DUMPSBOSS.

Administering a SQL Database Infrastructure

Microsoft 70-764

Version Demo

Total Demo Questions: 20

Total Premium Questions: 445

Buy Premium PDF

https://dumpsboss.co

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

QUESTION NO: 1

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

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

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:

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 index fragmentation and index width issue. What should you include in the recommendation? (Each correct answer presents part of the solution. Choose all that apply.)

A. Change the data type of the lastModified column to smalldatetime.

B. Remove the lastModified column from the clustered index.

C. Change the data type of the modifiedBy column to tinyint.

D. Change the data type of the id column to bigint.

E. Remove the modifiedBy column from the clustered index.

F. Remove the id column from the clustered index.

ANSWER: B E

Explanation:

Scenario: 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)

QUESTION NO: 2

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

QUESTION NO: 3

You configure the Microsoft SQL Server Agent service to start automatically. You create a SQL Server Agent operator. You receive email notifications from other SQL Server Agent alerts that have been defined.

You have a SQL Server Agent that is set to notify you through email when error number 50010 occurs. You do not receive any email notifications related to this alert. You verify that the error occurred recently.

You need to determine why this alert is not sending you email notifications.

Why is the alert not sending email notifications?

A. The fail-safe operator is disabled

B. The alert is enabled

C. The event is not written to the Windows Application log

D. The operator is disabled

ANSWER: D

Explanation:

The operator could be disabled.

Incorrect Answers:

A: A fail-safe operator is a user who receives the alert if the designated operator cannot be reached.

References:

https://www.mssqltips.com/sqlservertip/1523/how-to-setup-sql-server-alerts-and-email-operator-notifications/

QUESTION NO: 4

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: 5

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

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

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

A. A policy

B. A maintenance plan

C. An alert

D. A job

E. A trigger

ANSWER: D E

Explanation:

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

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

QUESTION NO: 6 - (DRAG DROP)

DRAG DROP

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

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

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

Select and Place:

ANSWER:

Explanation:

References:https://docs.microsoft.com/en-us/sql/database-engine/database-mirroring/upgrading-mirrored-instances?view=sql-server-2017

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.

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

Solution: You add the DatabaseMailUserRole to Mail1 in the master database.

Does the solution meet the goal?

A. Yes

B. No

ANSWER: B

Explanation:

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

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

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

QUESTION NO: 8

Background

Corporate Information

Fabrikam, Inc. is a retailer that sells electronics products on the Internet. The company has a headquarters site and one satellite sales office. You have been hired as the database administrator, and the company wants you to change the architecture of the Fabrikam ecommerce site to optimize performance and reduce downtime while keeping capital expenditures to a minimum. To help with the solution, Fabrikam has decided to use cloud resources as well as on-premise servers.

Physical Locations

All of the corporate executives, product managers, and support staff are stationed at the headquarters office. Half of the sales force works at this location. There is also a satellite sales office. The other half of the sales force works at the satellite office in order to have sales people closer to clients in that area. Only sales people work at the satellite location.

Problem Statement

To be successful, Fabrikam needs a website that is fast and has a high degree of system uptime. The current system operates on a single server and the company is not happy with the single point of failure this presents. The current nightly backups have been failing due to insufficient space on the available drives and manual drive cleanup often needing to happen to get past the errors. Additional space will not be made available for backups on the HQ or satellite servers. During your investigation, you discover that the sales force reports are causing significant contention.

Configuration

Windows Logins

The network administrators have set up Windows groups to make it easier to manage security. Users may belong to more than one group depending on their role. The groups have been set up as shown in the following table:

Group	Members
OurDomain\Management	All corporate executives
OurDomain\SalesStaff	All sales people
OurDomain\ProductionStaff	All product managers and support staff
OurDomain\AllUsers	Everyone
OurDomain\CustomerSupport	Customer support representatives

Server Configuration The IT department has configured two physical servers with Microsoft Windows Server 2012 R2 and SQL Server 2014 Enterprise Edition and one Windows Azure Server. There are two tiers of storage available for use by database files only a fast tier and a slower tier. Currently the data and log files are stored on the fast tier of storage only. If a possible use case exists, management would like to utilize the slower tier storage for data files. The servers are configured as shown in the following table:

Location	Server
Company headquarters	HQ_Server
Satellite sales office	Satellite_Server
Microsoft Windows Azure (cloud)	Cloud_File Server

Database

Currently all information is stored in a single database called ProdDB, created with the following script:

```
CREATE DATABASE ProdDB
GO
ALTER DATABASE ProdDB SET RECOVERY SIMPLE
GO
```

The Product table is in the Production schema owned by the ProductionStaff Windows group. It is the main table in the system so access to information in the Product table should be as fast as possible. The columns in the Product table are defined as shown in the following table:

Column	Data type
ProductID	INT
ProductName	VARCHAR(100)
ProductDescription	VARCHAR(MAX)
ProductPrice	SMALLMONEY
QuantityOnHand	INT
ProductCost	SMALLMONEY
ProductSupplierID	INT

The SalesOrderDetail table holds the details about each sale. It is in the Sales schema owned by the SalesStaff Windows group. This table is constantly being updated, inserted into, and read. The columns in the SalesOrderDetail table are defined as shown in the following table:

Column	Data type
SalesOrderDetailID	INT
ProductID	INT
SalePrice	SMALLMONEY
SaleQuantity	INT

Database Issues

The current database does not perform well. Additionally, a recent disk problem caused the system to go down, resulting in lost sales revenue. In reviewing the current system, you found that there are no automated maintenance procedures. The database is severely fragmented, and everyone has read and write access.

