



Microsoft

Exam Questions 70-764

Administering a SQL Database Infrastructure (beta)

NEW QUESTION 1

- (Exam Topic 1)

You have a database that stores information for a shipping company. You create a table named Customers by running the following Transact-SQL statement. (Line numbers are included for reference only.)

```

01 CREATE TABLE dbo.Customers (
02     customerId int,
03     customerName varchar(200),
04     salesPerson varchar(20)
05 )
06 CREATE FUNCTION fn_securitypredicateSalesPerson (@salesPerson sysname)
07
08 AS
09 RETURN SELECT 1 AS [fn_securityPredicateOrder_result]
10 FROM dbo.Customers
11 WHERE @salesPerson = user_name()

```

You need to ensure that salespeople can view data only for the customers that are assigned to them. Which Transact-SQL segment should you insert at line 07?

- A. RETURNS varchar(20)WITH Schemabinding
- B. RETURNS dbo.CustomersORDER BY @salesPerson
- C. RETURNS tableORDER BY @salesPerson
- D. RETURNS tableWITH Schemabinding

Answer: D

Explanation:

The return value can either be a scalar (single) value or a table.

SELECT 1 just selects a 1 for every row, of course. What it's used for in this case is testing whether any rows exist that match the criteria: if a row exists that matches the WHERE clause, then it returns 1, otherwise it returns nothing.

Specify the WITH SCHEMABINDING clause when you are creating the function. This ensures that the objects referenced in the function definition cannot be modified unless the function is also modified.

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

NEW QUESTION 2

- (Exam Topic 1)

You administer a Microsoft SQL Server 2016 database named Orders.

Users report that during peak usage periods, certain operations are taking more time than expected. Your initial analysis suggests that blocking is the cause. You need to gather more data to be able to determine which processes are being blocked and to identify the root cause.

What should you do?

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

Answer: B

NEW QUESTION 3

- (Exam Topic 1)

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

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

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

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

- A. In the Database Mail Configuration Wizard, configure the Database Mail profile as a private profile for the AgentPortalUser account.
- B. Disable the guest user in the msdb database.
- C. Use the sysmail_help_profileaccount_sp stored procedure to add accounts to the Database Mail profile.
- D. In the Database Mail Configuration Wizard, create an email account for each recipient's email address in the Database Mail profile.

Answer: A

Explanation:

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

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

NEW QUESTION 4

- (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 5

- (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 6

- (Exam Topic 1)

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

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

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

Answer: F

NEW QUESTION 7

- (Exam Topic 1)

You are the database administrator of a Microsoft SQL Server instance. Developers are writing stored procedures to send emails using sp_send_dbmail. Database Mail is enabled.

You need to configure each account's profile security and meet the following requirements:

Account SMTP1_Account must only be usable by logins that have been given explicit permissions to use the SMTP1_profile.

Account SMTP2_Account must only be usable by logins who are a member of the [DatabaseMailUserRole] role in msdb.

In the table below, identify the profile type that must be used for each account. NOTE: Make only one selection in each column.

Answer Area

Profile type	SMTP1_Account	SMTP2_Account
Private Profile	<input type="radio"/>	<input type="radio"/>
Public Profile	<input type="radio"/>	<input type="radio"/>
Default Profile	<input type="radio"/>	<input type="radio"/>

A. Mastered

B. Not Mastered

Answer: A

Explanation:

SMTP1_Account1: Private Profile

When no profile_name is specified, sp_send_dbmail uses the default private profile for the current user. If user does not have a default private profile, sp_send_dbmail uses the default public profile for the msdb database.

SMTP1_Account2: Default Profile

Execute permissions for sp_send_dbmail default to all members of the DatabaseMailUser database role in the msdb database.

References:

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

NEW QUESTION 8

- (Exam Topic 1)

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

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

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.idevelopment.info/data/SQLServer/DBA_tips/Database_Administration/DBA_20.shtml

NEW QUESTION 9

- (Exam Topic 1)

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

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

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

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

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

Answer: G

Explanation:

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


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

NEW QUESTION 10

- (Exam Topic 1)

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

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



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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Transact-SQL segments

