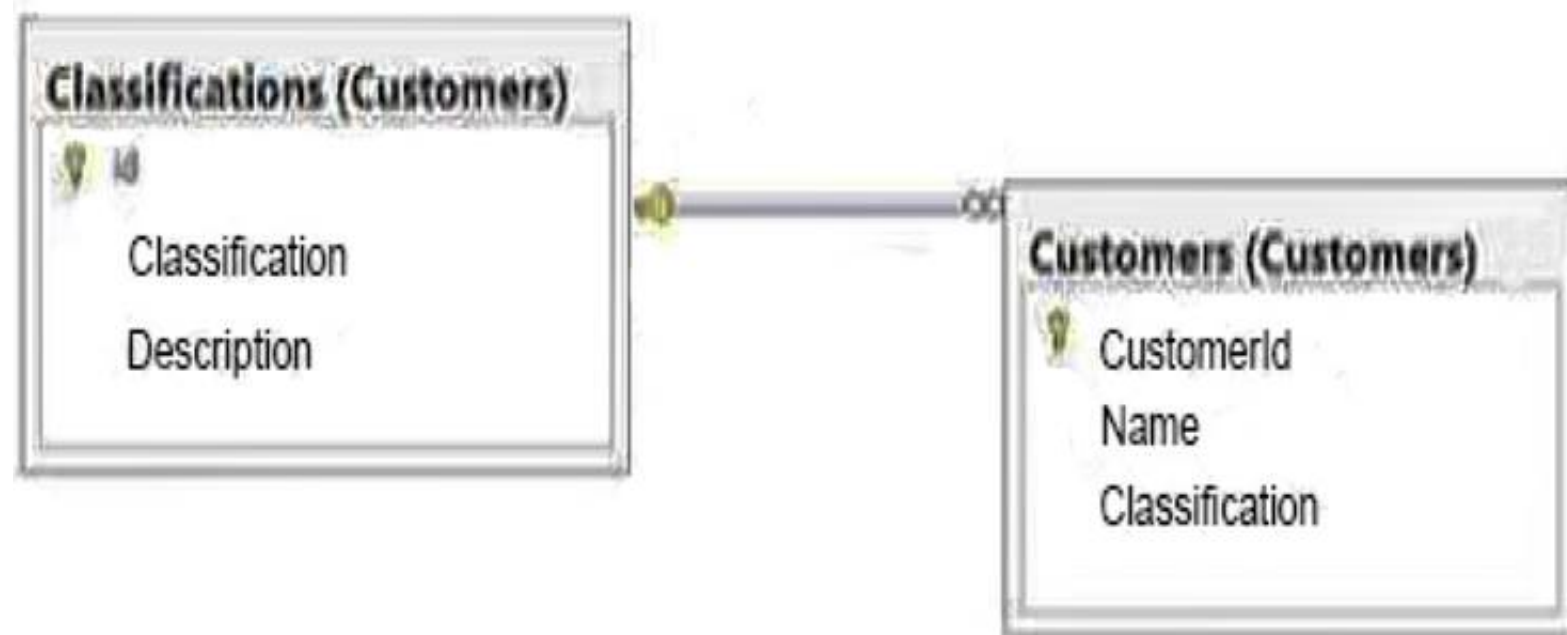
ADatum Corporation has offices in Miami and Montreal.

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

Databases

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

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



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

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

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

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

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

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

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

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

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

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

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

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

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

The nested stored procedures are never called directly. Design Requirements

Data Recovery

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

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

Classification Changes

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

Storage

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

Error Handling

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

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

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

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

Answer: D

Explanation:

Scenario:

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

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

NEW QUESTION 153

- (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 recommend a solution to back up DB1. What should you include in the recommendation?

- A. Azure Table Storage
- B. Azure Queue storage
- C. Azure Blob storage
- D. Azure Document DB

Answer: C

Explanation:

For SQL Server the Azure Blob Storage service offers a better alternative to the often used tape option to archive backups. Tape storage might require physical transportation to an off-site facility and measures to protect the media. Storing your backups in Azure Blob Storage provides an instant, highly available, and a durable archiving option.

References:

<https://azure.microsoft.com/en-us/documentation/articles/storage-use-storage-sql-server-backup/restore>

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 create the CustomerAccess table to support the reporting and performance requirements. Develop the solution by selecting and arranging the required code blocks in the correct order. You may not need all of the code blocks.