Requirements

Database

The database should be configured to maximize uptime and to ensure that very little data is lost in the event of a server failure. To help with performance, the database needs to be modified so that it can support in-memory data, specifically for the Product table, which the CIO has indicated should be a memoryoptimized table. The auto-update statistics option is set off on this database. Only product managers are allowed to add products or to make changes to the name, description, price, cost, and supplier. The changes are made in an internal database and pushed to the Product table in ProdDB during system maintenance time. Product managers and others working at the headquarters location also should be able to generate reports that include supplier and cost information.

Customer data access

Customers access the company's website to order products, so they must be able to read product information such asname, description, and price from the Product table. When customers place orders, stored procedures calledby the website update product quantityon-hand values. This means the product table is constantly updated at randomtimes.

Customer support data access

Customer support representatives need to be able to view and not update or change product information. Management does not want the customer support representatives to be able to see the product cost or any supplier information.

Sales force data access

Sales people at both the headquarters office and the satellite office must generate reports that read from the Product and SalesOrderDetail tables. No updates or inserts are ever made by sales people. These reports are run at random times and there can be no reporting downtime to refresh the data set except during the monthly maintenance window. The reports that run from the satellite office are process intensive queries with large data sets. Regardless of which office runs a sales force report, the SalesOrderDetail table should only return valid, committed order data; any orders not yet committed should be ignored.

Historical Data

The system should keep historical information about customers who access the site so that sales people can see how frequently customers log in and how long they stay on the site.

The information should be stored in a table called Customer Access. Supporting this requirement should have minimal impact on production website performance.

Backups The recovery strategy for Fabrikam needs to include the ability to do point in time restores and minimize the risk of data loss by performing transaction log backups every 15 minutes.

Database Maintenance

The company has defined a maintenance window every month when the server can be unavailable. Any maintenance functions that require exclusive access should be accomplished during that window.

Project milestones completed

Revoked all existing read and write access to the database, leaving the schema ownership in place.

Configured an Azure storage container secured with the storage account name MyStorageAccount with the primary access key StorageAccountKey on the cloud file server.

SQL Server 2014 has been configured on the satellite server and is ready for use.

On each database server, the fast storage has been assigned to drive letter F:, and the slow storage has been assigned to drive letter D:.

You need to implement changes to the system to reduce contention and improve performance of the SalesOrderDetail table.

Which three actions should you perform? Each correct answer presents part of the solution. (Choose three.)

A. Use (SNAPSHOT] hints in the report queries

B. ALTER DATABASE [ProdDB] SET READ_COMMITTED_SNAPSHOT ON

C. ALTER DATABASE [ProdDB] SET READ_COMMITTED_SNAPSHOT OFF

D. SET TRANSACTION ISOLATION LEVEL SNAPSHOT

E. Use (TABLOCK) hints in the report queries

F. SET TRANSACTION ISOLATION LEVEL SERIALIZABLE

G. ALTER DATABASE [ProdDB] SET ALLOW.SNAPSHOT ISOLATION ON

H. Use (SNAPSHOT] hints in the update statements

ANSWER: A B F

Explanation:

- Scenario: The SalesOrderDetail table holds the details about each sale. It is in the Sales schema owned by the SalesStaff Windows group. This table is constantly being updated, inserted into,and read.

- Regardless of which office runs a sales force report, the SalesOrderDetail table should only return valid, committed order data; any orders not yet committed should be ignored.

- READ_COMMITTED_SNAPSHOT { ON | OFF } ON Enables Read-Committed Snapshot option at the database level. When it is enabled, DML statements start generating row versions evenwhen no transaction uses Snapshot Isolation. Once this option is enabled, the transactions specifying the read committed isolation level use row versioning instead of locking.

When a transaction runs at the read committed isolation level, all statements see a snapshot of data as it exists at the start of the statement. OFF Turns off Read-Committed Snapshot option at the database level. Transactions specifying the READ COMMITTED isolation level use locking. ALTER DATABASE SET Options (Transact-SQL) SET Statements (Transact-SQL)

QUESTION NO: 9

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:

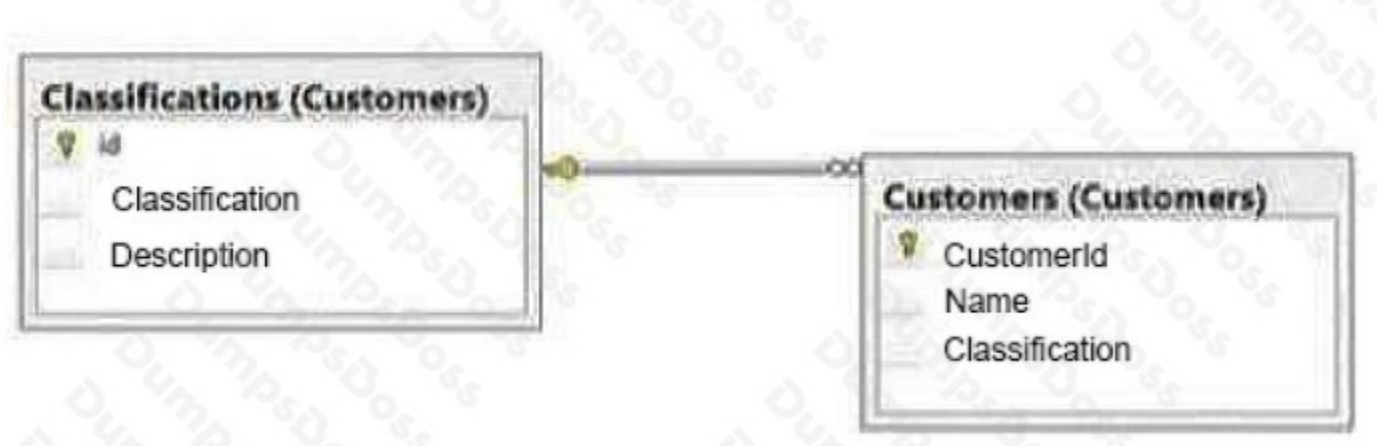

The following table shows the current data in the Classifications table:

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

QUESTION NO: 10

You are a database administrator for Contoso, Ltd. You configure a Microsoft SQL Server failover cluster with four nodes by using Windows Server 2012 R2 Datacenter Edition and SQL Server 2016 Enterprise edition. A server in the datacenter needs to be replaced. The server is part of the SQL Server Failover Instance (FCI).