DATABASE

REPLICA

SYNCHRONOUS_COMMIT

ASYNCHRONOUS_COMMIT

PRIMARY

MANUAL

AUTOMATIC

SECONDARY_ONLY

● ● ● ●

Answer Area

```
ALTER AVAILABILITY GROUP AG_1 MODIFY Transact-SQL segment ON 'INSTANCE6'
WITH (AVAILABILITY_MODE = Transact-SQL segment );
ALTER AVAILABILITY GROUP AG_1 MODIFY Transact-SQL segment ON 'INSTANCE6'
WITH (FAILOVER_MODE = Transact-SQL segment );
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

Box 1: REPLICA

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

You must minimize latency between the nodes in AG1

AVAILABILITY_MODE = { SYNCHRONOUS_COMMIT | ASYNCHRONOUS_COMMIT }

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

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

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

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

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

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

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

AVAILABILITY_MODE = SYNCHRONOUS_COMMIT.

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

NEW QUESTION 10

- (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 12

- (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 15

- (Exam Topic 1)

You administer a Microsoft SQL Server 2016 environment.

One of the SQL Server 2016 instances contains a database named Sales. You plan to migrate Sales to Windows Azure SQL Database.

To do so, you need to implement a contained database.

What should you do? (Each correct answer presents part of the solution. Choose all that apply.)

- A. Set database containment to AZURE.
- B. Enable server property contained database authentication.
- C. Disable server property cross db ownership chaining.
- D. Set database containment to PARTIAL.
- E. Disable server property contained database authentication.
- F. Set database containment to FULL.

Answer: BD

NEW QUESTION 19

- (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 24

- (Exam Topic 1)

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

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

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

Answer Area

Actions	Options
User1 server role assignment	<div style="border: 1px solid black; padding: 5px;"> <div style="background-color: #f0f0f0; padding: 2px; text-align: right;">▼</div> <div style="padding: 2px;"> diskadmin serveradmin securityadmin setupadmin </div> </div>
Transact-SQL syntax	<div style="border: 1px solid black; padding: 5px;"> <div style="background-color: #f0f0f0; padding: 2px; text-align: right;">▼</div> <div style="padding: 2px;"> sys.server_file_audits sys.server_audit_specifications sys.server_file_permissions sys.server_principals </div> </div>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: securityadmin

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

Box 2: sys.server_audit_specifications

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

NEW QUESTION 26

- (Exam Topic 1)

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

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

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

You need to address the performance issues with each table.

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

Answer Area	Action	TableA	TableB
	Run the following Transact-SQL statement: SET AUTO_UPDATE_STATISTICS_ASYNC ON	<input type="radio"/>	<input type="radio"/>
	Run the following Transact-SQL statement: SET AUTO_UPDATE_STATISTICS OFF	<input type="radio"/>	<input type="radio"/>
	Run the following Transact-SQL statement and then recreate all indexes and statistics using the INCREMENTAL keyword: SET AUTO_CREATE_STATISTICS ON (INCREMENTAL = ON)	<input type="radio"/>	<input type="radio"/>
	Run the sp_updatestats procedure instead of the following Transact-SQL statement: UPDATE STATISTICS	<input type="radio"/>	<input type="radio"/>

- A. Mastered

B. Not Mastered

Answer: A

Explanation:

Table A: Auto_update statistics off

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

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

SET AUTO_UPDATE_STATISTICS_ASYNC ON

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

References:

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

NEW QUESTION 31

- (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 34

- (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 style="border: 1px solid gray; padding: 5px;"> <div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> ▼ </div> <ul style="list-style-type: none"> Primary server instance Secondary server instance Monitor server instance Backup share file server </div>
Copy the backup job.	<div style="border: 1px solid gray; padding: 5px;"> <div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> ▼ </div> <ul style="list-style-type: none"> Primary server instance Secondary server instance Monitor server instance Backup share file server </div>
Restore the backup.	<div style="border: 1px solid gray; padding: 5px;"> <div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> ▼ </div> <ul style="list-style-type: none"> 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 37

- (Exam Topic 1)

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

You need to ensure that data changes are sent to a non-SQL Server database server in near real time. You also need to ensure that data on the primary server is unaffected.

Which configuration should you use?

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

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

Answer: A

Explanation:

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

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

Publishing data to and from Oracle has some restrictions.

References:

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

NEW QUESTION 38

- (Exam Topic 1)

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

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

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

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

Answer: B

Explanation:

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

References:

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

NEW QUESTION 41

- (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 Secondary
- B. Set the backup priority of Server02/HA to 20. Set the backup priority of Server03/HA to 10.
- C. Set the backup preference of HaContoso to Secondary only
- D. Set the backup priority of Server02/HA to 20. Set the backup priority of Server03/HA to 10.
- E. Set the backup preference of HaContoso to Secondary only
- F. Set the backup priority of Server02/HA to 10. Set the backup priority of Server03/HA to 20.
- G. set the exclude replica of Server01/HA to true
- 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 46

- (Exam Topic 1)

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

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

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

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1: Launch the Failover Cluster Manager and..

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

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

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

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

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

