DUMPSARENA

Administering a SQL Database Infrastructure

Microsoft 70-764

Version Demo

Total Demo Questions: 20

Total Premium Questions: 445

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

QUESTION NO: 1

You install a Microsoft SQL Server 2016 instance.

The instance will store data extracted from two databases running on Windows Azure SQL Database.

You hire a data steward to perform interactive data cleansing and ad hoc querying and updating of the database.

You need to ensure that the data steward is given the correct client tools to perform these tasks.

Which set of tools should you install?

A. SQL Server Management Studio and Distributed Replay Client

B. Master Data Services and Data Quality Client

C. Data Quality Client and Distributed Replay Client

D. Data Quality Client and SQL Server Management Studio

ANSWER: B

QUESTION NO: 2

You have a Microsoft SQL Server database named DB1. Nested and recursive triggers for DB1 are disabled. There is an existing DDL trigger named TableAudit in the database. The trigger captures the name of each table as the table is created.

You define a policy for the database by using SQL Server policy-based management. The policy requires that all table names use the prefix tbl. You set the evaluation mode for the policy to On change: prevent.

You observe that developers are able to add new tables that do not include the required prefix in the table name.

You need to ensure the policy is enforced.

What should you do?

A. Enable recursive triggers

B. Change the policy evaluation mode to On Schedule

C. Disable the TableAudit trigger

D. Enable nested triggers

ANSWER: D

Explanation:

If the nested triggers server configuration option is disabled, On change: prevent will not work correctly. Policy-Based Management relies on DDL triggers to detect and roll back DDL operations that do not comply with policies that use this evaluation mode. Removing the Policy-Based Management DDL triggers or disabling nest triggers, will cause this evaluation mode to fail or perform unexpectedly.

References:

https://docs.microsoft.com/en-us/sql/relational-databases/policy-based-management/administer-servers-by-using-policy-based-management

QUESTION NO: 3 - (DRAG DROP)

DRAG DROP

You have a SQL Azure database named Database1.

You need to design the schema for a table named table1.

Table1 will have less than one million rows.

Table1 will contain the following information for each row:

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

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

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

Select and Place:

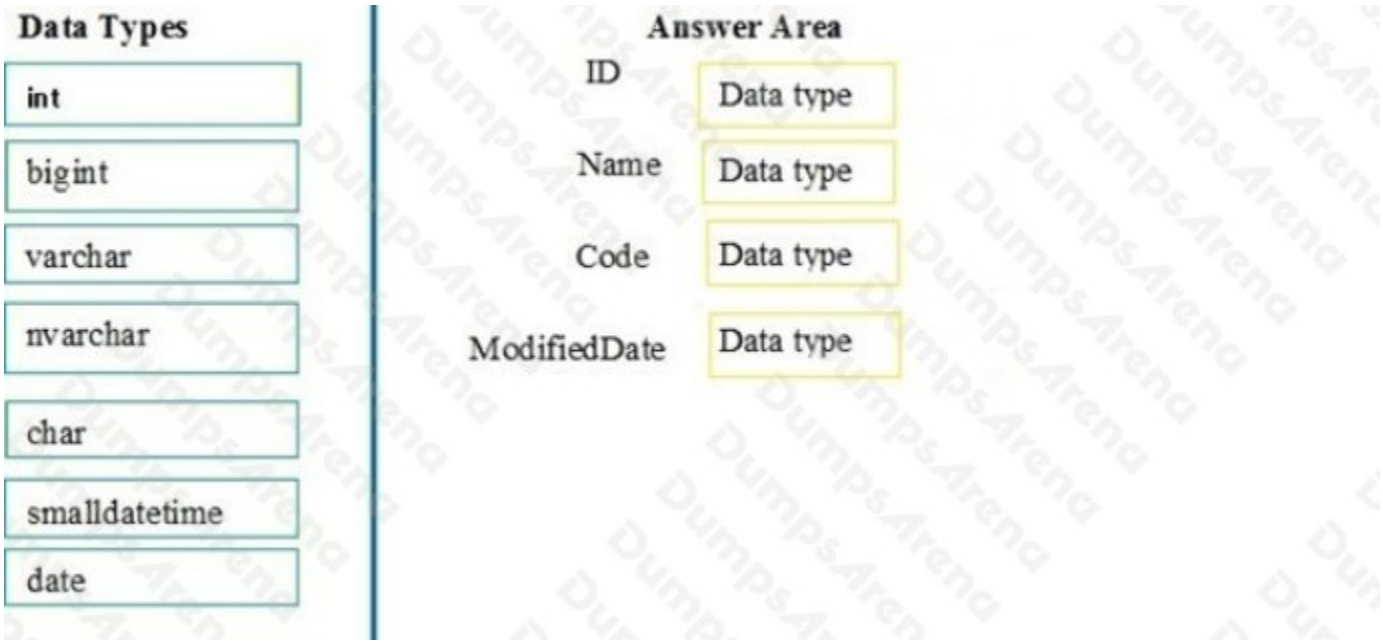

ANSWER:

