

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

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

What should you use?

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

Answer: D

Explanation:

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

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

Missing indexes can be found with the execution plans.

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

NEW QUESTION 2

- (Exam Topic 1)

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 [max worker threads] , 1
GO
RECONFIGURE WITH OVERRIDE
GO
sp_configure [max worker threads] , 0
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 3

- (Exam Topic 1)

You administer a Microsoft SQL Server 2016 database that includes a table named Application.Events. Application.Events contains millions of records about user activity in an application.

Records in Application.Events that are more than 90 days old are purged nightly.

When records are purged, table locks are causing contention with inserts.

You need to be able to modify Application.Events without requiring any changes to the applications that utilize Application.Events.

Which type of solution should you use?

- A. Partitioned tables
- B. Online index rebuild
- C. Change data capture
- D. Change tracking

Answer: A

NEW QUESTION 4

- (Exam Topic 1)

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

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

What should you do?

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

Answer: C

Explanation:

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

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

NEW QUESTION 5

- (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 6

- (Exam Topic 1)

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

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

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

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

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

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

Answer: D

NEW QUESTION 7

- (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 8

- (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.

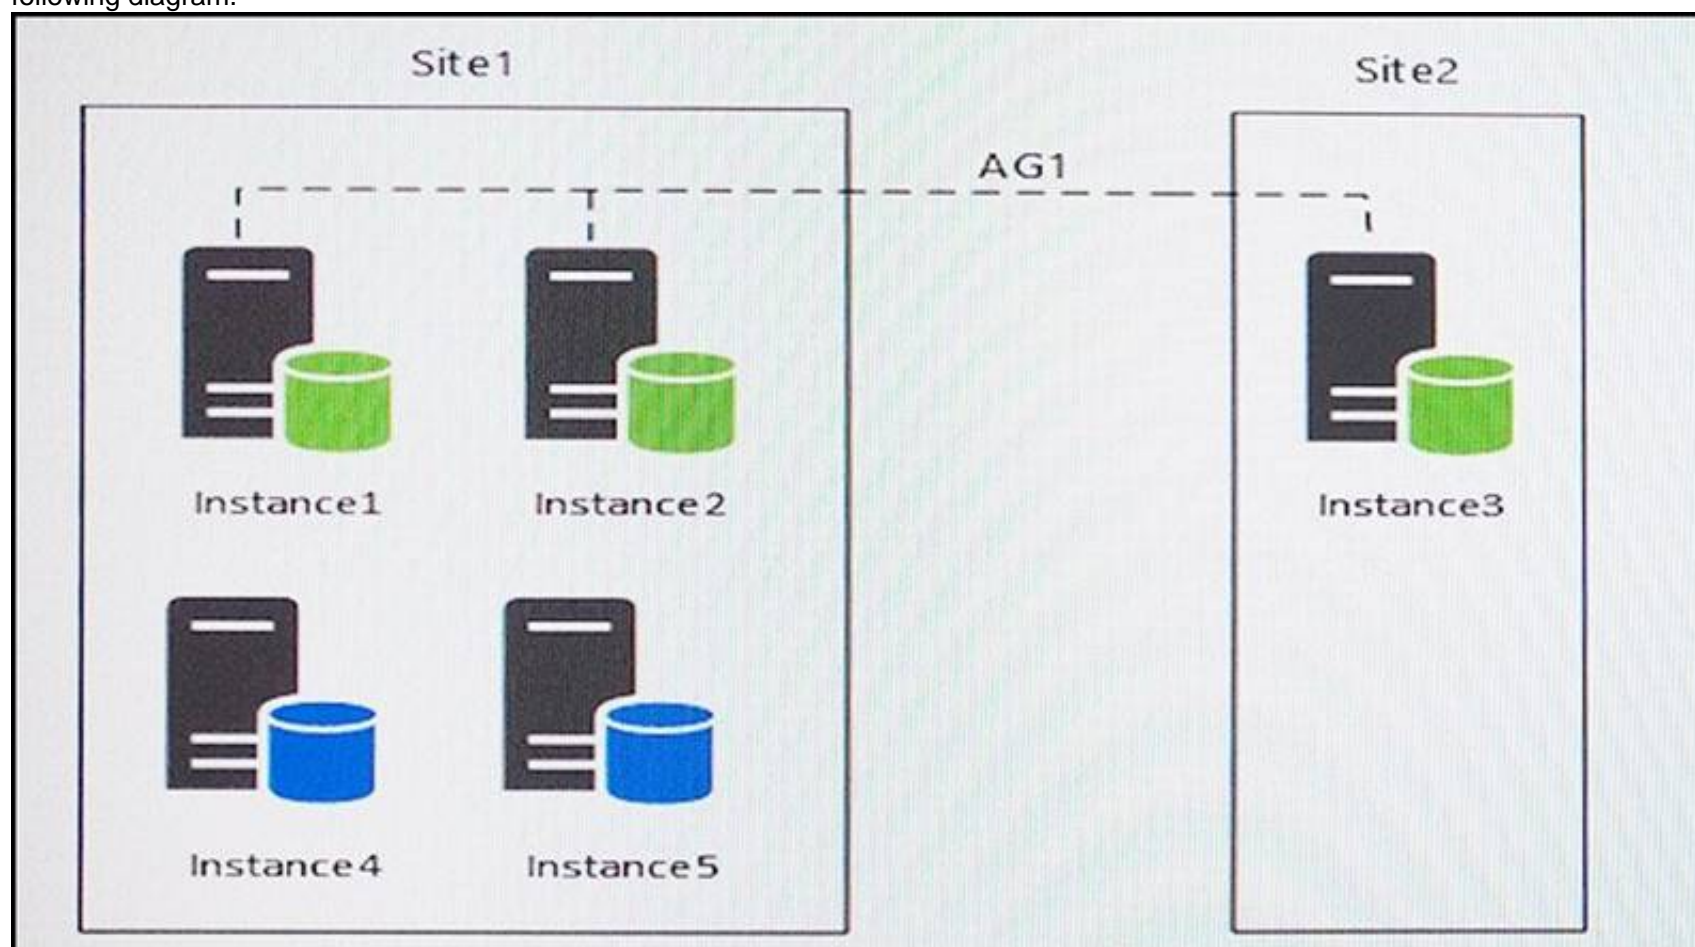
References: <https://docs.microsoft.com/en-us/sql/t-sql/statements/create-function-transact-sql>

NEW QUESTION 9

- (Exam Topic 1)