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

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

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

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

NEW QUESTION 49

- (Exam Topic 1)

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

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

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

You need to perform the following actions:

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

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

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

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

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

NEW QUESTION 52

- (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 54

- (Exam Topic 1)

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

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

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

You need to perform the following actions:

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

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

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

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

Does the solution meet the goal?

- A. Yes
- B. NO

Answer: B

NEW QUESTION 56

- (Exam Topic 1)

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

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

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

What should you use?

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

Answer: D

Explanation:

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

NEW QUESTION 58

- (Exam Topic 1)

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

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

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

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

Answer: H

Explanation:

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

References:

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

NEW QUESTION 60

- (Exam Topic 1)

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

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

What should you specify when you run the recovery command?

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

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

Answer: A

NEW QUESTION 65

- (Exam Topic 1)

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

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

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

Answer: A

Explanation:

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

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

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

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

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

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

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

NEW QUESTION 68

- (Exam Topic 2)

You are migrating a database named Orders to a new server that runs Microsoft SQL Server 2016. You attempt to add the [Corpnet\User1] login to the database. However, you receive the following error message: "User already exists in current database."

You need to configure the [Corpnet\User1] login to be able to access the Orders database and retain the original permissions.

You need to achieve this goal by using the minimum required permissions. Which Transact-SQL statement should you use?

- A. DROP USER [User1]; CREATE USER [Corpnet\User1] FOR LOGIN [Corpnet\User1]; ALTER ROLE [db_owner] ADD MEMBER [Corpnet\User1];
- B. ALTER SERVER RCL [sysadmin] ADD MEMBER [Corpnet\User1];
- C. ALTER USER [Corpnet\User1] WITH LOGIN [Corpnet\User1];
- D. ALTER ROLE [db_owner] ADD MEMBER [Corpnet\User1];

Answer: C

NEW QUESTION 71

- (Exam Topic 2)

You are designing a monitoring application for a new SQL Server 2014 instance.

You need to recommend a solution to generate a report that displays the 10 most frequent wait types that occur for the instance.

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

- A. The SQL Server error log
- B. The sys.dm_os_wait_stats dynamic management view
- C. The DBCC SQLPERF(WAITSTATS) command
- D. SQL Server Profiler

Answer: B

Explanation:

sys.dm_os_wait_stats

Returns information about all the waits encountered by threads that executed. You can use this aggregated view to diagnose performance issues with SQL Server and also with specific queries and batches.

Columns include: waiting_tasks_count

Number of waits on this wait type.

This counter is incremented at the start of each wait.

NEW QUESTION 75

- (Exam Topic 2)

Overview

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

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

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

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

Requirements Planned Changes

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

All database files will be stored in a highly available SAN. Database1 will contain two tables named Orders and OrderDetails. Database1 will also contain a stored procedure named usp_UpdateOrderDetails. The stored procedure is used to update order information. The stored procedure queries the Orders table twice each time the procedure executes. The rows returned from the first query must be returned on the second query unchanged along with any rows added to the table between the two read operations. Database1 will contain several queries that access data in the Database2 tables. Database2 will contain a table named Inventory. Inventory will contain over 100 GB of data. The Inventory table will have two indexes: a clustered index on the primary key and a nonclustered index. The column that is used as the primary key will use the identity property. Database2 will contain a stored procedure named usp_UpdateInventory. usp_UpdateInventory will manipulate a table that contains a self-join that has an unlimited number of hierarchies. All data in Database2 is recreated each day and does not change until the next data creation process. Data from Database2 will be accessed periodically by an external application named Application1. The data from Database2 will be sent to a database named Appl_Dbl as soon as changes occur to the data in Database2. Litware plans to use offsite storage for all SQL Server 2014 backups.

Business Requirements
 You have the following requirements:

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

You must encrypt the backup files to meet regulatory compliance requirements. The encryption strategy must minimize changes to the databases and to the applications. You need to recommend a solution to improve the performance of usp_UpdateInventory. The solution must minimize the amount of development effort. What should you include in the recommendation?

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

Answer: A

Explanation:

- Scenario: Database2 will contain a stored procedure named usp_UpdateInventory. Usp_UpdateInventory will manipulate a table that contains a self-join that has an unlimited number of hierarchies.
- A table variable can be very useful to store temporary data and return the data in the table format.