You need to remove the FCI node for the server that will be replaced.

What should you do?

A. Evict the node from Failover Cluster Manager.

B. Run the Remove-ClusterResource Windows PowerShell cmdlet.

C. Run the Remove-Cluster Windows PowerShell cmdlet.

D. Remove the shared storage from Failover Cluster Manager.

ANSWER: A

Explanation:

Recover from an irreparable failure

Use the following steps to recover from an irreparable failure. The failure could be caused, for example, by the failure of a disk controller or the operating system. In this case, failure is caused by hardware failure in Node 1 of a two-node cluster.

1. After Node 1 fails, the SQL Server FCI fails over to Node 2.

2. Evict Node 1 from the FCI. To do this, from Node 2, open the Failover Cluster Manager snap-in, right-click Node1, click Move Actions, and then click Evict Node.

3. Verify that Node 1 has been evicted from the cluster definition.

4. Install new hardware to replace the failed hardware in Node 1.

5. Using the Failover Cluster Manager snap-in, add Node 1 to the existing cluster. For more information, see Before Installing Failover Clustering.

6. Ensure that the administrator accounts are the same on all cluster nodes.

7. Run SQL Server Setup to add Node 1 to the FCI. For more information, see Add or Remove Nodes in a SQL Server Failover Cluster (Setup).

References: https://docs.microsoft.com/en-us/sql/sql-server/failover-clusters/windows/recover-from-failover-cluster-instance-failure

QUESTION NO: 11

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: 12

You administer a Microsoft SQL Server 2016 instance. The instance contains a database that supports a retail sales application.

The application generates hundreds of transactions per second and is online 24 hours per day and 7 days per week. You plan to define a backup strategy for the database.

You need to ensure that the following requirements are met:

- No more than 5 minutes worth of transactions are lost.
- Data can be recovered by using the minimum amount of administrative effort.

What should you do? Choose all that apply.

A. Configure the database to use the SIMPLE recovery model.

B. Create a DIFFERENTIAL database backup every 4 hours.

C. Create a LOG backup every 5 minutes.

D. Configure the database to use the FULL recovery model.

E. Create a FULL database backup every 24 hours.

F. Create a DIFFERENTIAL database backup every 24 hours.

ANSWER: B C D E

Explanation:

If there are only three options, the CDE (exclude differential backup), is the best answer.

QUESTION NO: 13

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

QUESTION NO: 14

Background

Corporate Information

Fabrikam, Inc. is a retailer that sells electronics products on the Internet. The company has a headquarters site and one satellite sales office. You have been hired as the database administrator, and the company wants you to change the architecture of the Fabrikam ecommerce site to optimize performance and reduce downtime while keeping capital expenditures to a minimum. To help with the solution, Fabrikam has decided to use cloud resources as well as on-premise servers.

Physical Locations

All of the corporate executives, product managers, and support staff are stationed at the headquarters office. Half of the sales force works at this location. There is also a satellite sales office. The other half of the sales force works at the satellite office in order to have sales people closer to clients in that area. Only sales people work at the satellite location.

Problem Statement

To be successful, Fabrikam needs a website that is fast and has a high degree of system uptime. The current system operates on a single server and the company is not happy with the single point of failure this presents. The current nightly backups have been failing due to insufficient space on the available drives and manual drive cleanup often needing to happen to get past the errors. Additional space will not be made available for backups on the HQ or satellite servers. During your investigation, you discover that the sales force reports are causing significant contention.

Configuration

Windows Logins

The network administrators have set up Windows groups to make it easier to manage security. Users may belong to more than one group depending on their role. The groups have been set up as shown in the following table:

Group	Members
OurDomain\Management	All corporate executives
OurDomain\SalesStaff	All sales people
OurDomain\ProductionStaff	All product managers and support staff
OurDomain\AllUsers	Everyone
OurDomain\CustomerSupport	Customer support representatives

Server Configuration The IT department has configured two physical servers with Microsoft Windows Server 2012 R2 and SQL Server 2014 Enterprise Edition and one Windows Azure Server. There are two tiers of storage available for use by database files only a fast tier and a slower tier. Currently the data and log files are stored on the fast tier of storage only. If a possible use case exists, management would like to utilize the slower tier storage for data files. The servers are configured as shown in the following table:

Location	Server
Company headquarters	HQ_Server
Satellite sales office	Satellite_Server
Microsoft Windows Azure (cloud)	Cloud_File Server

Database

Currently all information is stored in a single database called ProdDB, created with the following script:

```
CREATE DATABASE ProdDB
GO
ALTER DATABASE ProdDB SET RECOVERY SIMPLE
GO
```

The Product table is in the Production schema owned by the ProductionStaff Windows group. It is the main table in the system so access to information in the Product table should be as fast as possible. The columns in the Product table are defined as shown in the following table:

Column	Data type
ProductID	INT
ProductName	VARCHAR(100)
ProductDescription	VARCHAR(MAX)
ProductPrice	SMALLMONEY
QuantityOnHand	INT
ProductCost	SMALLMONEY
ProductSupplierID	INT

The SalesOrderDetail table holds the details about each sale. It is in the Sales schema owned by the SalesStaff Windows group. This table is constantly being updated, inserted into, and read. The columns in the SalesOrderDetail table are defined as shown in the following table:

Column	Data type
SalesOrderDetailID	INT
ProductID	INT
SalePrice	SMALLMONEY
SaleQuantity	INT

Database Issues

The current database does not perform well. Additionally, a recent disk problem caused the system to go down, resulting in lost sales revenue. In reviewing the current system, you found that there are no automated maintenance procedures. The database is severely fragmented, and everyone has read and write access.