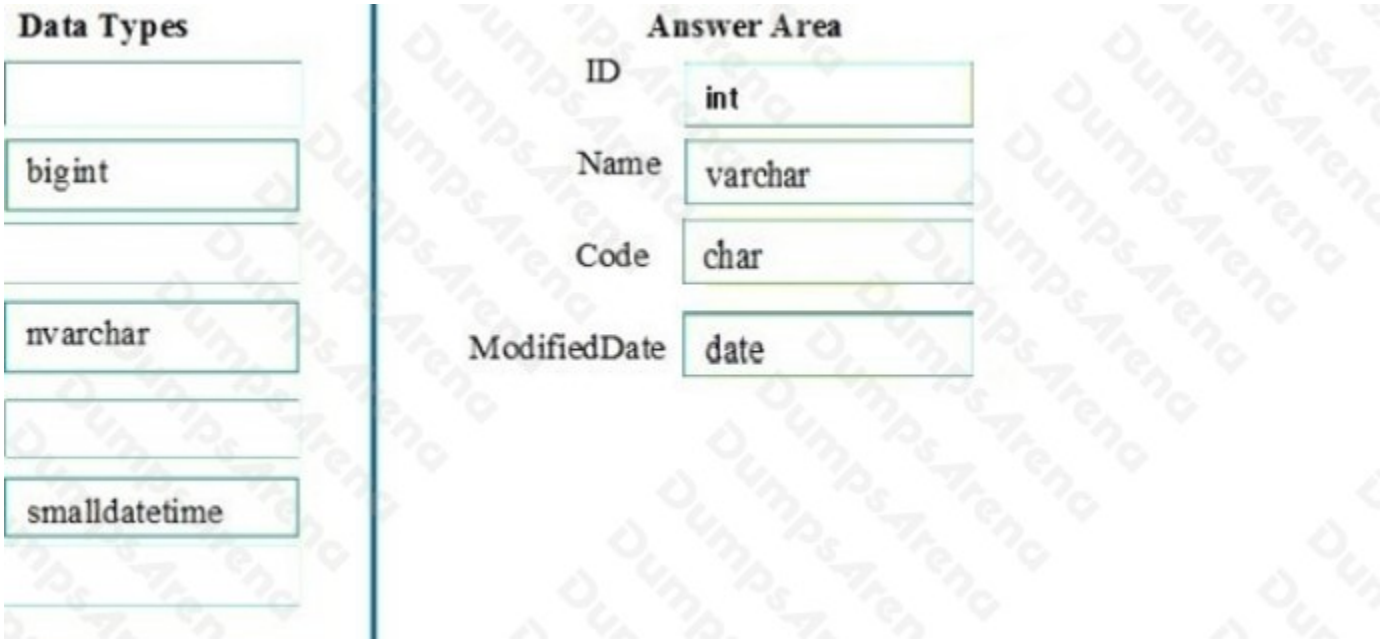

Explanation:

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

QUESTION NO: 4

You have a query that is used by a reporting dashboard. Users report that the query sometimes takes a long time to run. You need to recommend a solution to identify what is causing the issue.

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

A. Set the blocked process threshold, and then run SQL Server Profiler.

B. Set the blocked process threshold, and then create an alert.

C. Enable trace flag 1204, and then create an alert.

D. Create a job that queries the sys.dm_os_waiting_tasks dynamic management view.

ANSWER: B

Explanation:

Step 1: Turn on the blocked process report. This will look for any blocking taking 20 seconds or longer. - Make sure you don't have any pending changes

```
SELECT *
FROM sys.configurations
WHERE value <> value_in_use;
GO
exec sp_configure 'show advanced options', 1;
GO
RECONFIGURE
GO
exec sp_configure 'blocked process threshold (s)', 20;
GO
RECONFIGURE
GO
```

Step 2: Set up a trace to capture the blocked process report. Run it as a server side trace.

QUESTION NO: 5

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:

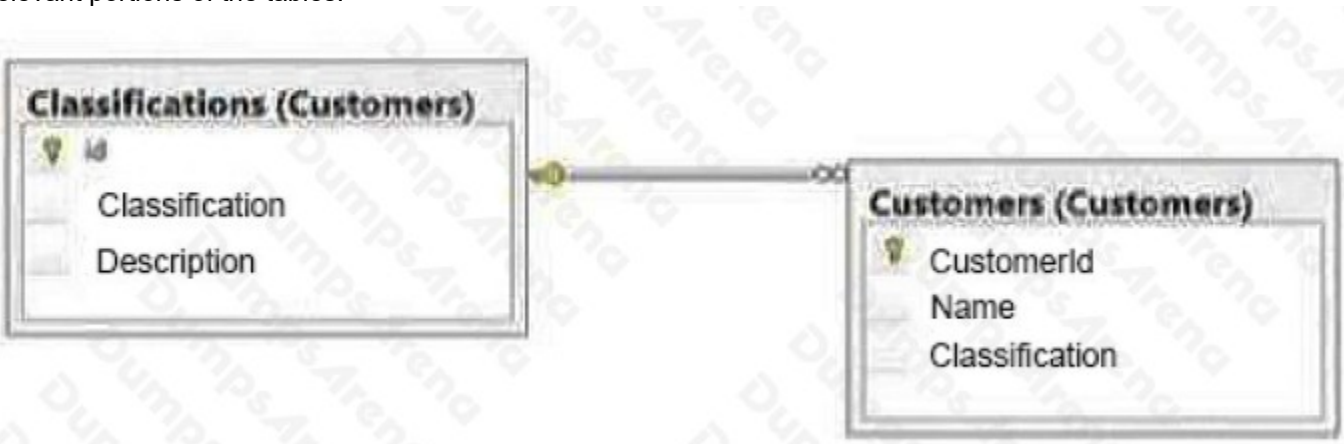

The following table shows the current data in the Classifications table:

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.
Datum wants to track which users run each stored procedure.

Storage
ADatum has limited storage. Whenever possible, all storage space should be minimized for all databases and all backups.
Error Handling
There is currently no error handling code in any stored procedure.
You plan to log errors in called stored procedures and nested stored procedures. Nested stored procedures are never called directly.
You need to recommend a solution for the planned changes to the customer classifications. What should you recommend? (Each correct answer presents part of the solution. Choose all that apply.)
Add a row to the Customers table each time a classification changes.

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

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

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

E. Implement change data capture.

ANSWER: C D

Explanation:

Scenario:

You plan to change the way customers are classified.

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

Incorrect Answers:

E: Change data capture provides information about DML changes on a table and a database. By using change data capture, you eliminate expensive techniques such as user triggers, timestamp columns, and join queries.

QUESTION NO: 6 - (HOTSPOT)

HOTSPOT

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.

Hot 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)
```

ANSWER:

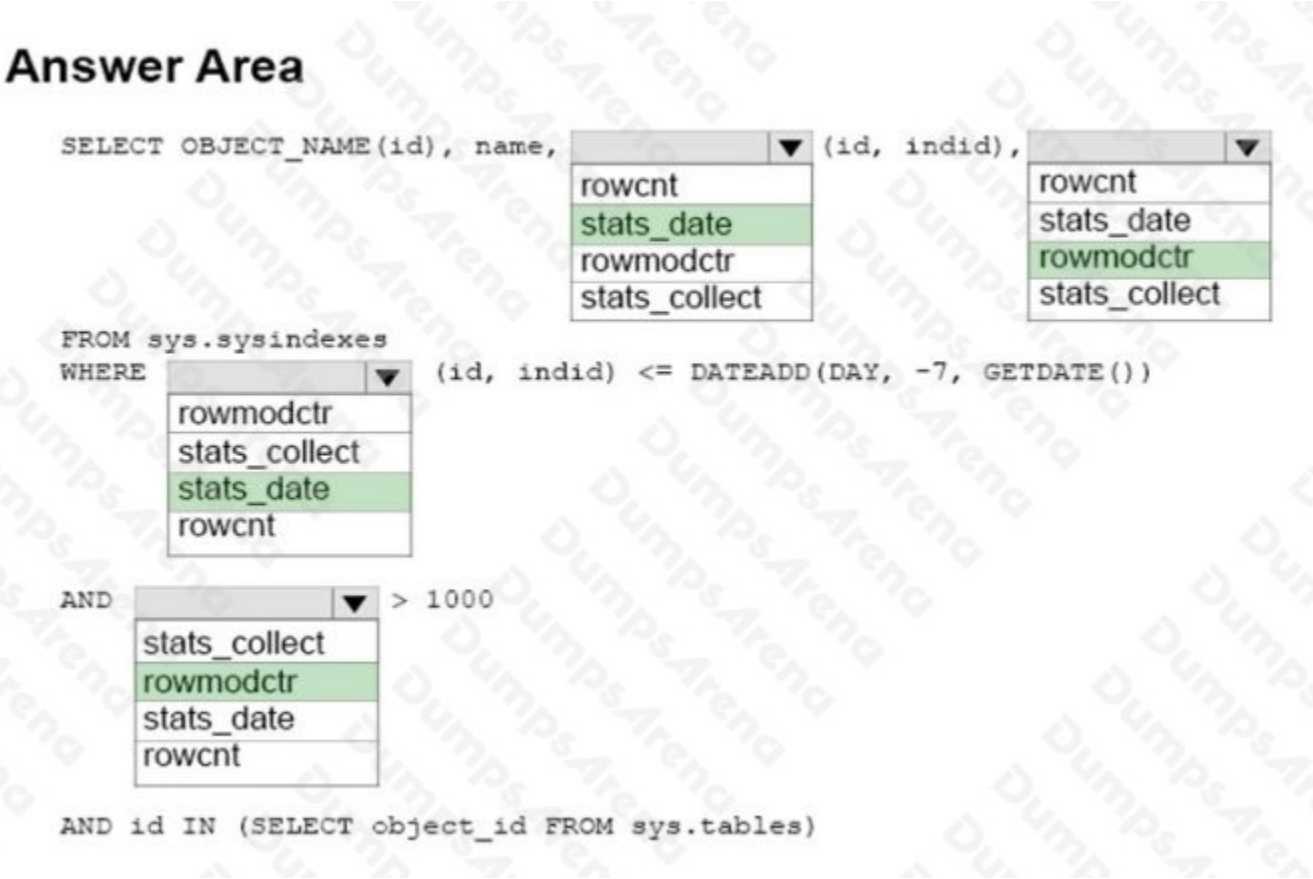

Explanation:

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

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

QUESTION NO: 7

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 is developing a new business intelligence application that will access data in a Microsoft Azure SQL Database instance. All objects in the instance have the same owner.

A new security principal named BI_User requires permission to run stored procedures in the database. The stored procedures read from and write to tables in the database. None of the stored procedures perform IDENTIFY_INSERT operations or dynamic SQL commands.

The scope of permissions and authentication of BI_User should be limited to the database. When granting permissions, you should use the principle of least privilege.

You need to create the required security principals and grant the appropriate permissions.

Solution: You run the following Transact-SQL statement in the database:

```
CREATE USER BI_User WITH PASSWORD = 'Pa$$wørd'
GRANT EXECUTE TO BI_User
EXEC sp_addrolemember 'db_datawriter', 'BI_user'
```

Does the solution meet the goal?

A. Yes

B. No

ANSWER: B

Explanation:

We need to add a login.

Secondly, it is enough to grant EXECUTE permissions on the stored procedures for database roles you want to be able to access the data. We do not need to add roles to this user.

Note: One method of creating multiple lines of defense around your database is to implement all data access using stored procedures or user-defined functions. You revoke or deny all permissions to underlying objects, such as tables, and grant EXECUTE permissions on stored procedures. This effectively creates a security perimeter around your data and database objects.

Best Practices

Simply writing stored procedures isn't enough to adequately secure your application. You should also consider the following potential security holes.

- Grant EXECUTE permissions on the stored procedures for database roles you want to be able to access the data.

- Revoke or deny all permissions to the underlying tables for all roles and users in the database, including the public role. All users inherit permissions from public. Therefore denying permissions to public means that only owners and sysadmin members have access; all other users will be unable to inherit permissions from membership in other roles. - Do not add users or roles to the sysadmin or db_owner roles. System administrators and database owners can access all database objects.

References: https://docs.microsoft.com/en-us/dotnet/framework/data/adonet/sql/managing-permissions-with-stored-procedures-in-sql-server

QUESTION NO: 8

General Overview

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

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

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

Databases

You plan to implement databases named Customers, Sales, Products, Current_Inventory, and TempReporting.

The Sales database contains a table named OrderTotals and a table named SalesInfo.

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

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

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

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

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

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

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

The GenInfo table is used for reports.

When the information in GenInfo is generated, a reporting process reads data from the Inv_Current table and queries information in the GenInfo table based on that data. The GenInfo table is deleted after the reporting process completes. The Products database contains tables named ProductNames and ProductTypes.

Current System

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

SQL Servers

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

Design Requirements

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

- Confidential information in the Current_ Inventory database that is accessed by ExternalApp1 must be securely stored.
- Direct access to database tables by developers or applications must be denied.
- The account used to generate reports must have restrictions on the hours when it is allowed to make a connection.
- Deadlocks must be analyzed with the use of Deadlock Graphs.
- In the event of a SQL Server failure, the databases must remain available.
- Software licensing and database storage costs must be minimized.
- Development effort must be minimized.
- The Tempdb databases must be monitored for insufficient free space.
- Failed authentication requests must be logged.