- Example: The following example uses a self-join to find the products that are supplied by more than one vendor. Because this query involves a join of the ProductVendor table with itself, the ProductVendor table appears in two roles. To distinguish these roles, you must give the ProductVendor table two different aliases (pv1 and pv2) in the FROM clause. These aliases are used to qualify the column names in the rest of the query. This is an example of the self-join Transact-SQL statement:

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

NEW QUESTION 76

- (Exam Topic 2)

Overview

General Overview


ADatum Corporation has offices in Miami and Montreal.

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

Databases

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

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



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

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

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

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

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

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

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

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

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

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

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

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

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

The nested stored procedures are never called directly. Design Requirements

Data Recovery

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

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

Classification Changes

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

Storage

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

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

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

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

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

Answer: A

Explanation:

Scenario:

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

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

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

NEW QUESTION 79

- (Exam Topic 2)

Overview

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

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

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

will use the application.

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

Requirements Planned Changes

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

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

Database1 will also contain a stored procedure named usp_UpdateOrderDetails.

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

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

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

Inventory will contain over 100 GB of data.

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

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

Business Requirements

You have the following requirements:

Costs for new licenses must be minimized.

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

Development effort must be minimized whenever possible.

The storage requirements for databases must be minimized.

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

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

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

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

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

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

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

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

Answer: C

Explanation:

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

NEW QUESTION 81

- (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. What should you recommend for the updates to Sales.TransactionHistory?

- A. a REPEATABLE READ isolation level
- B. implicit transactions
- C. query hints
- D. a SNAPSHOT isolation level

Answer: A

NEW QUESTION 82

- (Exam Topic 2)

You are designing a database named DB1.

Changes will be deployed to DB1 every Wednesday night.

You need to recommend a strategy to deploy the changes to DB1. The strategy must meet the following requirements:

The strategy must not disrupt backup operations.

DB1 must be unavailable to users while the changes are deployed.

You must be able to undo quickly the entire operation.

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

- A. Perform a copy-only database backup before the changes are deployed.If the deployment fails, restore the database to another server and recover the original.Objects from the restored database.
- B. Create a database snapshot.If the deployment fails, recover the objects from the database snapshot.
- C. Create a database snapshot.If the deployment fails, revert the database to the database snapshot.
- D. Perform a full database backup before the changes are deployed.If the deployment fails, restore the database to another server and recover the original objects from the restored database.

Answer: C

NEW QUESTION 84

- (Exam Topic 2)

You have a SQL Azure database named Database1.

You need to design the schema for a table named table1. Table1 will have less than one million rows.

Table1 will contain the following information for each row:

Column	Description
ID	An incremental numeric value used to identify the row
Name	A string in English
Code	An alphanumeric code that has five characters
ModifiedDate	The date of the last modification

The solution must minimize the amount of space used to store each row.

Which data types should you recommend for each column? To answer, drag the appropriate data type to the correct column in the answer area.

Data Types

- int
- bigint
- varchar
- nvarchar
- char
- smalldatetime
- date

Answer Area

ID	Data type
Name	Data type
Code	Data type
ModifiedDate	Data type

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

References:

<http://msdn.microsoft.com/en-US/library/ms187752.aspx>

NEW QUESTION 89

- (Exam Topic 2)

You plan to deploy a database by using SQL Server 2014. Your company identifies the following requirements for the database:

The name of all stored procedures must start with "usp_"always.

All distribution statistics must be updated daily

You need to identify which feature must be used to meet each database requirement.

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

Features	Answer Area
Change data capture	The name of all stored procedures must start with "usp_" always.
The CHECK constraint	
Extended Event	All distribution statistics must be updated daily.
A maintenance plan	
Policy-Based Management	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

- Policy-Based Management Each Stored Procedure that are created and that will be created has to have prefix "USP_".
- Maintenance plans create a workflow of the tasks required to make sure that your database is optimized, regularly backed up, and free of inconsistencies.

NEW QUESTION 94

- (Exam Topic 2)

Your company has offices in Seattle and Montreal.

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

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