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

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



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

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

Instance1 experiences heavy read-write traffic. The instance hosts a database named OperationsMain that is four terabytes (TB) in size. The database has

multiple data files and filegroups. One of the filegroups is read_only and is half of the total database size. Instance4 and Instance5 are not part of AG1. Instance4 is engaged in heavy read-write I/O. Instance5 hosts a database named StagedExternal. A nightly BULK INSERT process loads data into an empty table that has a rowstore clustered index and two nonclustered rowstore indexes. You must minimize the growth of the StagedExternal database log file during the BULK INSERT operations and perform point-in-time recovery after the BULK INSERT transaction. Changes made must not interrupt the log backup chain. You plan to add a new instance named Instance6 to a datacenter that is geographically distant from Site1 and Site2. You must minimize latency between the nodes in AG1. All databases use the full recovery model. All backups are written to the network location \\SQLBackup\\. A separate process copies backups to an offsite location. You should minimize both the time required to restore the databases and the space required to store backups. The recovery point objective (RPO) for each instance is shown in the following table.

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

Full backups of OperationsMain take longer than six hours to complete. All SQL Server backups use the keyword COMPRESSION. You plan to deploy the following solutions to the environment. The solutions will access a database named DB1 that is part of AG1. Reporting system: This solution accesses data inDB1with a login that is mapped to a database user that is a member of the db_datareader role. The user has EXECUTE permissions on the database. Queries make no changes to the data. The queries must be load balanced over variable read-only replicas. Operations system: This solution accesses data inDB1with a login that is mapped to a database user that is a member of the db_datareader and db_datawriter roles. The user has EXECUTE permissions on the database. Queries from the operations system will perform both DDL and DML operations. The wait statistics monitoring requirements for the instances are described in the following table.

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

You need to propose a new process for the StagedExternal database. Which five actions should you recommended be performed in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions

Drop all nonclustered indexes on the target table.

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

Run the nightly import process.

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

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

Drop all clustered and nonclustered indexes on the target table.

Recreate any dropped indexes on the target table.

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

Answer Area



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

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

All databases use the full recovery model.

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

NEW QUESTION 10

- (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 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 msdb database. Does the solution meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

Database Mail is guarded by the database role DatabaseMailUserRole in the msdb database in order to prevent anyone from sending arbitrary emails. Database users or roles must be created in the msdb database and must also be a member of DatabaseMailUserRole in order to send emails with the exception of sysadmin who has all privileges.

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

References:

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

NEW QUESTION 15

- (Exam Topic 1)

You plan to install a Microsoft SQL Server 2016 instance.

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

Store Excel workbooks on the file system.

Access the workbooks through Transact-SQL.

Include the workbooks in database backups.

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

Which feature should you use?

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

Answer: B

NEW QUESTION 18

- (Exam Topic 1)

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

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

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

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

You need to optimize query performance.

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

- A. Yes
- B. No

Answer: B

Explanation:

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

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

NEW QUESTION 19

- (Exam Topic 1)

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

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

Your company has several Microsoft SQL Saver instance. Each instance hosts many database. You observe I/O corruption on some of the instance.

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 != 'CHECKSUM'
GO
```

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

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

Does the solution meet the goal?

- A. Yes
- B. No

Answer: A

NEW QUESTION 23

- (Exam Topic 1)

You have configured Resource Governor with three resource pools.

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

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

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

Answer: B

Explanation:

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

NEW QUESTION 27

- (Exam Topic 1)

You administer a Microsoft SQL Server 2016 default instance. The instance is hosted by a server that has a local firewall configured.

The firewall only allows inbound connections on port 1433. The server only hosts a single instance of SQL Server.

You need to ensure that the instance is configured to allow remote connections even if the SQL Server is unresponsive to client connections.

What should you do?

- A. Enable inbound connections on TCP port 1434 in the Windows Firewall on the server.
- B. Execute the following Transact-SQL command: sp_configure 'remote admin connections',
- C. Execute the Reconfigure command.
- D. Execute the following Transact-SQL command: sp_configure 'remote access', 1
- E. Restart the SQL Server Agent Service.
- F. Enable inbound connections on TCP port 135 in the Windows Firewall on the server.

Answer: ABC

Explanation:

SQL Server provides a dedicated administrator connection (DAC). The DAC lets an administrator access a running server to execute diagnostic functions or Transact-SQL statements, or to troubleshoot problems on the server, even when the server is locked or running in an abnormal state and not responding to a SQL Server Database Engine connection. By default, the DAC is only available from a client on the server. To enable client applications on remote computers to use the DAC, use the remote admin connections option of sp_configure.

By default, the DAC only listens on the loop-back IP address (127.0.0.1), port 1434 The following example enables the DAC from a remote computer.

