

## 70-764 Dumps

### Administering a SQL Database Infrastructure (beta)

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



**NEW QUESTION 1**

- (Exam Topic 1)

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

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

What should you use?

- A. Activity Monitor
- B. Sp\_who3
- C. SQL Server Management Studio (SSMS) Object Explorer
- D. SQL Server Data Collector
- E. SQL Server Data Tools (SSDT)
- F. SQL Server Configuration Manager

**Answer:** D

**Explanation:**

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

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

Missing indexes can be found with the execution plans.

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

**NEW QUESTION 2**

- (Exam Topic 1)

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

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

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

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

Which permission should you grant?

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

**Answer:** A

**Explanation:**

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

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

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

**NEW QUESTION 3**

- (Exam Topic 1)

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

You manage a Microsoft SQL Server environment. You implement Transparent Data Encryption (TDE). A user will assist in managing TDE.

You need to ensure that the user can view the TDE metadata while following the principle of least privilege. Which permission should you grant?

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

**Answer:** G

**Explanation:**

Viewing the metadata involved with TDE requires the VIEW DEFINITION permission on the certificate. References:

<https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/transparent-data-encryption-tde>

**NEW QUESTION 4**

- (Exam Topic 1)

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

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

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

databases.

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

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

- A. Yes
- B. No

**Answer: B**

**Explanation:**

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

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

**NEW QUESTION 5**

- (Exam Topic 1)

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

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

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

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

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

Does the solution meet the goal?

- A. Yes
- B. No

**Answer: B**

**Explanation:**

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

References:

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

**NEW QUESTION 6**

- (Exam Topic 1)

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

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

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

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

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

- A. ALTER DATABASE DB1 SET QUERY\_STORE (QUERY\_CAPTURE\_MODE = ALL)
- B. ALTER DATABASE DB1 SET QUERY\_STORE (MAX\_STORAGE\_SIZE\_MB = 50)
- C. ALTER DATABASE DB1 SET QUERY\_STORE (CLEANUP\_POLICY = (STALE\_QUERY\_THRESHOLD\_DAYS = 14));
- D. ALTER DATABASE DB1 SET QUERY\_STORE (QUERY\_CAPTURE\_MODE = NONE)

**Answer: C**

**Explanation:**

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

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

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

```
ALTER DATABASE [QueryStoreDB]
```

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

References:

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

**NEW QUESTION 7**

- (Exam Topic 1)

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

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

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

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

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

Does the solution meet the goal?

- A. Yes
- B. No

**Answer: A**

**Explanation:**

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

**NEW QUESTION 8**

- (Exam Topic 1)

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

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

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

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

- A. Yes
- B. No

**Answer: B**

**Explanation:**

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

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

References:

[http://www.iddevelopment.info/data/SQLServer/DBA\\_tips/Database\\_Administration/DBA\\_20.shtml](http://www.iddevelopment.info/data/SQLServer/DBA_tips/Database_Administration/DBA_20.shtml)

**NEW QUESTION 9**

- (Exam Topic 1)

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

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

**Answer: D**

**Explanation:**

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

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

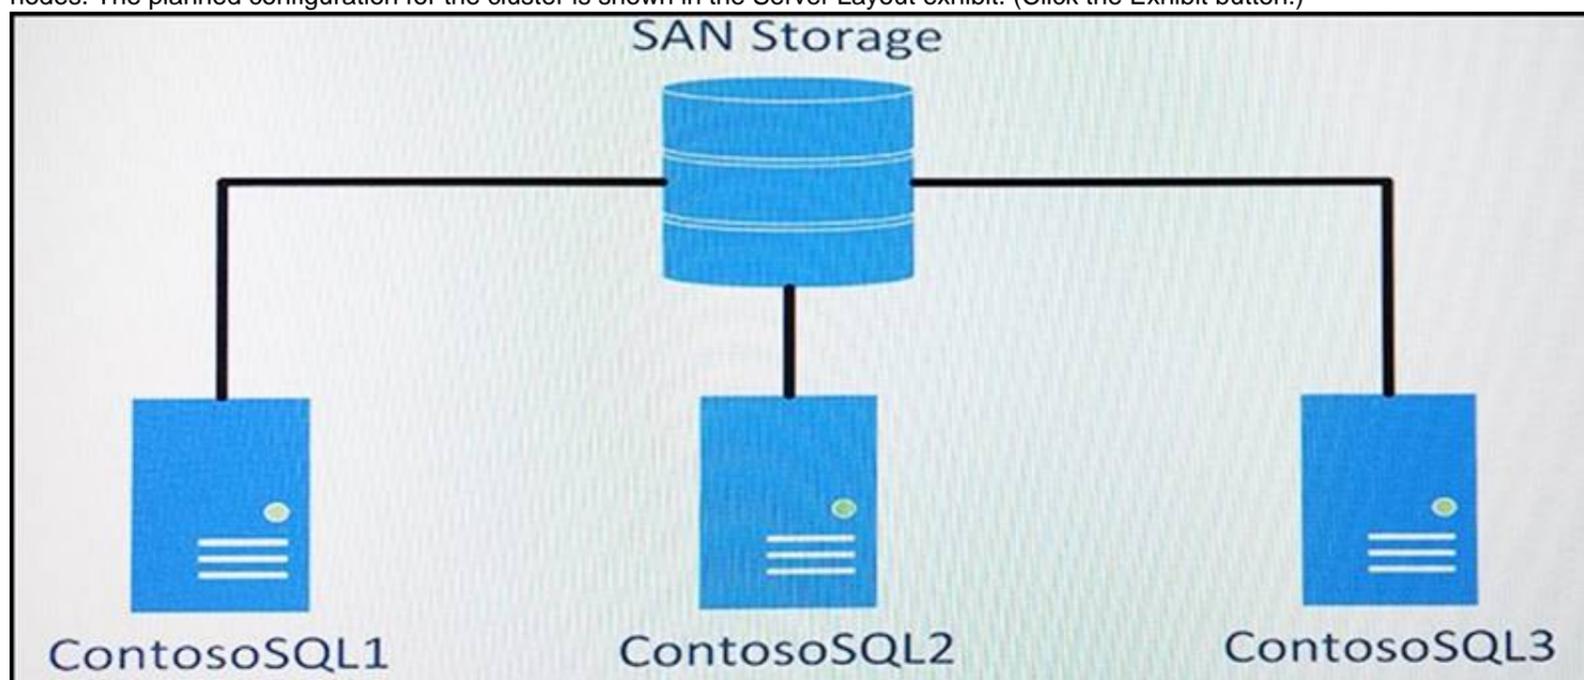
Encryptor: A certificate or asymmetric Key

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

**NEW QUESTION 10**

- (Exam Topic 1)

You are planning the deployment of two new Always On Failover Cluster Instances (FCIs) of Microsoft SQL Server to a single Windows Server Cluster with three nodes. The planned configuration for the cluster is shown in the Server Layout exhibit. (Click the Exhibit button.)



The SAN team has configured storage for the cluster and sent the configuration to you in the email shown in the SAN Team Email exhibit. (Click the Exhibit button.)

Conversation      SAN Storage for new SQL Cluster

Subject            SAN Storage for new SQL Cluster

DBA Team,

The following LUNs have been presented to the three servers of the new SQL Cluster:

Two 1 TB drives

Two 500 GB drives

Two 75 GB drives

Contoso SAN Admins

Each node of the cluster has identical local storage available as shown in the Local Storage exhibit. (Click the Exhibit button.)



All local storage is on SSD.

You need to plan specific configurations for the new cluster.

For each of the following statement, select Yes if the statement is true. Otherwise, select No.

## Answer Area

Statements	Yes	No
The Tempdb database for each cluster instance can be placed on the D: drive for the instance.	<input type="radio"/>	<input type="radio"/>
One virtual network name for each SQL Server instance must be configured in the cluster.	<input type="radio"/>	<input type="radio"/>
The shared storage has been formatted and configured on ContosoSQL1.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

**Explanation:**

Box 1: Yes

tempdb on local storage. FCIs now support placement of tempdb on local non-shared storage, such as a local solid-state-drive, potentially offloading a significant amount of I/O from a shared SAN.

