

## 70-764 Dumps

### Administering a SQL Database Infrastructure (beta)

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



**NEW QUESTION 1**

- (Exam Topic 1)

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

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

What should you use?

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

**Answer:** D

**Explanation:**

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

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

Missing indexes can be found with the execution plans.

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

**NEW QUESTION 2**

- (Exam Topic 1)

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

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

Customer	Cloud Type	Description
AdventureWorks Cycles	Private	The environment includes a database named <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 monitor WingDB and gather information for troubleshooting issues. What should you use?

- A. sp\_updatestats
- B. sp\_lock
- C. sys.dm\_os\_waiting\_tasks
- D. sys.dm\_tran\_active\_snapshot\_database\_transactions
- E. Activity Monitor

**Answer: B**

**Explanation:**

The sp\_lock system stored procedure is packaged with SQL Server and will give you insight into the locks that are happening on your system. This procedure returns much of its information from the syslock info in the master database, which is a system table that contains information on all granted, converting, and waiting lock requests.

Note: sp\_lock will be removed in a future version of Microsoft SQL Server. Avoid using this feature in new development work, and plan to modify applications that currently use this feature. To obtain information about locks in the SQL Server Database Engine, use the sys.dm\_tran\_locks dynamic management view.

sys.dm\_tran\_locks returns information about currently active lock manager resources in SQL Server 2008 and later. Each row represents a currently active request to the lock manager for a lock that has been granted or is waiting to be granted.

References:

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

**NEW QUESTION 3**

- (Exam Topic 1)

You administer a Windows Azure SQL Database database named Human\_Resources. The database contains 2 tables named Employees and SalaryDetails. You add two Windows groups as logins for the server:

- CORP\Employees - All company employees
  - CORP\HRAdmins - HR administrators only
- HR Administrators are also company employees.

You need to grant users access according to the following requirements:  
CORP\Employees should have SELECT access to the Employees table.  
Only users in CORP\HRAdmins should have SELECT access to the SalaryDetails table.  
Logins are based only on Windows security groups.  
What should you do?

- A. Create a database role called Employees. Add CORP\Employees to the db\_datareader role. Add all company employees except HR administrators to the Employees role. Deny SELECT access to the SalaryDetails table to the Employees role.
- B. Create a database role called HRAdmins. Add all company employees except HR administrators to the db\_datareader role, Add all HR administrators to the HRAdmins role. Grant SELECT access to the SalaryDetails table to the HRAdmins role. Deny SELECT access to the SalaryDetails table to the db\_datareader role.
- C. Create two database roles: Employees and HRAdmins. Add all company employees to the Employees role. Add HR administrators to the HRAdmins role. Grant SELECT access to all tables except SalaryDetails to the Employees role. Grant SELECT access to the SalaryDetails table to the HRAdmins role. Deny SELECT access to the SalaryDetails table to the Employees role.
- D. Create a database role called Employees. Add all HR administrators to the db\_datareader role. Add all company employees to the Employees role. Grant SELECT access to all tables except the SalaryDetails table to the Employees role. Deny SELECT access to the SalaryDetails table to the Employees role.

**Answer: D**

**NEW QUESTION 4**

- (Exam Topic 1)

You administer a Microsoft SQL Server 2016 server that hosts a transactional database and a reporting database.  
The transactional database is updated through a web application and is operational throughout the day. The reporting database is only updated from the transactional database.  
The recovery model and backup schedule are configured as shown in the following table:

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

At 16:20 hours, you discover that pages 17, 137, and 205 on one of the database files are corrupted on the transactional database. You need to ensure that the transactional database is restored. You also need to ensure that data loss is minimal.  
What should you do?

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

- K. Restore the latest full backup
- L. Then, restore the latest differential backup.

**Answer:** F

**Explanation:**

The goal of a page restore is to restore one or more damaged pages without restoring the whole database. Typically, pages that are candidates for restore have been marked as "suspect" because of an error that is encountered when accessing the page.

Note: Requirements for Restoring Pages

A page restore is subject to the following requirements:

The databases must be using the full or bulk-logged recovery model. Etc.

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

**NEW QUESTION 5**

- (Exam Topic 1)

You administer a Microsoft SQL Server 2016 database named Orders.

Users report that during peak usage periods, certain operations are taking more time than expected. Your initial analysis suggests that blocking is the cause.

You need to gather more data to be able to determine which processes are being blocked and to identify the root cause.

What should you do?

- A. Start a trace using SQL Server Profiler to catch the Lock: Deadlock event.
- B. Use sp\_configure to set the blocked process threshold
- C. Start a trace using SQL Server Profiler to catch the Blocked Process Report event.
- D. Schedule a SQL Agent job to run every 60 seconds and insert the results of executing the sys.dm\_os\_wait\_stats DMV into a table.
- E. Use System Monitor to catch the Lock Waits/sec event.

**Answer:** B

**NEW QUESTION 6**

- (Exam Topic 1)

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

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

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

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

- A. Yes
- B. No

**Answer:** A

**Explanation:**

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

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

References:

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

**NEW QUESTION 7**

- (Exam Topic 1)

You plan to install a Microsoft SQL Server 2016 instance.

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

Store Excel workbooks on the file system.

Access the workbooks through Transact-SQL.

Include the workbooks in database backups.

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

Which feature should you use?

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

**Answer:** B

**NEW QUESTION 8**

- (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 9**

- (Exam Topic 1)

You administer a Microsoft SQL Server 2016 database.

Users report that a billing application becomes unresponsive during busy times of the day. While investigating, you notice large number of processes taking or waiting for table locks. You suspect that SQL Server is assigning stronger locks to queries.

You start a SQL Profiler trace. Which event should you select?

- A. Deadlock graph
- B. Lock: Escalation
- C. Lock: Timeout
- D. Lock: Deadlock

**Answer:** B

**NEW QUESTION 10**

- (Exam Topic 1)

You have a database named DB1 that is configured to use the full recovery model. You have a full daily backup job that runs at 02:00. The job backs up data from DB1 to the file B:\DB1.bak.

You need to restore the DB1 database to the point in time of May 25, 2016 at 02:23 and ensure that the database is functional and starts to accept connections. Which Transact-SQL statement should you run?

```

A.
BACKUP LOG [DB1] TO DISK = N'B:\DB1Log.bak' WITH RECOVERY
RESTORE DATABASE [DB1] FROM DISK = N'B:\DB1.bak' WITH NORECOVERY
RESTORE LOG [DB1] FROM DISK = N'B:\DB1Log.bak' WITH STOPAT = N'2016-05-25T02:23:00'

B.
BACKUP LOG [DB1] TO DISK = N'B:\DB1Log.bak' WITH NORECOVERY
RESTORE DATABASE [DB1] FROM DISK = N'B:\DB1.bak' WITH NORECOVERY
RESTORE LOG [DB1] FROM DISK = N'B:\DB1Log.bak' WITH STOPAT = N'2016-05-25T02:23:00'

C.
BACKUP LOG [DB1] TO DISK = N'B:\DB1Log.bak' WITH NORECOVERY
RESTORE DATABASE [DB1] FROM DISK = N'B:\DB1.bak' WITH NORECOVERY
RESTORE LOG [DB1] FROM DISK = N'B:\DB1Log.bak' WITH STOPAT = N'2016-05-25T02:23:00', NORECOVERY

D.
RESTORE DATABASE [DB1] FROM DISK = N'B:\DB1.bak' WITH STOPAT = N'2016-05-25T02:23:00', RECOVERY
    
```

- A. Option A
- B. Option B
- C. Option C
- D. Option D

**Answer:** B

**NEW QUESTION 10**

- (Exam Topic 1)

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

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

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

Solution: You add the DatabaseMailUserRole to Mail1 in the tempdb 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 tempdb 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 SQLServer 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 13**

- (Exam Topic 1)

You administer a Microsoft SQL Server 2016 environment.