Code Blocks

```
CREATE TABLE [CustomerAccess]
ON [CustomerAccessFG] (
```

```
CustomerAccessID INT IDENTITY(1,1)
NOT NULL PRIMARY KEY
, CustomerID INT NOT NULL
, LoginDate DATETIME NOT NULL
, LogoffDate DATETIME NULL
```

```
CREATE TABLE [CustomerAccess] (
```

```
ALTER DATABASE [ProdDB] ADD FILE
(NAME = N'ProdDB_CustomerAccess',
FILENAME = N'F:\Data
\ProdDB_CustomerAccess.ndf') TO
FILEGROUP [CustomerAccessFG]
```

```
ALTER DATABASE [ProdDB] ADD FILE
(NAME = N'ProdDB_CustomerAccess',
FILENAME = N'D:\Data
\ProdDB_CustomerAccess.ndf') TO
FILEGROUP [CustomerAccessFG]
```

```
) ON [ProdDB_CustomerAccess]
```

```
) ON [CustomerAccessFG]
```

```
ALTER DATABASE [ProdDB] ADD FILE
(NAME = N'ProdDB_CustomerAccess',
FILENAME = N'D:\Data
\ProdDB_CustomerAccess.ndf')
```

```
)
```

```
ALTER DATABASE [ProdDB] ADD
FILEGROUP [CustomerAccessFG]
```

Answer Area

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Code Blocks

```
CREATE TABLE [CustomerAccess]
ON [CustomerAccessFG] (
```

```
CustomerAccessID INT IDENTITY(1,1)
NOT NULL PRIMARY KEY
, CustomerID INT NOT NULL
, LoginDate DATETIME NOT NULL
, LogoffDate DATETIME NULL
```

```
CREATE TABLE [CustomerAccess] (
```

```
ALTER DATABASE [ProdDB] ADD FILE
(NAME = N'ProdDB_CustomerAccess',
FILENAME = N'F:\Data
\ProdDB_CustomerAccess.ndf') TO
FILEGROUP [CustomerAccessFG]
```

```
ALTER DATABASE [ProdDB] ADD FILE
(NAME = N'ProdDB_CustomerAccess',
FILENAME = N'D:\Data
\ProdDB_CustomerAccess.ndf') TO
FILEGROUP [CustomerAccessFG]
```

```
) ON [ProdDB_CustomerAccess]
```

```
) ON [CustomerAccessFG]
```

```
ALTER DATABASE [ProdDB] ADD FILE
(NAME = N'ProdDB_CustomerAccess',
FILENAME = N'D:\Data
\ProdDB_CustomerAccess.ndf')
```

```
)
```

```
ALTER DATABASE [ProdDB] ADD
FILEGROUP [CustomerAccessFG]
```

Answer Area

```
ALTER DATABASE [ProdDB] ADD
FILEGROUP [CustomerAccessFG]
```

```
ALTER DATABASE [ProdDB] ADD FILE
(NAME = N'ProdDB_CustomerAccess',
FILENAME = N'F:\Data
\ProdDB_CustomerAccess.ndf') TO
FILEGROUP [CustomerAccessFG]
```

```
CREATE TABLE [CustomerAccess]
ON [CustomerAccessFG] (
```

```
CustomerAccessID INT IDENTITY(1,1)
NOT NULL PRIMARY KEY
, CustomerID INT NOT NULL
, LoginDate DATETIME NOT NULL
, LogoffDate DATETIME NULL
```

```
) ON [CustomerAccessFG]
```

NEW QUESTION 160

- (Exam Topic 3)

You use SQL Server 2014. You create a table within a database by using the following DDL:


```
CREATE TABLE OrderData
(
OrderID INT IDENTITY(1,1) Primary Key Clustered,
OrderDate SMALLDATETIME NOT NULL DEFAULT getdate(),
CustomerID INT,
IsTaxable INT,
SubTotal SmallMoney DEFAULT (0),
TaxAmount AS (Case IsTaxable when 1 then SubTotal * .0875 else NULL
END),
Freight SmallMoney,
OrderReturnedDate DATE,
OrderReturnedCustReason TEXT,
OrderReturnedEval Varchar(MAX)
)
```