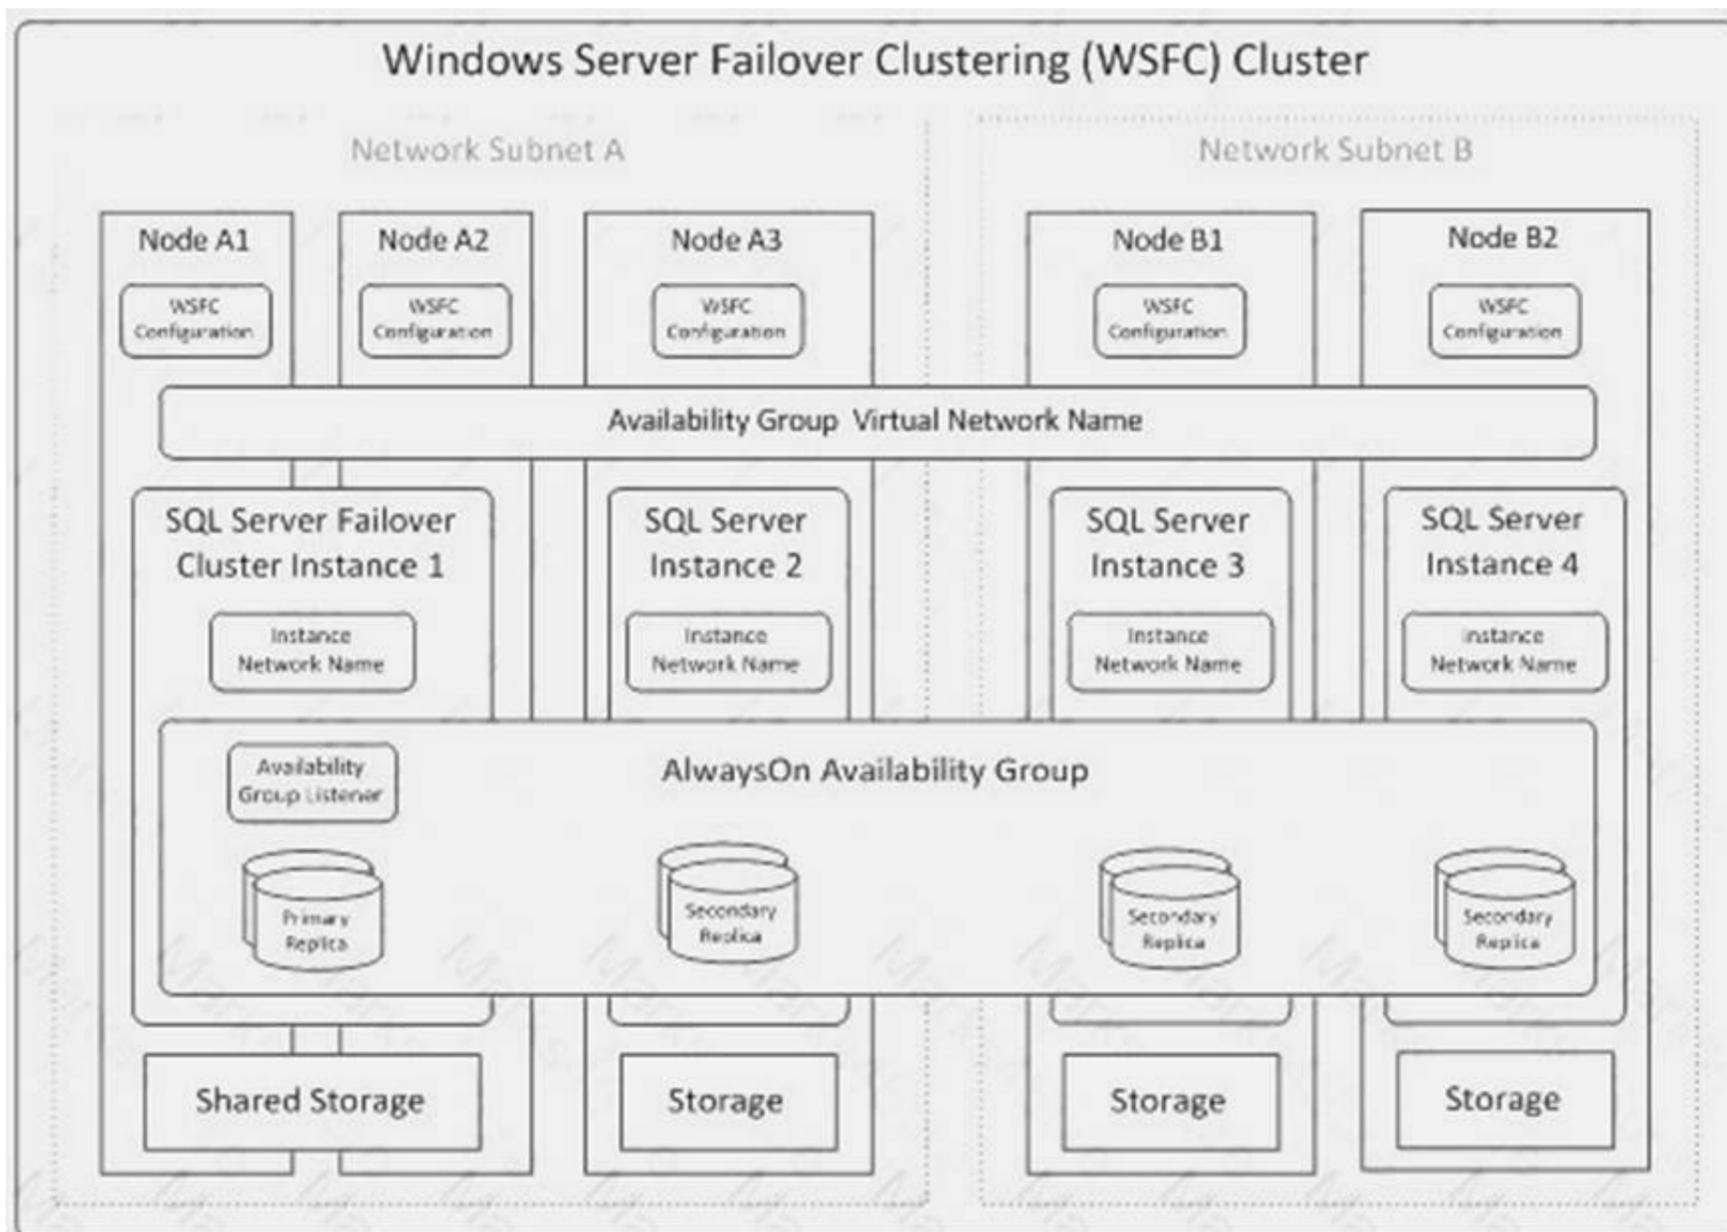
Prior to SQL Server 2012, FCIs required tempdb to be located on a symmetrical shared storage volume that failed over with other system databases.

Box 2: No

The VNN is set on the group level, not on the instance level.

Database client applications can connect directly to a SQL Server instance network name, or they may connect to a virtual network name (VNN) that is bound to an availability group listener. The VNN abstracts the WSFC cluster and availability group topology, logically redirecting connection requests to the appropriate SQL Server instance and database replica.

The logical topology of a representative AlwaysOn solution is illustrated in this diagram:



Box 3: No

You don't configure the SAN from a SQL Server, instead you can use a Microsoft Server server.

References:

<http://download.microsoft.com/download/d/2/0/d20e1c5f-72ea-4505-9f26-fef9550efd44/microsoft%20sql%20se>

**NEW QUESTION 10**

- (Exam Topic 1)

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

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

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

You need to configure auditing for the Adventure Works 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.

## Answer Area

```
USE master
```

```
GO
```

▼	AuditADUAccess
CREATE DATABASE AUDIT	
ALTER DATABASE AUDIT	
CREATE SERVER AUDIT	
ALTER SERVER AUDIT	

```
    TO FILE ( FILEPATH = 'C:\ADVAudit\' )
    WHERE object_name = 'SensitiveData'
```

```
GO
```

▼	AuditADUAccess WITH (STATE = ON)
CREATE DATABASE AUDIT	
ALTER DATABASE AUDIT	
CREATE SERVER AUDIT	
ALTER SERVER AUDIT	

```
GO
```

```
Use Adventureworks
```

▼	SPECIFICATION [FilterForSensitiveData]
CREATE DATABASE AUDIT	
ALTER DATABASE AUDIT	
CREATE SERVER AUDIT	
ALTER SERVER AUDIT	

▼	[AuditADUAccess]
FOR SERVER AUDIT	
FOR DATABASE AUDIT	
USE [AuditDataAcces]	
SELECT ID	

```
ADD (SELECT ON SCHEMA::[ADVSchema] BY [public])
WITH (STATE = ON)
GO
```

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: CREATE SERVER AUDIT

Create the server audit.