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

What should you recommend?

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

Answer: D

NEW QUESTION 96

- (Exam Topic 2)

You administer a Microsoft SQL Server 2016 instance.

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

What should you do?

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

Answer: B

NEW QUESTION 97

- (Exam Topic 2)

Overview

General Overview

ADatum Corporation has offices in Miami and Montreal.


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

Databases

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

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

Classifications (Customers)



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

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

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

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

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

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

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

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

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

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

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

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

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

The nested stored procedures are never called directly. Design Requirements

Data Recovery

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

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

Classification Changes

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

Storage

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

Error Handling

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

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

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

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

Answer: A

Explanation:

Scenario:

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

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

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

NEW QUESTION 102

- (Exam Topic 2)

You administer a Microsoft SQL Server 2016 failover cluster.

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

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

D. SqlDumperDumpPath

Answer: B

Explanation:

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

References:

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

NEW QUESTION 107

- (Exam Topic 2)

Overview

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

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

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

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

Requirements Planned Changes

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

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

Database1 will also contain a stored procedure named usp_UpdateOrderDetails.

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

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

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

Inventory will contain over 100 GB of data.

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

Database2 will contain a stored procedure named usp_UpdateInventory. usp_UpdateInventory will manipulate a table that contains a self-join that has an unlimited number of hierarchies. All data in Database2 is recreated each day and does not change until the next data creation process. Data from Database2 will be accessed periodically by an external application named Application1. The data from Database2 will be sent to a database named App1_Db1 as soon as changes occur to the data in Database2. Litware plans to use offsite storage for all SQL Server 2014 backups.

Business Requirements

You have the following requirements:

Costs for new licenses must be minimized.

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

Development effort must be minimized whenever possible.

The storage requirements for databases must be minimized.

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

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

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

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

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

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

You need to recommend a solution to synchronize Database2 to App1_Db1. What should you recommend?

- A. Change data capture
- B. Snapshot replication
- C. Master Data Services
- D. Transactional replication

Answer: D

Explanation:

Scenario:

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

Application1. The data from Database2 will be sent to a database named App1_Db1 as soon as changes occur to the data in Database2.

- All data in Database2 is recreated each day and does not change until the next data creation process.

NEW QUESTION 108

- (Exam Topic 2)

Overview

Application Overview

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

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

The new version will use SQL Server 2014.

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

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

Tables

The current database schema contains a table named OrderDetails.

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

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

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

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

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

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

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

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

Customer Problems Installation Issues

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

Index Fragmentation Issues

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

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

Backup Issues

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

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

Missing Data Issues

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

Query Performance Issues

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

Design Requirements

File Storage Requirements

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

Data Recovery Requirements

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

Security Requirements

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

Concurrency Requirements

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

What should you include in the recommendation?

- A. FileStream
- B. FileTable
- C. The varbinary data type
- D. The image data type

Answer: B

Explanation:

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

- FileTables remove a significant barrier to the use of SQL Server for the storage and management of unstructured data that is currently residing as files on file servers.

Enterprises can move this data from file servers into FileTables to take advantage of integrated administration and services provided by SQL Server. At the same time, they can maintain Windows application compatibility for their existing Windows applications that see this data as files in the file system.

