

70-764 Dumps

Administering a SQL Database Infrastructure (beta)

<https://www.certleader.com/70-764-dumps.html>



NEW QUESTION 1

- (Exam Topic 1)

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

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

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

You need to configure the Contoso instance.

How should you complete the Transact-SQL statement? To answer, select the appropriate Transact-SQL segments in the answer area.

Answer Area

```
sp_configure [dropdown] , 1
    'max worker threads'
    'show advanced options'
    'cost threshold for parallelism'
    'max degree of parallelism'

GO
RECONFIGURE WITH OVERRIDE
GO
sp_configure [dropdown] , 0
    'max worker threads'
    'show advanced options'
    'cost threshold for parallelism'
    'max degree of parallelism'

GO
RECONFIGURE
GO
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: show advanced options

Advanced configuration options are displayed by first setting show advanced option to 1. Box 2: max worker threads

SQL Server uses the native thread services of the operating systems so that one or more threads support each network that SQL Server supports simultaneously, another thread handles database checkpoints, and a pool of threads handles all users. The default value for max worker threads is 0. This enables SQL Server to automatically configure the number of worker threads at startup. The default setting is best for most systems.

References:

<https://docs.microsoft.com/en-us/sql/database-engine/configure-windows/configure-the-max-worker-threads-ser>

NEW QUESTION 2

- (Exam Topic 1)

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

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

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

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

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

Answer: B

Explanation:

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

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

References:

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

NEW QUESTION 3

- (Exam Topic 1)

You administer a Windows Azure SQL Database database named Human_Resources. The database contains 2 tables named Employees and SalaryDetails. You add two Windows groups as logins for the server:

- CORP\Employees - All company employees
- CORP\HRAdmins - HR administrators only
- HR Administrators are also company employees.

You need to grant users access according to the following requirements:
CORP\Employees should have SELECT access to the Employees table.
Only users in CORP\HRAdmins should have SELECT access to the SalaryDetails table.
Logins are based only on Windows security groups.
What should you do?

- A. Create a database role called Employees. Add CORP\Employees to the db_datareader role. Add all company employees except HR administrators to the Employees role. Deny SELECT access to the SalaryDetails table to the Employees role.
- B. Create a database role called HRAdmins. Add all company employees except HR administrators to the db_datareader role, Add all HR administrators to the HRAdmins role. Grant SELECT access to the SalaryDetails table to the HRAdmins role. Deny SELECT access to the SalaryDetails table to the db_datareader role.
- C. Create two database roles: Employees and HRAdmins. Add all company employees to the Employees role. Add HR administrators to the HRAdmins role. Grant SELECT access to all tables except SalaryDetails to the Employees role. Grant SELECT access to the SalaryDetails table to the HRAdmins role. Deny SELECT access to the SalaryDetails table to the Employees role.
- D. Create a database role called Employees. Add all HR administrators to the db_datareader role. Add all company employees to the Employees role. Grant SELECT access to all tables except the SalaryDetails table to the Employees role. Deny SELECT access to the SalaryDetails table to the Employees role.

Answer: D

NEW QUESTION 4

- (Exam Topic 1)

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

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

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

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

Which permission should you grant?

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

Answer: A

Explanation:

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

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

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

NEW QUESTION 5

- (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 6

- (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 threshol
- 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 7

- (Exam Topic 1)

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

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

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

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

Solution: You migrate DB1 to a named instance on a server than runs Microsoft SQL Server 2016 Standard edition.

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

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

References:

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

NEW QUESTION 8

- (Exam Topic 1)

You have a database named DB1 that stores more than 700 gigabyte (GB) of data and serves millions of requests per hour.

Queries on DB1 are taking longer than normal to complete. You run the following Transact-SQL statement:

```
SELECT * FROM sys.database_query_store_options
```

You determine that the Query Store is in Read-Only mode.

You need to maximize the time that the Query Store is in Read-Write mode. Which Transact-SQL statement should you run?

- A. ALTER DATABASE DB1SET QUERY_STORE (QUERY_CAPTURE_MODE = ALL)
- B. ALTER DATABASE DB1SET QUERY_STORE (MAX_STORAGE_SIZE_MB = 50)
- C. ALTER DATABASE DB1SET QUERY_STORE (CLEANUP_POLICY = (STALE_QUERY_THRESHOLD_DAYS = 14));
- D. ALTER DATABASE DB1SET QUERY_STORE (QUERY_CAPTURE_MODE = NONE)

Answer: C

Explanation:

Stale Query Threshold (Days): Time-based cleanup policy that controls the retention period of persisted runtime statistics and inactive queries.

By default, Query Store is configured to keep the data for 30 days which may be unnecessarily long for your scenario.

Avoid keeping historical data that you do not plan to use. This will reduce changes to read-only status. The size of Query Store data as well as the time to detect and mitigate the issue will be more predictable. Use Management Studio or the following script to configure time-based cleanup policy:

```
ALTER DATABASE [QueryStoreDB]
```

```
SET QUERY_STORE (CLEANUP_POLICY = (STALE_QUERY_THRESHOLD_DAYS = 14));
```

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/performance/best-practice-with-the-query-store>

NEW QUESTION 9

- (Exam Topic 1)

You administer a Microsoft SQL Server 2016 server.

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

What should you configure?

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

Answer: C

Explanation:

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

References:

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

NEW QUESTION 10

- (Exam Topic 1)

You install Microsoft SQL Server 2016 on a new server.

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

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

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

Answer: A

NEW QUESTION 10

- (Exam Topic 1)

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

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

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

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

- A. Yes
- B. No

Answer: B

Explanation:

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

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

References:

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

NEW QUESTION 12

- (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 16

- (Exam Topic 1)

You administer a Microsoft SQL Server 2016 database.

Users report that a billing application becomes unresponsive during busy times of the day. While investigating, you notice large number of processes taking or waiting for table locks. You suspect that SQL Server is assigning stronger locks to queries.

You start a SQL Profiler trace. Which event should you select?

- A. Deadlock graph
- B. Lock: Escalation
- C. Lock: Timeout
- D. Lock: Deadlock

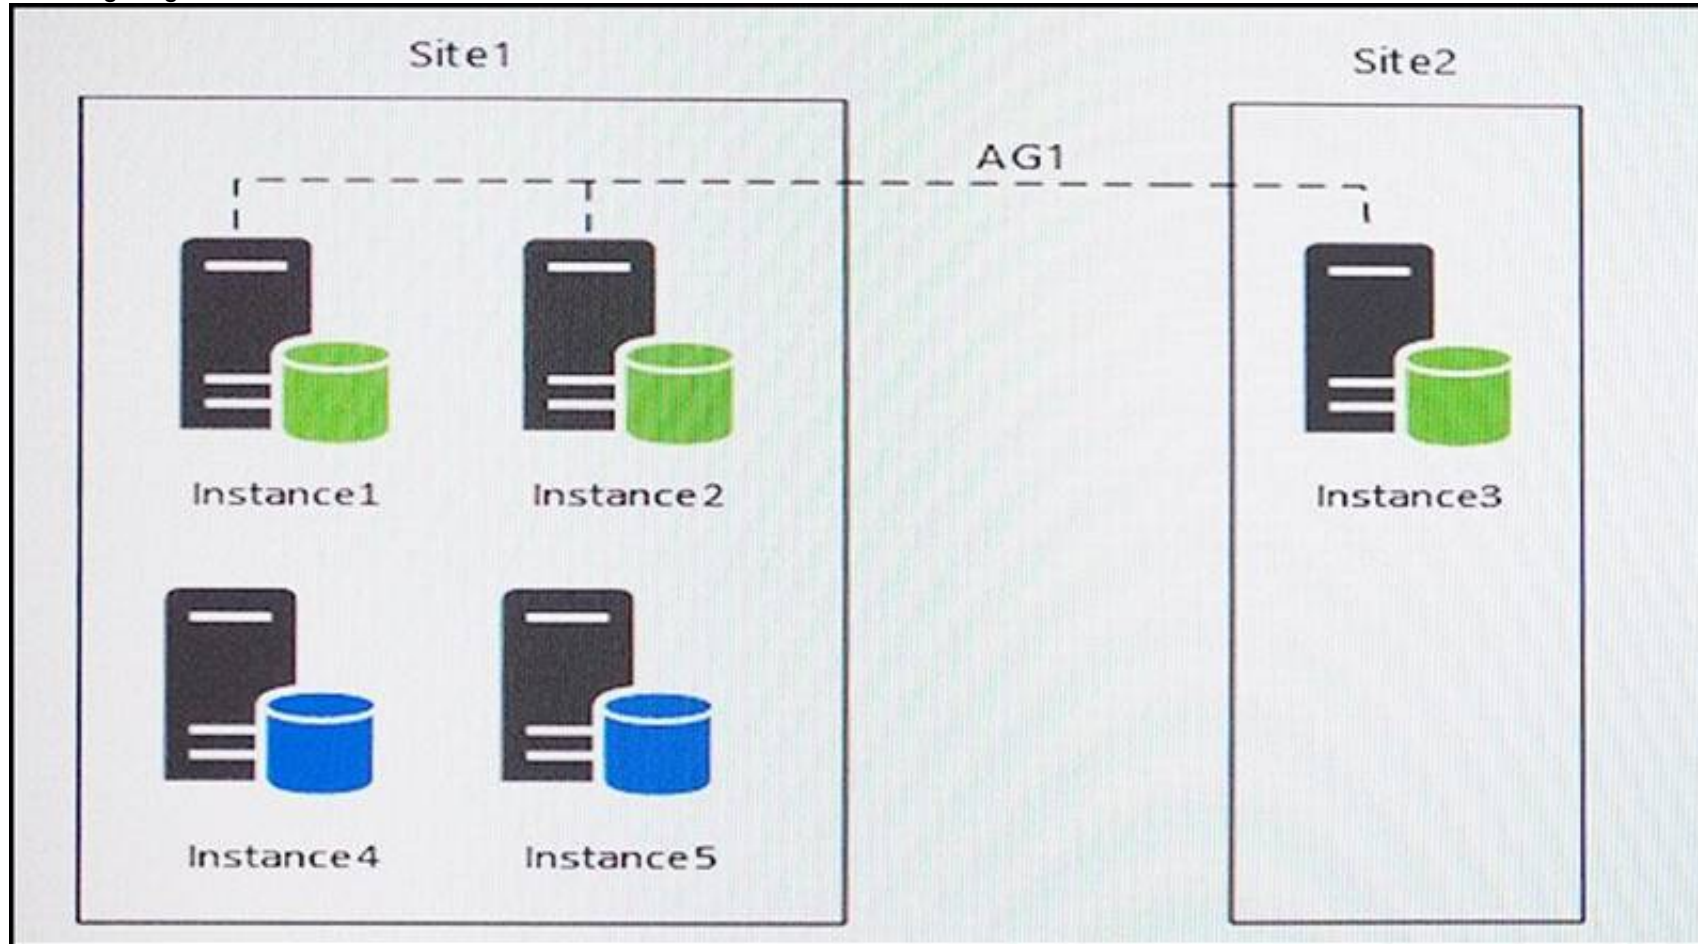
Answer: B

NEW QUESTION 19

- (Exam Topic 1)

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

You 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 20

- (Exam Topic 1)

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

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



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

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

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

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

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

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

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

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

You should minimize both the time required to restore

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

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

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

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

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

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

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

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

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

- A. Modify the backup script to use the keyword SKIP in the FILE_SNAPSHOT statement.
- B. Modify the backup script to use the keyword SKIP in the WITH statement
- C. Modify the backup script to use the keyword NO_COMPRESSION in the WITH statement.

D. Modify the full database backups script to stripe the backup across multiple backup files.

Answer: D

Explanation:

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

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

References:

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

NEW QUESTION 22

- (Exam Topic 1)

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

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

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

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

A. Yes

B. No

Answer: B

Explanation:

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

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

References:

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

NEW QUESTION 24

- (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 26

- (Exam Topic 1)

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

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

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

Answer Area

DBCC

CHECKDB
PHYSICAL_ONLY
REPAIR_FAST
REPAIR_REBUILD

 ('salesOrders',

CHECKDB
PHYSICAL_ONLY
REPAIR_FAST
REPAIR_REBUILD

)

A. Mastered

B. Not Mastered

Answer: A

Explanation:

Box 1: CHECKDB

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

DBCC CHECKDB

[(database_name | database_id | 0 [, NOINDEX

| , { REPAIR_ALLOW_DATA_LOSS | REPAIR_FAST | REPAIR_REBUILD }]

....

Box 2: REPAIR_REBUILD

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

DBCC CHECKDB repair the found errors.

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

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

NEW QUESTION 27

- (Exam Topic 1)

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

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

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

You need to configure auditing for WDWDB.

In the table below, identify the event type that you must audit for each activity.

Answer Area

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

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

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

NEW QUESTION 31

- (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 33

- (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

	▼
diskadmin	
serveradmin	
securityadmin	
setupadmin	

Transact-SQL syntax

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: securityadmin

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

Box 2: sys.server_audit_specifications

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

NEW QUESTION 38

- (Exam Topic 1)

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

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

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

You need to address the performance issues with each table.

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

Answer Area

Action

TableA

TableB

Run the following Transact-SQL statement:

```
SET AUTO_UPDATE_STATISTICS_ASYNC ON
```

☐☐

Run the following Transact-SQL statement:

```
SET AUTO_UPDATE_STATISTICS OFF
```

☐☐

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

```
SET AUTO_CREATE_STATISTICS on (INCREMENTAL = ON)
```

☐☐

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

```
UPDATE STATISTICS
```

☐☐

- A. Mastered

B. Not Mastered

Answer: A

Explanation:

Table A: Auto_update statistics off

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

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

SET AUTO_UPDATE_STATISTICS_ASYNC ON

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

References:

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

NEW QUESTION 40

- (Exam Topic 1)

You administer a Microsoft SQL Server 2016 database instance.

You plan to migrate the database to Windows Azure SQL Database.

You verify that all objects contained in the database are compatible with Windows Azure SQL Database. You need to ensure that database users and required server logins are migrated to Windows Azure SQL Database.

What should you do?

A. Use the Copy Database wizard.

B. Back up the database from the local server and restore it to Windows Azure SQL Database.

C. Use the Database Transfer wizard.

D. Use SQL Server Management Studio to deploy the database to Windows Azure SQL Database.

Answer: D

NEW QUESTION 43

- (Exam Topic 1)

You administer a Microsoft SQL Server 2016 database.

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

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

What should you do?

A. Create an Extended Events session by using the sqlserver.error_reported event.

B. Create a SQL Profiler session to capture all ErrorLog and EventLog events.

C. Flag all stored procedures for recompilation by using sp_recompile.

D. Execute sp_who.

Answer: A

Explanation:

Trapping SQL Server Errors with Extended Events

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

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

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

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

References:

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

NEW QUESTION 47

- (Exam Topic 1)

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

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

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

Answer Area

Action	Server instance
Complete the backup job.	<div>▼</div> <div>Primary server instance</div> <div>Secondary server instance</div> <div>Monitor server instance</div> <div>Backup share file server</div>
Copy the backup job.	<div>▼</div> <div>Primary server instance</div> <div>Secondary server instance</div> <div>Monitor server instance</div> <div>Backup share file server</div>
Restore the backup.	<div>▼</div> <div>Primary server instance</div> <div>Secondary server instance</div> <div>Monitor server instance</div> <div>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 52

- (Exam Topic 1)

You administer a SQL Server 2016 database instance.

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

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

Answer: A

Explanation:

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

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

Etc. References:

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

NEW QUESTION 53

- (Exam Topic 1)

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

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

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

Transact-SQL segments

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

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

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

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

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

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

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

Answer Area



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

NEW QUESTION 57

- (Exam Topic 1)

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

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

One instance hosts a user database named HRDB. The database contains sensitive human resources data. You need to grant an auditor permission to view the SQL Server audit logs while following the principle of least privilege.

Which permission should you grant?

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

H. sysadmin

Answer: F

Explanation:

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

Membership in the sysadmin fixed server role.

The CONTROL SERVER permission.

The VIEW SERVER STATE permission.

The ALTER ANY AUDIT permission.

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

References: [https://technet.microsoft.com/en-us/library/cc280386\(v=sql.110\).aspx](https://technet.microsoft.com/en-us/library/cc280386(v=sql.110).aspx)

NEW QUESTION 61

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

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

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

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

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

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

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

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

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

NEW QUESTION 66

- (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 69

- (Exam Topic 1)

You administer a Microsoft SQL Server 2016 database.

The database is currently configured to log ship to a secondary server.

You are preparing to cut over to the secondary server by stopping log-shipping and bringing the secondary database online.

You want to perform a tail-log backup. You need to leave the primary database in a restoring state. Which option of the BACKUP LOG command should you use?

- A. NO_TRUNCATE
- B. NORECOVERY
- C. STANDBY
- D. FORMAT

Answer: B

NEW QUESTION 74

- (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 plan to delegate encryption operations to a user.

You need to grant the user permission to implement cell-level encryption while following the principle of least privilege.

Which permission should you grant?

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

Answer: G

Explanation:

The following permissions are necessary to perform column-level encryption, or cell-level encryption.

CONTROL permission on the database.

CREATE CERTIFICATE permission on the database. Only Windows logins, SQL Server logins, and application roles can own certificates. Groups and roles cannot own certificates.

ALTER permission on the table.

Some permission on the key and must not have been denied VIEW DEFINITION permission. References:

<https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/encrypt-a-column-of-data>

NEW QUESTION 79

- (Exam Topic 1)

You have a Microsoft SQL Server instance that hosts a database named DB1 that contains 800 gigabyte (GB) of data. The database is used 24 hours each day.

You implement indexes and set the value of the Auto Update Statistics option set to True.

Users report that queries take a long time to complete.

You need to identify statistics that have not been updated for a week for tables where more than 1,000 rows changed.

How should you complete the Transact-SQL statement? To answer, configure the appropriate Transact-SQL segments in the answer area.

Answer Area