One of the SQL Server 2016 instances contains a database named Sales. You plan to migrate Sales to Windows Azure SQL Database.

To do so, you need to implement a contained database.

What should you do? (Each correct answer presents part of the solution. Choose all that apply.)

- A. Set database containment to AZURE.
- B. Enable server property contained database authentication.
- C. Disable server property cross db ownership chaining.
- D. Set database containment to PARTIAL.
- E. Disable server property contained database authentication.
- F. Set database containment to FULL.

**Answer:** BD

**NEW QUESTION 14**

- (Exam Topic 1)

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

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

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

You need to address the performance issues with each table.

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

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

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Table A: Auto\_update statistics off

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

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

SET AUTO\_UPDATE\_STATISTICS\_ASYNC ON

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

References:

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

**NEW QUESTION 18**

- (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 failure from the most recent full backup.
- K. Perform a page restore.
- L. Perform a partial restore.

**Answer:** C

**Explanation:**

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

**NEW QUESTION 19**

- (Exam Topic 1)

You are configuring log shipping for a Microsoft SQL Server database named salesOrders. You run the following Transact-SQL script:

```

DECLARE @LS_BackupJobId AS uniqueidentifier
DECLARE @LS_PrimaryId AS uniqueidentifier
DECLARE @SP_Add_RetCode As int
EXEC @SP_Add_RetCode = master.dbo.sp_add_log_shipping_primary_database
    @database = N'salesOrders'
    ,@backup_directory = N'C:\Backup'
    ,@backup_share = N'\\localhost\Backup'
    ,@backup_job_name = N'LSBackup_salesOrders'
    ,@backup_retention_period = 4320
    ,@backup_compression = 1
    ,@backup_threshold = 60
    ,@threshold_alert_enabled = 1
    ,@history_retention_period = 5760
    ,@backup_job_id = @LS_BackupJobId OUTPUT
    ,@primary_id = @LAS_PrimaryId OUTPUT
    ,@overwrite = 1
IF (@@ERROR = 0 AND @SP_Add_RetCode = 0)
    BEGIN
        DECLARE @LS_BackUpScheduleUID As uniqueidentifier
        DECLARE @LA_BackUpScheduleID AS int
        EXEC msdb.dbo.sp_add_schedule
            @schedule_name = N'LSBackupSchedule_ADATUM-SQL11'
            ,@enabled = 1
            ,@freq_type = 4
            ,@freq_interval = 1
            ,@freq_subday_type = 4
            ,@freq_subday_interval = 15
            ,@freq_recurrence_factor = 0
            ,@active_start_date = 20160720
            ,@active_end_date = 99991231
            ,@active_start_time = 0
            ,@active_end_time = 235900
            ,@schedule_uid = @LS_BackUpScheduleUID OUTPUT
            ,@schedule_id = @LS_BackupScheduleID OUTPUT
        EXEC msdb.dbo.sp_attach_schedule
            @job_id = @LS_BackupJobId
            ,@schedule_id = @LS_BackupScheduleID
        EXEC msdb.dbo.sp_update_job
            @job_id = @LS_BackupJobId
            ,@enabled = 1
    END
EXEC master.dbo.sp_add_log_shipping_alert_job

```

You need to determine the changes that the script has on the environment.

How does the script affect the environment? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

## Answer Area

A dedicated file share [answer choice] used to store the backups.

is
is not

A SQL Server monitor instance [answer choice] on a server named ADATUM-SQL11.

runs
does not run

Backup files will be deleted after [answer choice].

24 hours
48 hours
72 hours

The backup job will run every [answer choice].

15 minutes
60 minutes
24 hours

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: is

The dedicated backup file share is \\localhost\Backup Box 2: does not run

The only thing with a name related to ADATM-SQL11 is the schedule name. Box 3: 72 hours

4320 minutes equals 72 hours.

Note: @backup\_retention\_period= ] backup\_retention\_period

Isthe length of time, in minutes, to retain the log backup file in the backup directory on the primary server. backup\_retention\_period is int, with no default, and cannot be NULL.

Box 4: 15 minutes.

[ @freq\_subday\_type = ] freq\_subday\_type

Specifies the units for freq\_subday\_interval. freq\_subday\_typeis int, with a default of 0, and can be one of these values.

Here it is 4, which means minutes.

[ @freq\_subday\_interval = ] freq\_subday\_interval

The number of freq\_subday\_type periods to occur between eachexecution of a job. freq\_subday\_intervalis int, with a default of 0.

Note: Interval should be longer than 10 seconds. freq\_subday\_interval is ignored in those cases where freq\_subday\_type is equal to 1.

Here it is 15. References:

<https://docs.microsoft.com/en-us/sql/relational-databases/system-stored-procedures/sp-add-schedule-transact-sql> <https://docs.microsoft.com/en-us/sql/relational-databases/system-stored-procedures/sp-add-log-shipping-primary>

**NEW QUESTION 23**

- (Exam Topic 1)

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

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

"Executed as user CONTOSO\ServiceAccount. ...0.4035.00 for 64-bit Copyright (C) Microsoft Corp 1984-2011. All rights reserved. Started 63513 PM Error 2012-06-23 183536.87 Code 0XC001000E Source UserImport Description Code 0x00000000 Source Log Import Activity Descript... The package execution fa... The step failed."

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

- A. Expand agent logging to include information from all events.
- B. Disable the Limit size of job history log feature.
- C. Configure event forwarding.
- D. Configure output files.

**Answer:** D

**Explanation:**

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

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

**NEW QUESTION 27**

- (Exam Topic 1)

You administer a Microsoft SQL Server database named Contoso. You create a stored procedure named Sales.ReviewInvoice by running the following Transact-SQL statement:

```
CREATE PROCEDURE Sales.ReviewInvoice (@SaleID int)
AS
    DECLARE @tsql nvarchar(4000) = 'SELECT SaleID, CustomerID, TotalAmount FROM Sales.SalesInvoice WHERE SaleID = '
    SET @tsql = @tsql + CAST(@saleID AS varchar(20))
    EXEC sp_executesql @tsql
```

You need to create a Windows-authenticated login named ContosoSearch and ensure that ContosoSearch can run the Sales.ReviewInvoices stored procedure. Which three Transact-SQL segments should you use to develop the solution? To answer, move the appropriate Transact-SQL segments from the list of Transact-SQL segments to the answer area and arrange them in the correct order.