Requirements

Database

The database should be configured to maximize uptime and to ensure that very little data is lost in the event of a server failure. To help with performance, the database needs to be modified so that it can support in-memory data, specifically for the Product table, which the CIO has indicated should be a memoryoptimized table. The auto-update statistics option is set off on this database. Only product managers are allowed to add products or to make changes to the name, description, price, cost, and supplier. The changes are made in an internal database and pushed to the Product table in ProdDB during system maintenance time. Product managers and others working at the headquarters location also should be able to generate reports that include supplier and cost information.

Customer data access

Customers access the company's website to order products, so they must be able to read product information such asname, description, and price from the Product table. When customers place orders, stored procedures calledby the website update product quantityon-hand values. This means the product table is constantly updated at randomtimes.

Customer support data access

Customer support representatives need to be able to view and not update or change product information. Management does not want the customer support representatives to be able to see the product cost or any supplier information.

Sales force data access

Sales people at both the headquarters office and the satellite office must generate reports that read from the Product and SalesOrderDetail tables. No updates or inserts are ever made by sales people. These reports are run at random times and there can be no reporting downtime to refresh the data set except during the monthly maintenance window. The reports that run from the satellite office are process intensive queries with large data sets. Regardless of which office runs a sales force report, the SalesOrderDetail table should only return valid, committed order data; any orders not yet committed should be ignored.

Historical Data

The system should keep historical information about customers who access the site so that sales people can see how frequently customers log in and how long they stay on the site.

The information should be stored in a table called Customer Access. Supporting this requirement should have minimal impact on production website performance.

Backups The recovery strategy for Fabrikam needs to include the ability to do point in time restores and minimize the risk of data loss by performing transaction log backups every 15 minutes.

Database Maintenance

The company has defined a maintenance window every month when the server can be unavailable. Any maintenance functions that require exclusive access should be accomplished during that window.

Project milestones completed

- Revoked all existing read and write access to the database, leaving the schema ownership in place.
- Configured an Azure storage container secured with the storage account name MyStorageAccount with the primary access key StorageAccountKey on the cloud file server.

- SQL Server 2014 has been configured on the satellite server and is ready for use.
- On each database server, the fast storage has been assigned to drive letter F:, and the slow storage has been assigned to drive letter D:.

You need to configure security on the Product table for customer support representatives.

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

A. Create a view called CustProduct that includes columns ProductID. ProductName, Product Description, QuantityOnHand, ProductPrice, ProductCost, and ProductSupplierID.

B. GRANT ALL on CustProduct TO OurDomain\CustomerSupport

C. Create a user-defined data type called CustProduct that includes columns ProductID, ProductName, Product Description, and ProductPrice.

D. Create a view called CustProduct that includes columns ProductID, ProductName, Product Description, QuantityOnHand, and ProductPrice.

E. GRANT SELECT on CustProduct TO OurDomain\CustomerSupport.

F. GRANT SELECT on CustProduct TO public.

ANSWER: A E

Explanation:

Give access to CustomerSupport through a view. The view must include all these columns (refer to scenario). GRANT Object Permissions (Transact-SQL)

QUESTION NO: 15

You plan to migrate the db to azure.

You verify that all objects are valid for azure sql database. You need to ensure that users and logins are migrated to azure.

What should you do?

A. Use the Copy Database wizard

B. Use the Database Transfer wizard

C. Use the SQL Management Studio to deploy the db to azure

D. Back up the databasae from the local server and restore it to azure

ANSWER: C D

QUESTION NO: 16 - (HOTSPOT)

HOTSPOT

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.

Start of repeated scenario.

You are a database administrator for a company that has on-premises Microsoft SQL Server environment. There are two domains in separate forests. There are no trust relationships between the domains. The environment hosts several customer databases, and each customer uses a dedicated instance running SQL Server 2016 Standard edition. The customer environments are shown in the following table.

Customer	Domain	Description
AdventureWorks Cycles	DomainB	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	DomainA	Tailspin Toys has a database named **TSpinDB.** Tailspin Toys requires a custom application that monitors **TSpinDB** and captures information over time about which database objects are accessed and how frequently they are accessed.
Contoso, Ltd.	DomainB	The environment has a database named ConDB and is also running SQL Server Reporting Services (SSRS).
Wingtip Toys	DomainA	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 minimize data loss during the import process.
Wide World Importers	DomainB	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.

End of repeated scenario.