```
SELECT OBJECT_NAME(id), name, (id, indid), (id, indid),
FROM sys.sysindexes
WHERE (id, indid) <= DATEADD(DAY, -7, GETDATE())
AND > 1000
AND id IN (SELECT object_id FROM sys.tables)
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: stats_date See example below. Box 2: rowmodctr See examplebelow. Box 3: stats_date

You need to identify statistics that have not been updated for a week. Box 4: rowmodctr

You need to identify that more than 1,000 rows changed.

Rowmodctr counts the total number of inserted, deleted, or updated rows since the last time statistics were updated for the table.

Example: We will query every statistics object which was not updated in the last day and has rows modified since the last update. We will use the rowmodctr field of sys.sysindexes because it shows how many rows were inserted, updated or deleted since the last update occurred. Please note that it is not always 100% accurate in SQL Server 2005 and later, but it can be used to check if any rows were modified.

--Get the list of outdated statistics

```
SELECT OBJECT_NAME(id),name,STATS_DATE(id, indid),rowmodctr FROM sys.sysindexes
```

```
WHERE STATS_DATE (id, indid)<=DATEADD(DAY,-1,GETDATE())
```

```
AND rowmodctr>0
```

```
AND id IN (SELECT object_id FROM sys.tables) GO
```

After collecting this information, we can decide which statistics require an update.

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/system-compatibility-views/sys-sysindexes-transact-sq>

<https://www.mssqltips.com/sqlservertip/2628/how-to-find-outdated-statistics-in-sql-server-2008/>

NEW QUESTION 83

- (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-shippingevery 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:

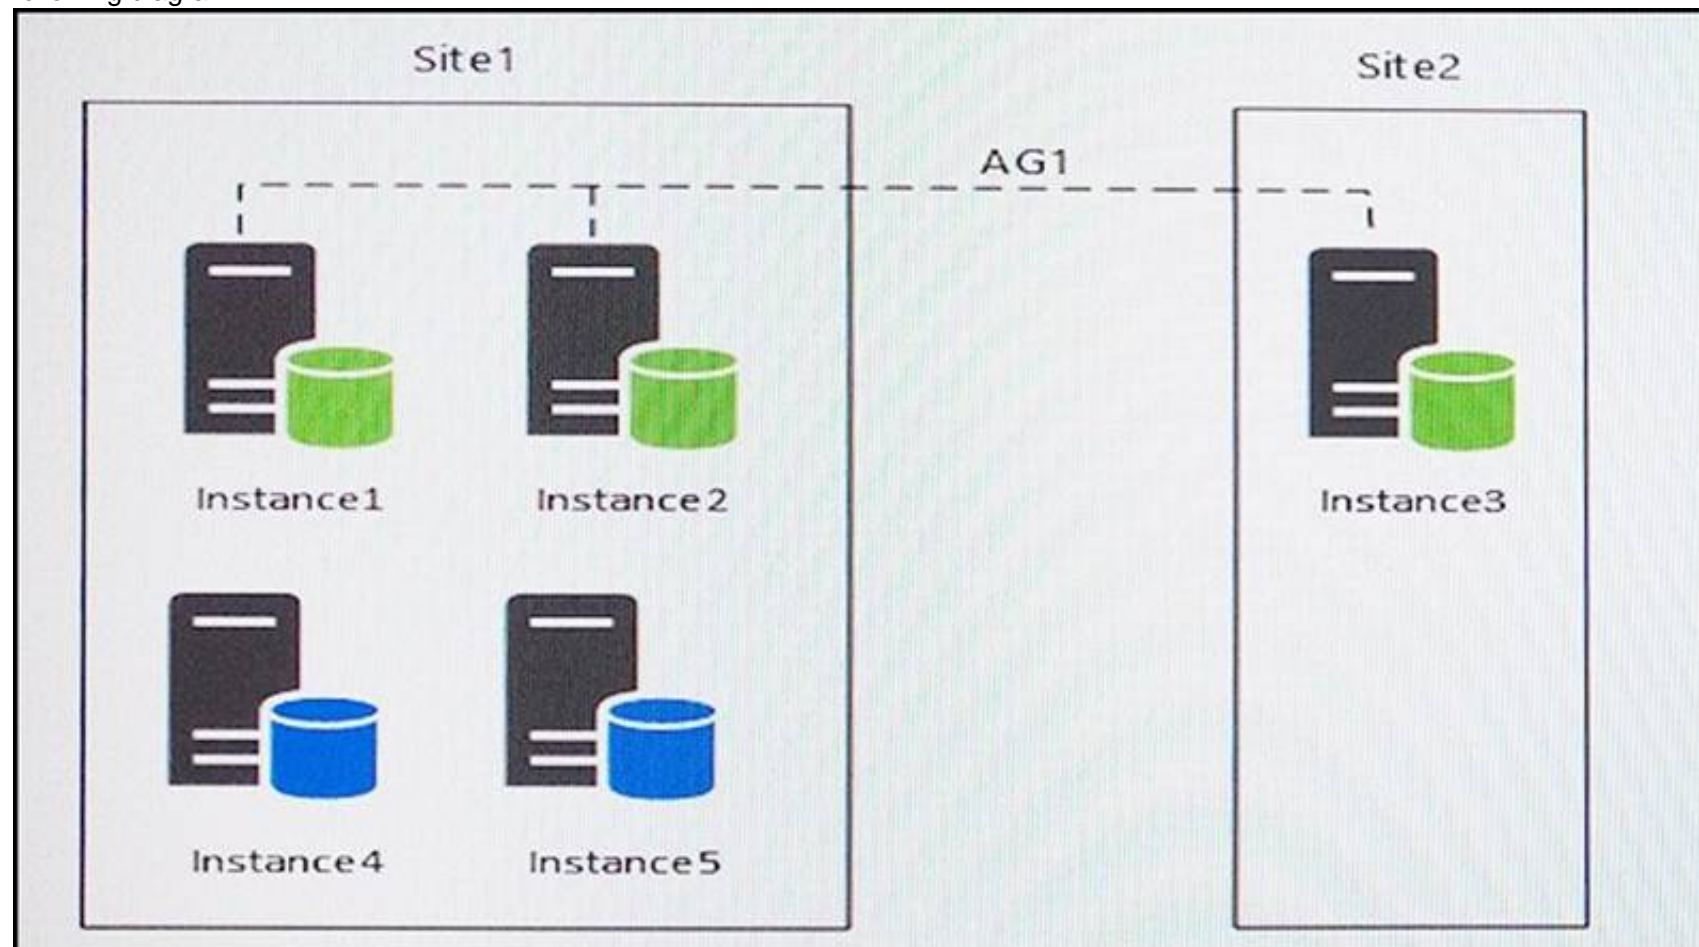
<https://docs.microsoft.com/en-us/sql/sql-server/failover-clusters/windows/windows-server-failover-clustering-ws>

NEW QUESTION 84

- (Exam Topic 1)

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

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



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

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

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

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

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

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

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

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

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

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

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

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

Reporting system: This solution accesses data in DB1 with 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 in DB1 with a login that is mapped to a database user that is a member of the db_datareader and db_datawriter roles. The user has EXECUTE permissions on the database. Queries from the operations system will perform both DDL and DML operations.

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

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

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

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

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

NOTE: Each correct selection is worth one point.

Objects

Sys.dm_os_wait_stats

Sys.dm_exec_connections

Sys.dm_exec_requests

Sys.dm_exec_procedure_stats

Sys.dm_exec_sessions

Sys.dm_exec_query_stats

Sys.dm_exec_query_re-
source_semaphores

Sys.dm_exec_ses-
sion_wait_stats

Answer Area

Instance Object

Instance1

Object

Instance4

Object

Instance5

Object

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Instance 1: sys.dm_exec_query_stats

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

Instance 4: sys.dm_os_wait_stats

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

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

Instance 5:sys.dm_exec_session_wait_stats

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

NEW QUESTION 85

- (Exam Topic 1)

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

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

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

You enabled log shipping for DB1 and configure backup and restore to occur every 30 minutes. You discover that the disks on the data server are almost full.

You need to reduce the amount of disk space used by the log shipping process. Solution: You enable compression for the transaction log backups:

Does this meet the goal?

- A. Yes

B. No

Answer: A

NEW QUESTION 87

- (Exam Topic 2)

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

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

Which solution should you use?