The following table illustrates a representative sample of data:

| OrderID | OrderDate | CustomerID | IsTaxable | SubTotal | TaxAmount | Freight |
|---------|------------------|------------|-----------|----------|-----------|---------|
| 1 | 11/13/2013 11:22 | 58465 | NULL | \$ 25.99 | NULL | \$ 5.40 |
| 2 | 11/15/2013 9:34 | 12588 | NULL | \$ 42.00 | NULL | NULL |
| 3 | 12/1/2013 14:34 | 85477 | NULL | \$ 23.99 | NULL | \$ 4.85 |
| 4 | 12/17/2013 4:31 | 58742 | NULL | \$ 19.00 | NULL | NULL |
| 5 | 1/3/2014 8:22 | 12477 | NULL | \$ 13.50 | NULL | \$ 5.40 |
| 6 | 1/5/2014 18:39 | 63214 | NULL | \$ 5.69 | NULL | NULL |
| 7 | 1/15/2014 14:22 | 85471 | NULL | \$ 18.99 | NULL | \$ 7.85 |
| 8 | 1/19/2014 3:20 | 85412 | NULL | \$ 65.77 | NULL | NULL |
| 9 | 1/22/2014 13:44 | 12588 | NULL | \$ 22.38 | NULL | \$ 7.35 |
| 10 | 1/28/2014 10:14 | 85471 | 1 | \$ 24.99 | \$ 2.19 | \$ 5.40 |

The system is expected to handle 50 million orders a month over the next five years.

You have been instructed by your Team Lead to follow best practices for storage and performance in the utilization of SPARSE columns.

Which columns should you designate as SPARSE? To answer, mark each column as SPARSE or NOT SPARSE in the answer area.

Answer Area

| Column Names | Sparse | Not Sparse |
|--------------|-----------------------|-----------------------|
| OrderID | <input type="radio"/> | <input type="radio"/> |
| OrderDate | <input type="radio"/> | <input type="radio"/> |
| CustomerID | <input type="radio"/> | <input type="radio"/> |
| IsTaxable | <input type="radio"/> | <input type="radio"/> |
| SubTotal | <input type="radio"/> | <input type="radio"/> |
| TaxAmount | <input type="radio"/> | <input type="radio"/> |
| Freight | <input type="radio"/> | <input type="radio"/> |

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Sparse columns are ordinary columns that have an optimized storage for null values. Sparse columns reduce the space requirements for null values at the cost of more overhead to retrieve nonnull values. Consider using sparse columns when the space saved is at least 20 percent to 40 percent.

NEW QUESTION 161

- (Exam Topic 4)

You plan to implement a fault tolerance solution for a Microsoft SQL Server database. The solution must provide failover storage on the local network. You need to ensure that the solution can route traffic to failover storage by using SMB 3.0. Which storage option should you use?

- A. Cluster Shared Volumes
- B. Microsoft Azure Blob storage
- C. Always On availability group
- D. Stretch Database

Answer: A

Explanation:

Clustered Shared Volumes (CSV) is a new clustered file system in Windows Server that is a layer of abstraction above the NTFS file system in a WSFC environment. It allows all Nodes in the failover cluster to read and write to the CSV volume. CSV leverages the investments Microsoft have made in SMB 3.0, such as SMB Direct and SMB Multichannel. SQL Server 2014 was the first version of SQL Server to support CSVs. References: <https://www.microsoftpressstore.com/articles/article.aspx?p=2832586&seqNum=5>

NEW QUESTION 162

- (Exam Topic 4)

A company has the following Microsoft SQL Server instances Instance1 and Instance2. You plan to enable Always Encrypted for both instances. You need to configure the instances to meet the following requirements:

- ▶ Instance1 must use an initialization vector that is different each time the instance is initiated.
- ▶ Instance2 must use an initialization vector that is derived from an algorithm.

In the table below, identify the encryption type that must be used for each instance. NOTE: Make only one selection in each column. Each correct selection is worth one point.

| | Instance1 | Instance2 |
|--------------------------|-----------------------|-----------------------|
| Randomized encryption | <input type="radio"/> | <input type="radio"/> |
| Deterministic encryption | <input type="radio"/> | <input type="radio"/> |
| Call-level encryption | <input type="radio"/> | <input type="radio"/> |

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Always Encrypted supports two types of encryption: randomized encryption and deterministic encryption.

- ▶ Randomized encryption uses a method that encrypts data in a less predictable manner. Randomized encryption is more secure, but prevents searching, grouping, indexing, and joining on encrypted columns.
- ▶ Deterministic encryption always generates the same encrypted value for any given plain text value. Using deterministic encryption allows point lookups, equality joins, grouping and indexing on encrypted columns. However, but may also allow unauthorized users to guess information about encrypted values by examining patterns in the encrypted column, especially if there is a small set of possible encrypted values, such as True/False, or North/South/East/West region.

References: <https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/always-encrypted-database-engine>

NEW QUESTION 163

- (Exam Topic 4)

You have an application that queries a database. Users report that the application is slower than expected. You discover that several server process identifiers (SPIDs) have PAGELATCH_UP and PAGELATCH_EX waits. The resource descriptions of the SPIDs contains 2:1:1. You need to resolve the issue. What should you do?