Transact-SQL segments	Answer Area
<pre>Use Contoso GO CREATE USER Contoso\SalesGroup FOR LOGIN Contoso\SalesGroup</pre>	
<pre>ALTER ROLE db_ddladmin ADD MEMBER Contoso\SalesGroup GRANT VIEW SEFINITION ON Sales.- SalesInvoice TO Contoso\SalesGroup</pre>	
<pre>use master CREATE LOGIN Contoso\SalesGroup FROM WINDOWS GO</pre>	
<pre>GRANT EXECUTE ON Sales.ReviewInvoice TO Contoso\SalesGroup GRANT SELECT ON Sales.SalesInvoice TO Contoso\SalesGroup</pre>	
<pre>use master CREATE LOGIN Contoso\ContosoSearch WITH PASSWORD=N'Pa\$\$w0rd' GO</pre>	
<pre>GRANT EXECUTE ON Sales.ReviewInvoice TO Contoso\SalesGroup GRANT VIEW DEFINITION ON Sales.SalesIn- voice TO Contoso\SalesGroup</pre>	
<pre>GRANT EXECUTE, SELECT ON Sales.Review- Invoice TO Contoso\SalesGroup</pre>	

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Transact-SQL segments	Answer Area
<pre>Use Contoso GO CREATE USER Contoso\SalesGroup FOR LOGIN Contoso\SalesGroup</pre>	<pre>use master CREATE LOGIN Contoso\ContosoSearch WITH PASSWORD=N' Pa\$\$w0rd' GO</pre>
<pre>ALTER ROLE db_ddladmin ADD MEMBER Contoso\SalesGroup GRANT VIEW SEFINITION ON Sales.- SalesInvoice TO Contoso\SalesGroup</pre>	<pre>Use Contoso GO CREATE USER Contoso\SalesGroup FOR LOGIN Contoso\SalesGroup</pre>
<pre>use master CREATE LOGIN Contoso\SalesGroup FROM WINDOWS GO</pre>	<pre>GRANT EXECUTE, SELECT ON Sales.Review- Invoice TO Contoso\SalesGroup</pre>
<pre>GRANT EXECUTE ON Sales.ReviewInvoice TO Contoso\SalesGroup GRANT SELECT ON Sales.SalesInvoice TO Contoso\SalesGroup</pre>	
<pre>use master CREATE LOGIN Contoso\ContosoSearch WITH PASSWORD=N' Pa\$\$w0rd' GO</pre>	
<pre>GRANT EXECUTE ON Sales.ReviewInvoice TO Contoso\SalesGroup GRANT VIEW DEFINITION ON Sales.SalesIn- voice TO Contoso\SalesGroup</pre>	
<pre>GRANT EXECUTE, SELECT ON Sales.Review- Invoice TO Contoso\SalesGroup</pre>	

**NEW QUESTION 28**

- (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 Database Engine instance.

You need to return information about processes that are not idle, that belong to a specific user, or that belong to a specific session.

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: B**

**Explanation:**

Use sp\_who3 to first view the current system load and to identify a session of interest. You should execute the query several times to identify which session id is most consuming the system resources.

Parameters

sp\_who3 null - who is active;  
 sp\_who3 1 or 'memory' - who is consuming the memory;  
 sp\_who3 2 or 'cpu' - who has cached plans that consumed the most cumulative CPU (top 10); sp\_who3 3 or 'count' - who is connected and how many sessions it has;  
 sp\_who3 4 or 'idle' - who is idle that has open transactions;  
 sp\_who3 5 or 'tempdb' - who is running tasks that use tempdb (top 5); and, sp\_who3 6 or 'block' - who is blocking.

**NEW QUESTION 33**

- (Exam Topic 1)

You are configuring a new Microsoft SQL Server Always On Availability Group. You plan to configure a shared network location at \\DATA-C11\SQL. You need to create an availability group listener named AGL1 on port 1433.

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

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

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Step 1: Launch the Failover Cluster Manager and..

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

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

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

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

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

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

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

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

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

**NEW QUESTION 36**

- (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 40**

- (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
<pre>WHERE ((sqlserver.data- base_id&gt;(4)) AND (severity&gt; (15)))</pre>	<div style="display: flex; justify-content: space-around; align-items: center;"> <div style="text-align: center;"> <span>⬅</span> <span>➡</span> </div> <div style="text-align: center;"> <span>⬆</span> <span>⬇</span> </div> </div>
<pre>(ACTION(sqlserver.client_ap- p_name, sqlserver.data- base_id,sqlserver.session_id)</pre>	
<pre>ALTER EVENT SESSION Contoso1 ON SERVER</pre>	
<pre>) GO</pre>	
<pre>ADD EVENT sqlserver.error_re- ported</pre>	
<pre>ADD TARGET sqlserver.er- ror_reported</pre>	

- 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 47

- (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 plan to delegate encryption operations to a user.

You need to grant the user permission to implement cell-level encryption 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 ServerState
- G. View Definition
- H. sysadmin

**Answer:** G

#### Explanation:

The following permissions are necessary to perform column-level encryption, or cell-level encryption.

CONTROL permission on the database.

CREATE CERTIFICATE permission on the database. Only Windows logins, SQL Server logins, and application roles can own certificates. Groups and roles cannot own certificates.

ALTER permission on the table.

Some permission on the key and must not have been denied VIEW DEFINITION permission. References:

<https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/encrypt-a-column-of-data>

#### NEW QUESTION 51

- (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 reorganize 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 53

- (Exam Topic 1)

You administer several Microsoft SQL Server 2016 database servers.

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

Occasionally, merge replication fails due to timeout errors. You need to reduce the occurrence of these timeout errors. What should you do?

- A. Set the Merge agent on the problem subscribers to use the slow link agent profile.
- B. Create a snapshot publication, and reconfigure the problem subscribers to use the snapshot publication.
- C. Change the Merge agent on the problem subscribers to run continuously.
- D. Set the Remote Connection Timeout on the Publisher to 0.

**Answer:** A

**Explanation:**

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

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

References:

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

**NEW QUESTION 57**

- (Exam Topic 1)

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

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

What should you use?

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

**Answer:** A

**Explanation:**

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

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

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

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

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

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

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

**NEW QUESTION 62**

- (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 63**

- (Exam Topic 2)

Overview

General Overview

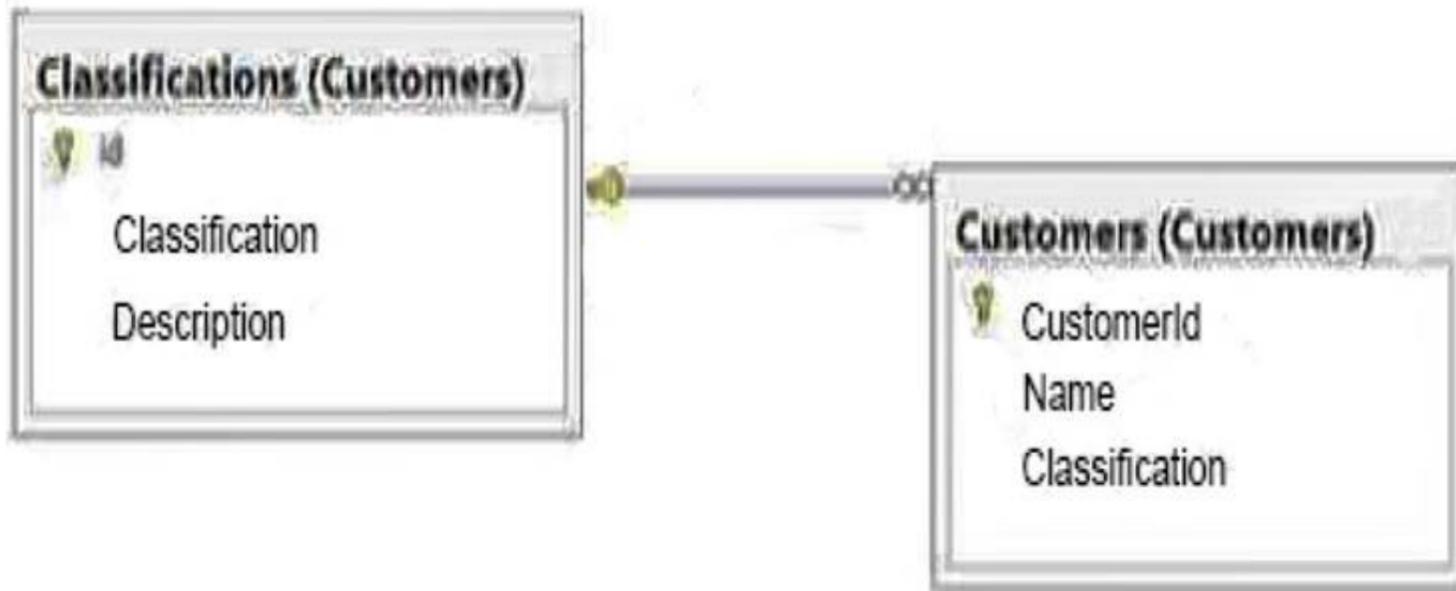
A Datum Corporation has offices in Miami and Montreal.