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

Answer: A

NEW QUESTION 88

- (Exam Topic 2)

You use SQL Server 2014 Enterprise Edition.

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

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

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

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

You may not need all of the actions.

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Note:

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

- SPLIT RANGE (boundary_value)

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

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

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

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

The following example creates a partition function to partition a table or index into four partitions. ALTER PARTITION FUNCTION splits one of the partitions into

two to create a total of five partitions. CREATE PARTITION FUNCTION myRangePF1 (int)
AS RANGE LEFT FOR VALUES (1, 100, 1000); GO
-Split the partition between boundary_values 100 and 1000
-to create two partitions between boundary_values 100 and 500
--and between boundary_values 500 and 1000. ALTER PARTITION FUNCTION myRangePF1 () SPLIT RANGE (500);

NEW QUESTION 91

- (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 ROLS [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 93

- (Exam Topic 2)

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

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

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

Fulfilled
Shipped
Ordered
Received

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

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

- A. To the Sales.Orders table, add a column named Status that will store the order status.
- B. Update the Status column as the order status changes.
- C. Create a new table named Sales.OrderStatus that contains three columns named OrderID, StatusDate, and Status.
- D. Insert new rows into the table as the order status changes.
- E. Implement change data capture on the Sales.Orders table.
- F. To the Sales.Orders table, add three columns named FulfilledDate, ShippedDate, and ReceivedDate. Update the value of each column from null to the appropriate date as the order status changes.

Answer: A

NEW QUESTION 94

- (Exam Topic 2)

Overview

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

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

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

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

Requirements Planned Changes

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

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

Database1 will also contain a stored procedure named usp_UpdateOrderDetails.

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

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

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

Inventory will contain over 100 GB of data.

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

Database2 will contain a stored procedure named usp_UpdateInventory. usp_UpdateInventory will manipulate a table that contains a self-join that has an unlimited number of hierarchies. All data in Database2 is recreated each day and does not change until the next data creation process. Data from Database2 will be accessed periodically by an external application named Application1. The data from Database2 will be sent to a database named Appl_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 98

- (Exam Topic 2)

You plan to create a database.

The database will be used by a Microsoft .NET application for a special event that will last for two days. During the event, data must be highly available. After the event, the database will be deleted. You need to recommend a solution to implement the database while minimizing costs. The solution must not affect any existing applications. What should you recommend? More than one answer choice may achieve the goal. Select the BEST answer.

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

Answer: D

NEW QUESTION 102

- (Exam Topic 2)

Overview

General Overview

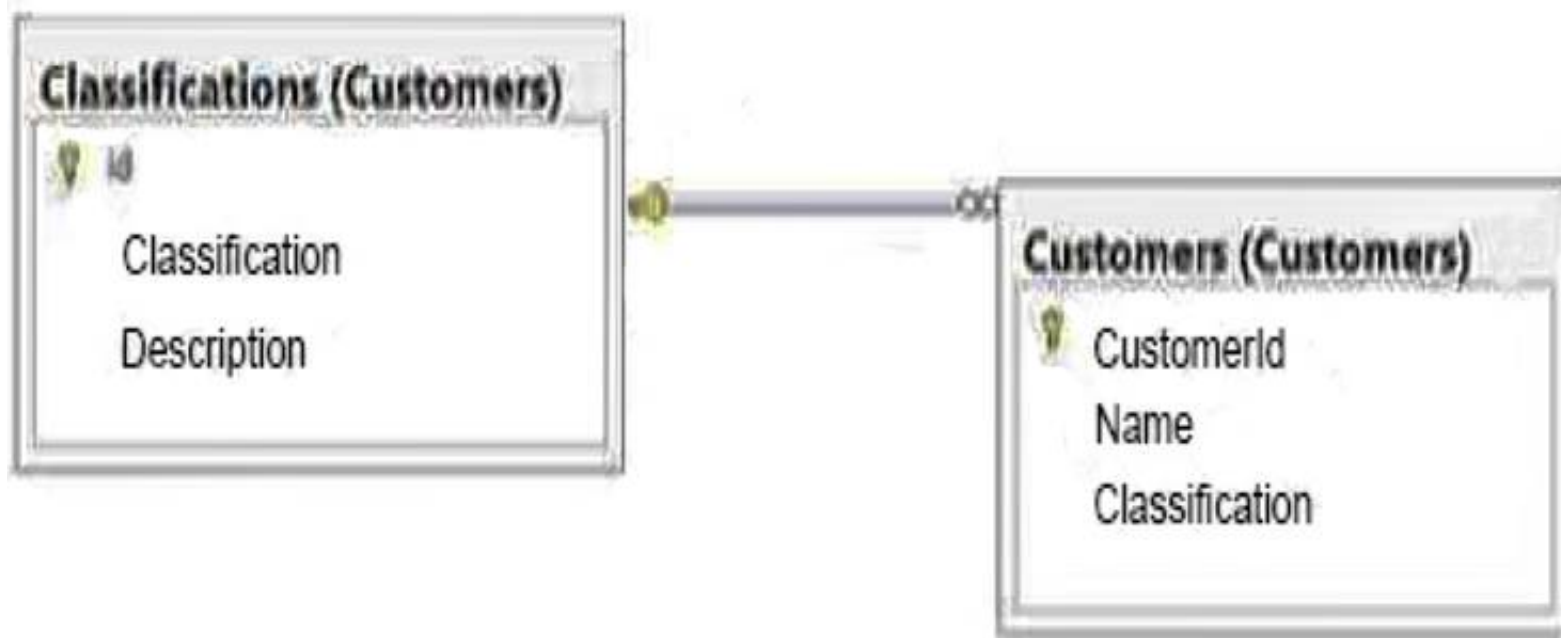
ADatum Corporation has offices in Miami and Montreal.