You must implement auditing to record access to data that is considered sensitive by the company. Create database 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

You must implement auditing for all objects in the ADVSchema.

References:

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

**NEW QUESTION 12**

- (Exam Topic 1)

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

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

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

You need to address the performance issues with each table.

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

## Answer Area

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

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Table A: Auto\_update statistics off

Table A does not change much. There is no need to update the statistics on this table. Table B: SET AUTO\_UPDATE\_STATISTICS\_ASYNC ON

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

SET AUTO\_UPDATE\_STATISTICS\_ASYNC ON

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

References:

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

**NEW QUESTION 13**

- (Exam Topic 1)

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

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

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

## Answer Area

Action	Server instance
Complete the backup job.	<div style="border: 1px solid black; padding: 5px;"> <div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> <span></span> <span>▼</span> </div> <div style="padding: 2px;">                     Primary server instance                      Secondary server instance                      Monitor server instance                      Backup share file server                 </div> </div>
Copy the backup job.	<div style="border: 1px solid black; padding: 5px;"> <div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> <span></span> <span>▼</span> </div> <div style="padding: 2px;">                     Primary server instance                      Secondary server instance                      Monitor server instance                      Backup share file server                 </div> </div>
Restore the backup.	<div style="border: 1px solid black; padding: 5px;"> <div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> <span></span> <span>▼</span> </div> <div style="padding: 2px;">                     Primary server instance                      Secondary server instance                      Monitor server instance                      Backup share file server                 </div> </div>

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Note: Before you configure log shipping, you must create a share to make the transaction log backups available to the secondary server. SQL Server Log shipping allows you to automatically send transaction log backups from a primary database on a primary server instance to one or more secondary databases on separate secondary server instances. The transaction log backups are applied to each of the secondary databases individually. An optional third server instance, known as the monitor server, records the history and status of backup and restore operations and, optionally, raises alerts if these operations fail to occur as scheduled.

Box 1: Primary server instance.

The primary server instance runs the backup job to back up the transaction log on the primary database. backup job: A SQL Server Agent job that performs the backup operation, logs history to the local server and the monitor server, and deletes old backup files and history information. When log shipping is enabled, the job category "Log Shipping Backup" is created on the primary server instance.

Box 2: Secondary server instance

Each of the three secondary server instances runs its own copy job to copy the primary log-backup file to its own local destination folder. copy job: A SQL Server Agent job that copies the backup files from the primary server to a configurable destination on the secondary server and logs history on the secondary server and the monitor server. When log shipping is enabled on a database, the job category "Log Shipping Copy" is created on each secondary server in a log shipping configuration.

Box 3: Secondary server instance.

Each secondary server instance runs its own restore job to restore the log backup from the local destination folder onto the local secondary database. restore job: A SQL Server Agent job that restores the copied backup files to the secondary databases. It logs history on the local server and the monitor server, and deletes old files and old history information. When log shipping is enabled on a database, the job category "Log Shipping Restore" is created on the secondary server instance.

References: <https://docs.microsoft.com/en-us/sql/database-engine/log-shipping/about-log-shipping-sql-server>

**NEW QUESTION 18**

- (Exam Topic 1)

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

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

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

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

**Answer:** B

**Explanation:**

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

Server.

References:

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

**NEW QUESTION 22**

- (Exam Topic 1)

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

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

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

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

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Step 1: Launch the Failover Cluster Manager and..

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

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

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

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

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

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

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

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

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

**NEW QUESTION 25**

- (Exam Topic 1)

You administer a Microsoft SQL Server 2016 server that hosts a transactional database and a reporting database. The transactional database is updated through a web application and is operational throughout the day.

The reporting database is only updated from the transactional database.

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

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

At 14:00 hours, you discover that pages 71, 520, and 713 on one of the database files are corrupted on the reporting database. You need to ensure that the databases are restored. You also need to ensure that data loss is minimal. What should you do?

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

**Answer:** H

**Explanation:**

At restore time, before you restore a differential backup, you must restore its base. Then, restore only the most recent differential backup to bring the database forward to the time when that differential backup was created. Typically, you would restore the most recent full backup followed by the most recent differential backup that is based on that full backup.

References: [https://technet.microsoft.com/en-us/library/ms345448\(v=sql.105\).aspx](https://technet.microsoft.com/en-us/library/ms345448(v=sql.105).aspx)

**NEW QUESTION 29**

- (Exam Topic 1)

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

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

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

You need to perform the following actions:

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

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

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

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

Does the solution meet the goal?

- A. Yes
- B. NO

**Answer: B**

#### NEW QUESTION 34

- (Exam Topic 1)

You administer a Microsoft SQL Server 2016 instance that contains a financial database hosted on a storage area network (SAN). The financial database has the following characteristics:

A data file of 2 terabytes is located on a dedicated LUN (drive D).

A transaction log of 10 GB is located on a dedicated LUN (drive E).

Drive D has 1 terabyte of free disk space.

Drive E has 5 GB of free disk space.

The database is continually modified by users during business hours from Monday through Friday between 09:00 hours and 17:00 hours.

Five percent of the existing data is modified each day.

The Finance department loads large CSV files into a number of tables each business day at 11:15 hours and 15:15 hours by using the BCP or BULK INSERT commands.

Each data load adds 3 GB of data to the database.

These data load operations must occur in the minimum amount of time. A full database backup is performed every Sunday at 10:00 hours.

Backup operations will be performed every two hours (11:00, 13:00, 15:00, and 17:00) during business hours. You implement log shipping of the financial database to another SQL Server 2016 instance. You decide to

failover to this secondary database.

You need to ensure that all transactions will be replicated to the secondary database. Which backup option should you use?

- A. Differential
- B. Transaction Log
- C. FULL
- D. SIMPLE
- E. SKIP
- F. RESTART
- G. STANDBY
- H. CHECKSUM
- I. DBO\_ONLY
- J. COPY\_ONLY
- K. NORECOVERY
- L. NO\_CHECKSUM
- M. CONTINUE\_AFTER\_ERROR
- N. BULK\_LOGGED

**Answer: K**

#### Explanation:

Roll back is controlled by the RESTORE statement through the [ RECOVERY | NORECOVERY ] options: NORECOVERY specifies that roll back not occur. This allows roll forward to continue with the next statement in the sequence.

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

#### NEW QUESTION 37

- (Exam Topic 1)

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

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

What should you use?

- A. Activity Monitor
- B. Sp\_who3
- C. SQL Server Management Studio (SSMS) Object Explorer
- D. SQL Server Data Collector
- E. SQL Server Data Tools (SSDT)
- F. SQL Server Configuration Manager

**Answer: A**

#### Explanation:

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

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

and Recent Expensive Queries.

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

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

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

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

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

**NEW QUESTION 41**

- (Exam Topic 2)

You are migrating a database named Orders to a new server that runs Microsoft SQL Server 2016. You attempt to add the [Corpnet\User1] login to the database.

However, you receive the following error message: "User already exists in current database."

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

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

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

**Answer: C**

**NEW QUESTION 42**

- (Exam Topic 2)

You plan to deploy SQL Server 2014.

You identify the following security requirements for the deployment:

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

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

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

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

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

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

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

**NEW QUESTION 47**

- (Exam Topic 2)

Overview

Application Overview

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

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

The new version will use SQL Server 2014.

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

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

Tables

The current database schema contains a table named OrderDetails.

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

The product price is stored in a table named Products. The Products table was defined by using the SQL\_Latin1\_General\_CP1\_CI\_AS collation.