```
sp_configure 'remote admin connections', 1; GO
```

```
RECONFIGURE; GO
```

References:

<https://docs.microsoft.com/en-us/sql/database-engine/configure-windows/remote-admin-connections-server-con>

NEW QUESTION 31

- (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 that runs Microsoft SQL Server 2016 Enterprise edition.

Does the solution meet the goal?

A. Yes

B. No

Answer: A

Explanation:

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

References:

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

NEW QUESTION 32

- (Exam Topic 1)

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 35

- (Exam Topic 1)

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

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

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

A. SERIALIZABLE

B. SNAPSHOT

C. READ COMMITTED SNAPSHOT

D. REPEATABLE READ

Answer: C

Explanation:

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

It consumes less tempdb space than snapshot isolation. Etc.

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


NEW QUESTION 39

- (Exam Topic 1)

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

MSG_AUDIT_FORCED_SHUTDOWN

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



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

Answer Area

In which user mode will the SQL Server instance start?

	▼
single-user	
multi-user	
restricted-user	

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

	▼
public	
serveradmin	
sysadmin	
setupadmin	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: single-user

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

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

References:

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

NEW QUESTION 42

- (Exam Topic 1)

You administer a Windows Azure SQL Database database named Inventory that contains a stored procedure named p_AddInventory.

Users need to be able to SELECT from all tables in the database and execute the stored procedure. You need to grant only the necessary permissions.

What should you do?

- A. Grant EXECUTE permission on p_AddInventory to all user
- B. Grant VIEW DEFINITION to all users.
- C. Grant EXECUTE permission on p_AddInventory to all user
- D. Add all users to the db_datawriter role.
- E. Add all users to the db_owner role.

F. Grant EXECUTE permission on p_AddInventory to all user
G. Add all users to the db_datareader role.

Answer: D

NEW QUESTION 43

- (Exam Topic 1)

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

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

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

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

You need to optimize query performance. Solution: You rebuild all indexes.

Does the solution meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

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

NEW QUESTION 47

- (Exam Topic 1)

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

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

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

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

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

Answer Area

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

SMTP1_Account1: Private Profile

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

SMTP1_Account2: Default Profile

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

References:

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

NEW QUESTION 48

- (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 49

- (Exam Topic 1)

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

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

Answer: D

Explanation:

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

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

Encryptor: A certificate or asymmetric Key

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

NEW QUESTION 50

- (Exam Topic 1)

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

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

You purchase a new server. You must restore the database to the new server. You need to restore the data to the most recent time possible.

Which three files should you restore in sequence? To answer, move the appropriate files from the list of files to the answer area and arrange them in the correct order.

Files

Log_20160503_0600.trn

Log_20160503_1500.trn

Full_20160502_2100.bak

Log_20160503_1900.trn

Log_20160503_0900.trn

Diff_20160503_1200.bak

Diff_20160503_1700.bak

Answer Area

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1: Full.

Start with the full backup.

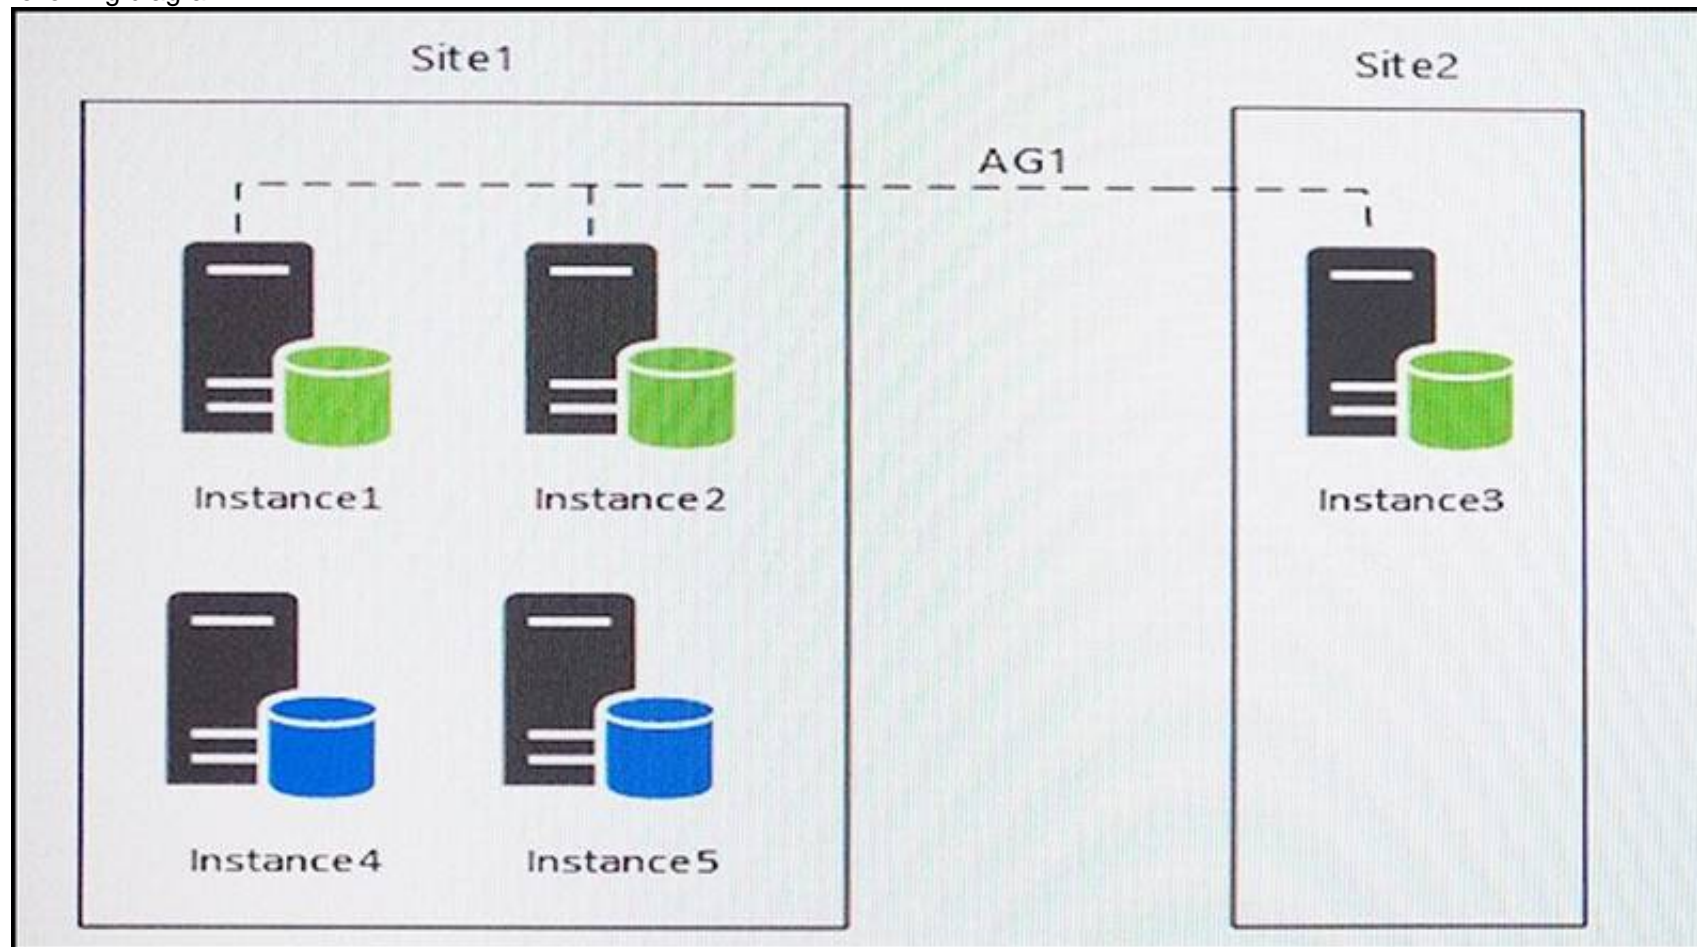
Step 2: Diff_20160503_1700.bak

Followed by the most recent differential backup. Step 3: Log_20160503_1900.bak
And finally the most recent log backup (the only log backup done after the most recent differential backup).
References:
<https://docs.microsoft.com/en-us/sql/relational-databases/backup-restore/differential-backups-sql-server>

NEW QUESTION 51

- (Exam Topic 1)

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



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

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

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

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

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

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

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

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

You should minimize both the time required to restore

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

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

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

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

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

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

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

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

You need to create a backup plan for Instance4. Which backup plan should you create?

- A. Weekly full backups, nightly differentia
- B. No transaction log backups are necessary.
- C. Weekly full backups, nightly differential backups, transaction log backups every 5 minutes.
- D. Weekly full backups, nightly differential backups, transaction log backups every 12 hours.
- E. Full backups every 60 minutes, transaction log backups every 30 minutes.

Answer: B

Explanation:

From scenario: Instance4 and Instance5 are not part of AG1. Instance4 is engaged in heavy read-write I/O. The recovery point objective of Instance4 is 60 minutes. RecoveryPoint Objectives are commonly described as the amount of data that was lost during the outage and recovery period. You should minimize both the time required to restore the databases and the space required to store backups.

References:

<http://sqlmag.com/blog/sql-server-recovery-time-objectives-and-recovery-point-objectives>

NEW QUESTION 52

- (Exam Topic 1)

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

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

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

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

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

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

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

References:

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

NEW QUESTION 57

- (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 60

- (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	<div><div></div><div>CHECKDB</div><div>PHYSICAL_ONLY</div><div>REPAIR_FAST</div><div>REPAIR_REBUILD</div></div>	('salesOrders' ,	<div><div></div><div>CHECKDB</div><div>PHYSICAL_ONLY</div><div>REPAIR_FAST</div><div>REPAIR_REBUILD</div></div>)
------	---	-------------------	---	---

- 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 61

- (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 65

- (Exam Topic 1)

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

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

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