- Every time a new row is added to the ProductTypes table in the Products database, a user defined function that validates the row must be called before the row is added to the table.
- When SPUpdateSalesInfo queries data in the OrderTotals table the first time, the same rows must be returned along with any newly added rows when SPUpdateSalesInfo queries data in the OrderTotals table the second time.

You need to plan the SQL Server 2012 deployment that meets the design requirements. Which of the following steps should you perform?

A. Upgrade the existing SQL Server 2005 server to SQL Server 2012.

B. Install one new server running SQL Server 2012.

C. Install two new servers running SQL Server 2012

D. Configure Failover Clustering

E. Configure AllwaysOn

ANSWER: A B E

QUESTION NO: 9

You maintain Microsoft SQL Server instances named SVR1 and SVR2 that are hosted on two different servers. You configure log shipping between the two instances as follows:

- DB1 on SVR1 is configured as the primary database.
- DB1 on SVR2 is configured as the secondary database for DB1 on SVR1.
- No monitoring server is configured.

You need to monitor error log messages about the copy job.

What are two possible ways to achieve this goal? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A. On SVR1, run the following Transact-SQL statement:
SELECT * FROM msdb.dbo.log_shipping_monitor_error_detail.

B. Use the Job Activity Monitor in SQL Server Management Studio by connecting to SVR1

C. View the Log Shipping Report in SQL Server Management Studio by connecting SVR1.

D. Use the Job Activity Monitor in SQL Server Management Studio by connecting to SVR2.

E. On SVR2 run the following Transact-SQL statement:
SELECT * FROM msdb.dbo.log_shipping_monitor_error_detail.

ANSWER: C E

Explanation:

C: The Log Shipping Report displays the status of any log shipping activity whose status is available from the server instance to which you are connected.

E: The history and status of the backup operation are stored at the primary server, and the history and status of the copy and restore operations are stored at the secondary server.

The log_shipping_monitor_error_detail table stores error details for log shipping jobs. You can query this table see the errors for an agent session. Optionally, you can sort the errors by the date and time at which each was logged. Each error is logged as a sequence of exceptions, and multiple errors (sequences) can per agent session.

References: https://docs.microsoft.com/en-us/sql/database-engine/log-shipping/view-the-log-shipping-report-sql-server-management-studio?view=sql-server-2017 https://docs.microsoft.com/en-us/sql/database-engine/log-shipping/monitor-log-shipping-transact-sql

QUESTION NO: 10 - (DRAG DROP)

DRAG DROP

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Select and Place:

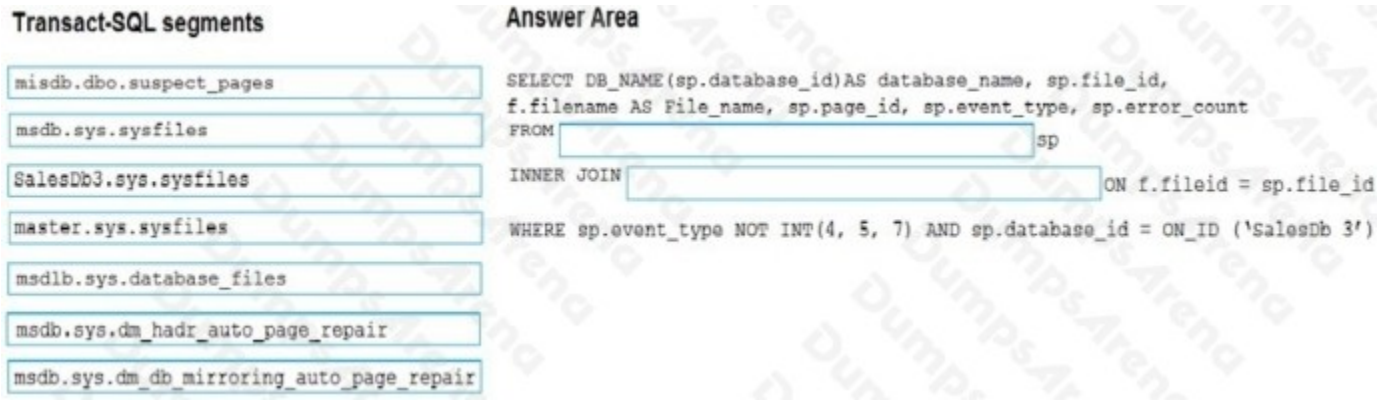

ANSWER:

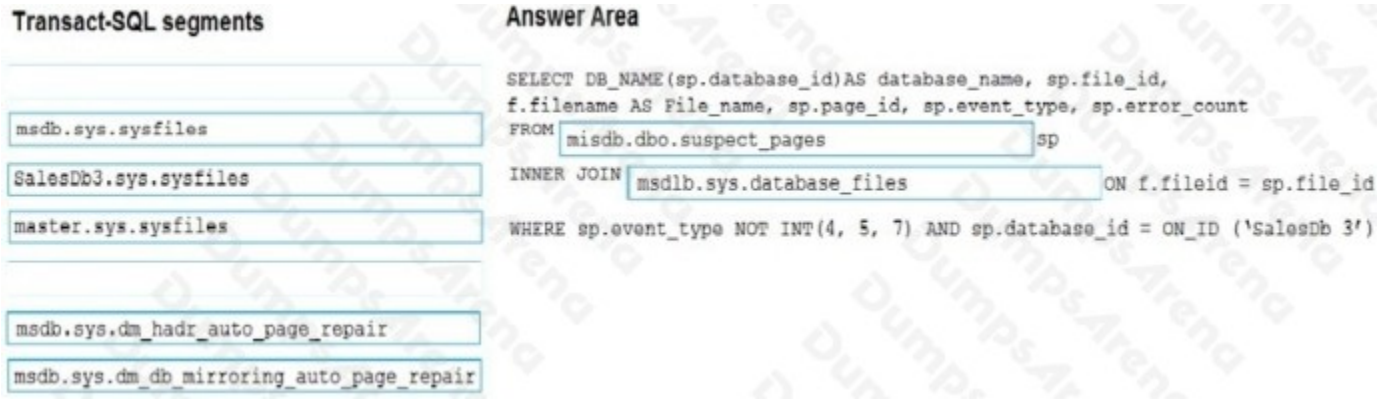

Explanation:

Box 1: msdb.dbo.suspect_pages

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

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

Box 2: msdb.sys.database_files

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

SELECT d.name AS databaseName, mf.name AS logicalFileName, mf.physical_name AS physicalFileName, sp.page_id, case sp.event_type

when 1 then N'823 or 824 error' when 2 then N'Bad Checksum' when 3 then N'Torn Page' when 4 then N'Restored' when 5 then N'Repaired' when 7 then N'Deallocated' end AS eventType, sp.error_count, sp.last_update_date

from msdb.dbo.suspect_pages as sp

join sys.databases as d ON sp.database_id = d.database_id join sys.master_files as mf on sp.[file_id] = mf.[file_id] and d.database_id = mf.database_id;