A column named ProductName was created by using the varchar data type. The database contains a table named Orders. Orders contains all of the purchase orders from the last 12 months. Purchase orders that are older than 12 months are stored in a table named OrdersOld. The previous version of the ERP application relied on table-level security. Stored Procedures  
The current version of the database contains stored procedures that change two tables. The following shows the relevant portions of the two stored procedures:

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

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

**Customer Problems Installation Issues**

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

**Index Fragmentation Issues**

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

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

**Backup Issues**

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

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

**Missing Data Issues**

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

**Query Performance Issues**

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

**Design Requirements**

**File Storage Requirements**

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

**Data Recovery Requirements**

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

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

**Concurrency Requirements**

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

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

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

**Answer: B**

**Explanation:**

- Concurrency Requirements

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

**NEW QUESTION 48**

- (Exam Topic 2)

Overview

General Overview

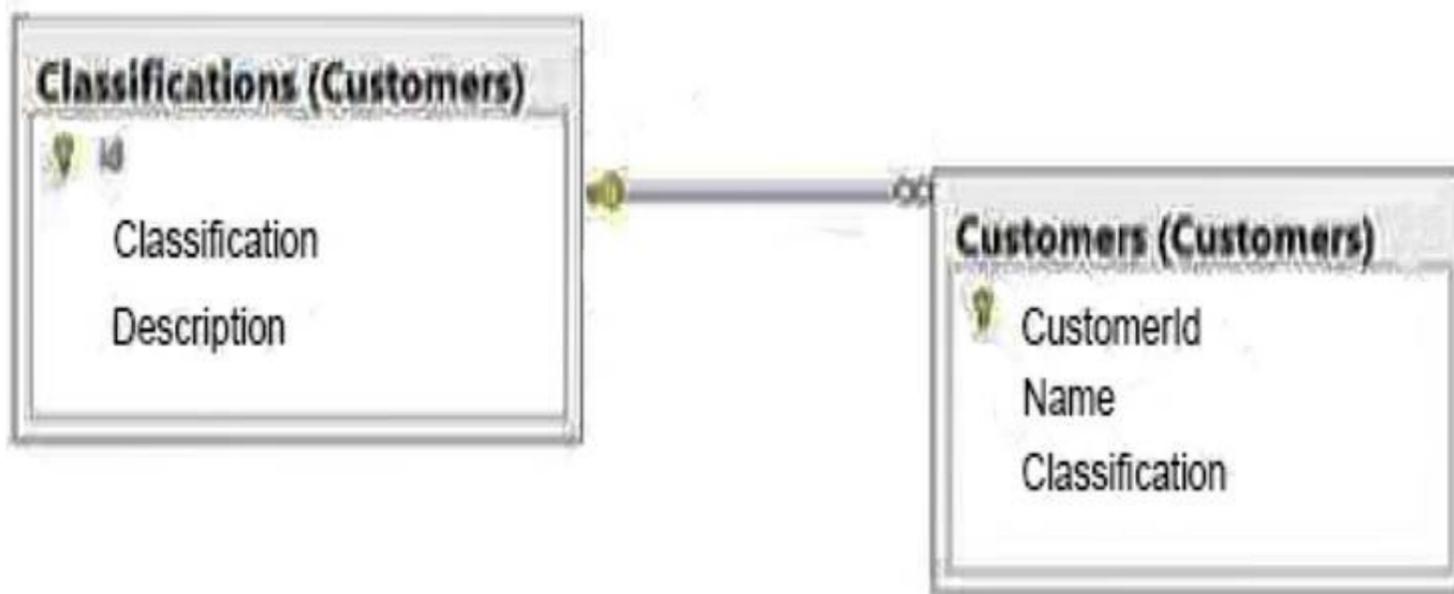
ADatum Corporation has offices in Miami and Montreal.

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

Databases

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

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



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

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

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

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

A stored procedure named USP\_1 generates millions of rows of data for multiple reports. USP\_1 combines data from five different tables from the Sales and Customers databases in a table named Table1.

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

A stored procedure named USP\_2 is used to generate a product list. The product list contains the names of products grouped by category.

USP\_2 takes several minutes to run due to locks on the tables the procedure accesses. The locks are caused by USP\_1 and USP\_3.

A stored procedure named USP\_3 is used to update prices. USP\_3 is composed of several UPDATE statements called in sequence from within a transaction.

Currently, if one of the UPDATE statements fails, the stored procedure fails. A stored procedure named USP\_4 calls stored procedures in the Sales, Customers, and Inventory databases.

The nested stored procedures read tables from the Sales, Customers, and Inventory databases. USP\_4 uses an EXECUTE AS clause.

All nested stored procedures handle errors by using structured exception handling. A stored procedure named USP\_5 calls several stored procedures in the same database. Security checks are performed each time USP\_5 calls a stored procedure.

You suspect that the security checks are slowing down the performance of USP\_5. All stored procedures accessed by user applications call nested stored procedures.

The nested stored procedures are never called directly. Design Requirements

Data Recovery

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

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

Classification Changes

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

Storage

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

Error Handling

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

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

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

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

**Answer:** CD

**Explanation:**

Scenario:

You plan to change the way customers are classified.

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

**NEW QUESTION 52**

- (Exam Topic 2)

Overview

General Overview

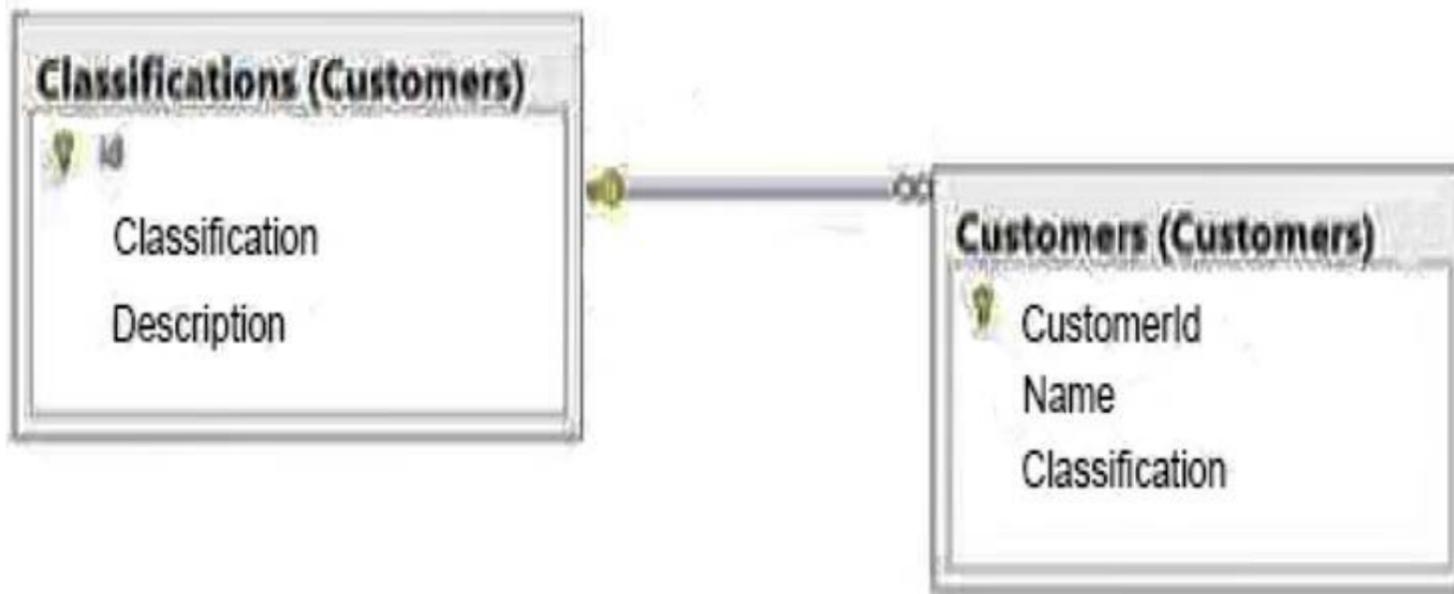
ADatum Corporation has offices in Miami and Montreal.

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

Databases

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

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



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

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

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

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

A stored procedure named USP\_1 generates millions of rows of data for multiple reports. USP\_1 combines data from five different tables from the Sales and Customers databases in a table named Table1.

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

A stored procedure named USP\_2 is used to generate a product list. The product list contains the names of products grouped by category.

USP\_2 takes several minutes to run due to locks on the tables the procedure accesses. The locks are caused by USP\_1 and USP\_3.

A stored procedure named USP\_3 is used to update prices. USP\_3 is composed of several UPDATE statements called in sequence from within a transaction.