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

Databases

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

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



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

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

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

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

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

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

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

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

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

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

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

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

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

The nested stored procedures are never called directly. Design Requirements

Data Recovery

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

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

Classification Changes

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

Storage

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

Error Handling

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

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

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

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

**Answer:** CD

**Explanation:**

Scenario:

You plan to change the way customers are classified.

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

**NEW QUESTION 64**

- (Exam Topic 2)

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

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

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

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

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

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

**Answer:** BC

**Explanation:**

References:

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

**NEW QUESTION 67**

- (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 72**

- (Exam Topic 2)

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

What should you create?

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

**Answer:** D

**NEW QUESTION 75**

- (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 78**

- (Exam Topic 2)

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

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

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

What should you include in the design?

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

**Answer:** D

**Explanation:**

References:

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

**NEW QUESTION 79**

- (Exam Topic 2)

You administer a Microsoft SQL Server 2016 failover cluster.

You need to ensure that a failover occurs when the server diagnostics returns query\_processing error. Which server configuration property should you set?

- A. SqlOumperDumpFlags
- B. FailureConditionLevel
- C. HealthCheckTimeout
- D. SqlDumperDumpPath

**Answer:** B

**Explanation:**

The SQL Server Database Engine resource DLL determines whether the detected health status is a condition for failure using the FailureConditionLevel property. The FailureConditionLevel property defines which detected health statuses cause restarts or failovers. Multiple levels of options are available, ranging from no automatic restart or failover to all possible failure conditions resulting in an automatic restart or failover.

References:

<https://docs.microsoft.com/en-us/sql/sql-server/failover-clusters/windows/failover-policy-for-failover-cluster-ins>

**NEW QUESTION 83**

- (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 86**

- (Exam Topic 2)

You plan to deploy a database to SQL Azure. You are designing two stored procedures named USP\_1 and USP\_2 that have the following requirements:

Prevent data read by USP\_1 from being modified by other active processes.

Allow USP\_2 to perform dirty reads.

You need to recommend the isolation level for the stored procedures. 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 uncommitted Note:

- SP1: repeatable read a repeatable read scan retains locks on every row it touches until the end of the transaction. Even rows that do not qualify for the query result remain locked. These locks ensure that the rows touched by the query cannot be updated or deleted by a concurrent session until the current transaction completes (whether it is committed or rolled back).
- SP2: read uncommitted permits repeatable reads

**NEW QUESTION 91**

- (Exam Topic 2)

You plan to deploy SQL Server 2012. You must create two tables named Table 1 and Table 2 that will have the following specifications:

Table1 will contain a date column named Column1 that will contain a null value approximately 80 percent of the time.

Table2 will contain a column named Column2 that is the product of two other columns in Table2. Both Table1 and Table2 will contain more than 1 million rows.

You need to recommend which options must be defined for the columns. The solution must minimize the storage requirements for the tables.

Which options should you recommend? To answer, drag the appropriate options to the correct column in the answer area.

Options		Answer Area
Sparse	Column1	Option
Computed	Column2	Option
Persisted computed		

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Column1 – Sparse; Column2 - Computed

- Sparse columns are ordinary columns that have an optimized storage for null values. Sparse columns reduce the space requirements for null values at the cost of more overhead to retrieve nonnull values. Consider using sparse columns when the space saved is at least 20 percent to 40 percent.

- A Persisted column would be faster to retrieve.

- A computed column is computed from an expression that can use other columns in the same table. The expression can be a noncomputed column name, constant, function, and any combination of these connected by one or more operators. Unless otherwise specified, computed columns are virtual columns that are not physically stored in the table. Their values are recalculated every time they are referenced in a query. The Database Engine uses the PERSISTED keyword in the CREATE TABLE and ALTER TABLE statements to physically store computed columns in the table. Their values are updated when any columns that are part of their calculation change.

References:

<http://msdn.microsoft.com/en-us/library/cc280604.aspx> <http://msdn.microsoft.com/en-us/library/ms186241.aspx>

**NEW QUESTION 95**

- (Exam Topic 2)

You are creating a database that will store usernames and passwords for an application. You need to recommend a solution to store the passwords in the database.

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

- A. One-way encryption
- B. Transparent Data Encryption (TDE)
- C. Encrypting File System (EFS)
- D. Reversible encryption

**Answer:** B

**Explanation:**

Transparent Data Encryption (TDE) is a special case of encryption using a symmetric key. TDE encrypts an entire database using that symmetric key called the database encryption key. The database encryption key is protected by other keys or certificates which are protected either by the database master key or by an asymmetric key stored in an EKM module.

SQL Server provides the following mechanisms for encryption:

- Transact-SQL functions
- Asymmetric keys
- Symmetric keys
- Certificates
- Transparent Data Encryption

**NEW QUESTION 98**

- (Exam Topic 2)

You need to recommend the actions that are required to partition a table.

In which order should the four actions be performed? To answer, move the actions from the list of actions to the answer area and arrange them in the correct order.

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

References:

<http://technet.microsoft.com/en-us/library/ms188730.aspx>

**NEW QUESTION 100**

- (Exam Topic 2)

You are designing a database for a university.

The database will contain two tables named Classes and Enrollment that have the following specifications:

Classes will store brochures in the XPS format. The brochures must be structured in folders and must be accessible by using UNC paths.

Enrollment will store information about students and their classes. Performance must be enhanced for queries of the current enrollments.

You need to identify which SQL Server technology meets the specifications of each table. Which technologies should you identify? To answer, drag the appropriate technology to the correct table in the answer area.

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

The FileTable feature brings support for the Windows file namespace and compatibility with Windows applications to the file data stored in SQL Server. FileTable lets an application integrate its storage and data management components, and provides integrated SQL Server services including full-text search and semantic search-over unstructured data and metadata. In other words, you can store files and documents in special tables in SQL Server called FileTables, but access them from Windows applications as if they were stored in the file system, without making any changes to your client applications.

**NEW QUESTION 103**

- (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 <input type="text" value="Index Storage Type"/>
Nonaligned	IX_Account <input type="text" value="Index Storage Type"/>
	IX_Date <input type="text" value="Index Storage Type"/>
	IX_Amount <input type="text" value="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 107**

- (Exam Topic 2)

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

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

You need to ensure that the following requirements are met:

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

What should you do? Choose all that apply.