The result of this query will give you a high level view of where you have potential corruption in your databases, from here it is important to use tools such as DBCC CHECKDB and your backups to recover from in line with your RPO and RTO.

References:

https://docs.microsoft.com/en-us/sql/relational-databases/backup-restore/manage-the-suspect-pages-table-sql-server
https://blogs.sentryone.com/johnmartin/monitoring-for-suspect-pages/

QUESTION NO: 11

You maintain three datacenters in different geographical regions. You have a four-node failover cluster that hosts a Microsoft SQL Server Failover Cluster Instance (FCI) in Datacenter1.

You must extend FCI to add four nodes in Datacenter2.

You need to configure a node majority quorum mode for the cluster.

What are the two possible ways to achieve the goal? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A. Use a cloud quorum witness that is hosted in Microsoft Azure

B. Use no witness

C. Use a disk witness that is hosted in Datacenter2

D. Use a disk witness that is hosted in Datacenter1

E. Use a file share witness that is hosted in Datacenter3

ANSWER: A E

Explanation:

A: Cloud Witness is a type of Failover Cluster quorum witness that uses Microsoft Azure to provide a vote on cluster quorum.

Cloud Witness is a new type of Failover Cluster quorum witness that leverages Microsoft Azure as the arbitration point (figure below).

Figure: Multi-site stretched clusters with Cloud Witness as a quorum witness

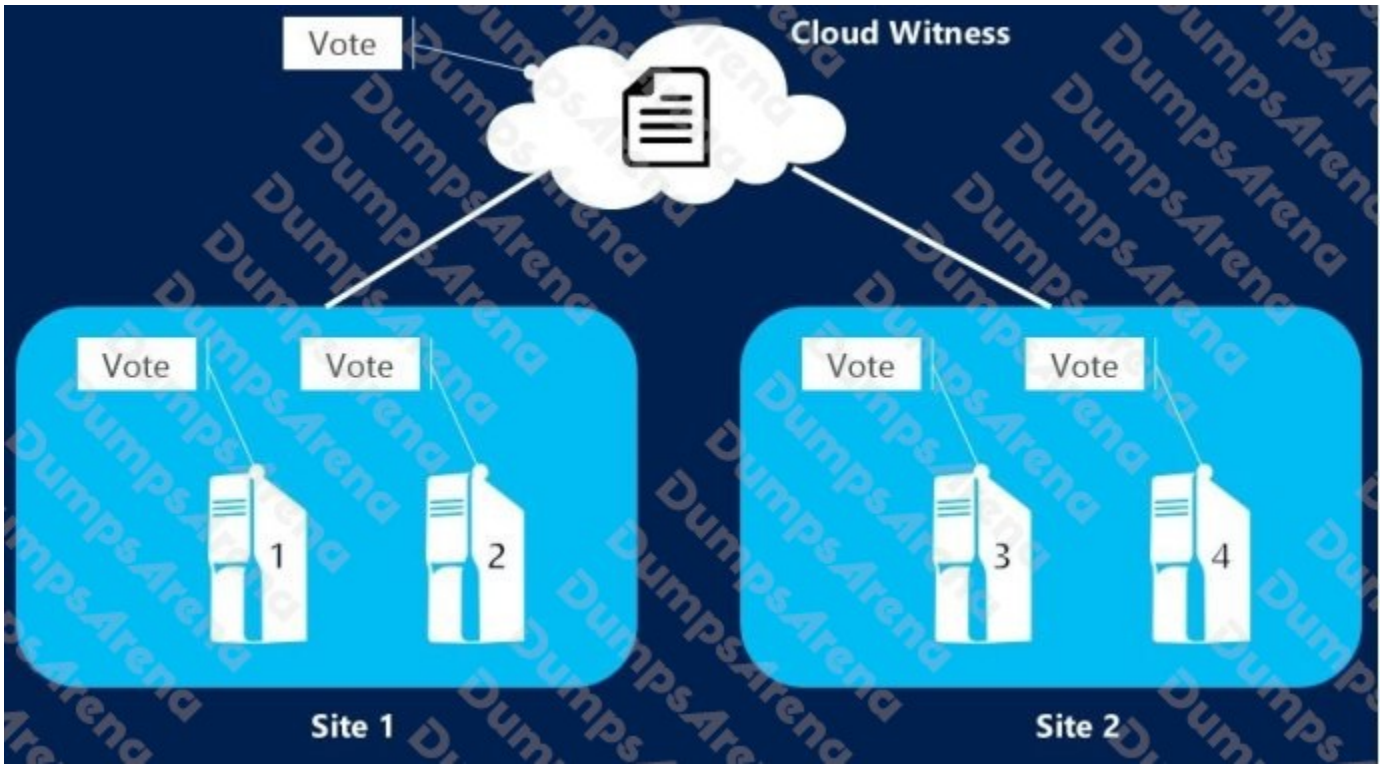

As shown in the figure , there is no third separate site that is required. Cloud Witness, like any other quorum witness, gets a vote and can participate in quorum calculations.

E: The figure below illustrates a multi-site stretched Failover Cluster quorum configuration with Windows Server 2016.

Figure: Using a File Share Witness as a quorum witness

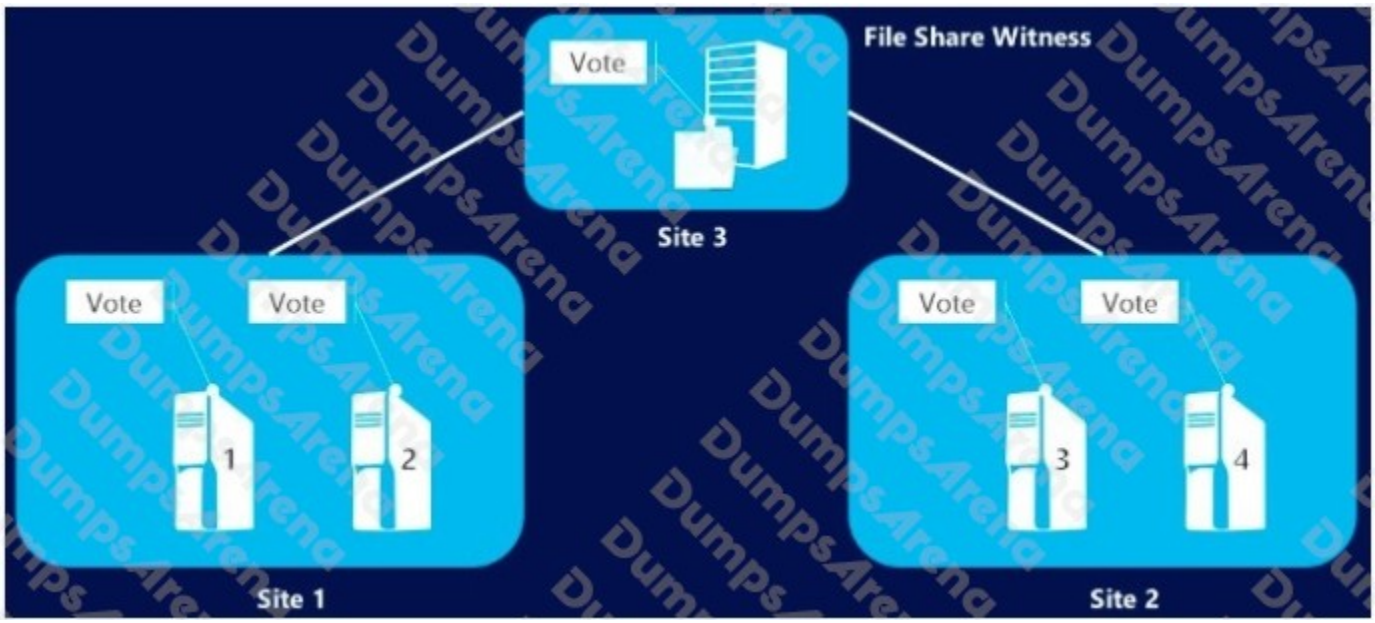

In case of power outage in one datacenter, to give equal opportunity for the cluster in other datacenter to keep it running, it is recommended to host the quorum witness in a location other than the two datacenters. This typically means requiring a third

separate datacenter (site) to host a File Server that is backing the File Share which is used as the quorum witness (File Share Witness).

References:

https://docs.microsoft.com/en-us/windows-server/failover-clustering/deploy-cloud-witness

QUESTION NO: 12

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/

QUESTION NO: 13

General Overview

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

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

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

Databases

You plan to implement databases named Customers, Sales, Products, Current_Inventory, and TempReporting.

The Sales database contains a table named OrderTotals and a table named SalesInfo.

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

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

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

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

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

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

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

The GenInfo table is used for reports.

When the information in GenInfo is generated, a reporting process reads data from the Inv_Current table and queries information in the GenInfo table based on that data. The GenInfo table is deleted after the reporting process completes. The Products database contains tables named ProductNames and ProductTypes.

Current System

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

