



# Microsoft

## Exam Questions 70-764

Administering a SQL Database Infrastructure (beta)

### NEW QUESTION 1

- (Exam Topic 1)

You administer a single server that contains a Microsoft SQL Server 2016 default instance on which several production databases have been deployed. You plan to install a new ticketing application that requires the deployment of a database on the server. The SQL login for this application requires sysadmin permissions. You need to ensure that the login for the ticketing application cannot access other production databases. What should you do?

- A. Use the SQL Server default instance and enable Contained Databases.
- B. Use the SQL Server default instance and configure a user-defined server rol
- C. Add the login for the ticketing application to this role.
- D. Install a new named SQL Server instance on the server.
- E. Install a new default SQL Server instance on the server.

**Answer:** C

#### Explanation:

SQL Server supports multiple instances of SQL Server on a single server or processor, but only one instance can be the default instance. All others must be named instances. A computer can run multiple instances of SQL Server concurrently, and each instance runs independently of other instances.  
References: [https://msdn.microsoft.com/en-us/library/ms143531\(v=SQL.105\).aspx](https://msdn.microsoft.com/en-us/library/ms143531(v=SQL.105).aspx)

### NEW QUESTION 2

- (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 backu
- 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 backu
- G. Then, restore the latest log backup.
- H. Perform a page restore.
- 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. Restore the latest full backu
- 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 3**

- (Exam Topic 1)

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

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

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

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

Which permission should you grant?

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

**Answer:** A

**Explanation:**

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

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

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

**NEW QUESTION 4**

- (Exam Topic 1)

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

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

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

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

You need to optimize query performance.

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

- A. Yes
- B. No

**Answer:** B

**Explanation:**

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

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

**NEW QUESTION 5**

- (Exam Topic 1)

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

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

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

You need to perform the following actions:

- identify databases where the PAGE verity option is not set.
- Configure full page protection for the identified databases. Solution: You run the following Transact-SQL Statement:

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

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

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

Does the solution meet the goal?

- A. Yes
- B. No

**Answer:** A

#### NEW QUESTION 6

- (Exam Topic 1)

You administer a Microsoft SQL Server 2016 server.

When transaction logs grow, SQL Server must send an email message to the database administrators. You need to configure SQL Server to send the email messages.

What should you configure?

- A. SQL Mail
- B. An Extended Events session
- C. Alerts and operators in SQL Server Agent
- D. Policies under Policy-Based Management

**Answer:** C

#### Explanation:

Operators are aliases for people or groups that can receive electronic notification when jobs have completed or alerts have been raised. The SQL Server Agent service supports the notification of administrators through operators. Operators enable notification and monitoring capabilities of SQL Server Agent.

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/database-mail/configure-sql-server-agent-mail-to-use-d>

#### NEW QUESTION 7

- (Exam Topic 1)

You administer a Windows Azure SQL Database database named Inventory that contains a stored procedure named p\_AddInventory.

Users need to be able to SELECT from all tables in the database and execute the stored procedure. You need to grant only the necessary permissions.

What should you do?

- A. Grant EXECUTE permission on p\_AddInventory to all user
- B. Grant VIEW DEFINITION to all users.
- C. Grant EXECUTE permission on p\_AddInventory to all user
- D. Add all users to the db\_datawriter role.
- E. Add all users to the db\_owner role.
- F. Grant EXECUTE permission on p\_AddInventory to all user
- G. Add all users to the db\_datareader role.

**Answer:** D

#### NEW QUESTION 8

- (Exam Topic 1)

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

You need to track all SELECT statements issued in the Contoso database only by users in a role named Sales. What should you create?

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

**Answer:** F

#### NEW QUESTION 9

- (Exam Topic 1)

You are the database administrator of a Microsoft SQL Server instance. Developers are writing stored procedures to send emails using sp\_send\_dbmail. Database Mail is enabled.

You need to configure each account's profile security and meet the following requirements:

Account SMTP1\_Account must only be usable by logins that have been given explicit permissions to use the SMTP1\_profile.

Account SMTP2\_Account must only be usable by logins who are a member of the [DatabaseMailUserRole] role in msdb.

In the table below, identify the profile type that must be used for each account. NOTE: Make only one selection in each column.

## Answer Area

| Profile type    | SMTP1_Account         | SMTP2_Account         |
|-----------------|-----------------------|-----------------------|
| Private Profile | <input type="radio"/> | <input type="radio"/> |
| Public Profile  | <input type="radio"/> | <input type="radio"/> |
| Default Profile | <input type="radio"/> | <input type="radio"/> |

- A. Mastered  
B. Not Mastered

**Answer:** A

**Explanation:**

SMTP1\_Account1: Private Profile

When no profile\_name is specified, sp\_send\_dbmail uses the default private profile for the current user. I user does not have a default private profile, sp\_send\_dbmail uses the default public profile for the msdb database.

SMTP1\_Account2: Default Profile

Execute permissions for sp\_send\_dbmail default to all members of the DatabaseMailUser database role in the msdb database.

References:

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

### 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 master database. Does the solution meet the goal?

- A. Yes  
B. No

**Answer:** B

**Explanation:**

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

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

References:

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

### NEW QUESTION 10

- (Exam Topic 1)

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

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

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

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

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

**Answer:** G

**Explanation:**

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

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

**NEW QUESTION 13**

- (Exam Topic 1)

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

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

**Answer: C**

**NEW QUESTION 14**

- (Exam Topic 1)

You plan to migrate the db to azure.

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

What should you do?

- A. Use the Copy Database wizard
- B. Use the Database Transfer wizard
- C. Use the SQL Management Studio to deploy the db to azure
- D. Back up the databases from the local server and restore it to azure