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 106

- (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 111

- (Exam Topic 2)

Overview

Application Overview

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

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

The new version will use SQL Server 2014.

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

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

Tables

The current database schema contains a table named OrderDetails.

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

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

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

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

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

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

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

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

Customer Problems Installation Issues

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

Index Fragmentation Issues

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

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

Backup Issues

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

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

Missing Data Issues

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

Query Performance Issues

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

Design Requirements

File Storage Requirements

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

Data Recovery Requirements

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

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

Concurrency Requirements

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

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

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

Answer: B

Explanation:

- Concurrency Requirements

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

NEW QUESTION 114

- (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 119

- (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 120

- (Exam Topic 2)

Overview

General Overview

ADatum Corporation has offices in Miami and Montreal.

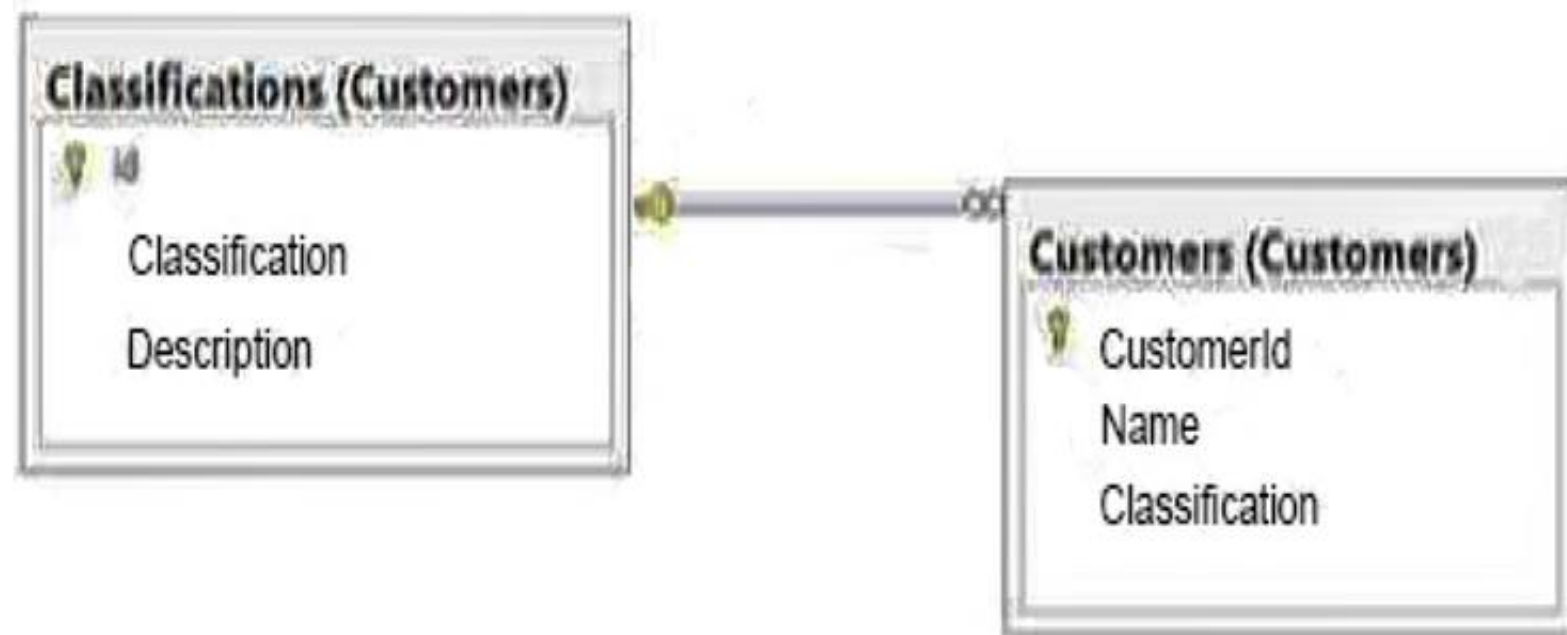
The network contains a single Active Directory forest named adatum.com. The offices connect to each other by using a WAN link that has 5-ms latency. A. Datum standardizes its database platform by using SQL Server 2014 Enterprise edition.

Databases

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

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

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



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

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

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

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

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

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

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

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

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

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

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

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

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

The nested stored procedures are never called directly. Design Requirements

Data Recovery

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

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

Classification Changes

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

Storage

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

Error Handling

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

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

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

- A. Add a row to the Customers table each time a classification changes.
- B. Add columns for each classification to the Customers table.
- C. Add a table to track any changes made to the classification of each customer.
- D. Add a column to the Classifications table to track the status of each classification.
- E. Implement change data capture.

Answer: CD

Explanation:

Scenario:

You plan to change the way customers are classified.

The new classifications will have four levels based on the number of orders. Classifications may be removed or added in the future.

NEW QUESTION 125

- (Exam Topic 2)

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

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

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

You need to achieve this goal by granting the minimum permissions required.

What should you do? (Each correct answer presents a complete solution. Choose all that apply.)