Currently, if one of the UPDATE statements fails, the stored procedure fails. A stored procedure named USP\_4 calls stored procedures in the Sales, Customers, and Inventory databases.

The nested stored procedures read tables from the Sales, Customers, and Inventory databases. USP\_4 uses an EXECUTE AS clause.

All nested stored procedures handle errors by using structured exception handling. A stored procedure named USP\_5 calls several stored procedures in the same database. Security checks are performed each time USP\_5 calls a stored procedure.

You suspect that the security checks are slowing down the performance of USP\_5. All stored procedures accessed by user applications call nested stored procedures.

The nested stored procedures are never called directly. Design Requirements

Data Recovery

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

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

Classification Changes

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

Storage

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

Error Handling

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

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

You need to recommend a change to USP\_3 to ensure that the procedure completes only if all of the UPDATE statements complete. Which change should you recommend?

- A. Set the XACT\_ABORT option to off
- B. Set the XACT\_ABORT option to on.
- C. Set the IMPLICIT\_TRANSACTIONS option to off.
- D. Set the IMPLICIT\_TRANSACTIONS option to on.

**Answer: B**

**Explanation:**

Scenario:

A stored procedure named USP\_3 is used to update prices. USP\_3 is composed of several UPDATE statements called in sequence from within a transaction. Currently, if one of the UPDATE statements fails, the stored procedure continues to execute.

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

**NEW QUESTION 53**

- (Exam Topic 2)

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

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

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

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

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

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

**Answer:** BC

**Explanation:**

References:

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

**NEW QUESTION 55**

- (Exam Topic 2)

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

The name of all stored procedures must start with "usp\_" always.

All distribution statistics must be updated daily

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

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

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

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

- Policy-Based Management Each Stored Procedure that are created and that will be created has to have prefix "USP\_".

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

**NEW QUESTION 57**

- (Exam Topic 2)

You administer a Microsoft SQL Server 2016 server. You plan to deploy new features to an application.

You need to evaluate existing and potential clustered and non-clustered indexes that will improve performance.

What should you do?

- A. Query the sys.dm\_db\_index\_usage\_stats DMV.
- B. Query the sys.dm\_db\_missing\_index\_details DMV.
- C. Use the Database Engine Tuning Advisor.
- D. Query the sys.dm\_db\_missing\_index\_columns DMV.

**Answer:** C

**Explanation:**

The Microsoft Database Engine Tuning Advisor (DTA) analyzes databases and makes recommendations that you can use to optimize query performance. You can use the Database Engine Tuning Advisor to select and create an optimal set of indexes, indexed views, or table partitions without having an expert understanding of the database structure or the internals of SQL Server. Using the DTA, you can perform the following tasks.

Troubleshoot the performance of a specific problem query  
Tune a large set of queries across one or more databases  
Perform an exploratory what-if analysis of potential physical design changes  
Manage storage space

References:  
<https://docs.microsoft.com/en-us/sql/relational-databases/performance/database-engine-tuning-advisor>

**NEW QUESTION 58**

- (Exam Topic 2)

Overview

General Overview

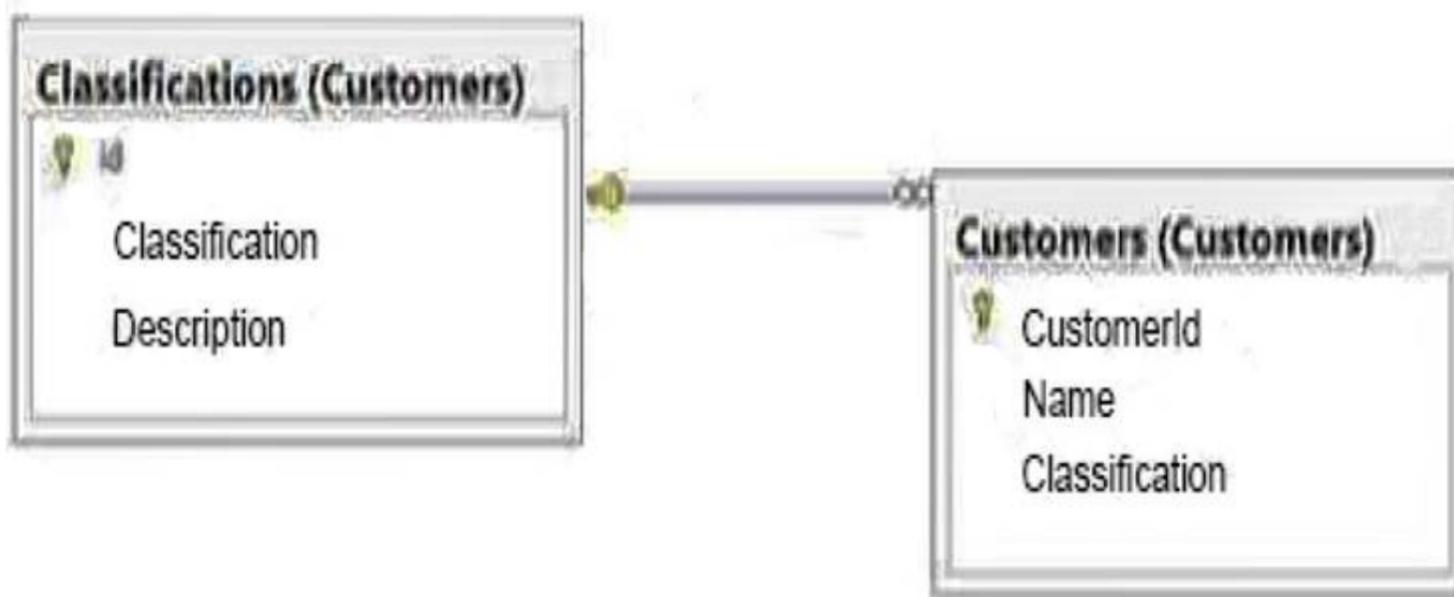
A Datum Corporation has offices in Miami and Montreal.

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

Databases

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

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



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

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

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

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

A stored procedure named USP\_1 generates millions of rows of data for multiple reports. USP\_1 combines data from five different tables from the Sales and Customers databases in a table named Table1.

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

A stored procedure named USP\_2 is used to generate a product list. The product list contains the names of products grouped by category.

USP\_2 takes several minutes to run due to locks on the tables the procedure accesses. The locks are caused by USP\_1 and USP\_3.

A stored procedure named USP\_3 is used to update prices. USP\_3 is composed of several UPDATE statements called in sequence from within a transaction.

Currently, if one of the UPDATE statements fails, the stored procedure fails. A stored procedure named USP\_4 calls stored procedures in the Sales, Customers, and Inventory databases.

The nested stored procedures read tables from the Sales, Customers, and Inventory databases. USP\_4 uses an EXECUTE AS clause.

All nested stored procedures handle errors by using structured exception handling. A stored procedure named USP\_5 calls several stored procedures in the same database. Security checks are performed each time USP\_5 calls a stored procedure.

You suspect that the security checks are slowing down the performance of USP\_5. All stored procedures accessed by user applications call nested stored procedures.

The nested stored procedures are never called directly. Design Requirements

Data Recovery

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

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

Classification Changes

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

Storage

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

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

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

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

- A. A server role
- B. A database role

- C. A credential
- D. A shared login

**Answer:** C

**Explanation:**

- Scenario: A group of junior database administrators must be able to view the server state of the SQL Server instance that hosts the Sales database. The junior database administrators will not have any other administrative rights.
- Credentials provide a way to allow SQL Server Authentication users to have an identity outside of SQL Server. Credentials can also be used when a SQL Server Authentication user needs access to a domain resource, such as a file location to store a backup.

**NEW QUESTION 60**

- (Exam Topic 2)

**Overview**

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

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

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

You have an existing desktop application that uses a SQL Server 2008 database named App1\_DB. App1\_DB will remain in production.

**Requirements Planned Changes**

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

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

Database1 will also contain a stored procedure named usp\_UpdateOrderDetails.

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

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

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

Inventory will contain over 100 GB of data.

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

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

**Business Requirements**

You have the following requirements:

Costs for new licenses must be minimized.

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

Development effort must be minimized whenever possible.

The storage requirements for databases must be minimized.

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

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

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

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

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

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

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

**Answer:** A

**Explanation:**

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

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

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

**NEW QUESTION 61**

- (Exam Topic 2)

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

If any of the new rows violates a table constraint, then no further additions must be attempted and all changes made by the stored procedure must be discarded.