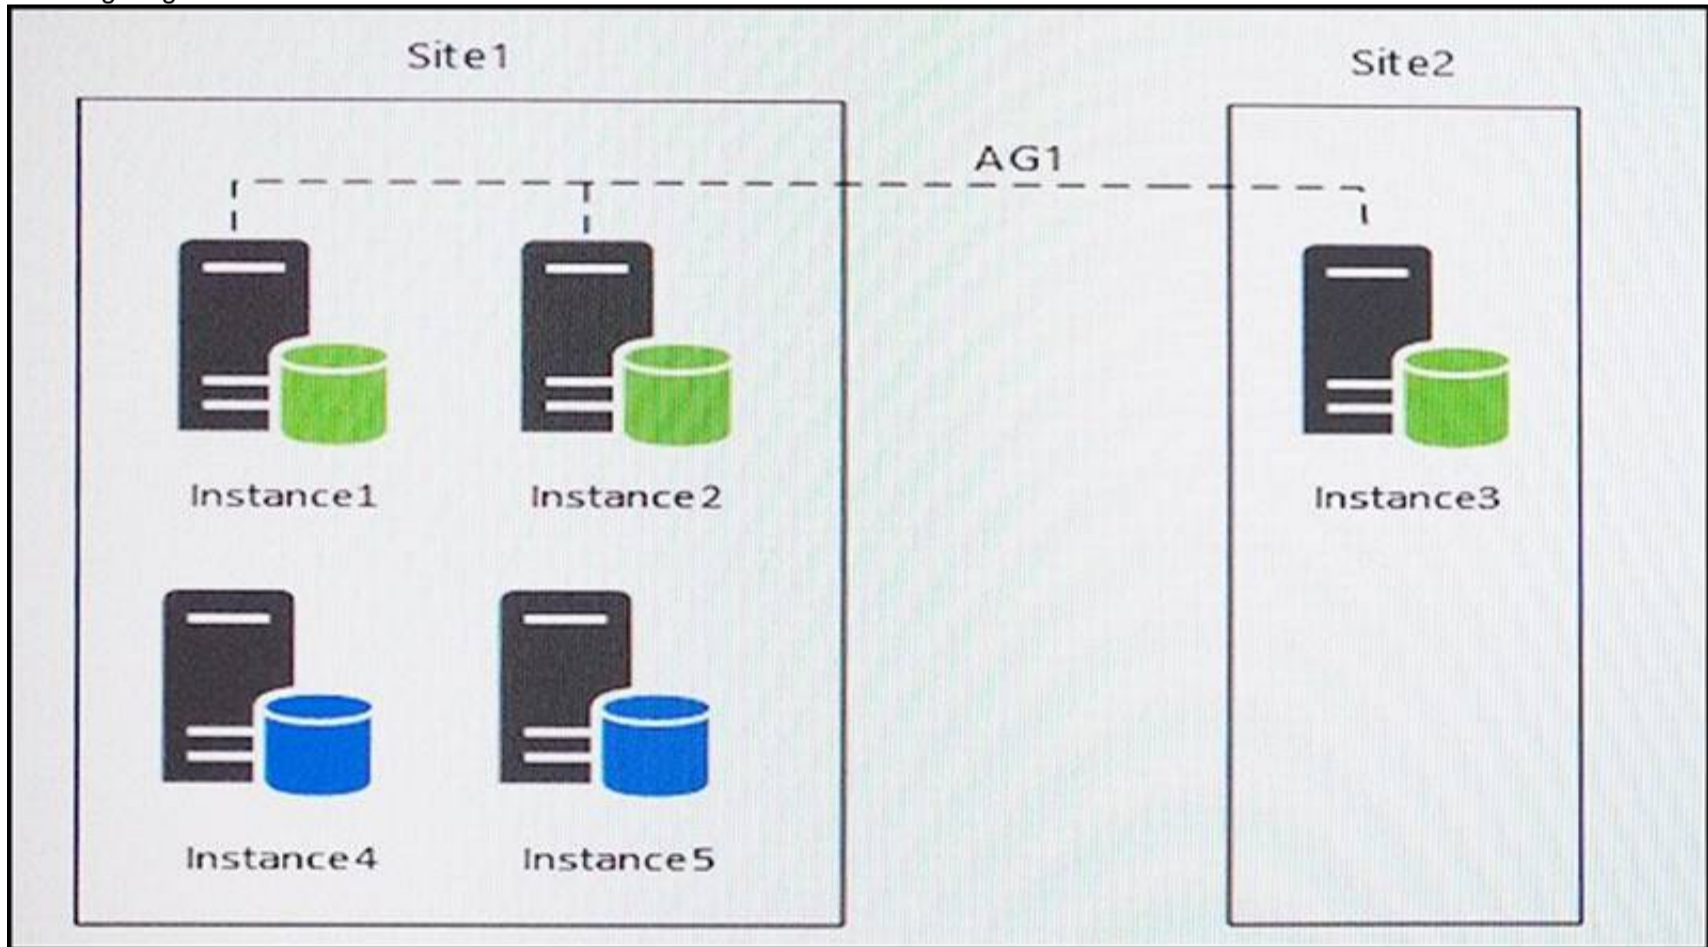
**Answer: CD**

**NEW QUESTION 17**

- (Exam Topic 1)

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

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



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

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

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

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

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

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

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

nodes in AG1.

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

You should minimize both the time required to restore

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

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

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

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

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

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

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

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

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

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

**Answer: D**

**Explanation:**

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

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

References:

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

**NEW QUESTION 22**

- (Exam Topic 1)

You administer a Microsoft SQL Server 2016 database instance.

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

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

What should you do?

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

**Answer: D**

**NEW QUESTION 24**

- (Exam Topic 1)

You are the administrator of a Microsoft SQL Server 2016 server. Some applications consume significant resources.

You need to manage the server workload by restricting resource-intensive applications. You need to dynamically limit resource consumption.

What should you do?

- A. Set up Service Broker to ensure that applications are not allowed to consume more than the specified amount of resources.
- B. Configure Resource Pools, Workload Groups, and Classifier Function, and then enable the Resource Governor.
- C. Configure Extended Events to monitor and restrict resource limits allowed by each application type.
- D. Create a new Plan Guide with a Scope Type of sql and define the resource limits for each application.

Answer: B

#### NEW QUESTION 28

- (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

Is the 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\_type is 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 each execution of a job. freq\_subday\_interval is 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 32

- (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 goals.

Some question sets might have more than one correct solution, while others might not have a correct solution.

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

You have a database named DB1 that is 640 GB and is updated frequently.

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

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

Solution: You increase the frequency of the transaction log backups to every 10 minutes. Does this meet the goal?

- A. Yes
- B. No

**Answer:** B

### NEW QUESTION 36

- (Exam Topic 1)

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

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

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

## Answer Area

| Action                   | Server instance  |
|--------------------------|--|
| Complete the backup job. | <div><div>▼</div><div>Primary server instance</div><div>Secondary server instance</div><div>Monitor server instance</div><div>Backup share file server</div></div> |
| Copy the backup job.     | <div><div>▼</div><div>Primary server instance</div><div>Secondary server instance</div><div>Monitor server instance</div><div>Backup share file server</div></div> |
| Restore the backup.      | <div><div>▼</div><div>Primary server instance</div><div>Secondary server instance</div><div>Monitor server instance</div><div>Backup share file server</div></div> |

- A. Mastered  
B. Not Mastered

**Answer:** A

**Explanation:**

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

Box 1: Primary server instance.

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

Box 2: Secondary server instance

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

Box 3: Secondary server instance.