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

The differential backup of the reporting database fails. Then, the reporting database fails at 14:00 hours.

You need to ensure that the reporting database is restored. You also need to ensure that data loss is minimal.

What should you do?

- A. Restore the latest full backup, and restore the latest differential backu

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

Answer: C

Explanation:

The differential backup of the reporting database has failed, so it can't be used.

NEW QUESTION 69

- (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 70

- (Exam Topic 1)

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

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

What should you do?

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

Answer: B

NEW QUESTION 71

- (Exam Topic 1)

You administer a Microsoft SQL Server 2016 database that includes a table named dbo.Log. This table contains millions of records about user activity in an application.

Records in dbo.Log that are more than 90 days old are purged nightly. When records are purged, table locks are causing contention with inserts.

You need to be able to modify dbo.Log without requiring any changes to the applications that utilize dbo.Log. Which type of solution should you use?

- A. Extended events
- B. Columnstore index
- C. Partitioned tables
- D. Read committed snapshot

Answer: C

NEW QUESTION 73

- (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 increase the frequency of the transaction log backups to every 10 minutes. Does this meet the goal?

- A. Yes
- B. No

Answer: B

NEW QUESTION 76

- (Exam Topic 1)

You have a test server that contains a database named DB1. Backups of the database are written to a single backup device. The backup device has a full, differential, and transaction log backup.

You discover that the database is damaged. You restore the database to the point at which the differential backup was taken.

You need to rebuild the database with data stored in the latest transaction logs.

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

Transact-SQL statements

NORECOVERY

RECOVERY

LOG

DBCC CHECKDB

CONTINUE_AFTER_ERROR

RESTORE

RESTORE VERIFYONLY

● ● ● ●

Answer Area

Transact-SQL segment

DB1 FROM DISK = N'Z:\Backups\Backup.bak WITH

Transact-SQL segment

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: RESTORE

Box 2: RECOVERY

The RESTORE ... WITH RECOVERY option puts the database into a useable state, so users can access a restored database.

References:

<https://www.mssqltips.com/sqlservertutorial/112/recovering-a-database-that-is-in-the-restoring-state/>

NEW QUESTION 77

- (Exam Topic 1)

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

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

Which configuration should you use?

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

Answer: A

Explanation:

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

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

Publishing data to and from Oracle has some restrictions.

References:

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

NEW QUESTION 80

- (Exam Topic 1)

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

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

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

A. Create SSIS configuration files for the packages.

B. Create an Integration Services catalog.

C. Install Data Quality Services.

D. Install Master Data services.

Answer: B

Explanation:

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

References:

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

NEW QUESTION 84

- (Exam Topic 1)

You administer a Microsoft SQL Server 2016 instance that has several SQL Server Agent jobs configured. SQL Server Agent jobs fail, the error messages returned by the job steps are truncated.

The following error message is an example of the truncated error message:

"Executed as user CONTOSO\ServiceAccount. ...0.4035.00 for 64-bit Copyright (C) Microsoft Corp

1984-2011. All rights reserved. Started 63513 PM Error 2012-06-23 183536.87 Code 0XC001000E Source UserImport Description Code 0x00000000 Source Log Import Activity Descript... The package execution fa... The step failed."

You need to ensure that all the details of the job step failures are retained for SQL Server Agent jobs. What should you do?

A. Expand agent logging to include information from all events.

B. Disable the Limit size of job history log feature.

C. Configure event forwarding.

D. Configure output files.

Answer: D

Explanation:

When you have a multiple-step job, then log all steps against a single file. Check the 'Append output to existing file' checkbox for all steps in the job that execute after the initial step. This results in a log file with all of the job steps from the last job execution. Each time the first step executes (each time the job is kicked-off) the file will be overwritten, so we have a record of the last set of output.

References: <https://www.mssqltips.com/sqlservertip/1411/verbose-sql-server-agent-logging/>

NEW QUESTION 87

- (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 90

- (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 91

- (Exam Topic 1)

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

Currently, Server01/HA is the primary replica.

You need to ensure that the following requirements are met:

Backup operations occur on Server02/HA.

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

Backup operations do not occur on Server01/HA.

How should you configure HaContoso?

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

Answer: B

Explanation:

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

Backup Priority (Lowest=1, Highest=100)

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

References:

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

NEW QUESTION 95

- (Exam Topic 1)

You administer two instances of Microsoft SQL Server 2016.

You deploy an application that uses a database on the named instance.

The application is unable to connect to the database on the named instance. You need to ensure that the application can connect to the named instance. What should you do?

- A. Use the Data Quality Client to configure the application.
- B. Start the SQL Server Browser Service.
- C. Use the Master Data Services Configuration Manager to configure the application.
- D. Start the SQL Server Integration Services Service.

Answer: B

Explanation:

The SQL ServerBrowser program runs as a Windows service. SQL Server Browser listens for incoming requests for Microsoft SQL Server resources and provides information about SQL Server instances installed on the computer. SQL Server Browser contributes to the following actions:

Browsing a list of available servers Connecting to the correct server instance Etc.

References: <https://docs.microsoft.com/en-us/sql/tools/configuration-manager/sql-server-browser-service>

NEW QUESTION 96

- (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 99

- (Exam Topic 1)

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

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

about the space used in the databases.

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

Answer Area

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

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

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

NEW QUESTION 104

- (Exam Topic 1)

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet goals.

Some question sets might have more than one correct solution, while others might not have a correct solution.

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

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

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

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

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

Does this meet the goal?

- A. Yes
B. No

Answer: B

NEW QUESTION 107

- (Exam Topic 1)

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

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

Your company has several Microsoft 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 != 'TORN_PAGE_DETECTION'
GO
```

