

## Exam Questions 70-764

Administering a SQL Database Infrastructure (beta)

<https://www.2passeasy.com/dumps/70-764/>



### NEW QUESTION 1

- (Exam Topic 1)

You administer a Microsoft SQL Server 2016 database.

You want to make a full backup of the database to a file on disk. In doing so, you need to output the progress of the backup. Which backup option should you use?

- A. STATS
- B. COMPRESSION
- C. CHECKSUM
- D. IN IT

**Answer:** A

### NEW QUESTION 2

- (Exam Topic 1)

You plan to install Microsoft SQL Server 2016 for a web hosting company.

The company plans to host multiple web sites, each supported by a SQL Server database.

You need to select an edition of SQL Server that features backup compression of databases, basic data integration features, and low total cost of ownership.

Which edition should you choose?

- A. Express Edition with Tools
- B. Standard Edition
- C. Web Edition
- D. Express Edition with Advanced Services

**Answer:** B

#### Explanation:

Backup compression is supported on SQL Server 2016 editions: Enterprise, Standard, and Developer. References: <https://docs.microsoft.com/en-us/sql/sql-server/editions-and-components-of-sql-server-2016>

### 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)

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 5

- (Exam Topic 1)

You administer a Microsoft SQL Server 2016 default instance. The instance is hosted by a server that has a local firewall configured. The firewall only allows inbound connections on port 1433. The server only hosts a single instance of SQL Server. You need to ensure that the instance is configured to allow remote connections even if the SQL Server is unresponsive to client connections. What should you do?

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

**Answer:** ABC

#### Explanation:

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

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

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

References:

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

### NEW QUESTION 6

- (Exam Topic 1)

You administer a Microsoft SQL Server 2016 database instance.

You create a new user named UserA. You need to ensure that UserA is able to create SQL Server Agent jobs and to execute SQL Server Agent jobs. To which role should you add UserA?

- A. Securityadmin
- B. RSExecRole
- C. SQLAgentUserRole
- D. DatabaseMailUserRole

**Answer:** C

### NEW QUESTION 7

- (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 ADVSchema. You must implement auditing for all objects in the ADVSchema schema. You must also implement auditing to record access to data that is considered sensitive by the company.
Tailspin Toys	Private	Tailspin Toys has a custom application that accesses a hosted database named <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 monitoring for Tailspin Toys.  
 In the table below, identify the monitoring tool that you must use for each activity.  
 NOTE: Make only one selection in each column.

**Answer Area**

Monitoring option	Monitoring from application	Trend analysis
Error logs	<input type="radio"/>	<input type="radio"/>
Transact-SQL	<input type="radio"/>	<input type="radio"/>
System Monitor	<input type="radio"/>	<input type="radio"/>
Distributed Replay	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Monitoring from application: Transact-SQL  
 Transact-SQL can be used to monitor a customized application. Trend analysis: System Monitor  
 System Monitor can provide trend analysis. From question:

Tailspin Toys has a custom application that accesses a hosted database named TSpinDB. The application will monitor TSpinDB and capture information over time about which database objects are accessed and how frequently they are accessed.

Tailspin Toys has a custom application that accesses a hosted database named TSpinDB. The application will monitor TSpinDB and capture information over time about which database objects are accessed and how frequently they are accessed.

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/performance/performance-monitoring-and-tuning-tools>

#### 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.

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 9

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

#### NEW QUESTION 10

- (Exam Topic 1)

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

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

You have a user database named HRDB that contains sensitive human resources data. The HRDB backup files must be encrypted.

You need to grant the correct permission to the service account that backs up the HRDB database. Which permission should you grant?

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

**Answer:** G

**Explanation:**

Restoring the encrypted backup: SQL Server restore does not require any encryption parameters to be specified during restores. It does require that the certificate or the asymmetric key used to encrypt the backup file be available on the instance that you are restoring to. The user account performing the restore must have VIEW DEFINITION permissions on the certificate or key.

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

**NEW QUESTION 10**

- (Exam Topic 1)

You administer a Microsoft SQL Server 2016 database named Contoso on a server named Server01. You need to collect data for a long period of time to troubleshoot wait statistics when querying Contoso. You also need to ensure minimum impact to the server. What should you create?

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

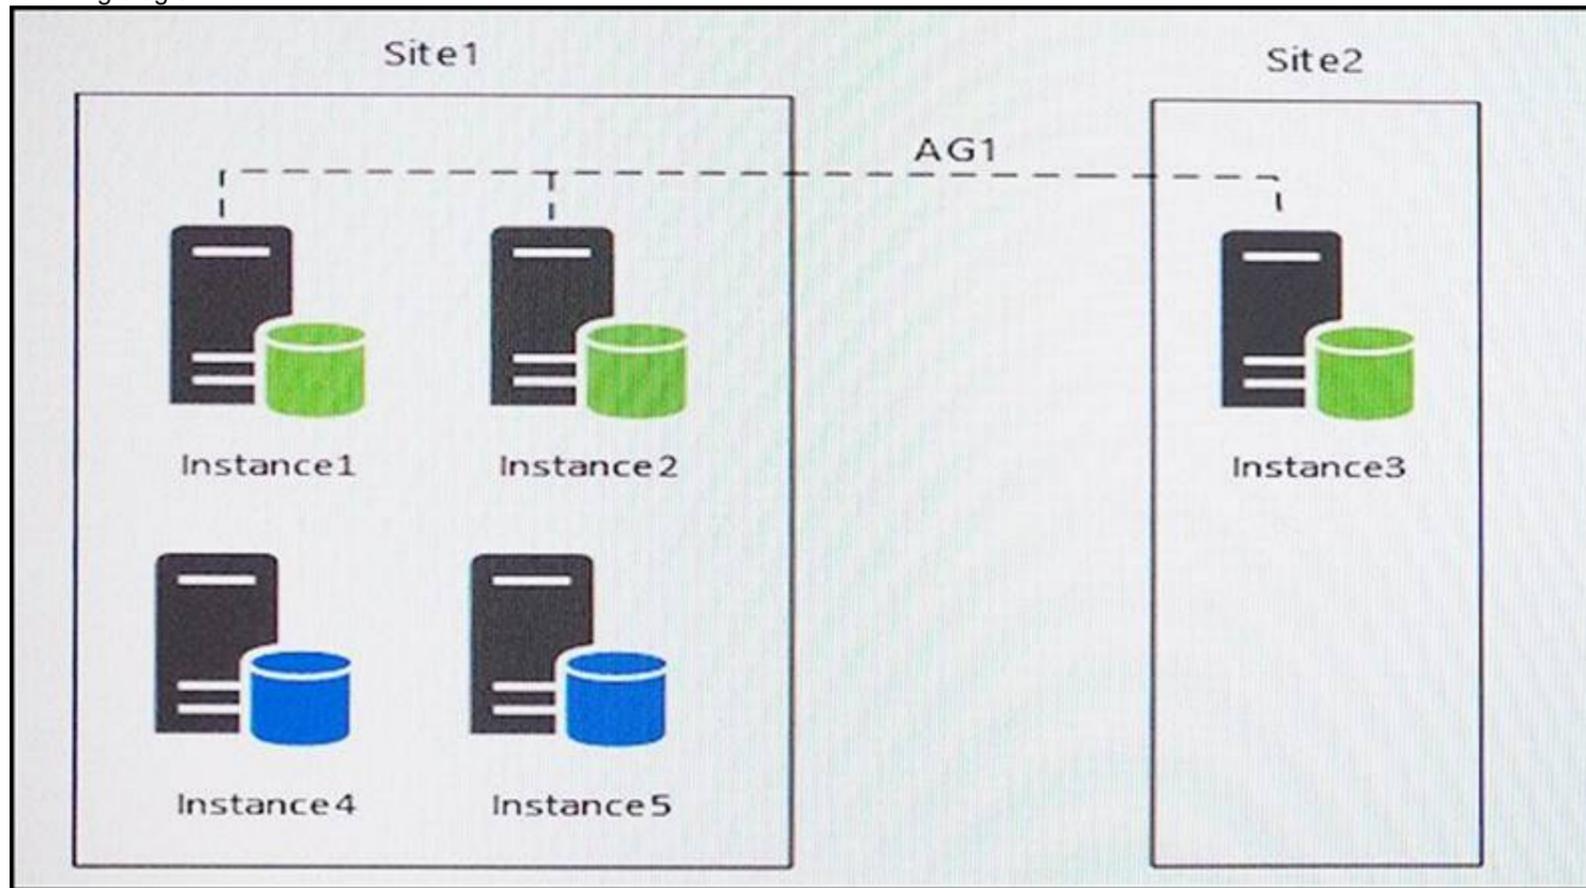
**Answer: C**

**NEW QUESTION 15**

- (Exam Topic 1)

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

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



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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

- A. Modify the backup script to use the keyword SKIP in the FILE\_SNAPSHOT statement.
- B. Modify the backup script to use the keyword SKIP in the WITH statement
- C. Modify the backup script to use the keyword NO\_COMPRESSION in the WITH statement.
- D. Modify the full database backups script to stripe the backup across multiple backup files.