NEW QUESTION 109

- (Exam Topic 2)

You plan to deploy a database to SQL Azure. You are designing two stored procedures named USP_1 and USP_2 that have the following requirements:
 Prevent data read by USP_1 from being modified by other active processes.
 Allow USP_2 to perform dirty reads.
 You need to recommend the isolation level for the stored procedures. The solution must maximize concurrency.
 Which isolation levels should you recommend? To answer, drag the appropriate isolation level to the correct stored procedure in the answer area.

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

SP1 – repeatable read; SP2 – read uncommitted Note:

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

NEW QUESTION 111

- (Exam Topic 2)

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

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

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

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

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

SP1 – repeatable read; SP2 – read committed

- REPEATABLE READ

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

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

NEW QUESTION 114

- (Exam Topic 2)

You are designing an authentication strategy for a new server that has SQL Server 2014 installed. The strategy must meet the following business requirements:

The account used to generate reports must be allowed to make a connection during certain hours only.

Failed authentication requests must be logged.

You need to recommend a technology that meets each business requirement. The solution must minimize the amount of events that are logged. Which technologies should you recommend? To answer, drag the appropriate solution to the correct business requirement in the answer area.

Isolation Levels	Answer area
Login auditing	The account used to generate reports must be allowed to make a connection during certain hours only. <input type="text" value="Technology"/>
Logon triggers	
C2 audit tracing	Failed authentication requests must be logged. <input type="text" value="Technology"/>
Policy-Based Management	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

1. Logon triggers fire stored procedures in response to a LOGON event. This event is raised when a user session is established with an instance of SQL Server. Logon triggers fire after the authentication phase of logging in finishes, but before the user session is actually established. You can use logon triggers to audit and control server sessions, such as by tracking login activity, restricting logins to SQL Server, or limiting the number of sessions for a specific login.
2. Login auditing can be configured to write to the error log on the following events.
 - Failed logins
 - Successful logins
 - Both failed and successful logins

NEW QUESTION 115

- (Exam Topic 2)

Overview

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

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

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

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

Requirements Planned Changes

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

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

Database1 will contain two tables named Orders and OrderDetails.

Database1 will also contain a stored procedure named usp_UpdateOrderDetails.

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

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

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

Inventory will contain over 100 GB of data.

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

Database2 will contain a stored procedure named usp_UpdateInventory. usp_UpdateInventory will manipulate a table that contains a self-join that has an unlimited number of hierarchies. All data in Database2 is recreated each day and does not change until the next data creation process. Data from Database2 will be accessed periodically by an external application named Application1. The data from Database2 will be sent to a database named App1_Db1 as soon as changes occur to the data in Database2. Litware plans to use offsite storage for all SQL Server 2014 backups.

Business Requirements

You have the following requirements:

Costs for new licenses must be minimized.

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

Development effort must be minimized whenever possible.

The storage requirements for databases must be minimized.

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

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

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

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

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

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

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

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

Answer: D

Explanation:

- Scenario:

- Data from Database2 will be accessed periodically by an external application named Application1
 - Application developers must be denied direct access to the database tables. Applications must be denied direct access to the tables.
- Tutorial: Signing Stored Procedures with a Certificate

NEW QUESTION 116

- (Exam Topic 2)

You plan to create a database.

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

After the event, the database will be deleted.

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

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

- A. SQL Server 2014 Enterprise
- B. SQL Server 2014 Standard
- C. SQL Azure
- D. SQL Server 2014 Express with Advanced Services

Answer: B

Explanation:

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

NEW QUESTION 118

- (Exam Topic 2)

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

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

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

Answer: B

Explanation:

Transparent Data Encryption (TDE) is a special case of encryption using a symmetric key. TDE encrypts an entire database using that symmetric key called the database encryption key. The database encryption key is protected by other keys or certificates which are protected either by the database master key or by an asymmetric key stored in an EKM module.

SQL Server provides the following mechanisms for encryption:

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

NEW QUESTION 120

- (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 changes to the ERP application to resolve the search issue.

The solution must minimize the impact on other queries generated from the ERP application. What should you recommend changing?

- A. The collation of the Products table
- B. The index on the ProductName column
- C. The collation of the ProductName column
- D. The data type of the ProductName column

Answer: C

Explanation:

References:

[http://technet.microsoft.com/en-us/library/aa214408\(v=sql.80\).aspx](http://technet.microsoft.com/en-us/library/aa214408(v=sql.80).aspx)

NEW QUESTION 121

- (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 123

- (Exam Topic 2)

Your company has offices in Seattle and Montreal.