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

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

Does the solution meet the goal?

- A. Yes
- B. NO

Answer: B

NEW QUESTION 109

- (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, 

|               |
|---------------|
| ▼             |
| rowcnt        |
| stats_date    |
| rowmodctr     |
| stats_collect |

 (id, indid), 

|               |
|---------------|
| ▼             |
| rowcnt        |
| stats_date    |
| rowmodctr     |
| stats_collect |



FROM sys.sysindexes
WHERE 

|               |
|---------------|
| ▼             |
| rowmodctr     |
| stats_collect |
| stats_date    |
| rowcnt        |

 (id, indid) <= DATEADD(DAY, -7, GETDATE())

AND 

|               |
|---------------|
| ▼             |
| stats_collect |
| rowmodctr     |
| stats_date    |
| rowcnt        |

 > 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 112

- (Exam Topic 1)

A Microsoft SQL Server database named DB1 has two filegroups named FG1 and FG2. You implement a backup strategy that creates backups for the filegroups. DB1 experiences a failure. You must restore FG1 and then FG2.
You need to ensure that the database remains in the RECOVERING state until the restoration of FG2 completes. After the restoration of FG2 completes, the database must be online.
What should you specify when you run the recovery command?

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

Answer: A

NEW QUESTION 113

- (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 116

- (Exam Topic 1)

You administer a database that is used for reporting purposes. The database has a large fact table that contains three hundred million rows. The table includes a clustered columnstore index and a nonclustered index on the ProductID column. New rows are inserted into the table every day. Performance of queries that filter the Product ID column have degraded significantly. You need to improve the performance of the queries. Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions	Answer Area
Drop the clustered columnstore index.	
Create a nonclustered index on ProductID.	
Drop and recreate the clustered columnstore index.	
Create a nonclustered columnstore index on ProductID.	
Recreate the clustered columnstore index using DROP EXISTING.	
Create a clustered rowstore index on ProductID.	
Rebuild the clustered columnstore index.	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1: Drop the clustered columnstore index

Step 2: Create a clustered rowstore index on ProductID.

Rowstore indexes perform best on queries that seek into the data, searching for a particular value, or for queries on a small range of values. Use rowstore indexes with transactional workloads since they tend to require mostly table seeks instead of table scans.

Step 3: Create a nonclustered index on ProductID

NEW QUESTION 119

- (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 122

- (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 RCLS [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 123

- (Exam Topic 2)

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

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

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

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

Answer: B

Explanation:

sys.dm_os_wait_stats

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

Columns include: waiting_tasks_count

Number of waits on this wait type.

This counter is incremented at the start of each wait.

NEW QUESTION 124

- (Exam Topic 2)

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

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

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

- Fulfilled
- Shipped
- Ordered

Received

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

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

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

Answer: A

NEW QUESTION 126

- (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 127

- (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 132

- (Exam Topic 2)

You plan to deploy SQL Server 2014.

You identify the following security requirements for the deployment:

Users must be prevented from intercepting and reading the T-SQL statements sent from the clients to the database engine.

All database files and log files must be encrypted if the files are moved to another disk on another server.

You need to identify which feature meets each security requirement. The solution must minimize processor overhead.

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

Features	Answer Area
Encrypting File System (EFS)	Users must be prevented from intercepting and reading the T-SQL statements sent from the clients to the database engine. <div>Feature</div>
Policy-Based Management	
Secure Socket Layer (SSL)	
Transparent Data Encryption (TDE)	All database files and log files must be encrypted if the files are moved to another disk on another server. <div>Feature</div>
Windows BitLocker Drive Encryption (BitLocker)	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

- Secure Sockets Layer (SSL) encryption enables transmitting encrypted data across the network between an instance of SQL Server and a client application.
- Transparent data encryption (TDE) performs real-time I/O encryption and decryption of the data and log files.

NEW QUESTION 136

- (Exam Topic 2)