**Answer:** D

**Explanation:**

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

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

References:

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

**NEW QUESTION 19**

- (Exam Topic 1)

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

You 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 ADVSchema. You must implement auditing for all objects in the ADVSchema schema. You must also implement auditing to record access to data that is considered sensitive by the company.
Tailspin Toys	Private	Tailspin Toys has a custom application that accesses a hosted database named <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 21**

- (Exam Topic 1)

You administer a Microsoft SQL Server 2016 database instance.

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

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

What should you do?

- A. Use the Copy Database wizard.
- B. Back up the database from the local server and restore it to Windows Azure SQL Database.
- C. Use the Database Transfer wizard.
- D. Use SQL Server Management Studio to deploy the database to Windows Azure SQL Database.

**Answer:** D

**NEW QUESTION 23**

- (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>

The differential backup of the reporting database fails. Then, the reporting database fails at 14:00 hours. You need to ensure that the reporting database is restored. You also need to ensure that data loss is minimal. What should you do?

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

**Answer:** C

**Explanation:**

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

**NEW QUESTION 25**

- (Exam Topic 1)

You plan to integrate an on-premises Microsoft SQL Server environment with Microsoft Azure. You need to create the authentication object so that you can connect to Azure.

Which Windows PowerShell command or commands should you run?

- A. Invoke-Sqlcmd "CREATE EXTERNAL DATA SOURCE MyAzureStorage WITH (LOCATION = 'wasbs://Azure@myaccount.blob.core.windows.net/', CREDENTIAL = Pa\$\$w0rd)"
- B. New-SqlAzureKeyVaultColumnMasterKeySettings-KeyUrihttps://myvault.vault.contoso.net:443/keys/C
- C. Invoke-Sqlcmd "CREATE CREDENTIAL AzureCred WITH IDENTITY = 'AzureKey', SECRET = 'Pa\$\$w0rd'"
- D. Invoke-Sqlcmd "CREATE LOGIN AzureCred WITH CREDENTIAL = 'AzureKey', PASSWORD = 'Pa\$\$w0rd'"

**Answer: C**

**Explanation:**

Invoke-Sqlcmd runs a script containing statements supported by the SQL Server SQLCMD utility.

The following example creates a SQL Server credential for the Database Engine to use when accessing the Azure Key Vault using the SQL Server Connector for Microsoft Azure Key Vault.

```
CREATE CREDENTIAL Azure_EKM_TDE_cred WITH IDENTITY = 'ContosoKeyVault',
SECRET = 'EF5C8E094D2A4A769998D93440D8115DSECRET_DBEngine'
FOR CRYPTOGRAPHIC PROVIDER AzureKeyVault_EKM_Prov ;
```

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

**NEW QUESTION 26**

- (Exam Topic 1)

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

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

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

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

### Transact-SQL statements

NORECOVERY

RECOVERY

LOG

DBCC CHECKDB

CONTINUE\_AFTER\_ERROR

RESTORE

RESTORE VERIFYONLY

● ● ● ●

---

### Answer Area

Transact-SQL segment DB1 FROM DISK = N'Z:\Backups\Backup.bak WITH

Transact-SQL segment

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

Box 1: RESTORE

Box 2: RECOVERY

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

References:

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

**NEW QUESTION 31**

- (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:**

You 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 33**

- (Exam Topic 1)

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

You are examining information about users, sessions, and processes in an on-premises Microsoft SQL Server 2016 Standard Edition server.

You need to identify waits for resources and return only the following information:

a list of all databases on the SQL Server instance, along with information about the database files, their paths, and names

a list of the queries recently executed that use most of memory, disk, and network resources

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:** E

**Explanation:**

SQL Server Data Tools (SSDT) is a Microsoft Visual Studio environment for creating business intelligence solutions. SSDT features the Report Designer authoring environment, where you can open, modify, preview, save, and deploy Reporting Services paginated report definitions, shared data sources, shared datasets, and report parts.

References: [https://msdn.microsoft.com/en-us/library/hh272686\(v=vs.103\).aspx](https://msdn.microsoft.com/en-us/library/hh272686(v=vs.103).aspx)

**NEW QUESTION 38**

- (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 observe that several indexes are fragmented. You need to rebuild the indexes.

What should you use?

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

**Answer:** C

**Explanation:**

How to: Rebuild an Index (SQL Server Management Studio) To rebuild an index

In Object Explorer, connect to an instance of the SQL Server Database Engine and then expand that instance.

Expand Databases, expand the database that contains the table with the specified index, and then expand Tables.

Expand the table in which the index belongs and then expand Indexes.

Right-click the index to rebuild and then click Rebuild.

To start the rebuild operation, click OK.

References: [https://technet.microsoft.com/en-us/library/ms187874\(v=sql.105\).aspx](https://technet.microsoft.com/en-us/library/ms187874(v=sql.105).aspx)

**NEW QUESTION 43**

- (Exam Topic 1)

You are the database administrator for a Microsoft SQL Server instance. You develop an Extended Events package to look for events related to application performance.

You need to change the event session to include SQL Server errors that are greater than error severity 15. Which five Transact-SQL segments should you use to develop the solution? To answer, move the appropriate

Transact-SQL segments from the list of Transact-SQL segments to the answer area and arrange them in the correct order.

Transact-SQL segments	Answer Area
WHERE ((sqlserver.data- base_id>(4)) AND (severity> (15)))	
(ACTION(sqlserver.client_ap- p_name, sqlserver.data- base_id,sqlserver.session_id)	
ALTER EVENT SESSION Contoso1 ON SERVER	
) GO	
ADD EVENT sqlserver.error_re- ported	
ADD TARGET sqlserver.er- ror_reported	

- A. Mastered  
B. Not Mastered

**Answer:** A

**Explanation:**

Step 1: ALTER EVENT SESSION Contoso1 ON SERVER

Step 2: ADD EVENT ... Step 3: (ACTION ... Step 4: WHERE...

Step 5: ) GO

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

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

```
ADD EVENT sqlserver.error_reported (
```

```
ACTION
```

```
(package0.collect_system_time,package0.last_error,sqlserver.client_app_name,sqlserver.client_hostname,sqlser
```

```
sqlserver.plan_handle,sqlserver.query_hash,sqlserver.session_id,sqlserver.sql_text,sqlserver.tsqf_frame,sqlserve
```

```
WHERE ([severity]>10)
```

```
)
```

```
ADD TARGET package0.event_file (
```

```
SET filename=N'D:\Program Files\Microsoft SQL Server\MSSQL11.MSSQLSERVER\MSSQL\XE\error_trap.xel'
```

```
) WITH (
```

```
STARTUP_STATE=OFF
```

```
) GO
```

References:

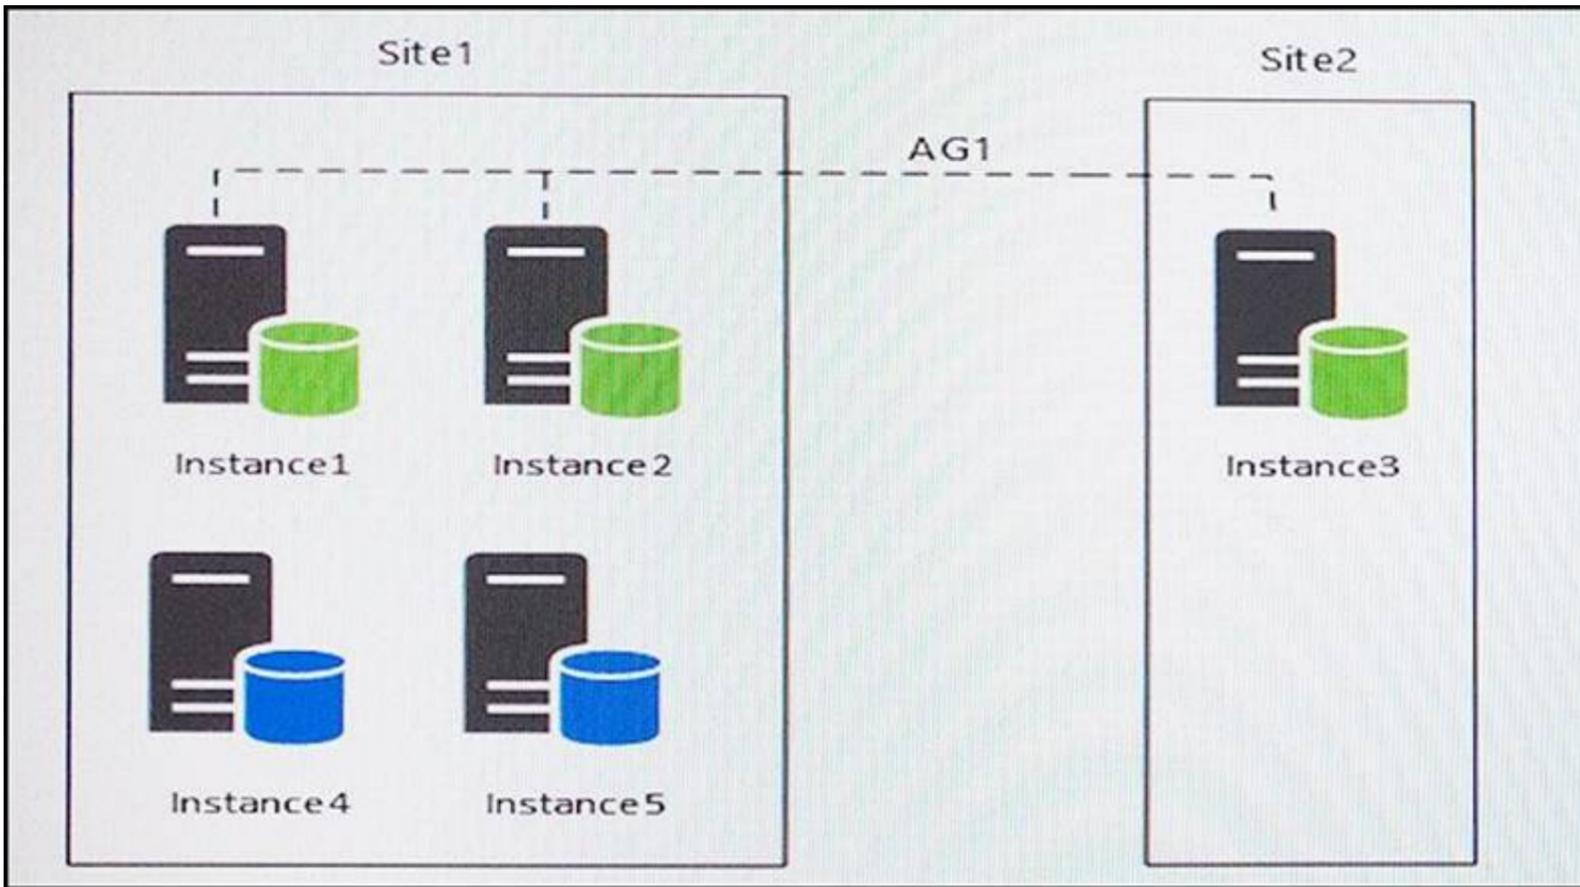
[http://sqlblog.com/blogs/davide\\_mauri/archive/2013/03/17/trapping-sql-server-errors-with-extended-events.aspx](http://sqlblog.com/blogs/davide_mauri/archive/2013/03/17/trapping-sql-server-errors-with-extended-events.aspx)

**NEW QUESTION 48**

- (Exam Topic 1)

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

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



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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

NOTE: Each correct selection is worth one point.

### Objects

- Sys.dm\_os\_wait\_stats
- Sys.dm\_exec\_connections
- Sys.dm\_exec\_requests
- Sys.dm\_exec\_procedure\_stats
- Sys.dm\_exec\_sessions
- Sys.dm\_exec\_query\_stats
- Sys.dm\_exec\_query\_resource\_semaphores
- Sys.dm\_exec\_session\_wait\_stats

### Answer Area

Instance	Object
Instance1	Object
Instance4	Object
Instance5	Object

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Instance 1: sys.dm\_exec\_query\_stats

From Scenario: Instance1 requirement: Aggregate statistics since last server restart. sys.dm\_exec\_query\_stats returns aggregate performance statistics for cached query plans in SQL Server.

Instance 4: sys.dm\_os\_wait\_stats

sys.dm\_os\_wait\_stats returns information about all the waits encountered by threads that executed. From Scenario: Instance4 requirement: Identify the most prominent wait types.

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

Instance 5: sys.dm\_exec\_session\_wait\_stats

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

**NEW QUESTION 51**

- (Exam Topic 1)

You administer a Microsoft SQL Server 2016 failover cluster that contains two nodes named Node A and Node B.

A single instance of SQL Server is installed on the cluster.

An additional node named Node C has been added to the existing cluster.

You need to ensure that the SQL Server instance can use all nodes of the cluster. What should you do?

- A. Create a ConfigurationFile.ini file from Node B, and then run the AddNode command-line tool on Node A.
- B. Use Node A to install SQL Server on Node C.
- C. Run the Add Node to SQL Server Failover Cluster Wizard on Node C.
- D. Use Cluster Administrator to add a new Resource Group to Node B.

**Answer:** C

**Explanation:**

To add a node to an existing SQL Server failover cluster

Insert the SQL Server installation media, and from the root folder, double-click Setup.exe. To install from a network share, navigate to the root folder on the share, and then double-click Setup.exe.

The Installation Wizard will launch the SQL Server Installation Center. To add a node to an existing failover cluster instance, click Installation in the left-hand pane. Then, select Add node to a SQL Server failover cluster.

Etc.

References:

<https://docs.microsoft.com/en-us/sql/sql-server/failover-clusters/install/add-or-remove-nodes-in-a-sql-server-fail>

### NEW QUESTION 52

- (Exam Topic 2)

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

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

Which solution should you use?

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

**Answer:** A

### NEW QUESTION 54

- (Exam Topic 2)

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

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

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

- A. The SQL Server error log
- B. The sys.dm\_os\_wait\_stats dynamic management view
- C. The DBCC SQLPERF(WAITSTATS) command
- D. SQL Server Profiler

**Answer:** B

#### Explanation:

sys.dm\_os\_wait\_stats

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

Columns include: waiting\_tasks\_count

Number of waits on this wait type.

This counter is incremented at the start of each wait.

### NEW QUESTION 55

- (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 solution to improve the performance of usp\_UpdateInventory.

The solution must minimize the amount of development effort. What should you include in the recommendation?

- A. A table variable
- B. A common table expression
- C. A subquery
- D. A cursor

**Answer:** A

**Explanation:**

- Scenario: Database2 will contain a stored procedure named usp\_UpdateInventory. Usp\_UpdateInventory will manipulate a table that contains a self-join that has an unlimited number of hierarchies.
- A table variable can be very useful to store temporary data and return the data in the table format.
- Example: The following example uses a self-join to find the products that are supplied by more than one vendor. Because this query involves a join of the ProductVendor table with itself, the ProductVendor table appears in two roles. To distinguish these roles, you must give the ProductVendor table two different aliases (pv1 and pv2) in the FROM clause. These aliases are used to qualify the column names in the rest of the query. This is an example of the self-join Transact-SQL statement:

```
USE AdventureWorks2008R2;
GO
SELECT DISTINCT pv1.ProductID, pv1.VendorID
FROM Purchasing.ProductVendor pv1
INNER JOIN Purchasing.ProductVendor pv2
ON pv1.ProductID = pv2.ProductID
AND pv1.VendorID <> pv2.VendorID
ORDER BY pv1.ProductID
```

**NEW QUESTION 56**

- (Exam Topic 2)

Overview

General Overview

A Datum Corporation has offices in Miami and Montreal.