If any errors occur, a row must be added to an audit table, and the original error must be returned to the caller of the stored procedure.

What should you include in the design?

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

**Answer:** D

**Explanation:**

References:

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

**NEW QUESTION 66**

- (Exam Topic 2)

You administer a Microsoft SQL Server 2016 instance.

You need to stop a blocking process that has an SPID of 64 without stopping other processes. What should you do?

- A. Execute the following Transact-SQL statement: EXECUTE sp\_KillSPID 64
- B. Restart the SQL Server service.
- C. Execute the following Transact-SQL statement: KILL 64
- D. Execute the following Transact-SQL statement: ALTER SESSION KILL '64'

**Answer: C**

**Explanation:**

KILL can be used to terminate a normal connection, which internally terminates the transactions that are associated with the specified session ID.

References:

<http://msdn.microsoft.com/en-us/library/ms173730.aspx>

**NEW QUESTION 69**

- (Exam Topic 2)

Overview

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

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

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

You have an existing desktop application that uses a SQL Server 2008 database named App1\_DB. App1\_DB will remain in production.

Requirements Planned Changes

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

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

Database1 will also contain a stored procedure named usp\_UpdateOrderDetails.

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

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

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

Inventory will contain over 100 GB of data.

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

Database2 will contain a stored procedure named usp\_UpdateInventory. usp\_UpdateInventory will manipulate a table that contains a self-join that has an unlimited number of hierarchies. All data in Database2 is recreated each day and does not change until the next data creation process. Data from Database2 will be accessed periodically by an external application named Application1. The data from Database2 will be sent to a database named App1\_Db1 as soon as changes occur to the data in Database2. Litware plans to use offsite storage for all SQL Server 2014 backups.

Business Requirements

You have the following requirements:

Costs for new licenses must be minimized.

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

Development effort must be minimized whenever possible.

The storage requirements for databases must be minimized.

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

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

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

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

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

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

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

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

**Answer: A**

**Explanation:**

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

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

Transparent Data Encryption (TDE)

**NEW QUESTION 72**

- (Exam Topic 2)

You have two SQL Server 2012 instances named SQLDev and SQLProd.

You plan to create a new database on SQLProd that will use SQL Server Authentication.

You need to ensure that when the new database is copied from SQLProd to SQLDev, users can connect to the database on SQLDev if they do not have a login on the SQLDev instance.

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

- A. SQL Server Integration Services (SSIS) scripts
- B. Extended Events
- C. Triggers
- D. SQL Server Analysis Services (SSAS) scripts
- E. Contained database

Answer: E

**Explanation:**

A fully contained database includes all the settings and metadata required to define the database and has no configuration dependencies on the instance of the SQL Server Database Engine where the database is installed.

**NEW QUESTION 75**

- (Exam Topic 2)

Overview

Application Overview

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

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

The new version will use SQL Server 2014.

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

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

Tables

The current database schema contains a table named OrderDetails.

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

The product price is stored in a table named Products. The Products table was defined by using the SQL\_Latin1\_General\_CP1\_CI\_AS collation.

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

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

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

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

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

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

Customer Problems Installation Issues

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

Index Fragmentation Issues

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

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

Backup Issues

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

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

Missing Data Issues

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

Query Performance Issues

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

Design Requirements

File Storage Requirements

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

#### Data Recovery Requirements

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

#### Security Requirements

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

#### Concurrency Requirements

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

What should you include in the recommendation?

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

**Answer: B**

#### Explanation:

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

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

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

#### NEW QUESTION 80

- (Exam Topic 2)

You administer a Microsoft SQL Server 2016 database.

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

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

**Answer: D**

#### NEW QUESTION 81

- (Exam Topic 2)

Overview

Application Overview

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

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

The new version will use SQL Server 2014.

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

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

Tables

The current database schema contains a table named OrderDetails.

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

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

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

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

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

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

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

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

**Customer Problems Installation Issues**

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

**Index Fragmentation Issues**

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

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

**Backup Issues**

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

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

**Missing Data Issues**

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

**Query Performance Issues**

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

**Design Requirements**

**File Storage Requirements**

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

**Data Recovery Requirements**

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

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

**Concurrency Requirements**

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

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

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

**Answer: C**

**Explanation:**

- Security Requirements

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

**NEW QUESTION 85**

- (Exam Topic 2)

Overview

You are a database administrator for a company named Litware, Inc. Litware is a book publishing house. Litware has a main office and a branch office. You are designing the database infrastructure to support a new web-based application that is being developed. The web application will be accessed at [www.litwareinc.com](http://www.litwareinc.com). Both internal employees and external partners will use the application. You have an existing desktop application that uses a SQL Server 2008 database named App1\_DB. App1\_DB will remain in production.

**Requirements Planned Changes**

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

**Business Requirements**

You have the following requirements:

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

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

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

**Answer:** D

**Explanation:**

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

**NEW QUESTION 90**

- (Exam Topic 2)

You have a server named SQL1 that has SQL Server 2012 installed. SQL1 hosts a database named Database1. Database1 contains a table named Table1. Table1 is partitioned across five filegroups based on the Date field. The schema of Table1 is configured as shown in the following table.

Column	Data type
ID	Bigint
Account	Bigint
Amount	Decimal
TransactionType	Int
TransactionDate	Date

Table1 contains the indexes shown in the following table.

Index	Type	Column
PK_Table1	Clustered, primary key	ID, TransactionType
IX_Account	Nonclustered	Account
IX_Type	Nonclustered	TransactionType
IX_Date	Nonclustered	TransactionDate
IX_Amount	Nonclustered	Amount

You need to recommend an index strategy to maximize performance for the queries that consume the indexes available to Table1.

Which type of index storage should you recommend? To answer, drag the appropriate index storage type to the correct index in the answer area.

Index Storage Types	Answer area
Aligned	IX_Type Index Storage Type
Nonaligned	IX_Account Index Storage Type
	IX_Date Index Storage Type
	IX_Amount Index Storage Type

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Index Storage Type

Designing a partitioned index independently (unaligned) of the base table can be useful in the following cases:

- The base table has not been partitioned.
- The index key is unique and it does not contain the partitioning column of the table.
- You want the base table to participate in collocated joins with more tables using different join columns.

**NEW QUESTION 91**

- (Exam Topic 2)

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:** BCDE

**Explanation:**

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

**NEW QUESTION 96**

- (Exam Topic 2)

You are planning on deploying a server that will be dedicated for ETL (Extraction, Transformation, and Loading) processes.

You want to ensure that SSIS (SQL Server Integration Services) packages will run on this dedicated ETL server and not on any other server on which they were started.

Which of the following features must you install on the ETL server in addition to SSIS to accomplish this goal?

- A. Database Engine
- B. SQL Server Reporting Services
- C. SQL Server Analysis Services
- D. Client Tools SDK

**Answer:** A

**NEW QUESTION 97**

- (Exam Topic 2)

Overview

Application Overview

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

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

The new version will use SQL Server 2014.

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

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

Tables

The current database schema contains a table named OrderDetails.

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

The product price is stored in a table named Products. The Products table was defined by using the SQL\_Latin1\_General\_CP1\_CI\_AS collation.

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

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

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

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

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

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

Customer Problems Installation Issues

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

Index Fragmentation Issues

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

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

Backup Issues

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

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

Missing Data Issues

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

Query Performance Issues

Customers report that query performance degrades very quickly. Additionally, the customers report that users cannot run queries when SQL Server runs maintenance tasks. Import Issues

During the monthly import process, database administrators receive many supports call from users who report that they cannot access the supplier data. The database administrators want to reduce the amount of time required to import the data.

Design Requirements

File Storage Requirements

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

Data Recovery Requirements

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

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

Concurrency Requirements

You must reduce the likelihood of deadlocks occurring when Sales.Prod and Sales.Proc2 execute. You need to recommend which statement should be used to update SalesOrder.

How should you recommend completing the statement? To answer, drag the appropriate elements to the correct locations. Each element 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.