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

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

**NEW QUESTION 39**

- (Exam Topic 1)

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

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

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

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

**Answer:** B

**Explanation:**

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

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

**NEW QUESTION 44**

- (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.  
One instance hosts a user database named HRDB. The database contains sensitive human resources data. You need to grant an auditor permission to view the SQL Server audit logs 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:** F

**Explanation:**  
Unless otherwise specified, viewing catalog views requires a principal to have one of the following:  
Membership in the sysadmin fixed server role.  
The CONTROL SERVER permission.  
The VIEW SERVER STATE permission.  
The ALTER ANY AUDIT permission.  
The VIEW AUDIT STATE permission (gives only the principal access to the sys.server\_audits catalog view).  
References: [https://technet.microsoft.com/en-us/library/cc280386\(v=sql.110\).aspx](https://technet.microsoft.com/en-us/library/cc280386(v=sql.110).aspx)

**NEW QUESTION 48**

- (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

WHERE ((sqlserver.data-base\_id>(4)) AND (severity>(15)))

(ACTION(sqlserver.client\_app\_name, sqlserver.data-base\_id,sqlserver.session\_id)

ALTER EVENT SESSION Contoso1 ON SERVER

)

GO

ADD EVENT sqlserver.error\_reported

ADD TARGET sqlserver.error\_reported

Answer Area

⏪

⏩

⏴

⏵

- A. Mastered
- B. Not Mastered

**Answer:** A

Your Partner of IT Exam

visit - <https://www.exambible.com>

**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.tsql_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 50**

- (Exam Topic 1)

You have a Microsoft SQL Server instance that hosts a database named DB1 that contains 800 gigabyte (GB) of data. The database is used 24 hours each day.

You implement indexes and set the value of the Auto Update Statistics option set to True.

Users report that queries take a long time to complete.

You need to identify statistics that have not been updated for a week for tables where more than 1,000 rows changed.

How should you complete the Transact-SQL statement? To answer, configure the appropriate Transact-SQL segments in the answer area.

## Answer Area

SELECT OBJECT\_NAME(id), name,

|               |   |
|---------------|---|
|               | ▼ |
| rowcnt        |   |
| stats_date    |   |
| rowmodctr     |   |
| stats_collect |   |

(id, indid),

|               |   |
|---------------|---|
|               | ▼ |
| rowcnt        |   |
| stats_date    |   |
| rowmodctr     |   |
| stats_collect |   |

FROM sys.sysindexes

WHERE

|               |   |
|---------------|---|
|               | ▼ |
| rowmodctr     |   |
| stats_collect |   |
| stats_date    |   |
| rowcnt        |   |

(id, indid) <= DATEADD(DAY, -7, GETDATE())

AND

|               |   |
|---------------|---|
|               | ▼ |
| stats_collect |   |
| rowmodctr     |   |
| stats_date    |   |
| rowcnt        |   |

> 1000

AND id IN (SELECT object\_id FROM sys.tables)

- A. Mastered  
B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: stats\_date See example below. Box 2: rowmodctr See examplebelow. Box 3: stats\_date

You need to identify statistics that have not been updated for a week. Box 4: rowmodctr

You need to identify that more than 1,000 rows changed.

Rowmodctr counts the total number of inserted, deleted, or updated rows since the last time statistics were updated for the table.

Example: We will query every statistics object which was not updated in the last day and has rows modified since the last update. We will use the rowmodctr field of sys.sysindexes because it shows how many rows were inserted, updated or deleted since the last update occurred. Please note that it is not always 100% accurate in SQL Server 2005 and later, but it can be used to check if any rows were modified.

--Get the list of outdated statistics

```
SELECT OBJECT_NAME(id),name,STATS_DATE(id, indid),rowmodctr FROM sys.sysindexes
```

```
WHERE STATS_DATE (id, indid)<=DATEADD(DAY,-1,GETDATE())
AND rowmodctr>0
AND id IN (SELECT object_id FROM sys.tables) GO
After collecting this information, we can decide which statistics require an update.
```

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/system-compatibility-views/sys-sysindexes-transact-sq>

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

#### NEW QUESTION 51

- (Exam Topic 1)

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

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

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

Drive D has 1 terabyte of free disk space.

Drive E has 5 GB of free disk space.

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

Five percent of the existing data is modified each day.

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

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

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

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

failover to this secondary database.

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

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

**Answer:** K

#### Explanation:

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

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

#### NEW QUESTION 55

- (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 60

- (Exam Topic 1)

You administer all the deployments of Microsoft SQL Server 2016 in your company.

You need to ensure that an OLTP database that uses a storage area network (SAN) remains available if any of the servers fail.

You also need to minimize the amount of storage used by the database. Which configuration should you use?

- A. Two servers configured in different data centers SQL Server Availability Group configured in Synchronous-Commit Availability Mode One server configured as an Active Secondary
- B. SQL Server that includes an application database configured to perform transactional replication
- C. Two servers configured in the same data center SQL Server Availability Group configured in AsynchronousCommit Availability Mode One server configured as an Active Secondary

- D. Two servers configured in different data centers SQL Server Availability Group configured in AsynchronousCommit Availability Mode
- E. Two servers configured in the same data center A primary server configured to perform log-shipping every 10 minutes A backup server configured as a warm standby
- F. Two servers configured on the same subnet SQL Server Availability Group configured in Synchronous-Commit Availability Mode
- G. SQL Server that includes an application database configured to perform snapshot replication
- H. Two servers configured in a Windows Failover Cluster in the same data center SQL Server configured as a clustered instance

**Answer:** H

**Explanation:**

A Windows Server Failover Cluster (WSFC) is a group of independent servers that work together to increase the availability of applications and services. SQL Server takes advantage of WSFC services and capabilities to support Always On availability groups and SQL Server Failover Cluster Instances.

References:

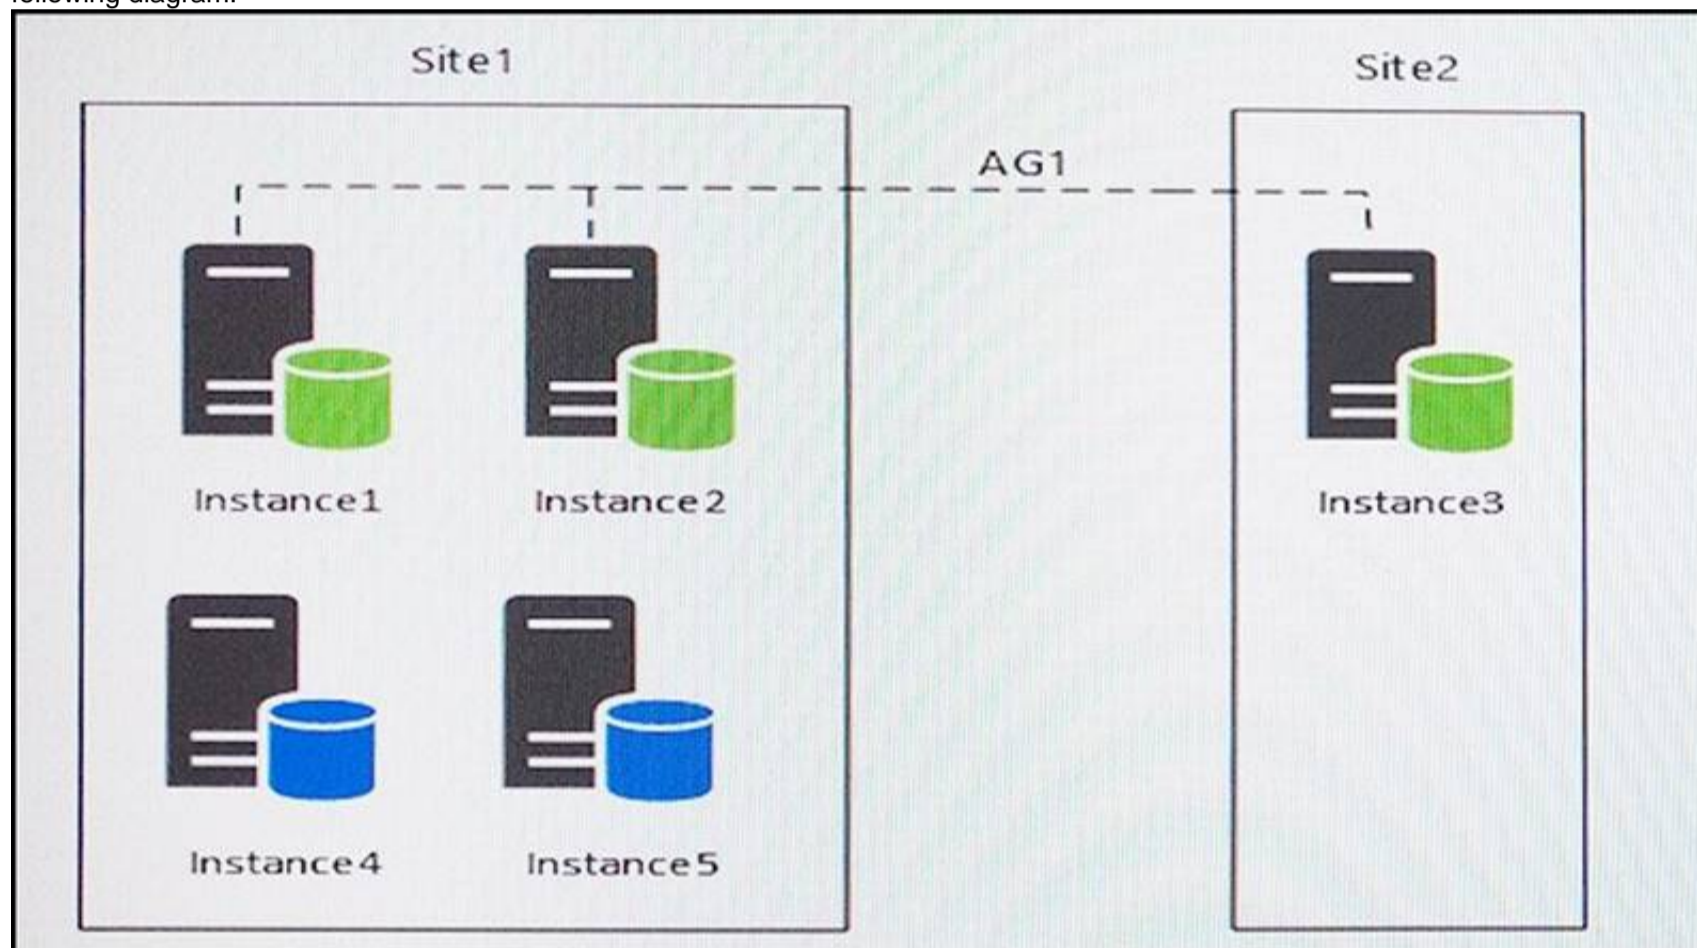
<https://docs.microsoft.com/en-us/sql/sql-server/failover-clusters/windows/windows-server-failover-clustering-ws>

**NEW QUESTION 65**

- (Exam Topic 1)

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

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



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

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

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

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

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

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

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

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

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

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

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

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

Reporting system: This solution accesses data in DB1 with a login that is mapped to a database user that is a member of the db\_datareader role. The user has EXECUTE permissions on the database. Queries make no changes to the data. The queries must be load balanced over variable read-only replicas.

Operations system: This solution accesses data in DB1 with a login that is mapped to a database user that is a member of the db\_datareader and db\_datawriter

roles. The user has EXECUTE permissions on the database. Queries from the operations system will perform both DDL and DML operations. The wait statistics monitoring requirements for the instances are described in the following table.

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

You need to analyze the wait type and statistics for specific instanced in the environment. Which object should you use to gather information about each instance? To answer, drag the appropriate objects to the correct instances. Each object may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.  
NOTE: Each correct selection is worth one point.

### Objects

Sys.dm\_os\_wait\_stats

Sys.dm\_exec\_connections

Sys.dm\_exec\_requests

Sys.dm\_exec\_procedure\_stats

Sys.dm\_exec\_sessions

Sys.dm\_exec\_query\_stats

Sys.dm\_exec\_query\_resource\_semaphores

Sys.dm\_exec\_session\_wait\_stats

### Answer Area

| Instance  | Object            |
|-----------|-------------------|
| Instance1 | <div>Object</div> |
| Instance4 | <div>Object</div> |
| Instance5 | <div>Object</div> |

- A. Mastered
- B. Not Mastered

Answer: A

**Explanation:**  
Instance 1: sys.dm\_exec\_query\_stats  
From Scenario: Instance1 requirement: Aggregate statistics since last server restart. sys.dm\_exec\_query\_stats returns aggregate performance statistics for cachedquery plans in SQL Server.  
Instance 4: sys.dm\_os\_wait\_stats  
sys.dm\_os\_wait\_statsreturns information about all the waits encountered by threads that executed. From Scenario: Instance4 requirement: Identify the most prominent wait types.

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

Instance 5:sys.dm\_exec\_session\_wait\_stats  
From Scenario: Instance5 requirement: Identify all wait types for queries currently running on the server. sys.dm\_exec\_session\_wait\_stats returns information about all the waits encountered by threads that executed for each session.

**NEW QUESTION 70**  
- (Exam Topic 1)

You administer a database that is used for reporting purposes. The database has a large fact table that contains three hundred million rows. The table includes a clustered columnstore index and a nonclustered index on the ProductID column. New rows are inserted into the table every day. Performance of queries that filter the Product ID column have degraded significantly. You need to improve the performance of the queries. Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

| Actions   | Answer Area |
|---|-------------|
| Drop the clustered columnstore index.                         |             |
| Create a nonclustered index on ProductID.                     |             |
| Drop and recreate the clustered columnstore index.            |             |
| Create a nonclustered columnstore index on ProductID.         |             |
| Recreate the clustered columnstore index using DROP EXISTING. |             |
| Create a clustered rowstore index on ProductID.               |             |
| Rebuild the clustered columnstore index.                      |             |

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Step 1: Drop the clustered columnstore index

Step 2: Create a clustered rowstore index on ProductID.