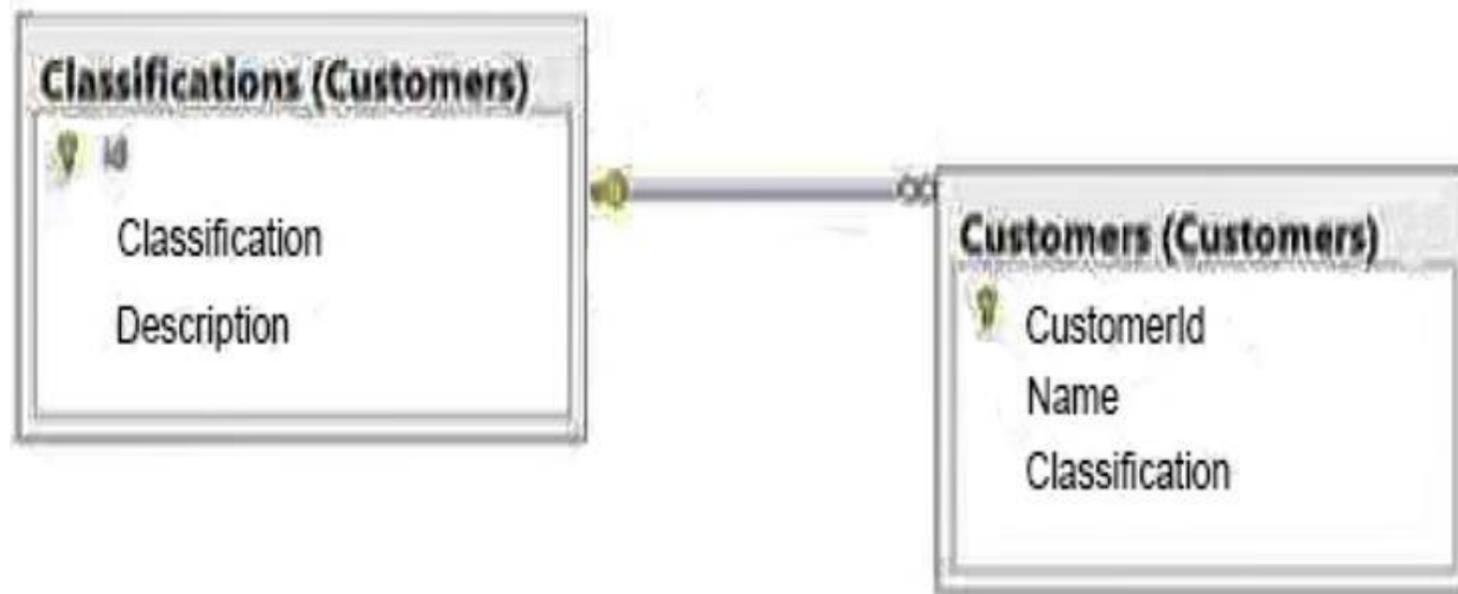
The network contains a single Active Directory forest named adatum.com. The offices connect to each other by using a WAN link that has 5-ms latency. A. Datum standardizes its database platform by using SQL Server 2014 Enterprise edition.

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 change to USP\_3 to ensure that the procedure continues to execute even if one of the UPDATE statements fails.

Which change should you recommend?

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

**Answer:** A

#### Explanation:

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

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

#### NEW QUESTION 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 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 solution for the deployment of SQL Server 2014. The solution must meet the business requirements. What should you include in the recommendation?

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

**Answer:** C

#### Explanation:

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

#### NEW QUESTION 64

- (Exam Topic 2)

You are designing a database named DB1.

Changes will be deployed to DB1 every Wednesday night.

You need to recommend a strategy to deploy the changes to DB1. The strategy must meet the following requirements:

The strategy must not disrupt backup operations.

DB1 must be unavailable to users while the changes are deployed.

You must be able to undo quickly the entire operation.

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

- A. Perform a copy-only database backup before the changes are deployed.If the deployment fails, restore the database to another server and recover the original.Objects from the restored database.
- B. Create a database snapshot.If the deployment fails, recover the objects from the database snapshot.
- C. Create a database snapshot.If the deployment fails, revert the database to the database snapshot.
- D. Perform a full database backup before the changes are deployed.If the deployment fails, restore the database to another server and recover the original objects from the restored database.

**Answer:** C

#### NEW QUESTION 69

- (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 74

- (Exam Topic 2)

You are troubleshooting an application that runs a query. The application frequently causes deadlocks. You need to identify which transaction causes the deadlock.

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

- A. Query the sys.dm\_exec\_requests dynamic management view.
- B. Create a trace in SQL Server Profiler that contains the Deadlock graph event.
- C. Query the sys.dm\_exec\_sessions dynamic management view.
- D. Create an extended events session to capture deadlock information.

**Answer:** D

#### Explanation:

Troubleshooting deadlocks

You have been receiving reports from users indicating that certain applications are returning deadlock errors. To maximize the effectiveness of troubleshooting these problems, you decide to focus on the deadlocks that are hit most frequently. You create an Extended Events session that:

Configures deadlock event tracking for the session.

Specifies a target that aggregates based on an identifier for the deadlock.

You run the Extended Events session, and after additional deadlocks are reported you are able to obtain aggregated deadlock information, along with the complete XML deadlock graph for each source. Using this information, you are able to pin point the most common deadlocks and start working on a solution.

#### NEW QUESTION 76

- (Exam Topic 2)

You want to reproduce the same SQL Server 2016 installation configuration across five servers. Which of the following files will you generate by using SQL Server Setup to accomplish this goal?

- A. Configuration.xml
- B. Setup.ini
- C. Setup.xml
- D. ConfigurationFile.ini

**Answer:** D

#### NEW QUESTION 80

- (Exam Topic 2)

You are planning to deploy a database to Windows Azure SQL Database.

You need to design a stored procedure to update rows. The stored procedure must meet the following requirements:

If more than one row is updated, an error must be raised to the application and the update must be discarded.

The stored procedure must be designed to maximize concurrency.

What should you include in the design? 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
Raise an error in a catch block	
Commit the transaction in a finally block	
Read the @@ROWCOUNT system variable	
Perform the update in a try block	
Raise an error and roll back the transaction if the row count is less than 1	
Issue a SELECT statement to count the number of rows	
Set the isolation level to serializable	
Begin an explicit transaction	

- A. Mastered  
B. Not Mastered

**Answer:** A

**Explanation:**

Note:

- Read Committed is SQL Server's default isolation level.
- @@ROWCOUNT returns the number of rows affected by the last statement.
- Using TRY...CATCH in a transaction

The following example shows how a TRY...CATCH block works inside a transaction. The statement inside the TRY block generates a constraint violation error.

```
- BEGIN TRANSACTION;  
- BEGIN TRY
```

- Generate a constraint violation error.

```
DELETE FROM Production.Product  
WHERE ProductID = 980;  
END TRY  
BEGIN CATCH  
SELECT  
ERROR_NUMBER() AS ErrorNumber  
,ERROR_SEVERITY() AS ErrorSeverity  
,ERROR_STATE() AS ErrorState  
,ERROR_PROCEDURE() AS ErrorProcedure  
,ERROR_LINE() AS ErrorLine  
,ERROR_MESSAGE() AS ErrorMessage;  
IF @@TRANCOUNT > 0  
ROLLBACK TRANSACTION;  
END CATCH;  
IF @@TRANCOUNT > 0  
COMMIT TRANSACTION;  
GO
```

**NEW QUESTION 84**

- (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\_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 89

- (Exam Topic 2)

You administer a Microsoft SQL Server 2016 database that contains a table named OrderDetail. You discover that the NCI\_OrderDetail\_CustomerID non-clustered index is fragmented.

You need to reduce fragmentation.