The network contains two servers named Server1 and Server2 that have SQL Server 2012 installed. Server1 is located in the Seattle office. Server2 is located in the Montreal office. The latency of the WAN link between the Montreal office and the Seattle office is more than 200 ms.

You plan to implement an AlwaysOn availability group on both servers.

You need to recommend a failover type for the availability group. What should you recommend?

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

Answer: D

NEW QUESTION 128

- (Exam Topic 2)

You have a SQL Server 2014 instance named SQL1. SQL1 creates error events in the Windows Application event log.

You need to recommend a solution that will run an application when SQL1 logs a specific error in the Application log.

Which SQL elements should you include in the recommendation? (Each correct answer presents part of the solution. Choose all that apply.)

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

Answer: DE

Explanation:

Use a trigger that starts a job which executes the application.

References:

<http://technet.microsoft.com/en-us/library/hh849759.aspx>

NEW QUESTION 129

- (Exam Topic 2)

You administer a Microsoft SQL Server 2016 database that has multiple tables in the Sales schema. Some users must be prevented from deleting records in any of the tables in the Sales schema. You need to

manage users who are prevented from deleting records in the Sales schema.

You need to achieve this goal by using the minimum amount of administrative effort. What should you do?

- A. Create a custom database role that includes the user
- B. Deny Delete permissions on the Sales schema for the custom database role.
- C. Include the Sales schema as an owned schema for the db_denydatawriter rol
- D. Add the users to the db_denydatawriter role.
- E. Deny Delete permissions on each table in the Sales schema for each user.
- F. Create a custom database role that includes the user
- G. Deny Delete permissions on each table in the Sales schema for the custom database role.

Answer: A

NEW QUESTION 131

- (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.

A. Mastered

B. Not Mastered

Answer: A

Explanation:

Policy-Based Management

A SQL Server Agent alert

SQL Server Integration Services (SSIS)

trace flags

triggers

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

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

NEW QUESTION 136

- (Exam Topic 2)

You administer a Microsoft SQL Server 2016 database named Contoso on a server named Server01. You need to write messages to the Application Log when users are added to or removed from a fixed server role in Server01. What should you create?

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

Answer: G

Explanation:

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

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

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

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

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

NEW QUESTION 138

- (Exam Topic 2)

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

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

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

Answer: A

Explanation:

References:

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

NEW QUESTION 140

- (Exam Topic 3)

You need to ensure that a stored procedure fails if an INSERT statement within the stored procedure fails. What action should you take?

- A. THROW 51000, 'Abort!'
- B. SET XACT_ABORT OFF
- C. SET XACT_ABORT ON
- D. TRY....CATCH

Answer: C

NEW QUESTION 144

- (Exam Topic 3)

You have an SQL Server 2014 server named SQL1. You are designing a performance monitoring solution. You need to monitor the following events on SQL1:

- A deadlock graph
- Missing column statistics
- CPU performance statistics
- A batch of completed Transact-SQL statements

Which two tools should you use? Each correct answer presents a complete solution.

- A. dynamic management views
- B. Database Engine Tuning Advisor
- C. SQL Server Profiler
- D. Activity Monitor
- E. Data Profile Viewer

Answer: BC

Explanation:

B: Database Engine Tuning Advisor examines how queries are processed in the databases you specify. When you run a Profiler Trace and feed it to the Database Engine Tuning Advisor, it also looks for missing column statistics, and it can automatically create them for you. C: 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.

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

NEW QUESTION 149

- (Exam Topic 3)

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

SQL Server auditing must be enabled on all server instances.

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

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

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

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

- A. Grant the manage auditing and security log permission to the SQL Server service account.
- B. Grant the generate security audits permission on the SQL Server service account.
- C. Update Windows security policy to audit object access.
- D. Restart the SQL Server Agent service.

Answer: BC

Explanation:

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

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

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

References:

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

NEW QUESTION 152

- (Exam Topic 3)

General Overview

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

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

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

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

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

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

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

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

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

SPUpdate2.

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

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

The GenInfo table is used for reports.

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

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

Current System

The current desktop application uses data stored in a SQL Server 2005 database named DesABCopAppDB.

This database will remain online and data from the Current_Inventory database will be copied to it as soon as data is changed in the Current_Inventory database.

SQL Servers

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

Design Requirements

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

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

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

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

Deadlocks must be analyzed with the use of Deadlock Graphs.

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

Software licensing and database storage costs must be minimized.

Development effort must be minimized.

The Tempdb databases must be monitored for insufficient free space.

Failed authentication requests must be logged.

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

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

SPUpdateSalesInfo queries data in the OrderTotals table the second time.

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

What should you configure?