Overview

General Overview

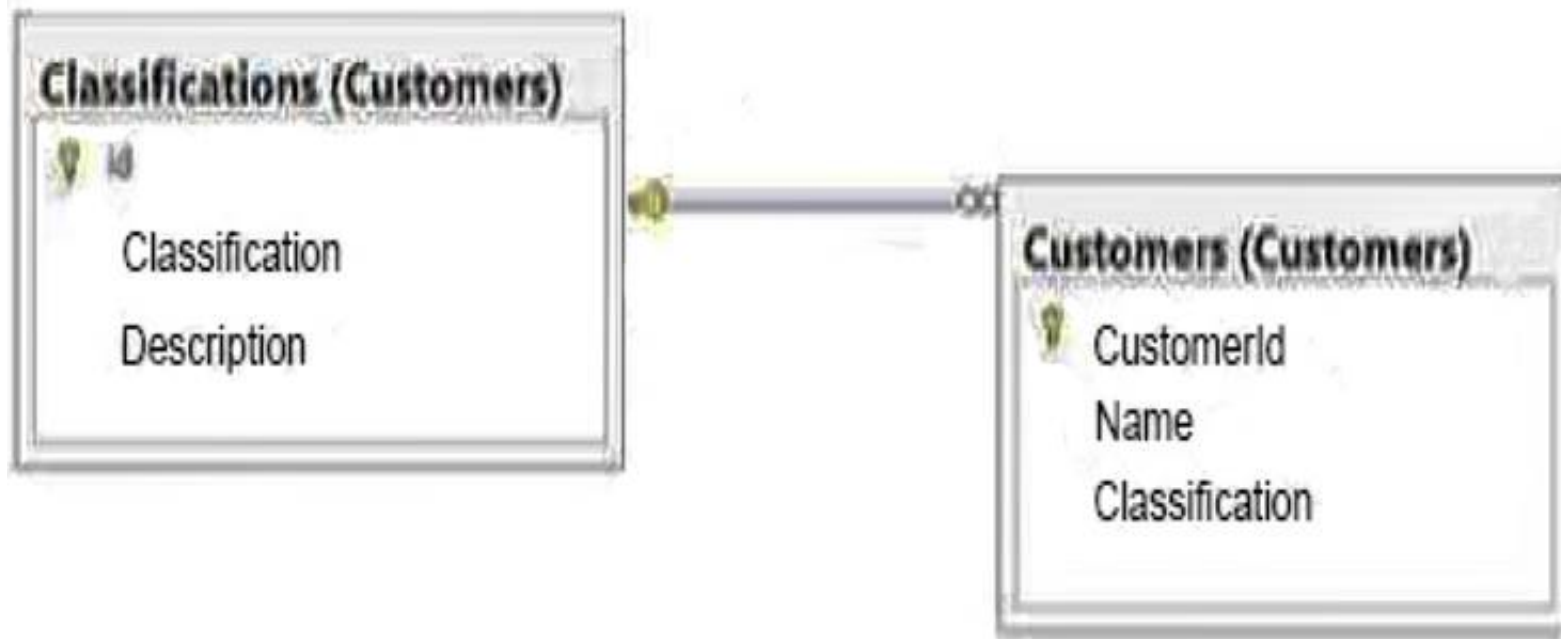
ADatum Corporation has offices in Miami and Montreal.

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 137

- (Exam Topic 2)

Overview

General Overview

ADatum Corporation has offices in Miami and Montreal.

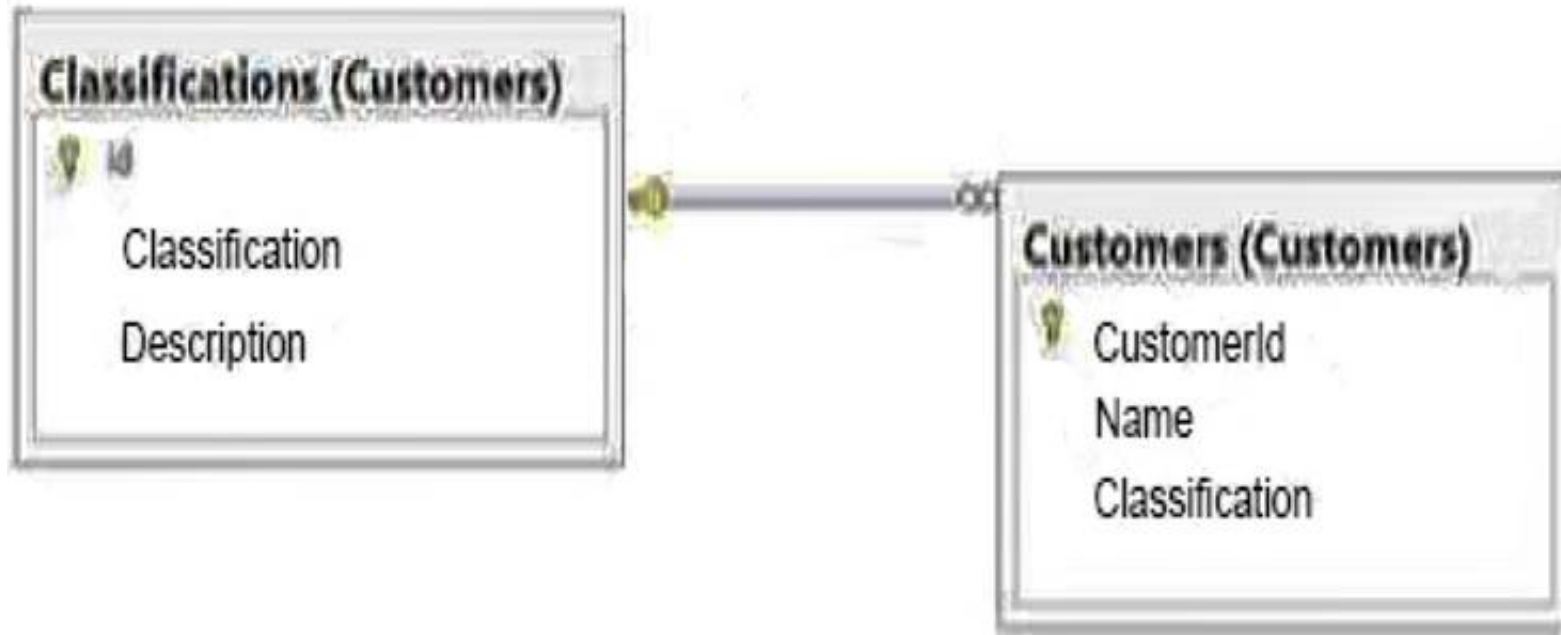
The network contains a single Active Directory forest named adatum.com. The offices connect to each other by using a WAN link that has 5-ms latency. A. Datum standardizes its database platform by using SQL Server 2014 Enterprise edition.

Databases

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

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

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



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

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

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

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

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

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

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

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

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

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

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

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

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

The nested stored procedures are never called directly. Design Requirements

Data Recovery

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

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

Classification Changes

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

Storage

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

Error Handling

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

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

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

Which change should you recommend?

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

Answer: A

Explanation:

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

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

NEW QUESTION 140

- (Exam Topic 2)

Overview

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

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

You are designing the database infrastructure to support a new web-based application that is being developed. The web application will be accessed at

www.litwareinc.com. Both internal employees and external partners will use the application.

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

Requirements Planned Changes

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

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

Database1 will also contain a stored procedure named usp_UpdateOrderDetails.

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

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

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

Inventory will contain over 100 GB of data.

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

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

Business Requirements

You have the following requirements:

Costs for new licenses must be minimized.

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

Development effort must be minimized whenever possible.

The storage requirements for databases must be minimized.

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

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

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

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

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

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

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

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

Answer: C

Explanation:

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

NEW QUESTION 144

- (Exam Topic 2)

Overview

General Overview


ADatum Corporation has offices in Miami and Montreal.