You need to achieve this goal without taking the index offline. Which Transact-SQL batch should you use?

- A. CREATE INDEX NCI\_OrderDetail\_CustomerID ON OrderDetail.CustomerID WITH DROP EXISTING
- B. ALTER INDEX NCI\_OrderDetail\_CustomerID ON OrderDetail.CustomerID REORGANIZE
- C. ALTER INDEX ALL ON OrderDetail REBUILD
- D. ALTER INDEX NCI\_OrderDetail\_CustomerID ON OrderDetail.CustomerID REBUILD

**Answer:** B

**Explanation:**

References:

<http://msdn.microsoft.com/en-us/library/ms188388.aspx>

#### NEW QUESTION 92

- (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

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

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

Index Fragmentation Issues

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

Backup Issues

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

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

Missing Data Issues

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

Query Performance Issues

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

Design Requirements

File Storage Requirements

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

Data Recovery Requirements

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

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

Concurrency Requirements

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

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

- A. Change the data type of the lastModified column to smalldatetime.
- B. Remove the lastModified column from the clustered index.
- C. Change the data type of the modifiedBy column to tinyint.
- D. Change the data type of the id column to bigint.
- E. Remove the modifiedBy column from the clustered index.
- F. Remove the id column from the clustered index.

**Answer:** BE

**Explanation:**

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

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

**NEW QUESTION 95**

- (Exam Topic 2)

You administer a single server that contains a Microsoft SQL Server 2016 default instance.

You plan to install a new application that requires the deployment of a database on the server. The application login requires sysadmin permissions.

You need to ensure that the application login is unable to access other production databases. What should you do?

- A. Use the SQL Server default instance and configure an affinity mask.
- B. Install a new named SQL Server instance on the server.
- C. Use the SQL Server default instance and enable Contained Databases.
- D. Install a new default SQL Server instance on the server.

**Answer:** B

**Explanation:**

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

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

**NEW QUESTION 100**

- (Exam Topic 2)

You administer a Microsoft SQL Server 2016 database.

You provide temporary securityadmin access to User1 to the database server. You need to know if User1 adds logins to securityadmin.

Which server-level audit action group should you use?

- A. SERVER\_STATE\_CHANGE\_GROUP
- B. SERVER\_PRINCIPAL\_IMPERSONATION\_GROUP
- C. SUCCESSFUL\_LOGIN\_GROUP
- D. SERVER\_ROLE\_MEMBER\_CHANGE\_GROUP

**Answer:** D

**Explanation:**

SERVER\_ROLE\_MEMBER\_CHANGE\_GROUP

This event is raised whenever a login is added or removed from a fixed server role. This event is raised for the sp\_addsrvrolemember and sp\_dropsrvrolemember stored procedures. Equivalent to the Audit Add Login to Server Role Event Class.

References:

<http://technet.microsoft.com/en-us/library/cc280663.aspx>

**NEW QUESTION 103**

- (Exam Topic 2)

You plan to deploy SQL Server 2014. You are designing two stored procedures named SP1 and SP2 that have the following requirements:

- Prevent data read by SP1 from being modified by other active processes.
- Prevent SP2 from performing dirty reads.

You need to recommend the isolation level for each stored procedure.

The solution must maximize concurrency. Which isolation levels should you recommend? To answer, drag the appropriate isolation level to the correct stored procedure in the answer area.

Isolation Levels		Answer area
Read committed	SP1	Isolation level
Read uncommitted	SP2	Isolation level
Repeatable read		
Serializable		

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

SP1 – repeatable read; SP2 – read committed  
 - REPEATABLE READ

This isolation level includes the guarantees given by SNAPSHOT isolation level. In addition, REPEATABLE READ guarantees that for any row that is read by the transaction, at the time the transaction commits the row has not been changed by any other transaction. Every read operation in the transaction is repeatable up to the end of the transaction.

- Committed Read is SQL Server's default isolation level. It ensures that an operation will never read data another application has changed but not yet committed.

**NEW QUESTION 108**

- (Exam Topic 2)

You are implementing a SQL Server 2016 five-node failover cluster. You need to choose a quorum configuration. Which configuration should you use?

- A. Distributed File System (DFS)
- B. Node Majority
- C. Cluster Shared Volume (CSV)
- D. Node and Disk Majority

**Answer: D**

**Explanation:**

Node and Disk Majority (recommended for clusters with an even number of nodes)

**NEW QUESTION 113**

- (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 116**

- (Exam Topic 2)

You install a Microsoft SQL Server 2016 instance.

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

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

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

**Answer: B**

**NEW QUESTION 121**

- (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 126**

- (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 131**

- (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.

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

The screenshot shows a drag-and-drop interface with two main sections: 'Elements' and 'Answer Area'.

**Elements:** A vertical list of eight boxes containing the following text: EXPLICIT, ISOLATION, READ UNCOMMITTED, ROLLBACK, SERIALIZABLE, SNAPSHOT, TABLOCK, and TRANSACTION.

**Answer Area:** A text area containing the following SQL code:
 

```
SET TRANSACTION ISOLATION
LEVEL SNAPSHOT
BEGIN TRANSACTION
UPDATE SalesOrder
...
COMMIT TRANSACTION;
```

 The words 'TRANSACTION' and 'ISOLATION' in the first line are highlighted with a red dashed border. 'SNAPSHOT' and 'TRANSACTION' in the second line are also highlighted with a red dashed border. 'UPDATE SalesOrder' is highlighted in pink. 'COMMIT TRANSACTION;' is highlighted in blue.

**NEW QUESTION 135**

- (Exam Topic 2)

You deploy a database by using SQL Server 2012. The database contains a table named Table1.

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

Stores the most recent data from Table1 by using the fastest storage solution possible.

Stores the historical data from Table1 by using a slower storage solution.

What should you recommend?

- A. partitioned views
- B. a database snapshot
- C. change data capture
- D. table partitioning

**Answer:** D

**NEW QUESTION 136**

- (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.

The screenshot shows a drag-and-drop interface with two main sections: 'Features' and 'Answer area'.

**Features:** A vertical list of four boxes containing the following text: Dynamic management view, Activity Monitor, Resource Governor, and SQL Trace.

**Answer area:** Two text boxes with corresponding empty boxes to their right for dragging a feature.
 

- Text box: "Tempdb must be monitored for insufficient free space." → Empty box: "Feature"
- Text box: "Deadlocks must be analyzed by using Deadlock graphs." → Empty box: "Feature"

- 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 138**

- (Exam Topic 2)

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

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

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

- A. A policy
- B. A maintenance plan
- C. An alert
- D. A job
- E. A trigger

**Answer:** DE

**Explanation:**

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

References:

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

**NEW QUESTION 141**

- (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 143**

- (Exam Topic 3)

You need to ensure that a stored procedure fails if an INSERT statement within the stored procedure fails. What action should you take?

- A. THROW 51000, 'Abort!'
- B. SET XACT\_ABORT OFF
- C. SET XACT\_ABORT ON
- D. TRY....CATCH

**Answer:** C

**NEW QUESTION 144**

- (Exam Topic 3)

You have a SQL Server instance on a server named Server1. You need to recommend a solution to perform the following tasks every week:

Rebuild the indexes by using a new fill factor.

Run a custom T-SQL command.

Back up the databases.

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

- A. A trigger
- B. An alert
- C. A maintenance plan
- D. Windows PowerShell
- E. A system policy

**Answer:** C

**Explanation:**

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

**NEW QUESTION 145**

- (Exam Topic 3)

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

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

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

- A. dynamic management views
- B. Database Engine Tuning Advisor
- C. SQL Server Profiler
- D. Activity Monitor
- E. Data Profile Viewer

**Answer:** BC

**Explanation:**

B: Database Engine Tuning Advisor examines how queries are processed in the databases you specify. When you run a Profiler Trace and feed it to the Database Engine Tuning Advisor, it also looks for missing column statistics, and it can automatically create them for you. C: Use SQL Server Profiler to identify the cause of a deadlock. A deadlock occurs when there is a cyclic dependency between two or more threads, or processes, for some set of resources within SQL Server. Using SQL Server Profiler, you can create a trace that records, replays, and displays deadlock events for analysis.