- A. You should configure a server role.
- B. You should configure a database role.
- C. You should configure functions that use the EXECUTE AS statement.
- D. You should configure stored procedures that use the EXECUTE AS statement.

Answer: D

NEW QUESTION 156

- (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.

You plan to create a stored procedure to access data in CorpDB.

You need to ensure that the stored procedure supports dirty reads. What should you do?

- A. You should configure the stored procedure to use the READ UNCOMMITTED isolation level.
- B. You should configure the stored procedure to use the READ COMMITTED isolation level.
- C. You should configure the stored procedure to use the READ SERIALIZABLE isolation level.
- D. You should configure the stored procedure to use the REPEATABLE READ isolation level.
- E. You should configure the stored procedure to use the SNAPSHOT isolation level.

Answer: A

NEW QUESTION 157

- (Exam Topic 3)


You have an SQL Server 2014 server. You plan to create four stored procedures that will use transactions. The stored procedures will be configured as shown in the following table.

Stored procedure name	Concurrency strategy
SP1	A transaction will only recognize data changes made before the start of the transaction. Other transactions are allowed to modify data.
SP2	A transaction can read data made from another transaction that has not been committed.
SP3	A transaction can only read data made from another transaction that has been committed.
SP4	During a transaction, no other transactions can modify data.

You need to recommend an isolation level for each stored procedure. The solution must support the concurrency strategy of each stored procedure and must minimize locks.

What should you recommend?

To answer, drag the appropriate isolation levels to the correct stored procedures. Each isolation level 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.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

SP1 – SNAPSHOT, SP2 – REPEATABLE READ, SP3 – READ COMMITTED, SP4 - SERIALIZABLE

References: SET TRANSACTION ISOLATION LEVEL (Transact-SQL) <https://msdn.microsoft.com/en-us/library/ms173763.aspx>

NEW QUESTION 159

- (Exam Topic 3)

You work as a Database Administrator (DBA) at ABC.com.

All databases are hosted on Windows Server 2012 servers running SQL Server 2012. The Sales department uses a database named SalesDB.

SalesDB contains a large table named Orders that lists every order ever received by the company. You want to improve the performance of SalesDB.

You want to configure the database to provide the fastest possible access to the most recent orders. Historical orders can be stored using a slower storage solution.

How can you achieve this goal?

- A. By configuring database mirroring.
- B. By configuring a failover cluster.
- C. By partitioning the Orders table.
- D. By partitioning a partitioned view of the Orders table.

Answer: C

NEW QUESTION 162

- (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 a table named SalesInfo.

You discover that some table has been deleted from the SalesInfo table. You are unable to find out who deleted the information.

You need to implement a solution to monitor the deletion of any further information from the SalesInfo table. You want to minimize the development effort required for the solution.

What should you configure?

- A. You should configure table permissions.
- B. You should configure a user role.
- C. You should configure change data capture.
- D. You should configure a trigger.

Answer: C

NEW QUESTION 163

- (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 166

- (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 a job to automate some database maintenance tasks. Which code fragment should you use in each location in the command to complete one of the commands you will need to include in the job? To answer, drag the appropriate lines of code to the correct locations in the command. Each line of code 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.

Lines of Code	Answer Area
STATISTICS	UPDATE Line of Code Production.Product
CONSTRAINT	Line of Code Line of Code
MEMBER	
WITH FULLSCAN	
WITH SAMPLE 50 PERCENT	
WITH SAMPLE 1000 ROWS	
RECOMPUTE	
NOCOMPUTE	
NORECOMPUTE	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Lines of Code	Answer Area
STATISTICS	UPDAT. STATISTICS Production.Product
CONSTRAINT	
MEMBER	
WITH FULLSCAN	WITH FULLSCAN NORECOMPUTE
WITH SAMPLE 50 PERCENT	
WITH SAMPLE 1000 ROWS	
RECOMPUTE	
NOCOMPUTE	
NORECOMPUTE	

NEW QUESTION 167

- (Exam Topic 3)

You have a server named Server1 that has 2 processors. You plan to deploy multiple instances of SQL Server 2014 to Server1. Each instance will have multiple databases.

You need to recommend a method to allocate processor time to each database.

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

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

Answer: A

Explanation:

SQL Server Resource Governor is a feature that you can use to manage SQL Server workload and system resource consumption. Resource Governor enables you to specify limits on the amount of CPU, physical IO, and memory that incoming application requests can use.

NEW QUESTION 171

- (Exam Topic 3)

General Overview

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

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

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

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

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

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

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

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

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

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

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

The GenInfo table is used for reports.

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

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

Current System

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

SQL Servers

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

Design Requirements

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

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

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

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

Deadlocks must be analyzed with the use of Deadlock Graphs.

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

Software licensing and database storage costs must be minimized.

Development effort must be minimized.

The Tempdb databases must be monitored for insufficient free space.

Failed authentication requests must be logged.

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

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

You need to meet the design requirement for the ProductTypes table in the Product database. Which of the following would be the best solution?

- A. A PRIMARY KEY constraint.
- B. A CHECK constraint.
- C. A UNIQUE constraint.
- D. A Data Definitions Language (DDL) trigger.
- E. A FOREIGN KEY constraint.

Answer: B

NEW QUESTION 176

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

70-764 Practice Exam Features:

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

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