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

Databases

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

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



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

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

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

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

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

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

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

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

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

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

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

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

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

The nested stored procedures are never called directly. Design Requirements

Data Recovery

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

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

Classification Changes

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

Storage

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

Error Handling

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

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

You need to recommend a change to USP_3 to ensure that the procedure completes only if all of the UPDATE statements complete. Which change should you recommend?

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

Answer: B

Explanation:

Scenario:

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

Currently, if one of the UPDATE statements fails, the stored procedure continues to execute.

- When SET XACT_ABORT is ON, if a Transact-SQL statement raises a run-time error, the entire transaction is terminated and rolled back.

NEW QUESTION 149

- (Exam Topic 2)

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

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

What should you recommend?

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

Answer: A

Explanation:

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

NEW QUESTION 154

- (Exam Topic 2)

You administer a Microsoft SQL Server 2016 database named Contoso on a server named Server01. You need to prevent users from disabling server audits in Server01.

What should you create?

- A. an Alert
- B. a Resource Pool
- C. an Extended Event session

- D. a Policy
- E. a Database Audit Specification
- F. a SQL Profiler Trace
- G. a Server Audit Specification

Answer: D

NEW QUESTION 158

- (Exam Topic 2)

You administer a Microsoft SQL Server 2016 server.

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

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

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

Answer: CD

Explanation:

How To Enable Instant File Initialization

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

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

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

References:

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

NEW QUESTION 161

- (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 disk monitoring solution that meets the business requirements. What should you include in the recommendation?

- A. a SQL Server Agent alert
- B. a dynamic management view
- C. a maintenance plan
- D. an audit

Answer: B

Explanation:

Dynamic Management Views and Functions (Transact-SQL)

NEW QUESTION 165

- (Exam Topic 2)

Overview

General Overview

ADatum Corporation has offices in Miami and Montreal.

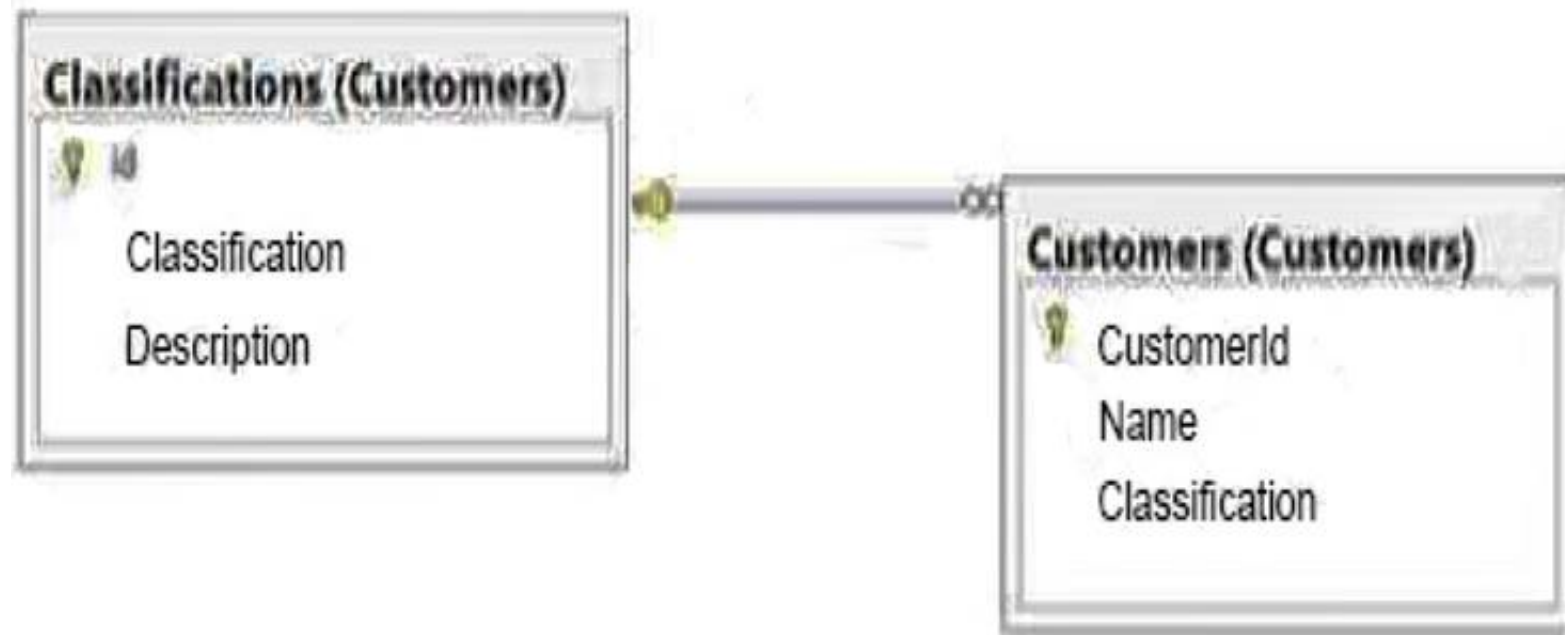
The network contains a single Active Directory forest named adatum.com. The offices connect to each other by using a WAN link that has 5-ms latency. A. Datum standardizes its database platform by using SQL Server 2014 Enterprise edition.

Databases

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

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

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



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

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

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

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

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

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

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

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

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

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

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

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

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

The nested stored procedures are never called directly. Design Requirements

Data Recovery

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

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

Classification Changes

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

Storage

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

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

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

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 170

- (Exam Topic 2)

You have two SQL Server instances named SQLDev and SQLProd that have access to various storage media. You plan to synchronize SQLDev and SQLProd.

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

The database schemas must be synchronized from SQLDev to SQLProd.

The database on SQLDev must be deployed to SQLProd by using a package.

The package must support being deployed to SQL Azure.

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

A. A database snapshot

B. A data-tier application

C. Change data capture

D. SQL Server Integration Services (SSIS)

Answer: B

Explanation:

*SIS supports connections to SQL Database by using the ADO.NET provider. OLEDB is not supported at this time. You can build the SSIS package connecting to SQL Database and create the data flow tasks the same way as you would against a typical on premise SQL Server.

<http://technet.microsoft.com/en-us/library/ee210546.aspx>

NEW QUESTION 173

- (Exam Topic 2)

Overview

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

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

You are designing the database infrastructure to support a new web-based application that is being developed. The web application will be accessed at

www.litwareinc.com. Both internal employees and external partners

will use the application.

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

Requirements Planned Changes

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

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

Database1 will also contain a stored procedure named usp_UpdateOrderDetails.