- A. Configure the database to use the SIMPLE recovery model.
- B. Create a DIFFERENTIAL database backup every 4 hours.
- C. Create a LOG backup every 5 minutes.
- D. Configure the database to use the FULL recovery model.
- E. Create a FULL database backup every 24 hours.
- F. Create a DIFFERENTIAL database backup every 24 hours.

**Answer:** BCDE

**Explanation:**

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

**NEW QUESTION 109**

- (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:**

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

**NEW QUESTION 114**

- (Exam Topic 2)

Overview

Application Overview

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

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

The new version will use SQL Server 2014.

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

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

Tables

The current database schema contains a table named OrderDetails.

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

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

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

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

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

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

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

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

**Customer Problems Installation Issues**

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

**Index Fragmentation Issues**

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

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

**Backup Issues**

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

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

**Missing Data Issues**

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

**Query Performance Issues**

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

**Design Requirements**

**File Storage Requirements**

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

**Data Recovery Requirements**

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

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

**Concurrency Requirements**

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

You need to recommend a solution that meets the data recovery requirement. What should you include in the recommendation?

- A. A differential backup
- B. A transaction log backup
- C. Snapshot isolation
- D. A database snapshot

**Answer: D**

**NEW QUESTION 118**

- (Exam Topic 2)

Your company has offices in Seattle and Montreal.

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

You plan to implement an AlwaysOn availability group on both servers.

You need to recommend a failover type for the availability group. What should you recommend?

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

D. Asynchronous manual failover

**Answer:** D

#### NEW QUESTION 122

- (Exam Topic 2)

You have a server that has SQL Server 2014 installed. The server contains 100 user databases.

You need to recommend a backup solution for the user databases. The solution must meet the following requirements:

Perform a transaction log backup every hour.

Perform a full backup of each database every week.

Perform a differential backup of each database every day.

Ensure that new user databases are added automatically to the backup solution.

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

A. A maintenance plan

B. SQL Server Agent jobs

C. Policy-Based Management

D. A Data Definition Language (DDL) trigger

**Answer:** A

#### Explanation:

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

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

#### NEW QUESTION 126

- (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 130

- (Exam Topic 2)

You administer a Microsoft SQL Server 2016 database named Contoso on a server named Server01.

You need to write messages to the Application Log when users are added to or removed from a fixed server role in Server01.

What should you create?

A. a Database Audit Specification

B. a Policy

C. an Alert

D. a SQL Profiler Trace

E. a Resource Pool

F. an Extended Event session

G. a Server Audit Specification

**Answer:** G

#### Explanation:

The SQL Server Audit feature enables you to audit server-level and database-level groups of events and individual events.

Audits can have the following categories of actions:

Server-level. These actions include server operations, such as management changes and logon and logoff operations.

Database-level. These actions encompass data manipulation languages (DML) and data definition language (DDL) operations.

Audit-level. These actions include actions in the auditing process.

References:

[http://technet.microsoft.com/en-us/library/cc280663\(v=sql.105\).aspx](http://technet.microsoft.com/en-us/library/cc280663(v=sql.105).aspx)

#### NEW QUESTION 134

- (Exam Topic 2)

You administer a Microsoft SQL Server database named Sales. The database is 3 terabytes in size.

The Sales database is configured as shown in the following table. You discover that all files except Sales\_2.ndf are corrupt.

You need to recover the corrupted data in the minimum amount of time. What should you do?

Filegroup	File
PRIMARY	<ul style="list-style-type: none"> <li>• Sales.mdf</li> </ul>
XACTIONS	<ul style="list-style-type: none"> <li>• Sales_1.ndf</li> <li>• Sales_2.ndf</li> <li>• Sales_3.ndf</li> </ul>
ARCHIVES	<ul style="list-style-type: none"> <li>• SalesArch_1.ndf</li> <li>• SalesArch_2.ndf</li> </ul>

- A. Perform a restore from a full backup.
- B. Perform a transaction log restore.
- C. Perform a file restore.
- D. Perform a filegroup restore.

**Answer:** A

**Explanation:**

References:

<http://technet.microsoft.com/en-us/library/ms187048.aspx>

**NEW QUESTION 135**

- (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 attempt to restore a database on a new SQL Server instance and receive the following error message: "Msg 33111, Level 16, State 3, Line 2

Cannot find server certificate with thumbprint '0x7315277C70764B1F252DC7A5101F6F66EFB1069D.'" You need to ensure that you can restore the database successfully.

Solution: You add the backup set password to the restore command. Does this meet the goal?

- A. Yes
- B. No

**Answer:** B

**Explanation:**

The error is related to the certificate.

References: <https://dba.stackexchange.com/questions/3388/restore-encrypted-database-to-another-server?rq=1>

**NEW QUESTION 138**

- (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 SPUdateSalesInfo 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 SPUdate2.

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 SPUdateGenInfo combines data from multiple databases and generates millions of rows of data in the GenInfo table.

The GenInfo table is used for reports.

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

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

Current System

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

SQL Servers

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

Design Requirements

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

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

Direct access to database tables by developers or applications must be denied.  
The account used to generate reports must have restrictions on the hours when it is allowed to make a connection.  
Deadlocks must be analyzed with the use of Deadlock Graphs.  
In the event of a SQL Server failure, the databases must remain available.  
Software licensing and database storage costs must be minimized.  
Development effort must be minimized.  
The Tempdb databases must be monitored for insufficient free space.  
Failed authentication requests must be logged.  
Every time a new row is added to the ProductTypes table in the Products database, a user defined function that validates the row must be called before the row is added to the table.  
When SPUpdateSalesInfo queries data in the OrderTotals table the first time, the same rows must be returned along with any newly added rows when SPUpdateSalesInfo queries data in the OrderTotals table the second time.  
The performance of the SPUpdate2 stored procedure needs to be improved. Your solution must meet the design requirements.  
What should your solution include?

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

**Answer:** A

#### NEW QUESTION 141

- (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 need to address the requirements for disc usage monitoring for the SQL Servers. What should you do?

- A. You should configure disc quotas.
- B. You should configure a Dynamic Management View.
- C. You should configure alerts sent by the SQL Server Agent.
- D. You should configure a SQL Server Maintenance Plan.

**Answer:** B

#### NEW QUESTION 149

- (Exam Topic 3)

You work as a Database Administrator (DBA) for a company named ABC.com. The company uses a Microsoft SQL Server 2012 infrastructure. You have a database named CorpDB. You plan to create a stored procedure to access data in CorpDB. You need to ensure that the stored procedure supports dirty reads. What should you do?

- A. You should configure the stored procedure to use the READ UNCOMMITTED isolation level.
- B. You should configure the stored procedure to use the READ COMMITTED isolation level.
- C. You should configure the stored procedure to use the READ SERIALIZABLE isolation level.
- D. You should configure the stored procedure to use the REPEATABLE READ isolation level.
- E. You should configure the stored procedure to use the SNAPSHOT isolation level.

**Answer:** A

#### NEW QUESTION 151

- (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 153

- (Exam Topic 3)

You create an availability group that has replicas named HA/Server01 and HA/Server02. Currently, HA/Server01 is the primary replica.

You have multiple queries that read data and produce reports from the database. You need to offload the reporting workload to the secondary replica when HA/Server01 is the primary replica. What should you do?