The screenshot shows a web-based exam interface. On the left, under the heading "Elements", there is a vertical list of eight transaction isolation levels, each in a rectangular box: EXPLICIT, ISOLATION, READ UNCOMMITTED, ROLLBACK, SERIALIZABLE, SNAPSHOT, TABLOCK, and TRANSACTION. To the right, under the heading "Answer Area", there is a code editor. The code in the editor is: SET [ ] [ ]  
LEVEL [ ]  
BEGIN [ ]  
UPDATE SalesOrder  
...  
COMMIT TRANSACTION;

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Elements	Answer Area
EXPLICIT	SET TRANSACTION ISOLATION
ISOLATION	LEVEL SNAPSHOT
READ UNCOMMITTED	BEGIN TRANSACTION
ROLLBACK	UPDATE SalesOrder
SERIALIZABLE	...
SNAPSHOT	COMMIT TRANSACTION;
TABLOCK	
TRANSACTION	

**NEW QUESTION 99**

- (Exam Topic 2)

Overview

Application Overview

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

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

The new version will use SQL Server 2014.

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

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

Tables

The current database schema contains a table named OrderDetails.

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

The product price is stored in a table named Products. The Products table was defined by using the SQL\_Latin1\_General\_CP1\_CI\_AS collation.

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

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

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

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

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

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

**Customer Problems Installation Issues**

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

**Index Fragmentation Issues**

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

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

**Backup Issues**

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

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

**Missing Data Issues**

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

**Query Performance Issues**

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

**Design Requirements**

**File Storage Requirements**

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

**Data Recovery Requirements**

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

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

**Concurrency Requirements**

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

You need to recommend a solution that reduces the time it takes to import the supplier data. What should you include in the recommendation?

- A. Enable instant file initialization.
- B. Reorganize the indexes.
- C. Disable Resource Governor.
- D. Enable Auto Update Statistics.

**Answer: C**

**Explanation:**

- 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.

**NEW QUESTION 102**

- (Exam Topic 2)

You plan to deploy SQL Server 2014. Your company identifies the following monitoring requirements:

Tempdb must be monitored for insufficient free space.

Deadlocks must be analyzed by using Deadlock graphs.

You need to identify which feature meets each monitoring requirement.

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

Features	Answer area
Dynamic management view	Tempdb must be monitored for insufficient free space.
Activity Monitor	Deadlocks must be analyzed by using Deadlock graphs.
Resource Governor	
SQL Trace	

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

You can use the sys.dm\_db\_file\_space\_usage dynamic management view to monitor the disk space used by the user objects, internal objects, and version stores in the tempdb files. Additionally, to monitor the page allocation or deallocation activity in tempdb at the session or task level, you can use the sys.dm\_db\_session\_space\_usage and sys.dm\_db\_task\_space\_usage dynamic management views. These views can be used to identify large queries, temporary tables, or table variables that are using a large amount of tempdb disk space.

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.

SQL Server Profiler and SQL Server Management Studio use a deadlock wait-for graph to describe a deadlock. The deadlock wait-for graph contains process nodes, resource nodes, and edges representing the relationships between the processes and the resources. References: Troubleshooting Insufficient Disk Space in tempdb

References: Analyze Deadlocks with SQL Server Profiler

**NEW QUESTION 104**

- (Exam Topic 2)

You have four databases that are accessed by using an Online Transaction Processing (OLTP) application. The databases are stored on a server named SQL1 that has SQL Server 2014 installed. You plan to deploy an additional server that has SQL Server 2014 installed. You need to design a high-availability solution for the databases that meets the following requirements:

If SQL1 fails, the databases must be available.

Users must be able to run reports against a secondary copy of the databases.

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

- A. AlwaysOn availability groups
- B. Database mirroring
- C. Log shipping
- D. Failover Clustering

**Answer:** A

**Explanation:**

The AlwaysOn Availability Groups feature is a high-availability and disaster-recovery solution that provides an enterprise-level alternative to database mirroring. Introduced in SQL Server 2012, AlwaysOn Availability Groups maximizes the availability of a set of user databases for an enterprise. An availability group supports a failover environment for a discrete set of user databases, known as availability databases, that fail over together.

**NEW QUESTION 109**

- (Exam Topic 2)

You administer a Microsoft SQL Server 2016 Enterprise Edition server that uses 64 cores.

You discover performance issues when large amounts of data are written to tables under heavy system load. You need to limit the number of cores that handle I/O.

What should you configure?

- A. Processor affinity
- B. Lightweight pooling
- C. Max worker threads
- D. I/O affinity

**Answer:** D

**Explanation:**

References:

<http://msdn.microsoft.com/en-us/library/ms189629.aspx>

**NEW QUESTION 111**

- (Exam Topic 2)

Your company has offices in Seattle and Montreal.

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

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

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

**Answer: D**

**NEW QUESTION 114**

- (Exam Topic 2)

You administer a Microsoft SQL Server 2016 database that has multiple tables in the Sales schema. Some users must be prevented from deleting records in any of the tables in the Sales schema. You need to manage users who are prevented from deleting records in the Sales schema. You need to achieve this goal by using the minimum amount of administrative effort. What should you do?

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

**Answer: A**

**NEW QUESTION 119**

- (Exam Topic 3)

You need to recommend a backup process for data warehouse database. The solution must meet the following requirements:  
Ensure that if a hardware failure occurs, you can bring the database online without losing more than 24 hours of transactions.  
Minimize the amount of administrative effort required to restore any lost data.  
Minimize the space used by transaction logs.  
What should you include in the recommendation? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions	Work Area
Perform a weekly full backup of the database	
Ship the logs to a secondary server	
Set the database to the simple recovery model	
Set the database to the full recovery model	
Create a database snapshot	
Perform a backup of the transaction log every hour	
Perform a differential backup of the database every night	

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

Set the database to the simple recovery model; Perform a weekly full backup of the database; Perform a differential backup of the database every night Note:  
- Simple recovery model No log backups.  
- Full recovery model Requires log backups

**NEW QUESTION 120**

- (Exam Topic 3)

Background Corporate Information

Fabrikam, Inc. is a retailer that sells electronics products on the Internet. The company has a headquarters site and one satellite sales office. You have been hired as the database administrator, and the company wants you to change the architecture of the Fabrikam ecommerce site to optimize performance and reduce

downtime while keeping capital expenditures to a minimum. To help with the solution, Fabrikam has decided to use cloud resources as well as on-premise servers.  
Physical Locations

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

**Problem Statement**

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

**Configuration Windows Logins**

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

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

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

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

**Database**

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

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

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

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

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

Column	Data type
SalesOrderDetailID	INT
ProductID	INT
SalePrice	SMALLMONEY
SaleQuantity	INT

**Database Issues**

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

**Requirements Database**

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

**Customer data access**

Customers access the company's website to order products, so they must be able to read product information such as name, description, and price from the Product table. When customers place orders, stored procedures called by the website update product quantity-on-hand values. This means the product table is constantly updated at random times.

**Customer support data access**

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

**Sales force data access**

Sales people at both the headquarters office and the satellite office must generate reports that read from the Product and SalesOrderDetail tables. No updates or inserts are ever made by sales people. These reports are run at random times and there can be no reporting downtime to refresh the data set except during the monthly

maintenance window. The reports that run from the satellite office are process-intensive queries with large data sets. Regardless of which office runs a sales force report, the SalesOrderDetail table should only return valid, committed order data; any orders not yet committed should be ignored.

**Historical Data**

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

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

**Backups**

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

**Database Maintenance**

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

**Project milestones completed**

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

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

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

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

You need to implement a backup strategy to support the requirements.

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

- A. Create a credential called MyCredential on SQL Server by using a Windows domain account and password.
- B. Schedule a full backup by using the command `BACKUP DATABASE ProdDB TO DISK...`
- C. Create a share on your Windows Azure site by using your Windows Azure storage account information, and grant permission to the SQL Server service login.
- D. Schedule a full backup by using the command `BACKUP DATABASE ProdDB TO URL ... WITH CREDENTIAL=N'MyCredential'`
- E. Create a share on the hot standby site and grant permission to the SQL Server service login.
- F. Create a credential called MyCredential on SQL Server, using MyStorageAccount for the storage account name and StorageAccountKey for the access key.
- G. Schedule a full backup by using the command `BACKUP DATABASE ProdDB TO SHARE ... WITH CREDENTIAL=N'MyCredential'`

**Answer:** CD

**Explanation:**

- Scenario: 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.

- Need to store files in the cloud.