You need to configure auditing for the AdventureWorks environment. How should you complete the Transact-SQL statement? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

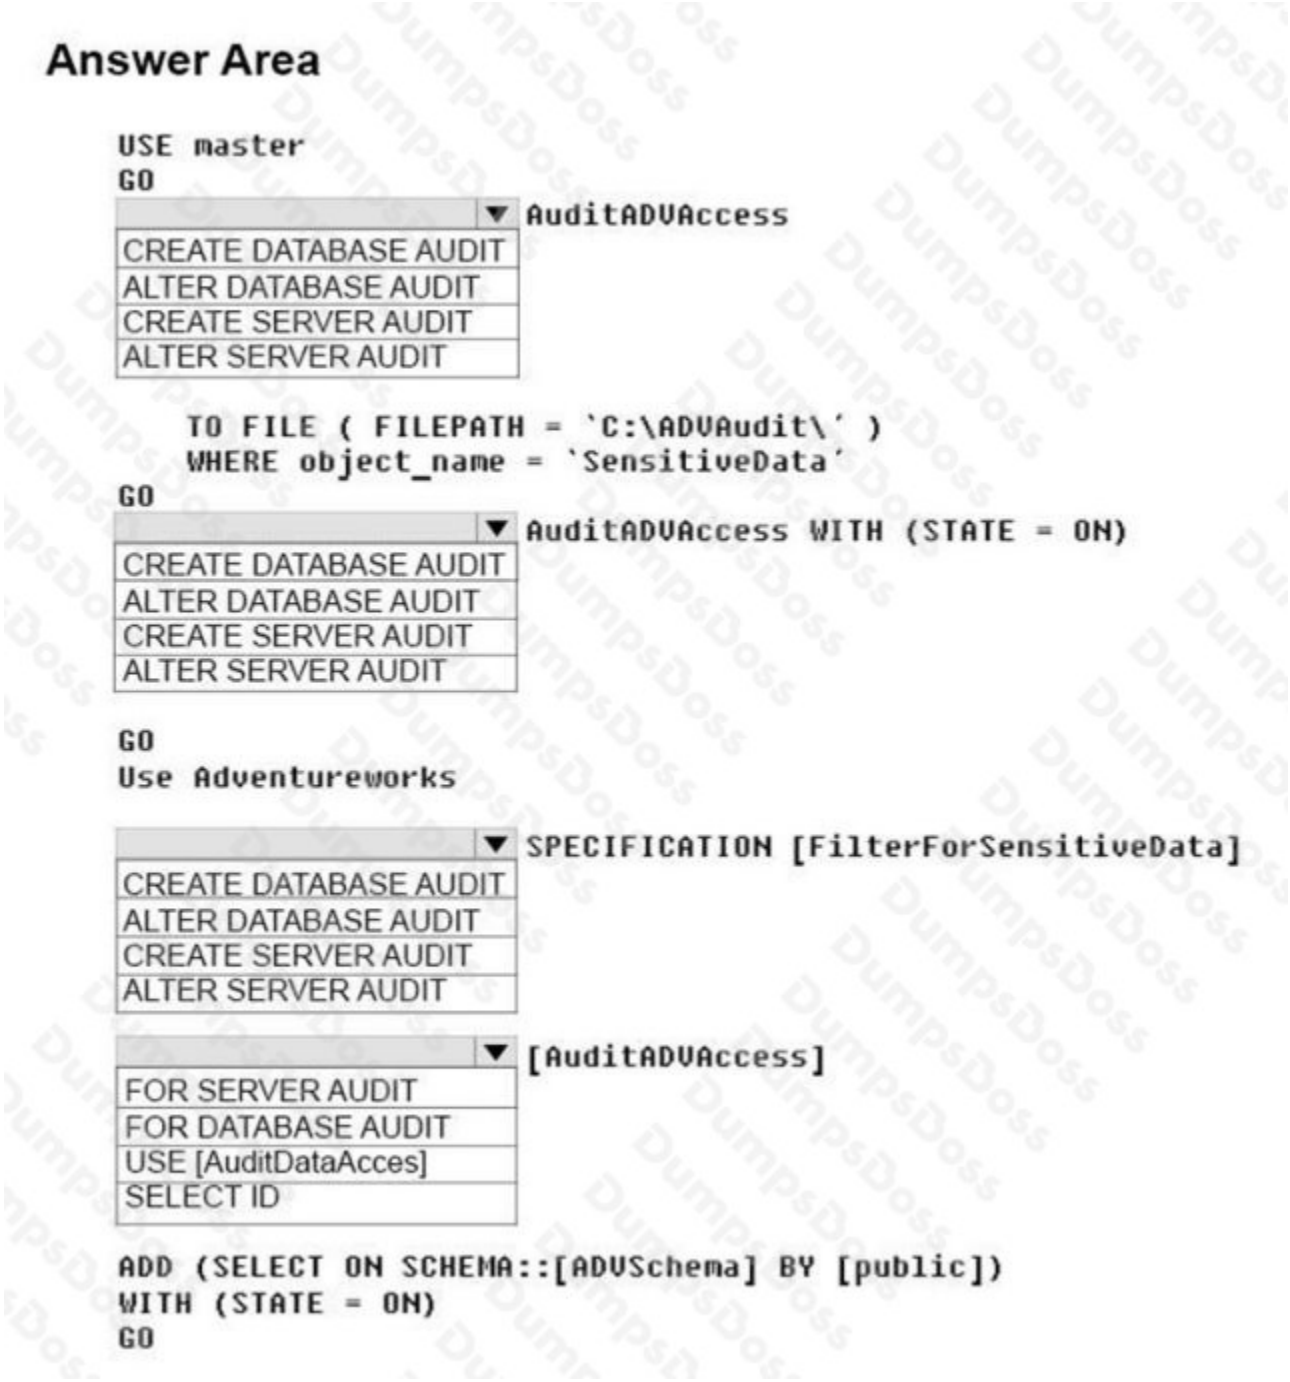

ANSWER:

Answer Area

```
USE master
GO
[ CREATE DATABASE AUDIT | ALTER DATABASE AUDIT | CREATE SERVER AUDIT | ALTER SERVER AUDIT ] AuditADVAccess
    TO FILE ( FILEPATH = 'C:\ADVAudit\' )
    WHERE object_name = 'SensitiveData'
GO
[ CREATE DATABASE AUDIT | ALTER DATABASE AUDIT | CREATE SERVER AUDIT | ALTER SERVER AUDIT ] AuditADVAccess WITH (STATE = ON)
GO
Use Adventureworks
[ CREATE DATABASE AUDIT | ALTER DATABASE AUDIT | CREATE SERVER AUDIT | ALTER SERVER AUDIT ] SPECIFICATION [FilterForSensitiveData]
[ FOR SERVER AUDIT | FOR DATABASE AUDIT | USE [AuditDataAcces] | SELECT ID ] [AuditADVAccess]
ADD (SELECT ON SCHEMA::[ADVSchema] BY [public])
WITH (STATE = ON)
GO
```

Explanation:

Scenario:

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.

You must implement auditing for all objects in the ADVSchema schema.

Box 1: CREATE SERVER AUDIT Create the server audit.

Box 2: ALTER SERVER AUDIT Enable the server audit.

Box 3: CREATE DATABASE AUDIT Create the database audit specification.

Box 4: FOR SERVER AUDIT

Example: The following example creates a server audit called Payrole_Security_Audit and then a database audit specification called Payrole_Security_Audit that audits SELECT and INSERT statements by the dbo user, for the HumanResources.EmployeePayHistory table in the AdventureWorks2012 database.

```
USE master ; GO

-- Create the server audit.

CREATE SERVER AUDIT Payrole_Security_Audit

TO FILE ( FILEPATH =

'C:\Program Files\Microsoft SQL Server\MSSQL13.MSSQLSERVER\MSSQL\DATA' ) ; GO

-- Enable the server audit.

ALTER SERVER AUDIT Payrole_Security_Audit

WITH (STATE = ON) ; GO

-- Move to the target database.

USE AdventureWorks2012 ; GO

-- Create the database audit specification.

CREATE DATABASE AUDIT SPECIFICATION Audit_Pay_Tables

FOR SERVER AUDIT Payrole_Security_Audit

ADD (SELECT , INSERT

ON HumanResources.EmployeePayHistory BY dbo )

WITH (STATE = ON) ; GO
```

References: https://docs.microsoft.com/en-us/sql/t-sql/statements/create-database-audit-specification-transact-sql?view=sql-server-2017

QUESTION NO: 17

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 attempt to restore a database on a new SQL Server instance and receive the following error message:

“Msg 33111, Level 16, State 3, Line 2

Cannot find server certificate with thumbprint ‘0x7315277C70764B1F252DC7A5101F6F66EFB1069D.”

You need to ensure that you can restore the database successfully.

Solution: You restore the certificate on the new instance.

Does this meet the goal?

A. Yes

B. No

ANSWER: B

Explanation:

In order to successfully make the restore in a different server you will need to create a master certificate in the detonation and transfer the certificates and backups in that order.

References: https://deibymarcos.wordpress.com/2017/11/15/how-to-restore-encrypted-databases-cannot-find-server-certificate-with-thumbprint/

QUESTION NO: 18

You have a server that has SQL Server 2014 installed.

The server contains 100 user databases.

You need to recommend a backup solution for the user databases.

The solution must meet the following requirements:

- Perform a transaction log backup every hour.
- Perform a full backup of each database every week.
- Perform a differential backup of each database every day.
- Ensure that new user databases are added automatically to the backup solution.

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

A. A maintenance plan

B. SQL Server Agent jobs

C. Policy-Based Management

D. A Data Definition Language (DDL) trigger

ANSWER: A

Explanation:

Maintenance plans create a workflow of the tasks required to make sure that your database is optimized, regularly backed up, and free of inconsistencies.

Maintenance plans can be created to perform the following task (among others): Back up the database and transaction log files. Database and log backups can be retained for a specified period. This lets you create a history of backups to be used if you have to restore the database to a time earlier than the last database backup. You can also perform differential backups.

QUESTION NO: 19 - (HOTSPOT)

HOTSPOT

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

You configure an Always On Availability Group for a customer. You must create log reports for the customer that detail when the log is flushed to disk on the primary and secondary replica.

You need to develop a report containing the requested information.

In the table below, identify the log type that you should use for each replica.

NOTE: Make only one selection in each column. Each correct selection is worth one point.

Hot Area:

Log type	Flush on primary	Flush on secondary
Log capture	○	○
Log hardened	○	○
Log receive/Log cache	○	○
Log flush	○	○

ANSWER:

Log type	Flush on primary	Flush on secondary
Log capture	○	○
Log hardened	○	◯
Log receive/Log cache	○	○
Log flush	◯	○

Explanation:

Flush on primary: Log flush

Log flush. Log data is generated and flushed to disk on the primary replica in preparation for replication to the secondary replica. It then enters the send queue.

Flush on secondary: Log hardened

The log is flushed on the secondary replica, and then a notification is sent to the primary replica to acknowledge completion of the transaction.

Incorrect Answers:

Not Log capture:

Log capture. Logs for each database are captured on the primary replica, compressed, and sent to the corresponding queue on the secondary replica. This process runs continuously as long as database replicas are connecting. If this process is not able to scan and enqueue the messages quickly enough, the log send queue continues to grow.

Not Log receive/Log cache

Log receive/Log cache. Each secondary replica gets messages from the primary replica and then caches the messages.

References: http://www.futas.net/ora/doc/SQL_Server_2016_Higher_availability_eBook_EN_US.pdf

QUESTION NO: 20 - (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.