- A. Set the Availability Mode property of HA/Server02 to Asynchronous commit.
- B. Set the Readable Secondary property of HA/Server02 to Read-intent only.
- C. Set the Connections in Primary Role property of HA/Server01 to Allow read/write connections.
- D. Set the Availability Mode property of HA/Server01 to Asynchronous commit.

**Answer: B**

**Explanation:**

References: <https://msdn.microsoft.com/en-us/library/jj542414.aspx>

**NEW QUESTION 156**

- (Exam Topic 3)

You are using dynamic management views to monitor an SQL Server server named SQL1. A database administrator named Dbal must monitor the health of SQL1.

You need to ensure that Dbal can access dynamic management views for SQL1.

The solution must use the principle of least privilege. Which permissions should you assign to Dbal?

- A. VIEW ANY DEFINITION
- B. VIEW SERVER STATE
- C. VIEW DEFINITION
- D. CONTROL SERVER

**Answer: B**

**Explanation:**

To query a dynamic management view or function requires SELECT permission on object and VIEW SERVER STATE or VIEW DATABASE STATE permission.

References: <https://msdn.microsoft.com/en-us/library/ms188754.aspx>

**NEW QUESTION 157**

- (Exam Topic 3)

Background Corporate Information

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

Physical Locations

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

Problem Statement

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

Configuration Windows Logins

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

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

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

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

Database

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

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

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

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

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

Column	Data type
SalesOrderDetailID	INT
ProductID	INT
SalePrice	SMALLMONEY
SaleQuantity	INT

**Database Issues**

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

**Requirements Database**

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

**Customer data access**

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

**Customer support data access**

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

**Sales force data access**

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

**Historical Data**

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

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

**Backups**

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

**Database Maintenance**

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

**Project milestones completed**

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

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

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

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

You need to create a job to automate some database maintenance tasks. Which code fragment should you use in each location in the command to complete one of the commands you will need to include in the job? To answer, drag the appropriate lines of code to the correct locations in the command. Each line of code 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.

Lines of Code	Answer Area	
STATISTICS	UPDATE	Production.Product
CONSTRAINT	Line of Code	Line of Code
MEMBER		
WITH FULLSCAN		
WITH SAMPLE 50 PERCENT		
WITH SAMPLE 1000 ROWS		
RECOMPUTE		
NOCOMPUTE		
NORECOMPUTE		

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Lines of Code	Answer Area	
STATISTICS	UPDAT	Production.Product
CONSTRAINT	STATISTICS	
MEMBER		
WITH FULLSCAN	WITH FULLSCAN	NORECOMPUTE
WITH SAMPLE 50 PERCENT		
WITH SAMPLE 1000 ROWS		
RECOMPUTE		
NOCOMPUTE		
NORECOMPUTE		

**NEW QUESTION 159**

- (Exam Topic 3)  
General Overview

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

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

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

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

A stored procedure named SPUpdateSalesInfo reads data in the OrderTotals table and modifies data in the SalesInfo table. The stored procedure then reads data in the OrderTotals table a second time and makes further changes to the information in the SalesInfo table. The Current\_Inventory database contains a large table named Inv\_Current. The Inv\_Current table has a clustered index for the primary key and a nonclustered index. The primary key column uses the identity property. The data in the Inv\_Current table is over 120GB in size. The tables in the Current\_Inventory database are accessed by multiple queries in the Sales database. Another table in the Current\_Inventory database contains a self-join with an unlimited number of hierarchies. This table is modified by a stored procedure named SPUpdate2. An external application named ExternalApp1 will periodically query the Current\_Inventory database to generate statistical information. The TempReporting database contains a single table named GenInfo. A stored procedure named SPUPdateGenInfo combines data from multiple databases and generates millions of rows of data in the GenInfo table. The GenInfo table is used for reports. When the information in GenInfo is generated, a reporting process reads data from the Inv\_Current table and queries information in the GenInfo table based on that data. The GenInfo table is deleted after the reporting process completes. The Products database contains tables named ProductNames and ProductTypes.

**Current System**

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

**SQL Servers**

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

**Design Requirements**

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

- Confidential information in the Current\_Inventory database that is accessed by ExternalApp1 must be securely stored.
- Direct access to database tables by developers or applications must be denied.
- The account used to generate reports must have restrictions on the hours when it is allowed to make a connection.
- Deadlocks must be analyzed with the use of Deadlock Graphs.
- In the event of a SQL Server failure, the databases must remain available.
- Software licensing and database storage costs must be minimized.
- Development effort must be minimized.
- The Tempdb databases must be monitored for insufficient free space.
- Failed authentication requests must be logged.

Every time a new row is added to the ProductTypes table in the Products database, a user defined function that validates the row must be called before the row is added to the table. When SPUpdateSalesInfo queries data in the OrderTotals table the first time, the same rows must be returned along with any newly added rows when SPUpdateSalesInfo queries data in the OrderTotals table the second time. The performance of the SPUpdate2 stored procedure needs to be improved. Your solution must meet the design requirements. What should your solution include?

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

**Answer:** A

#### **NEW QUESTION 160**

- (Exam Topic 3)

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution. After you answer a question in this sections, you will NOT be able to return to it. As a result, these questions will not appear in the review screen. You have a server named Server1 that has Microsoft SQL Server installed. Server1 has SQL Server Audit configured to send audit even records to a file. You need to ensure that a database user named User1 can review the audit data. Solution: You grant the VIEW ANY DEFINITION permission to User1. Does this meet the goal?

- A. Yes
- B. No

**Answer:** B

#### **Explanation:**

Each feature and command for SQL Server Audit has individual permission requirements. Unless otherwise specified, viewing catalog views requires a principal to have one of the following:

- The VIEW SERVER STATE permission.
- The VIEW AUDIT STATE permission (gives only the principal access to the sys.server\_audits catalog view).
- Membership in the sysadmin fixed server role.
- The CONTROL SERVER permission.
- The ALTER ANY AUDIT permission.

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

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

#### **NEW QUESTION 161**

- (Exam Topic 3)

You use a contained database named ContosoDb within a domain. You need to create a user who can log on to the ContosoDb database. You also need to ensure that you can port the database to different database servers within the domain without additional user account configurations. Which type of user should you create?

- A. User mapped to a certificate
- B. SQL user without login
- C. Domain user
- D. SQL user with login

Answer: C

**NEW QUESTION 162**

- (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 load the data in a heap table. Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

**Explanation:**

It is better to use a clustered index.

Columnstore indexes are the standard for storing and querying large data warehousing fact tables. It uses column-based data storage and query processing to achieve up to 10x query performance gains in your data warehouse over traditional row-oriented storage, and up to 10x data compression over the uncompressed data size.

In SQL Server, rowstore refers to table where the underlying data storage format is a heap, a clustered index, or a memory-optimized table.

References: <https://docs.microsoft.com/en-us/sql/relational-databases/indexes/columnstore-indexes-overview>

**NEW QUESTION 167**

- (Exam Topic 3)

Overview

General Overview

ADatum Corporation has offices in Miami and Montreal.