References: <https://msdn.microsoft.com/en-us/library/ms188246.aspx>

**NEW QUESTION 149**

- (Exam Topic 3)

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

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

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

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

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

- A. Grant the manage auditing and security log permission to the SQL Server service account.
- B. Grant the generate security audits permission on the SQL Server service account.
- C. Update Windows security policy to audit object access.
- D. Restart the SQL Server Agent service.

**Answer:** BC

**Explanation:**

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

The audit object access setting must be configured to capture the events.

The account that the SQL Server service is running under must have the generate security audits permission to write to the Windows Security log.

References:

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

**NEW QUESTION 154**

- (Exam Topic 3)

You are designing two stored procedures named Procedure1 and Procedure2. You identify the following requirements:

Procedure1 must take a parameter that ensures that multiple rows of data can pass into the stored procedure.

Procedure2 must use business logic that resides in a Microsoft .NET Framework assembly. You need to identify the appropriate technology for each stored procedure.

Which technologies should you identify? To answer, drag the appropriate technology to the correct stored procedure in the answer area. (Answer choices may be used once, more than once, or not at all.)

Technologies	Answer Area
Common language runtime (CLR)	Procedure 1 Technology
Extensible Markup Language (XML)	Procedure 2 Technology
A table-valued parameter (TVP)	

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Procedure 1 - A table-valued parameter (TVP); Procedure 2 - Common language runtime (CLR) References:  
<http://msdn.microsoft.com/en-us/library/ms131102.aspx> <http://msdn.microsoft.com/en-us/library/bb522446.aspx> <http://msdn.microsoft.com/en-us/library/bb510489.aspx>

**NEW QUESTION 158**

- (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 SERVER STATE permission to User1.

Does this meet the goal?

- A. Yes
- B. No

**Answer: A**

**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 159**

- (Exam Topic 3)

You need to address the Sales Director's requirements regarding the customer classification. You need to recommend a solution for changing the classifications. What should you recommend?

- A. Add each classification change to a new row in the Customers table.
- B. Record each change to the classification of each customer in a new row in the Customers table.
- C. Add a new row to the Customers table for each new classification.
- D. Record each change to the classification of each customer in a new table in the Customers database.

**Answer: D**

**NEW QUESTION 163**

- (Exam Topic 3)

You are the administrator for a SQL Server 2014 instance that stores the data for an online transaction processing sales system.

The company takes full backups every week; differential backups on the days with no full backups; and hourly transaction backups.

These backups are stored on a backup server in the company's data center. Every week, the company places the full backup on a tape and sends it to a third-party backup storage system. The company is worried that a disaster might occur that could destroy their computer center and cause them to lose orders.

You need to determine the best method for providing the smallest amount of data loss and downtime without leasing or purchasing additional physical locations.

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

- A. Set up SQL Server Always On with a SQL Azure database as a replica.
- B. Set up SQL Server Always On by using a SQL Server on a Windows Azure Virtual Machine.
- C. Put the differential backup on tape and send it to the third-party backup storage system.
- D. Use the Microsoft SQL Server Backup to Microsoft Windows Azure Tool to direct all backups to a different geographical location.

**Answer: D**

**Explanation:**

SQL Server 2012 was the first version to provide the ability to back up databases to the Cloud, and SQL Server 2014 improves on the process.

Microsoft SQL Server Backup to Windows Azure Tool enables backup to Windows Azure Blob Storage and encrypts and compresses SQL Server backups stored locally or in the cloud.

**NEW QUESTION 166**

- (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 data warehouse that stored sales data. One fact table has 100 million rows. You must reduce storage needs for the data warehouse.

You need to implement a solution that uses column-based storage and provides real-time analytics for the operational workload.

Solution: You generate a new certificate on new instance. Does the solution meet the goal?

- A. Yes
- B. No

**Answer: B**

**Explanation:**

Certificates are of no use in this scenario.

#### NEW QUESTION 168

- (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 171

- (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 172

- (Exam Topic 3)

You have a server named Server1 that has 16 processors.

You plan to deploy multiple instances of SQL Server 2014 to Server1. You need to recommend a method to allocate processors to each instance.

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

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

**Answer: A**

#### Explanation:

CPU affinity management through Windows System Resource Manager is not recommended for SQL Server multi-instance management. Instead, use the processor affinity settings in SQL Server.

#### NEW QUESTION 175

- (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.

You need to meet the design requirement for the ProductTypes table in the Product database. Which of the following would be the best solution?

- A. A PRIMARY KEY constraint.
- B. A CHECK constraint.
- C. A UNIQUE constraint.
- D. A Data Definitions Language (DDL) trigger.
- E. A FOREIGN KEY constraint.

Answer: B

**NEW QUESTION 176**

- (Exam Topic 3)

You administer a SQL Server 2014 instance.

Users report that the SQL Server has seemed slow today.

A large database was being restored for much of the day, which could be causing issues. You want to write a query of the system views that will report the following:

Number of users that have a connection to the server

Whether a user's connection is active

Whether any connections are blocked

What queries are being executed

Whether the database restore is still executing and, if it is, what percentage of the restore is complete. Which system objects should you use in your query to best achieve this task?

- A. sys.dm\_exec\_requests, sys.dm\_exec\_sessions, sys.objects
- B. sys.dm\_exec\_sessions, sys.dm\_exec\_query\_stats, sys.dm\_exec\_query\_text, sys.objects
- C. sys.sysprocesses, sys.dm\_exec\_query\_text, sys.objects
- D. sys.dm\_exec\_requests, sys.dm\_exec\_sessions, sys.dm\_exec\_query\_text

Answer: D

**Explanation:**

- sys.dm\_exec\_requests

Returns information about each request that is executing within SQL Server.

- sys.dm\_exec\_sessions

Returns one row per authenticated session on SQL Server. sys.dm\_exec\_sessions is a serverscope view that shows information about all active user connections and internal tasks. This information includes client version, client program name, client login time, login user, current session setting, and more.

- sys.dm\_exec\_query\_text

Returns the text of the SQL batch that is identified by the specified sql\_handle.

**NEW QUESTION 177**

- (Exam Topic 3)

You plan to install two SQL Server 2014 environments named Environment1 and Environment2. Your company identifies the following availability requirements for each environment:

Environment1 must have mirroring with automatic failover implemented.

Environment2 must have Always On with automatic failover implemented.

You need to identify the minimum number of SQL Server 2014 servers that must be deployed to each environment to ensure that all data remains available if a physical server fails.

How many servers should you identify? To answer, drag the appropriate number to the correct environment in the answer area.

Number of Servers		Answer Area
2	Environment1	Number of Server
3	Environment2	Number of Server
4		

- A. Mastered
- B. Not Mastered

Answer: A

**Explanation:**

Environment1 – 3; Environment2 - 2

**NEW QUESTION 180**

- (Exam Topic 3)