Users reports that they encounter the following error when they query SalesDb1: “SQL Server detected a logical consistency-based I/O error: incorrect checksum (expected: 0x5d672d9b; actual:

0xdd672d98). It occurred during a read of page (1:232) in database ID 12 at offset 0x000000001d0000 in file F:\Databases\MSSQLServer Databases\MSSQL13.MSSQL2016\MSSQL\DATA

\SalesDb1.mdt.”

You must restore the impacted page from SalesDb1Full_1.bak. A single backup set named SalesDb1Log.bak was created since the latest full backup operation.

You need to restore the impacted page.

Which four 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.

Select and Place:

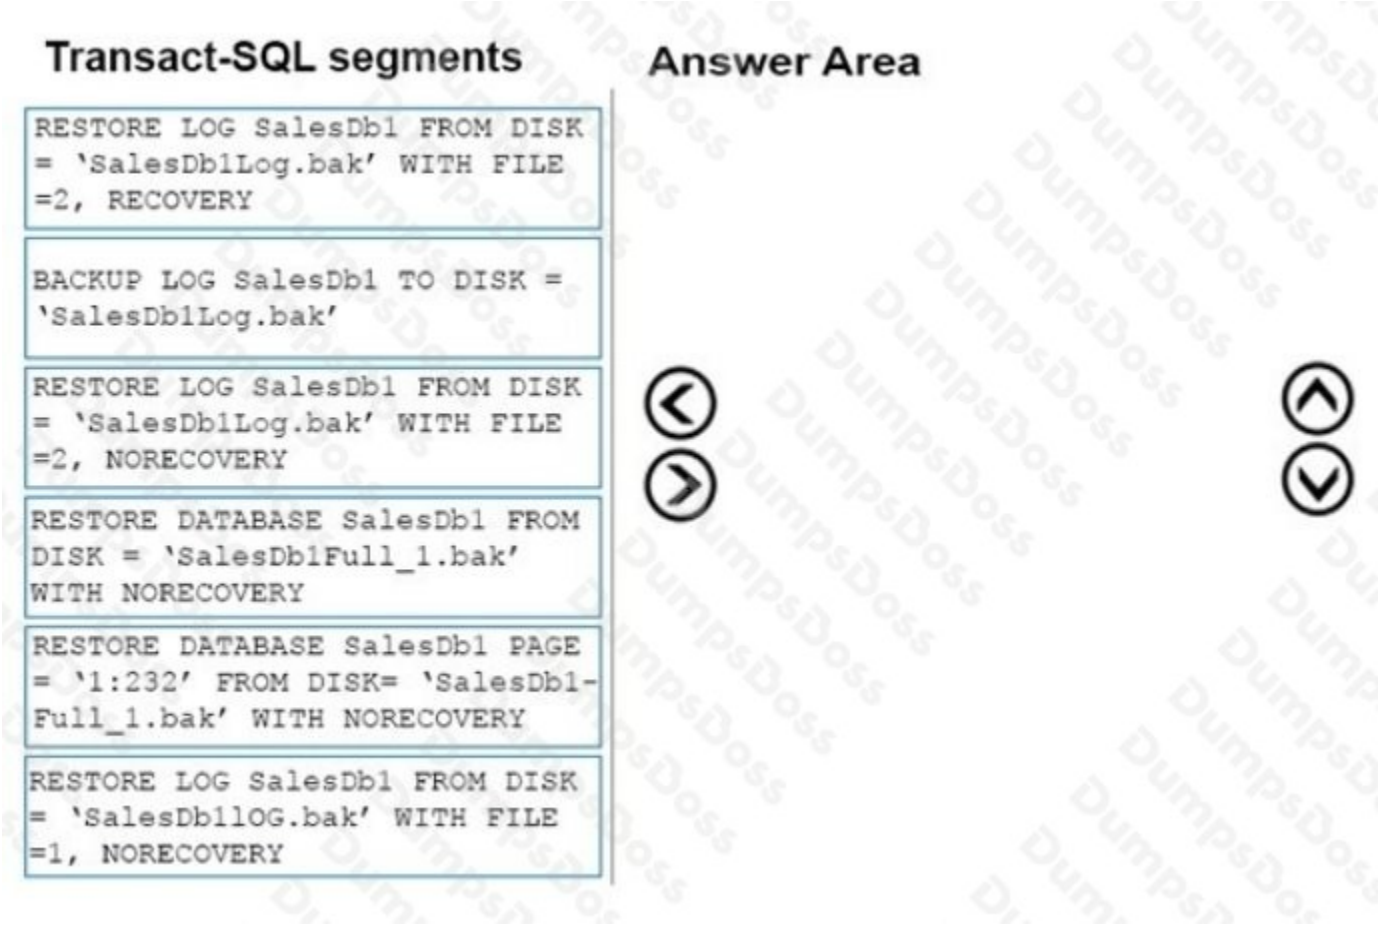
Transact-SQL segments
RESTORE LOG SalesDb1 FROM DISK = 'SalesDb1Log.bak' WITH FILE =2, RECOVERY
BACKUP LOG SalesDb1 TO DISK = 'SalesDb1Log.bak'
RESTORE LOG SalesDb1 FROM DISK = 'SalesDb1Log.bak' WITH FILE =2, NORECOVERY
RESTORE DATABASE SalesDb1 FROM DISK = 'SalesDb1Full_1.bak' WITH NORECOVERY
RESTORE DATABASE SalesDb1 PAGE = '1:232' FROM DISK= 'SalesDb1-Full_1.bak' WITH NORECOVERY
RESTORE LOG SalesDb1 FROM DISK = 'SalesDb1lOG.bak' WITH FILE =1, NORECOVERY
Answer Area