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

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

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

Inventory will contain over 100 GB of data.

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

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

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

Business Requirements

You have the following requirements:

Costs for new licenses must be minimized.

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

Development effort must be minimized whenever possible.

The storage requirements for databases must be minimized.

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

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

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

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

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

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

Which isolation level should recommend?

A. Read committed

B. Repeatable read

C. Read uncommitted

D. Serializable

Answer: B

Explanation:

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

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

NEW QUESTION 176

- (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 181

- (Exam Topic 2)

You are building a stored procedure for a SQL Azure database. The procedure will add multiple rows to a table. You need to design the stored procedure to meet the following requirements:

- If any of the new rows violates a table constraint, then no further additions must be attempted and all changes made by the stored procedure must be discarded.
 - If any errors occur, a row must be added to an audit table, and the original error must be returned to the caller of the stored procedure.
- What should you include in the design?

- A. An implicit transaction that has XACT_ABORT enabled
- B. An explicit transaction that has XACT_ABORT disabled
- C. An implicit transaction that has error handling enabled
- D. An explicit transaction that has error handling enabled

Answer: D

Explanation:

References:
[http://technet.microsoft.com/en-us/library/ms175127\(v=SQL.105\).aspx](http://technet.microsoft.com/en-us/library/ms175127(v=SQL.105).aspx)

NEW QUESTION 183

- (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 187

- (Exam Topic 2)

Overview

General Overview

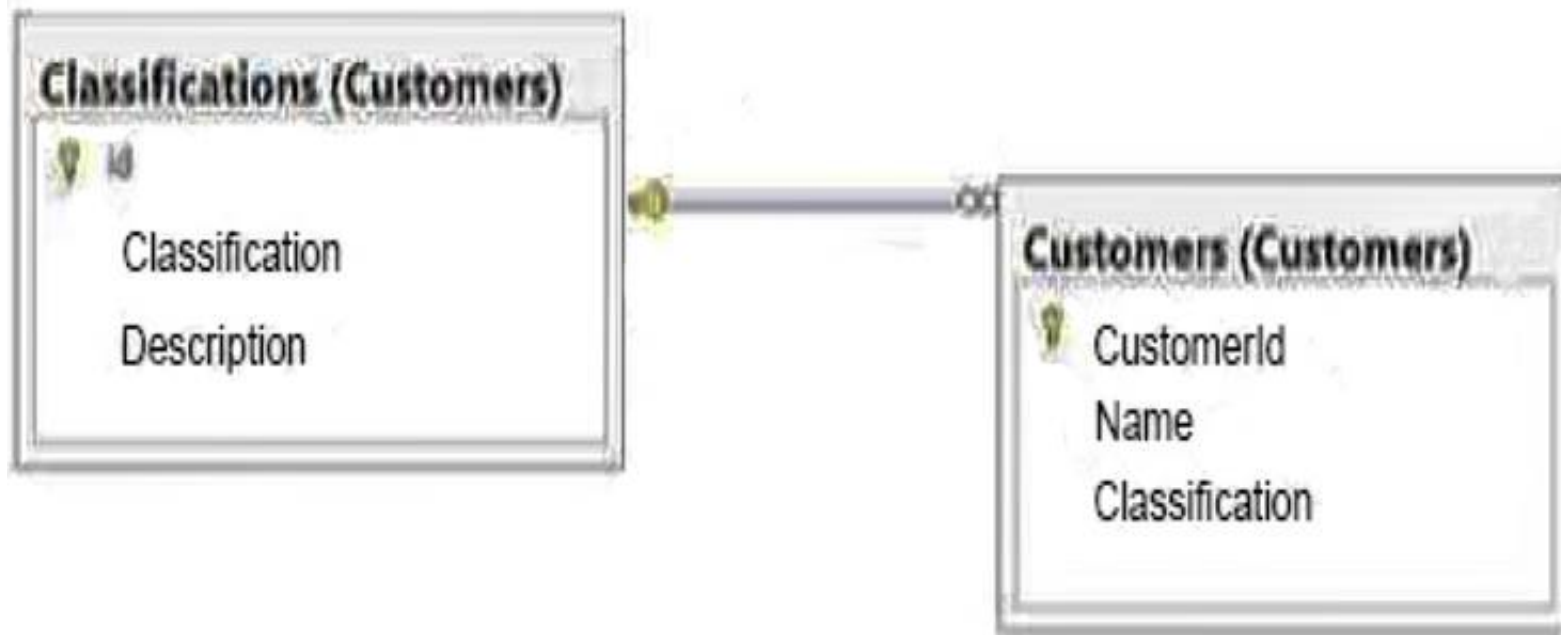
ADatum Corporation has offices in Miami and Montreal.

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

Databases

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

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



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

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

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

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

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

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

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

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

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

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

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

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

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

The nested stored procedures are never called directly. Design Requirements

Data Recovery

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

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

Classification Changes

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

Storage

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

Error Handling

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

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

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

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

Answer: A

Explanation:

Scenario:

- You must be able to recover data from the Inventory database if a storage failure occurs. You have a Recovery Point Objective (RPO) of one hour.
- A. Datum Corporation has offices in Miami and Montreal.
- SQL Server Log shipping allows you to automatically send transaction log backups from a primary database on a primary server instance to one or more secondary databases on separate secondary server instances. The transaction log backups are applied to each of the secondary databases individually.

NEW QUESTION 190

- (Exam Topic 2)

Overview

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

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

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

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

Requirements Planned Changes

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

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

Database1 will also contain a stored procedure named usp_UpdateOrderDetails.

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

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

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

Inventory will contain over 100 GB of data.

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

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

Business Requirements

You have the following requirements:

Costs for new licenses must be minimized.

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

Development effort must be minimized whenever possible.

The storage requirements for databases must be minimized.

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

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

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

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

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

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

You need to recommend a feature to support your backup solution. What should you include in the recommendation?

- A. Transparent Data Encryption (TDE)
- B. Column-level encryption
- C. An NTFS file permission
- D. A Secure Sockets Layer (SSL)

Answer: A

Explanation:

- Scenario: You must encrypt the backup files to meet regulatory compliance requirements. The encryption strategy must minimize changes to the databases and to the applications.

- Transparent data encryption (TDE) performs real-time I/O encryption and decryption of the data and log files. The encryption uses a database encryption key (DEK), which is stored in the database boot record for availability during recovery.

Transparent Data Encryption (TDE)

NEW QUESTION 193

.....

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>