Overview

General Overview

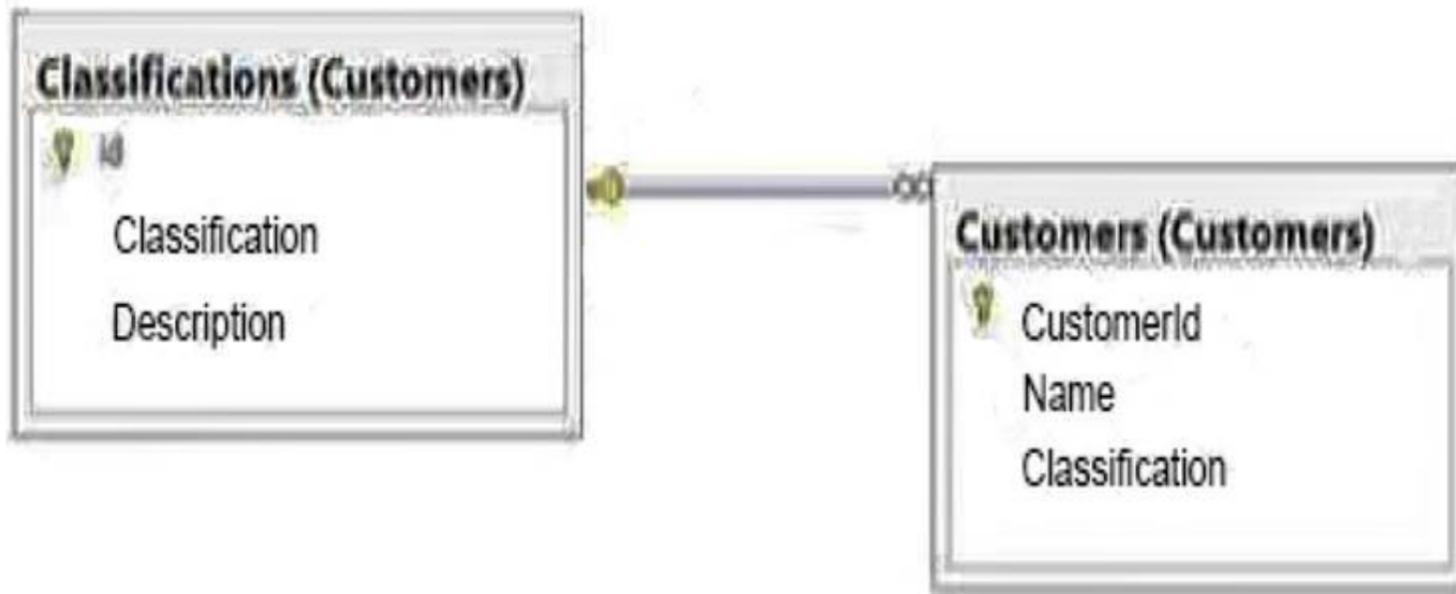
ADatum Corporation has offices in Miami and Montreal.

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

Databases

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

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



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

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

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

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

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

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

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

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

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

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

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

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

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

The nested stored procedures are never called directly. Design Requirements

Data Recovery

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

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

Classification Changes

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

Storage

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

Error Handling

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

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

You need to recommend a solution to minimize the amount of time it takes to execute USP\_1. With what should you recommend replacing Table1?

- A. An indexed view
- B. A function

- C. A table variable
- D. A temporary table

**Answer:** D

**Explanation:**

Scenario:

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 a table in the Products database and searches for information in Table1 based on input from the Products table. After the process is complete, Table1 is deleted.

**NEW QUESTION 185**

- (Exam Topic 3)

You administer two Microsoft SQL Server 2012 servers. Each server resides in a different, untrusted domain. You plan to configure database mirroring. You need to be able to create database mirroring endpoints on both servers. What should you do?

- A. Configure the SQL Server service account to use Network Service.
- B. Use a server certificate.
- C. Use a database certificate.
- D. Configure the SQL Server service account to use Local System.

**Answer:** B

**NEW QUESTION 189**

- (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 need to recommend a solution to back up DB1. What should you include in the recommendation?

- A. Azure Table Storage
- B. Azure Queue storage
- C. Azure Blob storage
- D. Azure Document DB

**Answer: C**

**Explanation:**

For SQL Server the Azure Blob Storage service offers a better alternative to the often used tape option to archive backups. Tape storage might require physical transportation to an off-site facility and measures to protect the media. Storing your backups in Azure Blob Storage provides an instant, highly available, and a durable archiving option.

References:

<https://azure.microsoft.com/en-us/documentation/articles/storage-use-storage-sql-server-backup/restore>

**NEW QUESTION 194**

- (Exam Topic 3)

You use SQL Server 2014.

You need to create a single object that inserts a provided value into Table1, and then returns a count of the records in Table1.

Develop the solution by selecting and arranging the required code blocks in the correct order. You may not need all of the code blocks.

Code Blocks	Answer Area
<pre>END</pre>	
<pre>CREATE FUNCTION dbo.Func_Table1 @InsertWord Varchar (10), @Var1 int OUTPUT AS BEGIN</pre>	
<pre>CREATE FUNCTION dbo.Func_Table1 (@InsertWord Varchar (10)), Returns INT AS BEGIN</pre>	
<pre>CREATE PROCEDURE dbo.Spr_Table1 @InsertWord Varchar (10), @Var1 INT AS BEGIN</pre>	
<pre>Select @Var1 = count(*) from TABLE1</pre>	
<pre>Declare @Var1 INT</pre>	
<pre>CREATE PROCEDURE dbo.Spr_Table1 @InsertWord Varchar (10), @Var1 int OUTPUT AS BEGIN</pre>	
<pre>INSERT INTO TABLE1 (FIELD1) values (@InsertWord)</pre>	

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Code Blocks	Answer Area
END	CREATE PROCEDURE dbo.Spr_Table1 @InsertWord Varchar (10), @Var1 int OUTPUT AS BEGIN
CREATE FUNCTION dbo.Func_Table1 @InsertWord Varchar (10), @Var1 int OUTPUT AS BEGIN	-----
CREATE FUNCTION dbo.Func_Table1 (@InsertWord Varchar (10)), Returns INT AS BEGIN	INSERT INTO TABLE1 (FIELD1) values (@InsertWord)
CREATE PROCEDURE dbo.Spr_Table1 @InsertWord Varchar (10), @Var1 INT AS BEGIN	-----
Select @Var1 = count(*) from TABLE1	Select @Var1 = count(*) from TABLE1
Declare @Var1 INT	-----
CREATE PROCEDURE dbo.Spr_Table1 @InsertWord Varchar (10), @Var1 int OUTPUT AS BEGIN	END
INSERT INTO TABLE1 (FIELD1) values (@InsertWord)	-----

**NEW QUESTION 198**

- (Exam Topic 3)

You have a database hosted on SQL Server 2012 R2. The database contains 5 million rows.

You need to recommend a repeatable method to migrate the database to SQL Azure.

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

- A. Create a SQL Server Integration Services (SSIS) package, and then run the package.
- B. Back up the database, and then restore the database.
- C. Extract a data-tier application, and then import the application.
- D. Generate scripts to create all of the all database objects and all of the data, and then execute the scripts by using SQL Azure.

**Answer:** A

**Explanation:**

SQL Server Integration Services Most flexibility

Data Transfer Efficiency: Good

- SSIS can be used to perform a broad range of data migration tasks. SSIS provides support for complex workflow and data transformation between the source and destination. It is a good choice to transfer of data for databases that require many changes to work on Microsoft Azure SQL Database.

You can use SSIS data transfer packages with another mechanism for transferring the database schema, such as a Data-tier Application package.

**NEW QUESTION 202**

- (Exam Topic 3)

You define a table by using the following statement.

```
CREATE TABLE School.dbo.Teachers
(
  ID INT PRIMARY KEY,
  Teacher XML
)
```

The Teachers table contains one million records. A sample record of the XML data output is as follows:

```
<teacher>
  <id>12345</id>
  <name>Elementary Teacher1</name>
  <grade>5</grade>
  <subjects>
    <subject>English</subject>
    <subject>Math</subject>
  </subjects>
</teacher>
```

You plan to create queries against the table based on grade.

You need to recommend which indexes should be applied to the XML fields. The solution must meet the following requirements:

Reduce the amount of time required to query the table based on grade.

Minimize the size of the index.

What should you recommend? To answer, drag the appropriate index types to the correct fields. Each index type 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

Index Types	Answer Area
No index	Grade: <input type="text"/>
Primary XML index	Id: <input type="text"/>
Secondary XML index	Subject: <input type="text"/>
Selective XML index	

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Grade – Secondary XML index; Id – Primary XML index; Subject – No index

**NEW QUESTION 206**

- (Exam Topic 3)

Overview

General Overview

ADatum Corporation has offices in Miami and Montreal.