ANSWER:

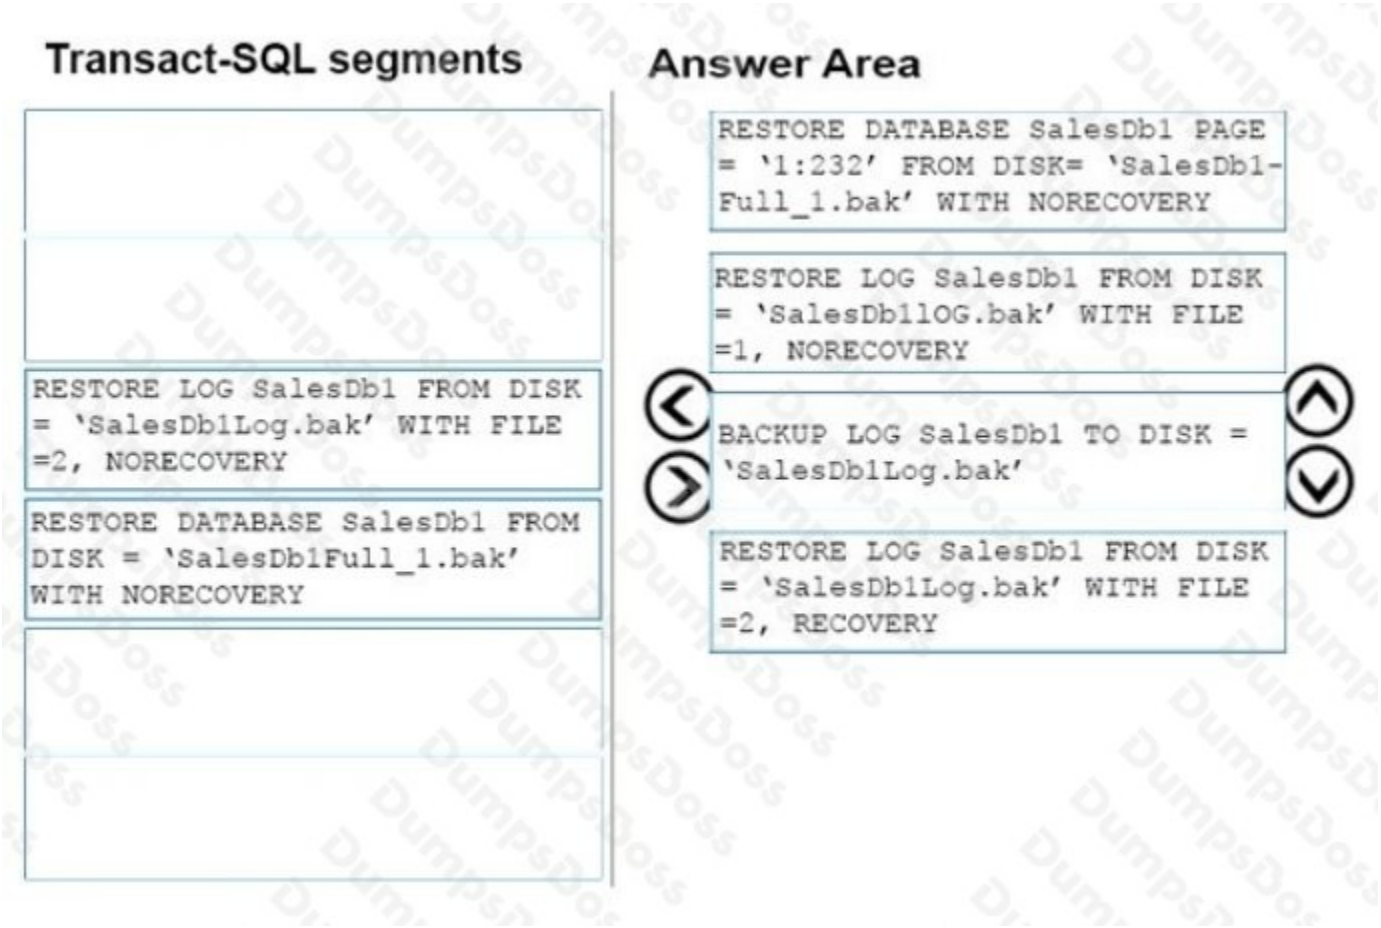

Explanation:

Step 1:

Restore page

Start a page restore with a full database, file, or filegroup backup that contains the page. In the RESTORE DATABASE statement, use the PAGE clause to list the page IDs of all of the pages to be restored.

Step 2:

Restore log file with norecovery. Use the first file (FILE = 1).

Step 3:

Backup the tail-end of the log.

Create a new log backup of the database that includes the final LSN of the restored pages, that is, the point at which the last restored page is taken offline.

Step 4:

Restore database with recovery. Use second file (FILE = 2).

Restore the new log backup. After this new log backup is applied, the page restore is completed and the pages are now usable. Example:

The following example restores four damaged pages of file B with NORECOVERY. Next, two log backups are applied with NORECOVERY, followed with the tail-log backup, which is restored with RECOVERY. This example performs an online restore. In the example, the file ID of file B is 1, and the page IDs of the damaged pages are 57, 202, 916, and 1016.

RESTORE DATABASE PAGE='1:57, 1:202, 1:916, 1:1016'

FROM

WITH NORECOVERY;

RESTORE LOG FROM

WITH NORECOVERY;

RESTORE LOG FROM

WITH NORECOVERY;

BACKUP LOG TO ;

RESTORE LOG FROM WITH RECOVERY;

References: https://docs.microsoft.com/en-us/sql/relational-databases/backup-restore/restore-pages-sql-server