- A. Allocate additional processor cores to the server.
- B. Add files to the file group of the application database.
- C. Reduce the fill factor of all clustered indexes.
- D. Add data files to tempdb.

Answer: D

Explanation:

PAGELATCH contention in tempdb is typically on allocation bitmaps and occurs with workloads with many concurrent connections creating and dropping small temporary tables (which are stored in tempdb). Assuming that the temporary tables are needed for performance, the trick is to have multiple data files for tempdb so that the allocations are done round-robin among the files, the contention is split over multiple PFS pages, and so the overall contention goes down. References: <https://sqlperformance.com/2015/10/sql-performance/knee-jerk-wait-statistics-pagelatch>

NEW QUESTION 167

- (Exam Topic 4)

You are implementing log shipping. You configure a secondary server in standby mode. You configure log shipping to occur hourly and keep up to seven days of log backups.

Users create reports by querying the secondary server. Transaction logs are not applied to the secondary server.

You need to ensure that the secondary server has current data at the beginning of each business day. What are two possible ways to achieve the goal? Each answer represents a complete solution.

- A. Create a snapshot database from a daily backup to use as the secondary databas
- B. Revert the snapshot database back to the previous business day.
- C. Disconnect users while transaction logs are applied.
- D. Configure the secondary server as a database mirroring target server.
- E. Allow restore operations to accumulat
- F. Run a Microsoft SQL Agent job that disconnects all users that are connected to the secondary server.

Answer: CD

NEW QUESTION 168

- (Exam Topic 4)

You are deploying a Microsoft Azure SQL Database environment:

You need to configure the environment to meet the following requirements:

- ▶ Access to databases must be logged.
- ▶ Email alerts must be sent to the database administrator if anomalous activity is detected.

What should you do? To answer, drag the appropriate solutions to the correct requirements. Each solution 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.

| Solutions | | Answer Area | |
|--------------------------------|---|-------------|--|
| | Requirement | Solution | |
| Enable database auditing | <div> <div>⬅</div> <div>Log database access</div> <div>➡</div> </div> | Solution | <div> <div>⬆</div> <div>⬇</div> </div> |
| Configure row-level security | | Solution | |
| Configure dynamic data masking | <div> <div>⬅</div> <div>Generate email alerts</div> <div>➡</div> </div> | | |
| Enable threat detection | | | |
| Create a custom role | | | |

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Enable database auditing Box 2: Enable threat detection

SQL Database Threat Detection runs multiple sets of algorithms which detect potential vulnerabilities and SQL injection attacks, as well as anomalous database access patterns (such as access from an unusual location or by an unfamiliar principal). Security officers or other designated administrators get email notification once a threat is detected on the database. Each notification provides details of the suspicious activity and recommends how to further investigate and mitigate the threat.

References:

NEW QUESTION 169

- (Exam Topic 4)

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

A company has a Microsoft SQL Server environment that has multiple databases. A database named DB1 has multiple online file groups. It is configured to use the full recovery model. A full backup is preformed nightly and transaction logs are performed on the hour. A large number of records are accidentally deleted at 17:20.

You need to perform a point-in-time recovery. Which option should you use first?

- A. backup compression
- B. backup encryption
- C. file snapshot backup
- D. mirrored backup media sets
- E. SQL Server backup to URL
- F. SQL Server Managed Backup to Azure

- G. tail-log backup
H. back up and truncate the transaction log

Answer: G

Explanation:

To back up the tail of the log (that is, the active log), check Back up the tail of the log, and leave database in the restoring state.

A tail-log backup is taken after a failure to back up the tail of the log in order to prevent work loss. Back up the active log (a tail-log backup) both after a failure, before beginning to restore the database, or when failing over to a secondary database. Selecting this option is equivalent to specifying the NORECOVERY option in the BACKUP LOG statement of Transact-SQL.