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

Answer: BC

Explanation:

References:

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

NEW QUESTION 129

- (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. <div>Feature</div>
The CHECK constraint	
Extended Event	All distribution statistics must be updated daily. <div>Feature</div>
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 130

- (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 131

- (Exam Topic 2)

Overview

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

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

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

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

Requirements Planned Changes

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

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

Database1 will also contain a stored procedure named usp_UpdateOrderDetails.

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

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

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

Inventory will contain over 100 GB of data.

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

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

Business Requirements

You have the following requirements:

Costs for new licenses must be minimized.

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

Development effort must be minimized whenever possible.

The storage requirements for databases must be minimized.

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

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

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

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

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

The encryption strategy must minimize changes to the databases and to the applications. You need to recommend a database reporting solution that meets the business requirements. What should you include in the recommendation?

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

Answer: A

Explanation:

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

2. The data collector provides an historical report for each of the System Data collection sets. Each of the following reports use data that is stored in the management data warehouse:

You can use these reports to obtain information for monitoring system capacity and troubleshooting system performance.

NEW QUESTION 136

- (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 installation issues.

What should you include in the recommendation?

- A. Windows logins
- B. Server roles
- C. Contained users
- D. Database roles

Answer: C

Explanation:

- Scenario: 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.

- Creating contained users enables the user to connect directly to the contained database. This is a very significant feature in high availability and disaster recovery scenarios such as in an AlwaysOn solution. If the users are contained users, in case of failover, people would be able to connect to the secondary without creating logins on the instance hosting the secondary. This provides an immediate benefit.

NEW QUESTION 137

- (Exam Topic 2)

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

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

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

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

Answer: B

Explanation:

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

NEW QUESTION 138

- (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 143

- (Exam Topic 2)

You administer a Microsoft SQL Server 2016 database named Contoso on a server named Server01. You need to be notified immediately when fatal errors occur on Server01.

What should you create?

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

Answer: A

Explanation:

References:

<http://www.sqlskills.com/blogs/glenn/creating-sql-server-agent-alerts-for-critical-errors/>

NEW QUESTION 147

- (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 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 150

- (Exam Topic 2)

Overview

Application Overview

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

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

The new version will use SQL Server 2014.

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

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

Tables

The current database schema contains a table named OrderDetails.

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

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

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

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

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

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

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

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

Customer Problems Installation Issues

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

Index Fragmentation Issues

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

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

Backup Issues

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

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

Missing Data Issues

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

Query Performance Issues

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

Design Requirements

File Storage Requirements

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

Data Recovery Requirements

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

Security Requirements

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

Concurrency Requirements

You must reduce the likelihood of deadlocks occurring when Sales.Prod and Sales.Proc2 execute. You need to recommend a solution that 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 153

- (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 155

- (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 156

- (Exam Topic 2)

You administer a Microsoft SQL Server 2016 database.

You need to ensure that the size of the transaction log file does not exceed 2 GB. What should you do?

- A. Execute sp_configure 'max log size', 2G.
- B. use the ALTER DATABASE...SET LOGFILE command along with the maxsize parameter.
- C. In SQL Server Management Studio, right-click the instance and select Database Setting
- D. Set the maximum size of the file for the transaction log.
- E. in SQL Server Management Studio, right-click the database, select Properties, and then click Files.Open the Transaction log Autogrowth window and set the maximum size of the file.

Answer: D

NEW QUESTION 161

- (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 security requirement. What should you recommend?

- A. Revoke user permissions on the table
- B. Create stored procedures that manipulate dat
- C. Grant the users the EXECUTE permission on the stored procedures.
- D. Grant the users the SELECT permission on the table
- E. Create views that retrieve data from the tables.Grant the users the SELECT permission on the views.
- F. Deny the users SELECT permission on the table
- G. Create views that retrieve data from the table
- H. Grant the users the SELECT permission on the views.
- I. Deny the users the SELECT permission on the table
- J. Create stored procedures that manipulate data.Grant the users the EXECUTE permission on the stored procedures.

Answer: C

Explanation:

- Security Requirements

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

NEW QUESTION 162

- (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 167

- (Exam Topic 2)

You install a Microsoft SQL Server 2016 instance.

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

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

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

Answer: B

NEW QUESTION 169

- (Exam Topic 2)

You administer a Microsoft SQL Server 2016 database.

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

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

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

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

Answer: D

NEW QUESTION 172

- (Exam Topic 2)

You have a SQL Server 2012 database named DB1.

You plan to import a large number of records from a SQL Azure database to DB1.

You need to recommend a solution to minimize the amount of space used in the transaction log during the import operation.

What should you include in the recommendation?

- A. a new log file
- B. a new filegroup
- C. the full recovery model
- D. a new partitioned table
- E. the bulk-logged recovery model

Answer: E

Explanation:

Compared to the full recovery model, which fully logs all transactions, the bulk-logged recovery model minimally logs bulk operations, although fully logging other transactions. The bulk-logged recovery model protects against media failure and, for bulk operations, provides the best performance and least log space usage.

Note:

The bulk-logged recovery model is a special-purpose recovery model that should be used only intermittently to improve the performance of certain large-scale bulk operations, such as bulk imports of large amounts of data. Recovery Models (SQL Server)

NEW QUESTION 175

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your 70-764 Exam with Our Prep Materials Via below:

<https://www.certleader.com/70-764-dumps.html>