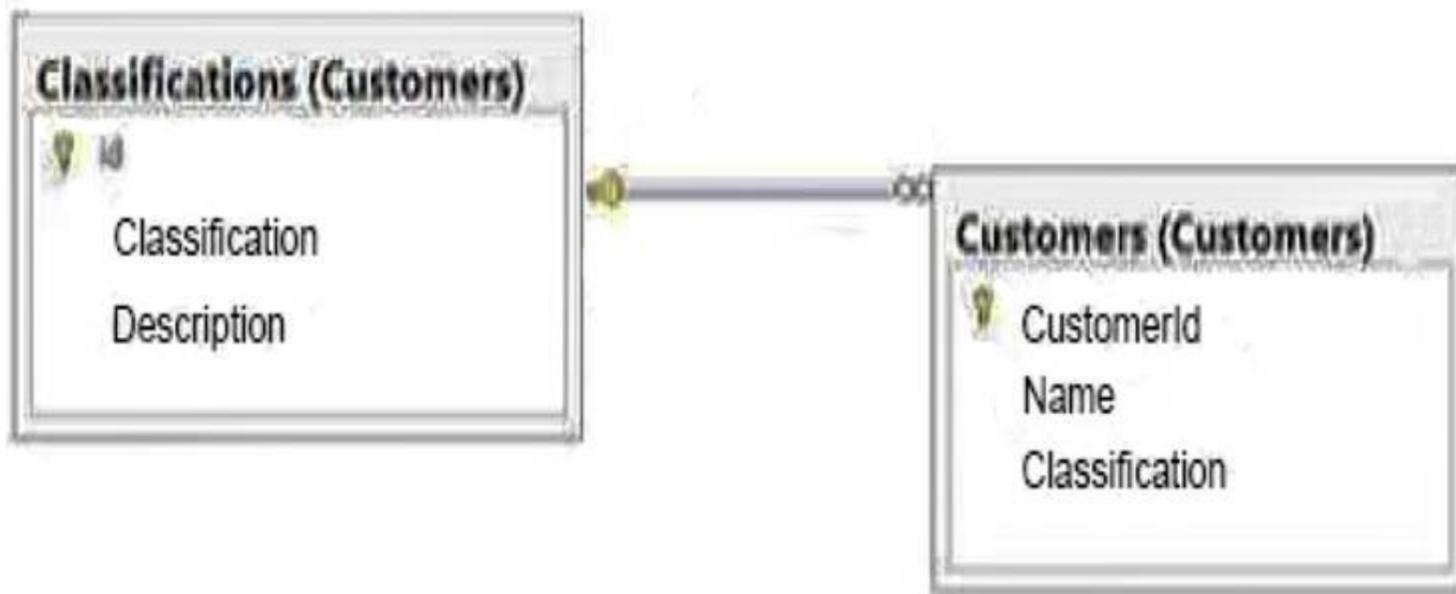
The network contains a single Active Directory forest named adatum.com. The offices connect to each other by using a WAN link that has 5-ms latency. A. Datum standardizes its database platform by using SQL Server 2014 Enterprise edition.

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 that meets the data recovery requirement. What should you include in the recommendation?

- A. A database snapshot
- B. A transaction log backup
- C. Snapshot isolation
- D. A differential backup

**Answer: D**

#### Explanation:

References:

<http://stackoverflow.com/questions/5299812/alternatives-to-snapshot-functionality-sql-serverstandard>

#### NEW QUESTION 170

- (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 172

- (Exam Topic 3)

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

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

You have a database that includes a table named Candidate.

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

Solution: You run the following query:

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

Does this meet the goal?

- A. Yes
- B. No

**Answer: B**

#### Explanation:

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

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

**NEW QUESTION 176**

- (Exam Topic 3)

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. A service account has the required permissions to backup all databases.

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 service account that backs up the HRDB database the permission necessary to access the encryption key. Your solution must use 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 ServerState
- G. View Definition
- H. sysadmin

**Answer:** G

**Explanation:**

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 179**

- (Exam Topic 3)

You have a SQL Server 2014 environment That includes four servers. The servers are configured as shown in the following table.

Server name	SQL Server 2014 edition	SQL Server version	Details
Server1	Enterprise	SQL Server 2014	A production Online Transaction Processing (OLTP) server
Server2	Web	SQL Server 2014	A test server
Server3	Standard	SQL Server 2012	A production report server
Server4	Express	SQL Server 2008 R2	A witness server

You plan to configure Policy-Based Management to enforce the following rules:

On Server1, enable SQL Server password policies and enable the default trace.

On Server3, ensure that the names of user-defined stored procedures begin with the prefix "usp\_" and ensure that all databases use a casesensitive collation.

You need to recommend which server you must configure as a Central Management Server.

Which server should you recommend? (Each correct answer presents a complete solution. Choose all that apply.)

- A. Server1
- B. Server2
- C. Server3
- D. Seiver4

**Answer:** AC

**Explanation:**

Need Standard or Enterprise edition of SQL Server.

**NEW QUESTION 182**

- (Exam Topic 3)

Overview

General Overview

ADatum Corporation has offices in Miami and Montreal.

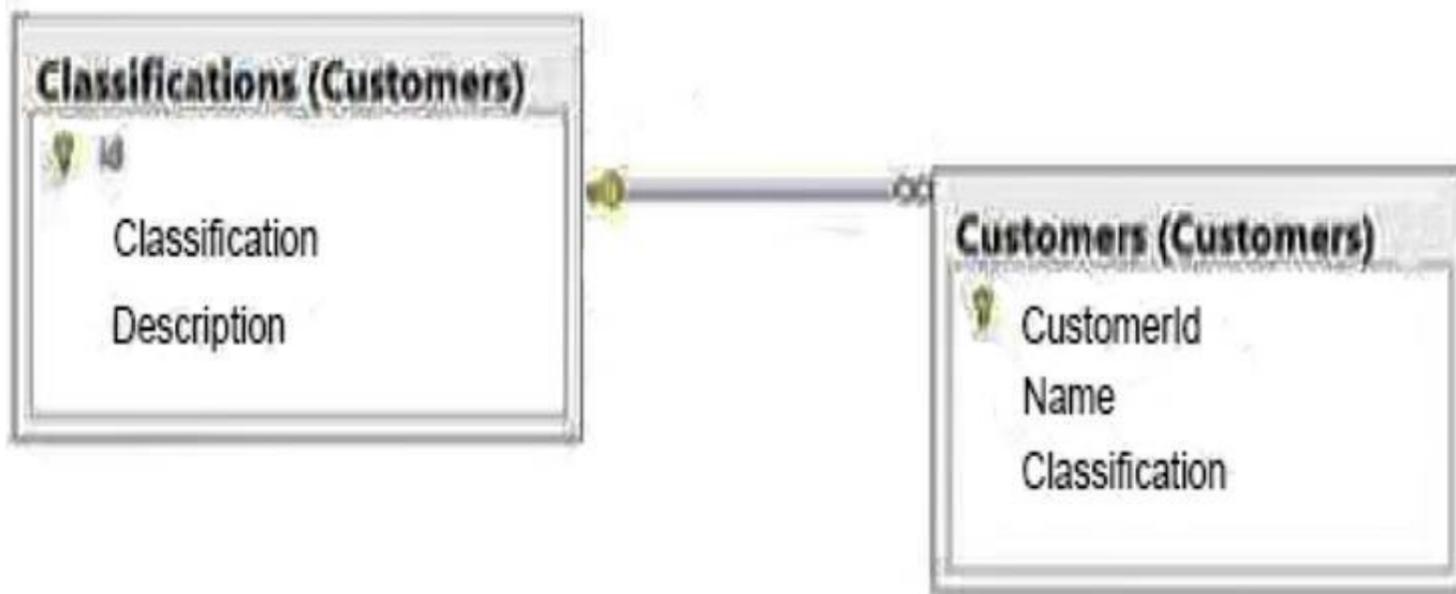
The network contains a single Active Directory forest named adatum.com. The offices connect to each other by using a WAN link that has 5-ms latency. A. Datum standardizes its database platform by using SQL Server 2014 Enterprise edition.