Rowstore indexes perform best on queries that seek into the data, searching for a particular value, or for queries on a small range of values. Use rowstore indexes with transactional workloads since they tend to require mostly table seeks instead of table scans.

Step 3: Create a nonclustered index on ProductID

**NEW QUESTION 73**

- (Exam Topic 2)

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

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

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

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

**Answer:** B

**Explanation:**

sys.dm\_os\_wait\_stats

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

Columns include: waiting\_tasks\_count

Number of waits on this wait type.

This counter is incremented at the start of each wait.

**NEW QUESTION 75**

- (Exam Topic 2)

Overview

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

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

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

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

Requirements Planned Changes

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

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

**Business Requirements**

You have the following requirements:

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

The encryption strategy must minimize changes to the databases and to the applications. You need to recommend a solution to improve the performance of usp\_UpdateInventory. The solution must minimize the amount of development effort. What should you include in the recommendation?

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

**Answer:** A

**Explanation:**

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

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

**NEW QUESTION 77**

- (Exam Topic 2)

You plan to create a database.

The database will be used by a Microsoft .NET application for a special event that will last for two days. During the event, data must be highly available. After the event, the database will be deleted. You need to recommend a solution to implement the database while minimizing costs. The solution must not affect any existing applications. What should you recommend? More than one answer choice may achieve the goal. Select the BEST answer.

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

**Answer:** D

**NEW QUESTION 81**

- (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 84**

- (Exam Topic 2)

You administer a Microsoft SQL Server 2016 server.

The MSSQLSERVER service uses a domain account named CONTOSO\SQLService. You plan to configure Instant File Initialization.

You need to ensure that Data File Autogrow operations use Instant File Initialization. What should you do? Choose all that apply.

- A. Restart the SQL Server Agent Service.
- B. Disable snapshot isolation.
- C. Restart the SQL Server Service.
- D. Add the CONTOSO\SQLService account to the Perform Volume Maintenance Tasks local security policy.
- E. Add the CONTOSO\SQLService account to the Server Operators fixed server role.
- F. Enable snapshot isolation.

**Answer:** CD

**Explanation:**

How To Enable Instant File Initialization

Open Local Security Policy and go to Local Policies → User Rights Assignment.

Double click Perform Volume Maintenance Tasks and add your SQL Server database engine service account.

Restart the SQL Server service using SQL Server Configuration Manager and this setting should now be enabled.

References:

<http://msdn.microsoft.com/en-us/library/ms175935.aspx>

**NEW QUESTION 89**

- (Exam Topic 2)

Overview

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

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

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

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

Requirements Planned Changes

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

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

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

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

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

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

Inventory will contain over 100 GB of data.

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

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

Business Requirements

You have the following requirements:

Costs for new licenses must be minimized.

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

Development effort must be minimized whenever possible.

The storage requirements for databases must be minimized.

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

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

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

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

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

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

- A. a SQL Server Agent alert
- B. a dynamic management view
- C. a maintenance plan
- D. an audit

**Answer:** B

**Explanation:**

Dynamic Management Views and Functions (Transact-SQL)

### NEW QUESTION 93

- (Exam Topic 2)

Overview

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

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

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

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

Requirements Planned Changes

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

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

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

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

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

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

Inventory will contain over 100 GB of data.

The Inventory table will have two indexes: a clustered index on the primary key and a nonclustered index.

The column that is used as the primary key will use the identity property.

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

Business Requirements

You have the following requirements:

Costs for new licenses must be minimized.

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

Development effort must be minimized whenever possible.

The storage requirements for databases must be minimized.

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

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

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

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

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

The encryption strategy must minimize changes to the databases and to the applications. You need to recommend an isolation level for usp\_UpdateOrderDetails.

Which isolation level should you recommend?

- A. Read committed
- B. Repeatable read
- C. Read uncommitted
- D. Serializable

**Answer: B**

**Explanation:**

- Scenario: 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.

- REPEATABLE READ Specifies that statements cannot read data that has been modified but not yet committed by other transactions and that no other transactions can modify data that has been read by the current transaction until the current transaction completes.

### NEW QUESTION 94

- (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 98

- (Exam Topic 2)

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

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

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

Answer: A

**Explanation:**

If we are going to encrypt credit card number for storage, then we should have Data Encryption Key(DEK) for encrypting the credit card number.

**NEW QUESTION 102**

- (Exam Topic 2)

Overview

Application Overview

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

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

The new version will use SQL Server 2014.

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

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

Tables

The current database schema contains a table named OrderDetails.

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

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

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

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

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

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

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

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

Customer Problems Installation Issues

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

Index Fragmentation Issues

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

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

Backup Issues

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

Users report that when they search product names, the search results exclude product names that contain

accents, unless the search string includes the accent. Missing Data Issues

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

Query Performance Issues

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

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

Design Requirements

File Storage Requirements

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

Data Recovery Requirements

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

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

Concurrency Requirements

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

What should you include in the recommendation?

- A. Windows logins
- B. Server roles
- C. Contained users
- D. Database roles

**Answer: C**

**Explanation:**

- Scenario: 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.

- Creating contained users enables the user to connect directly to the contained database. This is a very significant feature in high availability and disaster recovery scenarios such as in an AlwaysOn solution. If the users are contained users, in case of failover, people would be able to connect to the secondary without creating logins on the instance hosting the secondary. This provides an immediate benefit.

**NEW QUESTION 107**

- (Exam Topic 2)

Overview

General Overview

ADatum Corporation has offices in Miami and Montreal.