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/backup-restore/back-up-a-transaction-log-sql-server?vi>

NEW QUESTION 172

- (Exam Topic 4)

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 maintain a Microsoft SQL Server instance that contains the following databases SalesDb1, SalesDb2, and SalesDb3. Each database has table named Products and Sales. The following table shows the configuration of each database.

| Option of configuration | SalesDb1 | SalesDb2 | SalesDb3 |
|---------------------------------------|------------|---------------|------------|
| Recovery model | Full | Full | Simple |
| Query Store operation model | Read Write | Off | Off |
| Auto Update Statistics | True | False | False |
| Auto Update Statistics asynchronously | False | False | False |
| Sales data age | < 1 month | 1 to 6 months | > 6 months |

The backup strategies for each database are described in the following table.

| Database | Strategy | Backup file names |
|----------|---|---|
| SalesDb1 | Full database backups occur daily at 00:00. Log backups occur every hour. | SalesDb1Full_*.bak SalesDb1Log.bak |
| SalesDb2 | Full database backups occur every three months. Differential backups occur every month. Logs are not backed up. | SalesDb2Delta_*.bak SalesDb2Full_*.bak |
| SalesDb3 | Full database backups occur every five years. Differential backups occur every six months. | SalesDb3Delta_*.bak SalesDb3Full_*.bak |

Each full or differential backup operation writes into a new file and uses a different sequence number. You observe the following database corruption issues.

| Database | Error | Description |
|----------|-------|--|
| SalesDb2 | 824 | Some data pages that store table row data are torn. All backups for SalesDb2 are lost. |
| SalesDb3 | 823 | You observe bad checksum issues for data pages that store table row data. All backups are available. No new data has been added to the table since the latest differential backup. |

SalesDb3 reports a number of database corruption issues related to error 823 and 824 when reading data pages. You must display the following information about the corrupted pages:

- ▶ database name
- ▶ impacted file id
- ▶ impacted file physical name
- ▶ impacted page id
- ▶ event type that identifies the error type
- ▶ error count

Users report performance issues when they run queries against SalesDb2. You plan to monitor query statistics and execution plans for SalesDb2 by using Query Store. The monitoring strategy must meet the following requirements:

- ▶ Perform automatic data cleanup when query store disk usage reaches 500 megabyte (MB).
- ▶ Capture queries based on resource consumption.
- ▶ Use a stale query threshold value of 60 days.

The query optimizer generates suboptimal execution plans for a number of queries on the Sales table in SalesDb2. You will create a maintenance plan that updates statistics for the table. The plan should only update statistics that were automatically created and have not been updated for 30 days. The update should be based on all data in the table.

You need to view the information about the corrupted pages on SalesDb3.

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

Transact-SQL segments

Answer Area

misdb.dbo.suspect_pages

msdb.sys.sysfiles

SalesDb3.sys.sysfiles

master.sys.sysfiles

msdlb.sys.database_files

msdb.sys.dm_hadr_auto_page_repair

msdb.sys.dm_db_mirroring_auto_page_repair

```
SELECT DB_NAME(sp.database_id)AS database_name, sp.file_id,
f.filename AS File_name, sp.page_id, sp.event_type, sp.error_count
FROM [ ] sp
INNER JOIN [ ] ON f.fileid = sp.file_id
WHERE sp.event_type NOT IN(4, 5, 7) AND sp.database_id = ON_ID ('SalesDb 3')
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: msdb.dbo.suspect_pages

suspect_pages contains one row per page that failed with a minor 823 error or an 824 error. Pages are listed in this table because they are suspected of being bad, but they might actually be fine. When a suspect page is repaired, its status is updated in the event_type column.

The suspect_pages table resides in the msdb database. SalesDb3 has pages with checksum errors.

Box 2: msdb.sys.database_files

We want to identify these pages and which database they are in, this is easy enough to do when we join out to sys.databases and sys.master_files, as seen here:

```
SELECT d.name AS databaseName, mf.name AS logicalFileName, mf.physical_name AS physicalFileName, sp.page_id,
case sp.event_type
when 1 then N'823 or 824 error' when 2 then N'Bad Checksum' when 3 then N'Torn Page' when 4 then N'Restored'
when 5 then N'Repaired' when 7 then N'Deallocated' end AS eventType, sp.error_count, sp.last_update_date
```

from msdb.dbo.suspect_pages as sp
join sys.databases as d ON sp.database_id = d.database_id join sys.master_files as mf on sp.[file_id] = mf.[file_id] and d.database_id = mf.database_id;
The result of this query will give you a high level view of where you have potential corruption in your databases, from here it is important to use tools such as DBCC CHECKDB and your backups to recover from in line with your RPO and RTO.
References:
<https://docs.microsoft.com/en-us/sql/relational-databases/backup-restore/manage-the-suspect-pages-table-sql-ser>
<https://blogs.sentryone.com/johnmartin/monitoring-for-suspect-pages/>

NEW QUESTION 173

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual 70-764 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the 70-764 Product From:

<https://www.2passeasy.com/dumps/70-764/>

Money Back Guarantee

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 FirstTry
- * 70-764 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year