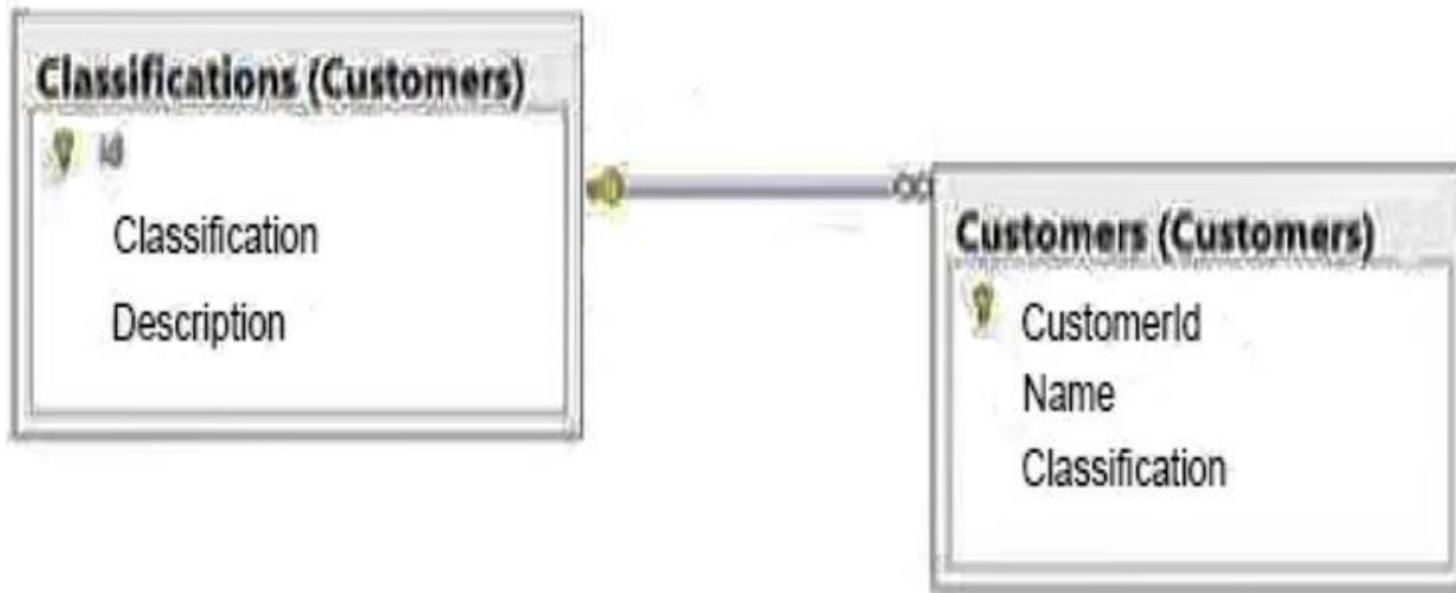
The network contains a single Active Directory forest named adatum.com. The offices connect to each other by using a WAN link that has 5-ms latency. A. Datum standardizes its database platform by using SQL Server 2014 Enterprise edition.

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 error handling of USP\_4.

The solution must handle errors for nested stored procedures in the code for USP\_4. What should you recommend?

- A. Use the @@ERROR variable in the nested stored procedures.
- B. Use the @@ERROR variable in USP\_4.
- C. Use the RAISERROR command in the nested stored procedures.
- D. Use the RAISERROR command in USP\_4.

**Answer: C**

**Explanation:**

- 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.

**NEW QUESTION 208**

- (Exam Topic 3)

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

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

- A. Set the blocked process threshold, and then run SQL Server Profiler.
- B. Set the blocked process threshold, and then create an alert.
- C. Enable trace flag 1204, and then create an alert.

D. Create a job that queries the sys.dm\_os\_waiting\_tasks dynamic management view.

**Answer: B**

**Explanation:**

Step 1: Turn on the blocked process report. This will look for any blocking taking 20 seconds or longer.

- Make sure you don't have any pending changes

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

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

**NEW QUESTION 211**

- (Exam Topic 3)

You create a stored procedure that retrieves all of the rows from a table named Table1.

You need to recommend a solution to ensure that all of the statements in the stored procedure can be executed if another transaction is modifying rows in Table1 simultaneously.

What should you recommend?

- A. Snapshot isolation
- B. A database snapshot
- C. Filegroups
- D. Indexes

**Answer: A**

**Explanation:**

Once snapshot isolation is enabled, updated row versions for each transaction are maintained in tempdb. A unique transaction sequence number identifies each transaction, and these unique numbers are recorded for each row version.

The transaction works with the most recent row versions having a sequence number before the sequence number of the transaction.

Newer row versions created after the transaction has begun are ignored by the transaction.

**NEW QUESTION 214**

- (Exam Topic 3)

You work as a Database Administrator (DBA) for a company named ABC.com.

The company has a Windows Azure subscription.

The company uses a cloud based SQL Server environment hosted on SQL Azure. Developers in your company are creating an ecommerce website.

You are designing a database for the website. The database will be hosted on SQL Azure.

The database will store and reuse web site login details and customer credit card numbers.

You need to ensure the username, passwords and credit card details are securely stored in the database. Which of the following would be the most suitable secure storage solution?

- A. Secure Sockets Layer (SSL)
- B. IPSec
- C. Data encryption
- D. Transparent Data Encryption (TDE)
- E. Encrypting File System (EFS)

**Answer: C**

**NEW QUESTION 215**

- (Exam Topic 3)

You are the new database administrator for a SQL Server 2014 instance.

You conduct an assessment on the instance and determine that the auto create statistics setting on the database named DB1 has been turned off.

You see no evidence that any maintenance has been occurring. You want to set up monitoring to see if query performance is being affected. You need to set up a monitoring process that will capture any cases where statistics could have been useful if they existed.

What should you do?

- A. Create a SQL Server Agent job to execute DBCC SHOWSTATISTICS on each of the primary key columns in the database.
- B. Use the missing\_column\_statistics extended event.
- C. Query the sys.statistics system view to see all cases where the statistics were last needed.
- D. Write a query using the sys.dm\_db\_missing\_index\_group\_stats DMV Joining to sys.indexes, filtering on is\_hypothetical.

**Answer: B**

**Explanation:**

The Missing Column Statistics event class indicates that column statistics that could have been useful for the optimizer are not available. By monitoring the Missing Column Statistics event class, you can determine if there are statistics missing for a column used by a query.

This can cause the optimizer to choose a less efficient query plan than expected.

#### NEW QUESTION 219

- (Exam Topic 3)

Your company has a SQL Azure subscription.

You implement a database named Database1. Database1 has two tables named Table1 and Table2. You create a stored procedure named sp1. Sp1 reads data from Table1 and inserts data into Table2. A user named User1 informs you that he is unable to run sp1.

You verify that User1 has the SELECT permission on Table1 and Table2. You need to ensure that User1 can run sp1.

The solution must minimize the number of permissions assigned to User1. What should you do?

- A. Grant User1 the INSERT permission on Table2.
- B. Add User1 to the db\_datawriter role.
- C. Change sp1 to run as the sa user.
- D. Grant User1 the EXECUTE permission on sp1.

**Answer:** D

#### Explanation:

References:

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

#### NEW QUESTION 223

- (Exam Topic 4)

You have 10 Microsoft SQL Server 2016 servers.

You deploy a management data warehouse named dwi. You configure dwi to gather all the performance data from the servers.

You configure a Data Collector on a SQL server named svi.

You query the data warehouse on dwi and discover that data from svi is unavailable.

You need to ensure that you can review the performance data from svi when you query dwi. What should you do?

- A. Start the SQL Server Agent service on SV1.
- B. Execute the msdb.sp\_syscollector\_enable\_collector stored procedure on DW1.
- C. Start the SQL Server Agent service on DW1.
- D. Execute the msd
- E. sp\_syscollector\_set\_warehouse\_connection\_user Stored procedure on SVI.

**Answer:** C

#### NEW QUESTION 228

- (Exam Topic 4)

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 named DB1 that is 640 GB and is updated frequently.

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

You need to reduce the amount of disk space used by the log shipping process.

Solution: You configure log shipping to backup and restore by using a single shared folder. Does this meet the goal?

- A. Yes
- B. No

**Answer:** B

#### Explanation:

You should compress the transaction log backups. References:

<https://docs.microsoft.com/en-us/sql/database-engine/log-shipping/configure-log-shipping-sql-server?view=sql->

#### NEW QUESTION 230

.....

## THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual 70-764 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the 70-764 Product From:

<https://www.2passeasy.com/dumps/70-764/>

### Money Back Guarantee

#### **70-764 Practice Exam Features:**

- \* 70-764 Questions and Answers Updated Frequently
- \* 70-764 Practice Questions Verified by Expert Senior Certified Staff
- \* 70-764 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- \* 70-764 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year