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

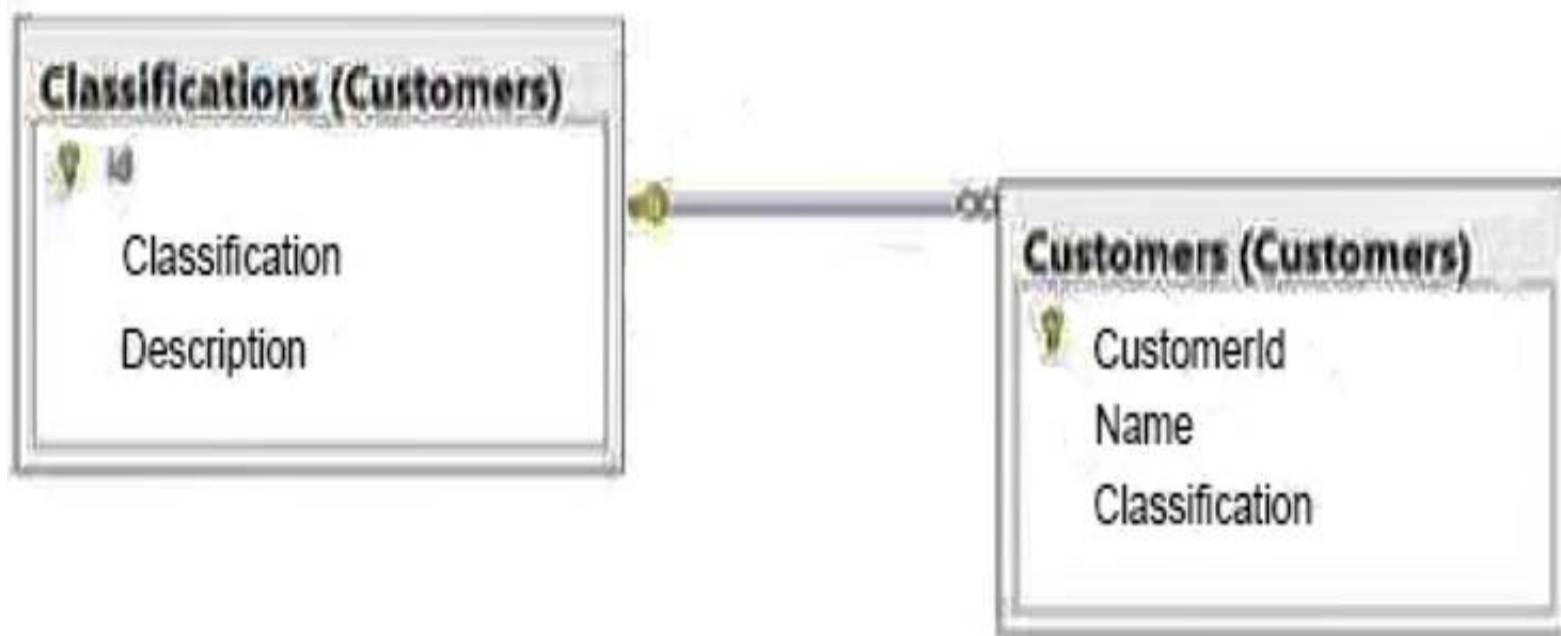
Databases

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

administrators have the same level of permissions on all of the servers and all of the databases.

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

Classifications (Customers)



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 disaster recovery strategy for the Inventory database. What should you include in the recommendation?

- A. Log shipping
- B. SQL Server Failover Clustering
- C. AlwaysOn availability groups
- D. Peer-to-peer replication

**Answer:** A

#### Explanation:

Scenario:

- You must be able to recover data from the Inventory database if a storage failure occurs. You have a Recovery Point Objective (RPO) of one hour.
- A. Datum Corporation has offices in Miami and Montreal.
- SQL Server Log shipping allows you to automatically send transaction log backups from a primary database on a primary server instance to one or more secondary databases on separate secondary server instances. The transaction log backups are applied to each of the secondary databases individually.

#### NEW QUESTION 108

- (Exam Topic 2)

You administer a Microsoft SQL Server 2016 instance.

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

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

**Answer:** C

#### Explanation:

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

References:

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

#### NEW QUESTION 112

- (Exam Topic 2)

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

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

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

The stored procedure must be designed to maximize concurrency.

What should you include in the design? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

| Actions  | Work Area |
|--|-----------|
| Raise an error in a catch block  |           |
| Commit the transaction in a finally block                                    |           |
| Read the @@ROWCOUNT system variable  |           |
| Perform the update in a try block  |           |
| Raise an error and roll back the transaction if the row count is less than 1 |           |
| Issue a SELECT statement to count the number of rows                         |           |
| Set the isolation level to serializable                                      |           |
| Begin an explicit transaction  |           |

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Note:

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

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

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

- Generate a constraint violation error.

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

**NEW QUESTION 116**

- (Exam Topic 2)

Overview

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

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

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

will use the application.

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

Requirements Planned Changes

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

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

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

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

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

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

Inventory will contain over 100 GB of data.

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

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

Business Requirements

You have the following requirements:

Costs for new licenses must be minimized.

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

Development effort must be minimized whenever possible.

The storage requirements for databases must be minimized.

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

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

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

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

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

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

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

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

**Answer: A**

**Explanation:**

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

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

Transparent Data Encryption (TDE)