SQL Servers

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

Design Requirements

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

- Confidential information in the Current_ Inventory database that is accessed by ExternalApp1 must be securely stored.
- Direct access to database tables by developers or applications must be denied.
- The account used to generate reports must have restrictions on the hours when it is allowed to make a connection.

- Deadlocks must be analyzed with the use of Deadlock Graphs.
- In the event of a SQL Server failure, the databases must remain available.
- Software licensing and database storage costs must be minimized.
- Development effort must be minimized.
- The Tempdb databases must be monitored for insufficient free space.
- Failed authentication requests must be logged.
- Every time a new row is added to the ProductTypes table in the Products database, a user defined function that validates the row must be called before the row is added to the table.
- When SPUpdateSalesInfo queries data in the OrderTotals table the first time, the same rows must be returned along with any newly added rows when SPUpdateSalesInfo queries data in the OrderTotals table the second time.

You need to ensure that the account used to generate reports can only connect during certain hours.

What should you configure?

A. A CHECK constraint.

B. Windows Server Resource Manager (WSRM).

C. Logon Triggers.

D. Login Auditing.

ANSWER: C

QUESTION NO: 14

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

- SQL Server auditing must be enabled on all server instances.
- Auditing results must be logged in the Windows Security even log.

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

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

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

A. Grant the manage auditing and security log permission to the SQL Server service account.

B. Grant the generate security audits permission on the SQL Server service account.

C. Update Windows security policy to audit object access.

D. Restart the SQL Server Agent service.

ANSWER: B C

Explanation:

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

- The audit object access setting must be configured to capture the events.
- The account that the SQL Server service is running under must have the generate security audits permission to write to the Windows Security log.

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

QUESTION NO: 15

You administer a Microsoft SQL Server 2012 instance.

You need to configure a new database to support FILETABLES.

What should you do? Choose all that apply.

A. Disable FILESTREAM on the Database.

B. Enable FILESTREAM on the Server Instance.

C. Configure the Database for Partial Containment.

D. Create a non-empty FILESTREAM file group.

E. Enable Contained Databases on the Server Instance.

F. Set the FILESTREAM directory name on the Database.

ANSWER: B D F

Explanation:

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

QUESTION NO: 16

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

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

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

Which feature should you use?

A. Database Mirroring

B. Peer-to-Peer Replication

C. Log Shipping

D. Availability Groups

ANSWER: B

QUESTION NO: 17 - (HOTSPOT)

HOTSPOT

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

You configure instances for a specific customer as an Always On Availability Group. The primary replica is located on-premises and the secondary replica is in Azure.

You need to configure the availability group for planned manual failovers and forced failovers.