- Manage your backups to Windows Azure: Using the same methods used to backup to DISK and TAPE, you can now back up to Windows Azure storage by specifying URL as the backup destination.

You can use this feature to manually backup or configure your own backup strategy like you would for a local storage or other off-site options.

This feature is also referred to as SQL Server Backup to URL. SQL Server Managed Backup to Windows Azure

**NEW QUESTION 122**

- (Exam Topic 3)

**Background**

**Corporate Information**

Fabrikam, Inc. is a retailer that sells electronics products on the Internet. The company has a headquarters site and one satellite sales office. You have been hired as the database administrator, and the company wants you to change the architecture of the Fabrikam e-commerce 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 memory-optimized table. The auto-update statistics option is set off on this database. Only product managers are allowed to add products or to make changes to the name, description, price, cost, and supplier. The changes are made in an internal database and pushed to the Product table in ProdDB during system maintenance time. Product managers and others working at the headquarters location also should be able to generate reports that include supplier and cost information.

#### Customer data access

Customers access the company's website to order products, so they must be able to read product information such as name, description, and price from the Product table. When customers place orders, stored procedures called by the website update product quantity-on-hand values. This means the product table is constantly updated at random times.

#### Customer support data access

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

#### Sales force data access

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

#### Historical Data

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

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

#### Backups

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

#### Database Maintenance

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

#### Project milestones completed

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

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

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

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

You need to 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, 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:** AE

#### Explanation:

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

#### NEW QUESTION 124

- (Exam Topic 3)

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

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

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

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

How can you achieve this goal?

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

**Answer:** C

#### NEW QUESTION 126

- (Exam Topic 3)

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 have a database that includes a table named Candidate.

You need to update the statistics for a column named Skills in the table and turn off automatic statistics updates for the column.

Solution: You run the following query:

```
USE CustomerDatabase
GO
UPDATE STATISTICS Person.Candidate(Skills)
WITH FULLSCAN, NORECOMPUTE
GO
```

Does this meet the goal?

- A. Yes
- B. No

**Answer:** A

**Explanation:**

The following example updates the Products statistics in the Product table, forces a full scan of all rows in the Product table, and turns off automatic statistics for the Products statistics.

```
USE AdventureWorks2012;
```

```
GO
```

```
UPDATE STATISTICS Production.Product(Products) WITH FULLSCAN, NORECOMPUTE;
```

Note: NORECOMPUTE

Disable the automatic statistics update option, AUTO\_UPDATE\_STATISTICS, for the specified statistics. If this option is specified, the query optimizer completes this statistics update and disables future updates.

To re-enable the AUTO\_UPDATE\_STATISTICS option behavior, run UPDATE STATISTICS again without the NORECOMPUTE option or run sp\_autostats.

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

**NEW QUESTION 128**

- (Exam Topic 3)

You deploy a SQL Server instance named SQLProd that uses SQL Server 2014.

You need to recommend a solution to monitor the transactions that are running currently against SQLProd. The solution must minimize the amount of custom code required.

What should you recommend?

- A. Statistics
- B. A dynamic management view
- C. A trigger
- D. User-defined views

**Answer:** B

**Explanation:**

Dynamic management views and functions return server state information that can be used to monitor the health of a server instance, diagnose problems, and tune performance.

Transactions can be monitored.

**NEW QUESTION 132**

- (Exam Topic 3)

General Overview

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

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

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

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

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

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

The Current\_Inventory database contains a large table named Inv\_Current. The Inv\_Current table has a clustered index for the primary key and a nonclustered index. The primary key column uses the identity property.

The data in the Inv\_Current table is over 120GB in size. The tables in the Current\_Inventory database are accessed by multiple queries in the Sales database.

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

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

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

The GenInfo table is used for reports.

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

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

Current System

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

SQL Servers

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

Design Requirements

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

Confidential information in the Current\_Inventory database that is accessed by ExternalApp1 must be securely stored.

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

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

Deadlocks must be analyzed with the use of Deadlock Graphs.

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

Software licensing and database storage costs must be minimized.

Development effort must be minimized.

The Tempdb databases must be monitored for insufficient free space.

Failed authentication requests must be logged.

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

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

SPUpdateSalesInfo queries data in the OrderTotals table the second time.  
The performance of the SPUpdate2 stored procedure needs to be improved. Your solution must meet the design requirements.  
What should your solution include?

- A. A common table expression.
- B. A derived table.
- C. A Cursor.
- D. A table variable.

**Answer:** A

#### NEW QUESTION 136

- (Exam Topic 3)

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 sections, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You have a server named Server1 that has Microsoft SQL Server installed. Server1 has SQL Server Audit configured to send audit even records to a file. You need to ensure that a database user named User1 can review the audit data. Solution: You grant the VIEW ANY DEFINITION permission to User1.

Does this meet the goal?

- A. Yes
- B. No

**Answer:** B

#### Explanation:

Each feature and command for SQL Server Audit has individual permission requirements.

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

The VIEW SERVER STATE permission.

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

Membership in the sysadmin fixed server role.

The CONTROL SERVER permission.

The ALTER ANY AUDIT permission.

A principal must have the VIEW SERVER STATE or ALTER ANY AUDIT permission to use the Dynamic Management Views.

References: [https://technet.microsoft.com/en-us/library/cc280665\(v=sql.105\).aspx](https://technet.microsoft.com/en-us/library/cc280665(v=sql.105).aspx)

#### NEW QUESTION 137

- (Exam Topic 3)

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 sections, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You have a database that includes a table named Candidate.

You need to update the statistics for a column named Skills in the table and turn off automatic statistics updates for the column.

Solution: You run the following query:

```
USE CustomerDatabase
GO
UPDATE STATISTICS Person.Candidate(Skills)
RESAMPLE, NORECOMPUTE
GO
```

Does this meet the goal?

- A. Yes
- B. No

**Answer:** B

#### Explanation:

RESAMPLE: Update each statistic using its most recent sample rate.

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

#### NEW QUESTION 139

- (Exam Topic 3)

Background Corporate Information

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

Physical Locations

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

Problem Statement

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

Configuration Windows Logins

The network administrators have set up Windows groups to make it easier to manage security. Users may belong to more than one group depending on their role.

The groups have been set up as shown in the following table:

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

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

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

Database

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

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

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

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

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

Column	Data type
SalesOrderDetailID	INT
ProductID	INT
SalePrice	SMALLMONEY
SaleQuantity	INT

Database Issues

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

Requirements Database

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

Customer data access

Customers access the company's website to order products, so they must be able to read product information such as name, description, and price from the Product table. When customers place orders, stored procedures called by the website update product quantity-on-hand values. This means the product table is constantly updated at random times.

#### Customer support data access

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

#### Sales force data access

Sales people at both the headquarters office and the satellite office must generate reports that read from the Product and SalesOrderDetail tables. No updates or inserts are ever made by sales people. These reports are run at random times and there can be no reporting downtime to refresh the data set except during the monthly

maintenance window. The reports that run from the satellite office are process intensive queries with large data sets. Regardless of which office runs a sales force report, the SalesOrderDetail table should only return valid, committed order data; any orders not yet committed should be ignored.

#### Historical Data

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

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

#### Backups

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

#### Database Maintenance

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

#### Project milestones completed

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

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

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

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

You are designing your maintenance plan.

Which command should you use only during the monthly maintenance window?

- A. DBCC INDEXDEFRAG (ProdDB, SalesOrderDetail, SODIndex)
- B. ALTER INDEX SODIndex ON SalesOrderDetail REORGANIZE
- C. ALTER INDEX SODIndex ON SalesOrderDetail REBUILD
- D. ALTER INDEX SODIndex ON SalesOrderDetail REBUILD WITH (ONLINE \* ON)

**Answer: D**

#### Explanation:

- Scenario: 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.

- After the degree of fragmentation is known, use the following table to determine the best method to correct the fragmentation.

```
avg_fragmentation_in_percent value
/ > 5% and < = 30%
then use
ALTER INDEX REORGANIZE
/ > 30%
then use
ALTER INDEX REBUILD WITH (ONLINE = ON)
ALTER INDEX (Transact-SQL)
```

#### NEW QUESTION 143

.....

## Thank You for Trying Our Product

\* 100% Pass or Money Back

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

\* One year free update

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

\* Trusted by Millions

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

\* Shop Securely

All transactions are protected by VeriSign!

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

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