## NEW QUESTION 119

- (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 120

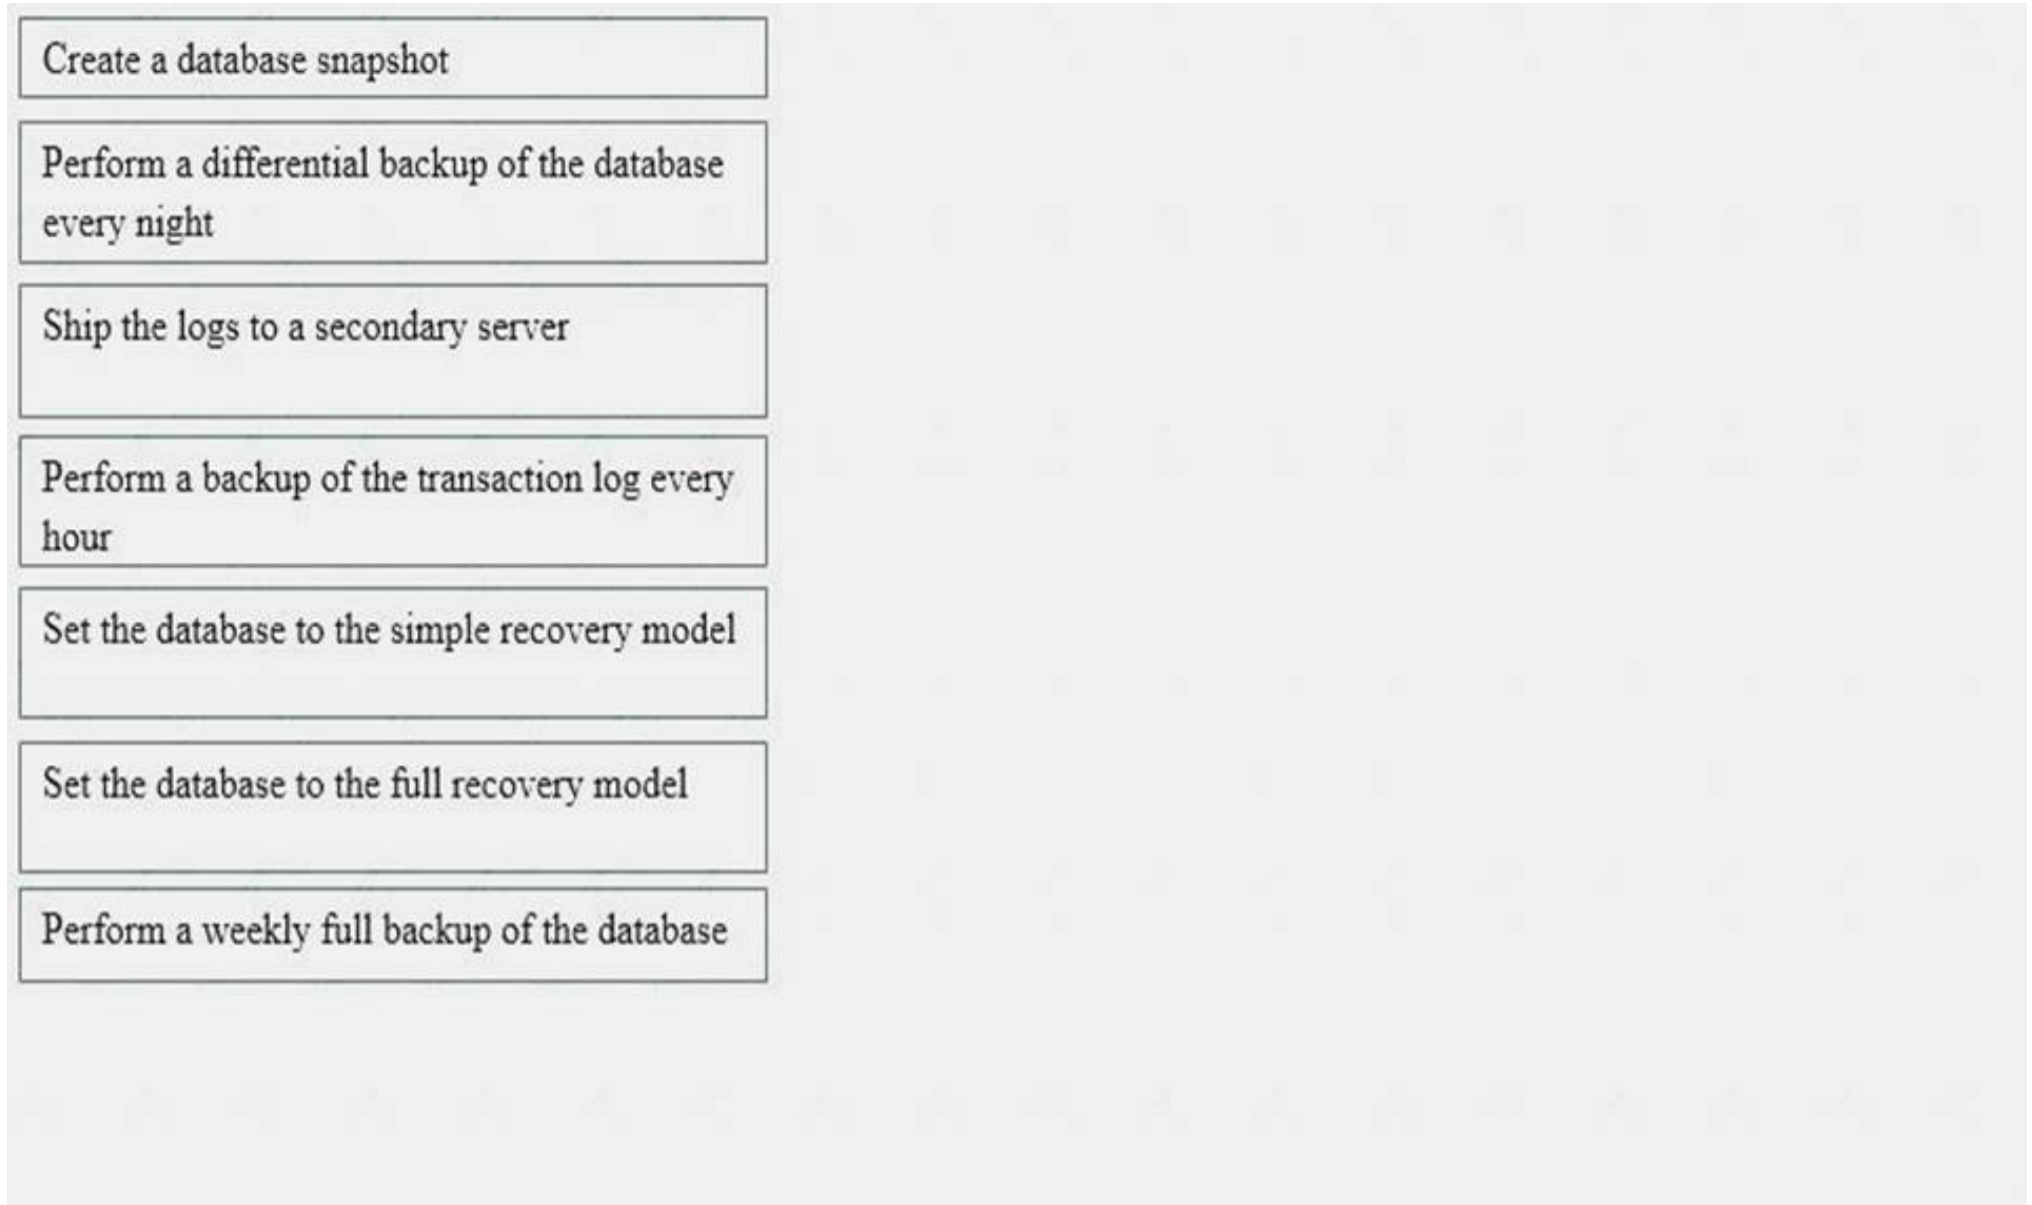
- (Exam Topic 2)

You need to recommend a backup process for an Online Transaction Processing (OLTP) database. The process must meet the following requirements:

Ensure that if a hardware failure occurs, you can bring the database online with a minimum amount of data loss.

Minimize the amount of administrative effort required to restore any lost data.

What should you include in the recommendation? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.