In the table below, identify the failover mode that you must use for each failover type.

NOTE: Make only one selection in each column.

Hot Area:

ANSWER:

Explanation:

Column 1: Asynchronous-commit

Asynchronous-commit replicas support only the manual failover mode.

Column 2: Synchronous-commit

Synchronous-commit replicas support two settings—automatic or manual. The "automatic" setting supports both automatic failover and manual failover.

Three forms of failover exist: automatic failover (without data loss), planned manual failover (without data loss), and forced manual failover (with possible data loss), typically called forced failover.

References:

https://docs.microsoft.com/en-us/sql/database-engine/availability-groups/windows/failover-and-failover-modes-always-on-availability-groups?view=sql-server-2017

QUESTION NO: 18

You have a Microsoft SQL Server instance that has a database named DB1. The database is used for reporting purposes. You plan to capture all queries for a specific table and save the data as a text file.

You need to ensure that queries are captured and that a failure to capture a query will shut down the SQL Server instance.

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

NOTE: Each correct selection is worth one point.

A. Enable SQL Server Audit.

B. Use Extended Events.

C. Use a logon trigger.

D. Create a SQL Server Profiler trace and disable c2 audit tracing.

E. Enable SQL Server Query Store.

ANSWER: A D

Explanation:

The general process for creating and using an audit is as follows.

Create an audit and define the target.

Create either a server audit specification or database audit specification that maps to the audit. Enable the audit specification.

Enable the audit.

Read the audit events by using the Windows Event Viewer, Log File Viewer, or the fn_get_audit_file function.

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

QUESTION NO: 19

You administer several Microsoft SQL Server 2016 database servers.

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

Occasionally, merge replication fails due to timeout errors.

You need to reduce the occurrence of these timeout errors.

What should you do?

A. Set the Merge agent on the problem subscribers to use the slow link agent profile.

B. Create a snapshot publication, and reconfigure the problem subscribers to use the snapshot publication.

C. Change the Merge agent on the problem subscribers to run continuously.

D. Set the Remote Connection Timeout on the Publisher to 0.

ANSWER: A

Explanation:

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

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

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

QUESTION NO: 20

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

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

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

A. dynamic management views

B. Database Engine Tuning Advisor

C. SQL Server Profiler

D. Activity Monitor

E. Data Profile Viewer

ANSWER: B C

Explanation:

B: Database Engine Tuning Advisor examines how queries are processed in the databases you specify.

When you run a Profiler Trace and feed it to the Database Engine Tuning Advisor, it also looks for missing column statistics, and it can automatically create them for you. C: Use SQL Server Profiler to identify the cause of a deadlock. A deadlock occurs when there is a cyclic dependency between two or more threads, or processes, for some set of resources within SQL Server. Using SQL Server Profiler, you can create a trace that records, replays, and displays deadlock events for analysis. References: https://msdn.microsoft.com/en-us/library/ms188246.aspx