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\_2. What should you recommend?

- A. A database snapshot
- B. A table variable
- C. A temporary table
- D. Snapshot isolation

**Answer: C**

**Explanation:**

Scenario:

A stored procedure named USP\_2 is used to generate a product list.

USP\_2 takes several minutes to run due to locks on the tables the procedure accesses.

**NEW QUESTION 184**

- (Exam Topic 3)

Overview

General Overview

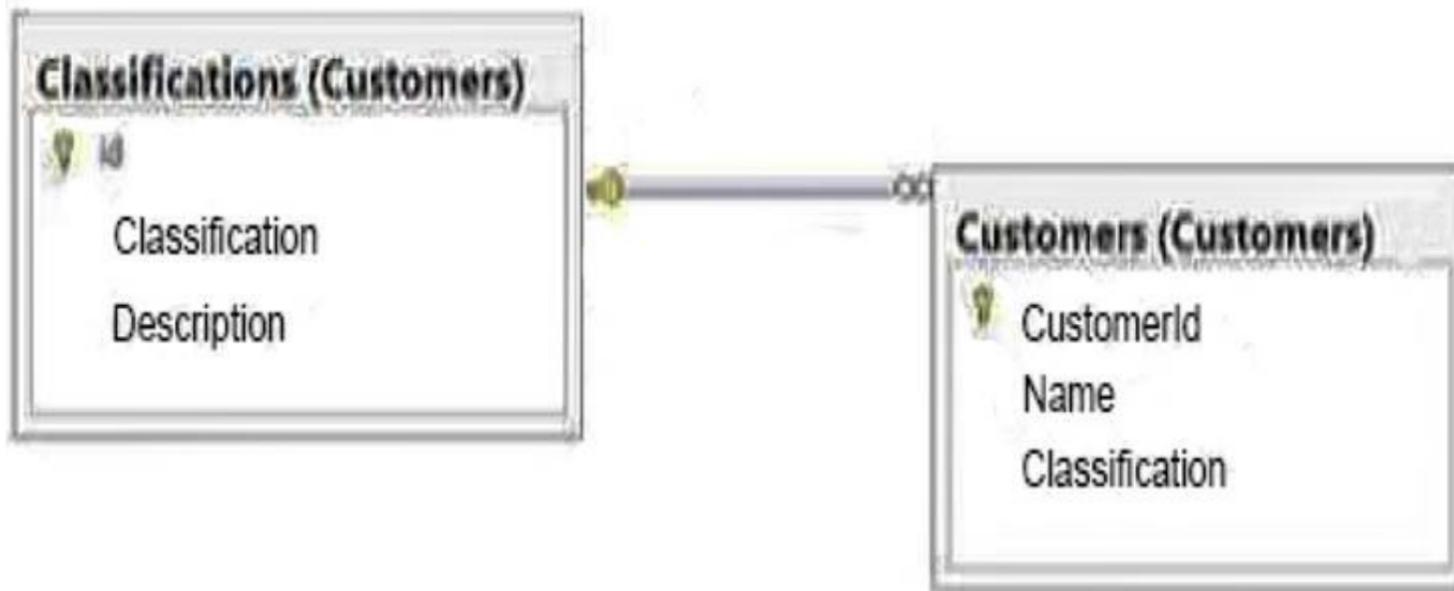
ADatum Corporation has offices in Miami and Montreal.

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 186**

- (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:.

What should you create in Azure to support the creation of the backups for DB1?

- A. an Azure Content Delivery Network (CDN) endpoint
- B. a Service Bus namespace
- C. a storage account
- D. a cloud service

**Answer: C**

**NEW QUESTION 191**

- (Exam Topic 3)

**Overview**

**General Overview**

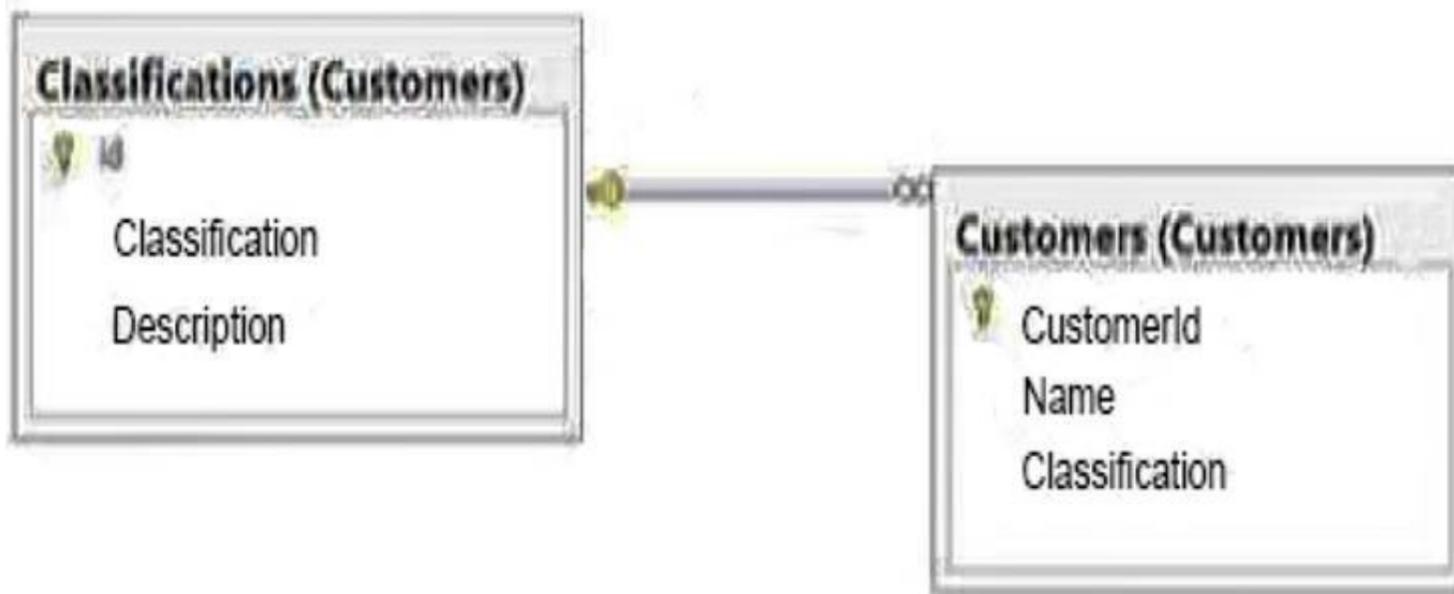
ADatum Corporation has offices in Miami and Montreal.

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 ensure that USP\_4 adheres to the security requirements. What should you include in the recommendation?

- A. Enable SQL Server Audit.
- B. Enable trace flags.
- C. Configure data manipulation language (DML) triggers.
- D. Enable C2 audit tracing.

**Answer:** A

**Explanation:**

Scenario:

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.

Beginning in SQL Server 2008 Enterprise, you can set up automatic auditing by using SQL Server Audit.

**NEW QUESTION 193**

.....

## Thank You for Trying Our Product

\* 100% Pass or Money Back

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

\* One year free update

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

\* Trusted by Millions

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

\* Shop Securely

All transactions are protected by VeriSign!

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

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