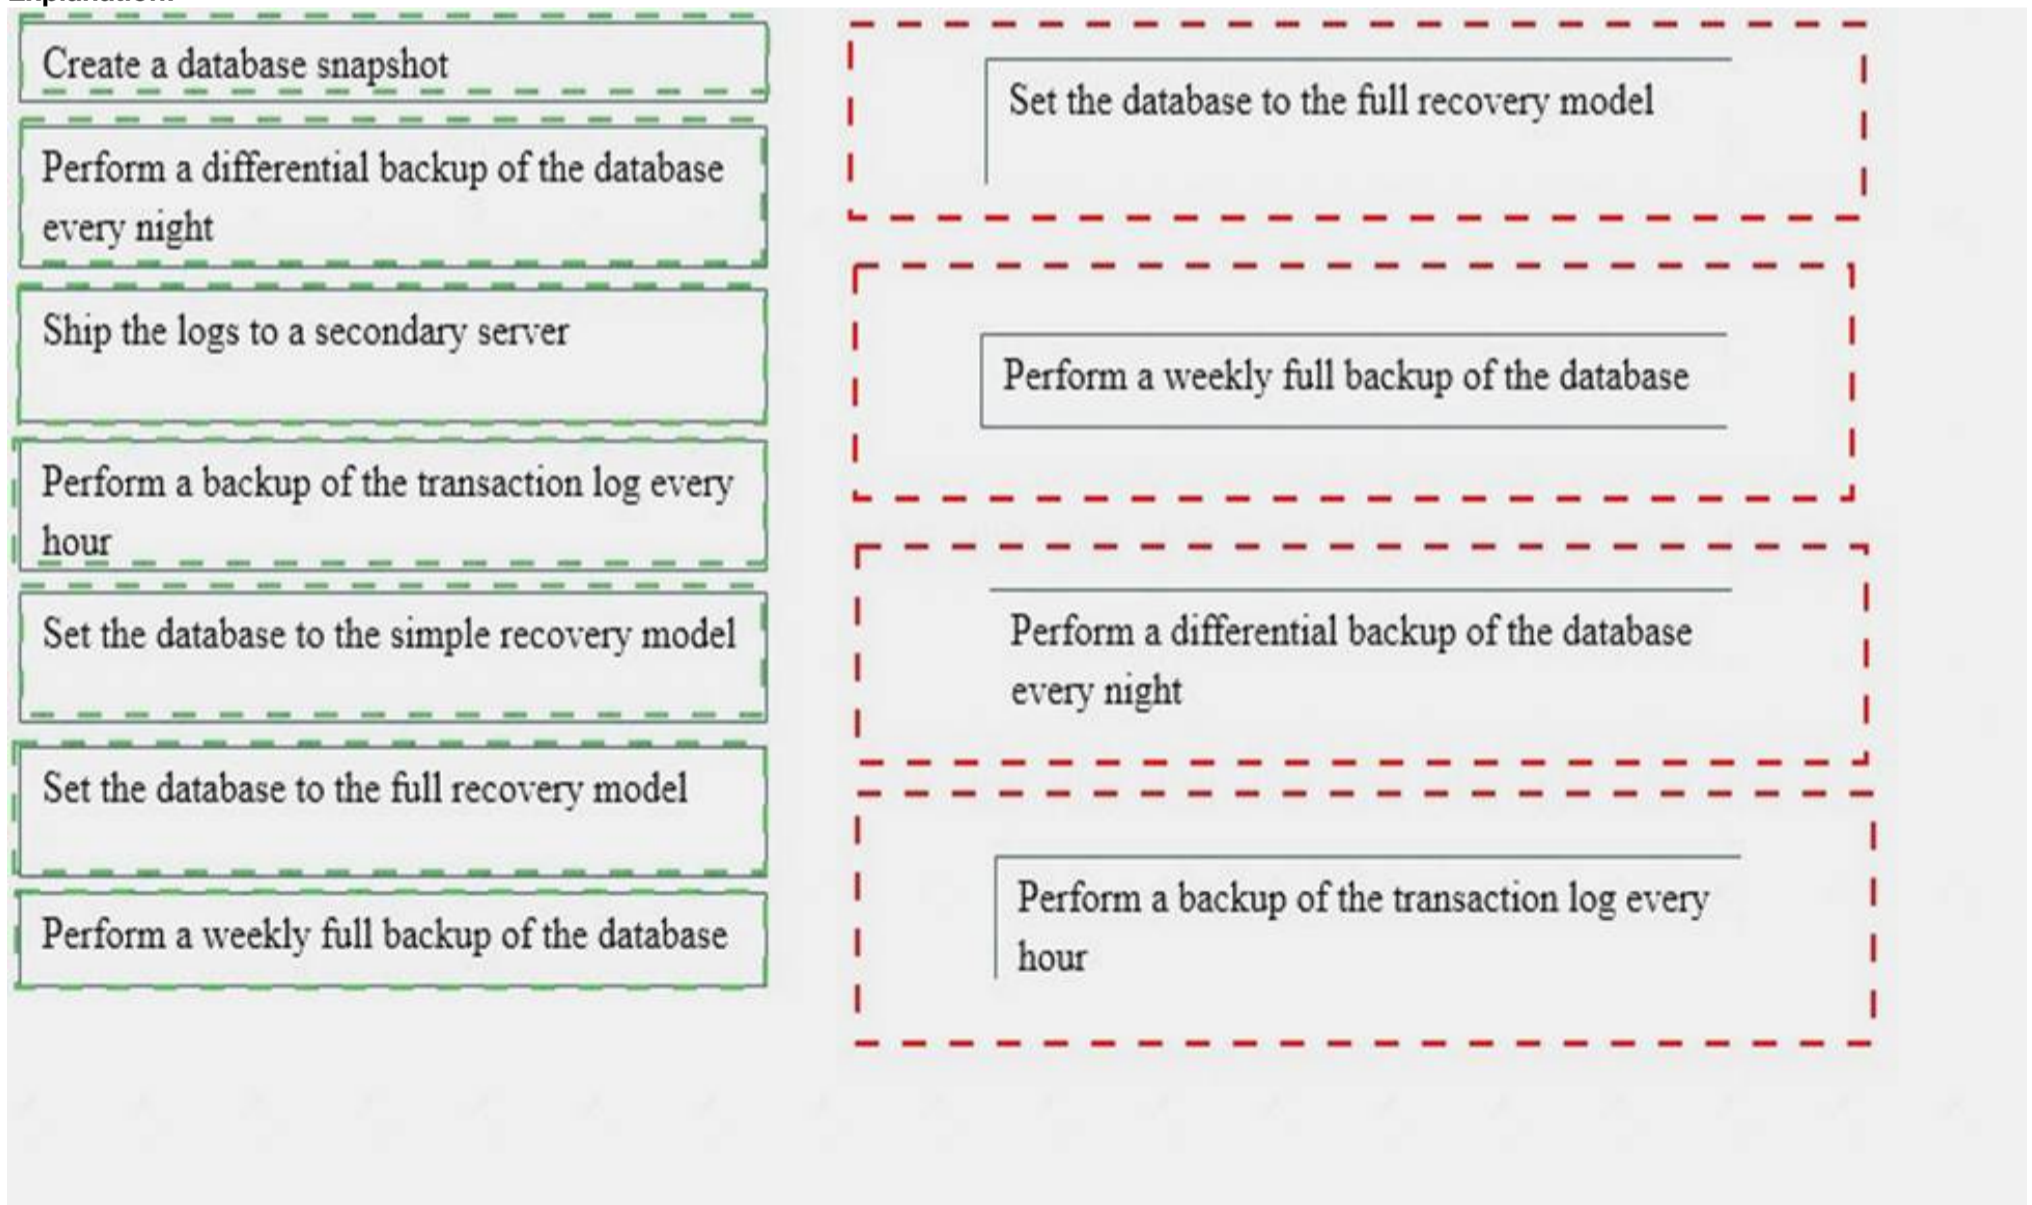
The screenshot shows a list of tasks in a SQL Server Enterprise console. The tasks are:

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

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**



The screenshot shows a list of tasks in a SQL Server Enterprise console. The tasks are:

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

A red dashed box highlights the task "Set the database to the full recovery model".

#### NEW QUESTION 124

- (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 129

- (Exam Topic 2)

Overview

Application Overview

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

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

The new version will use SQL Server 2014.

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

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

Tables

The current database schema contains a table named OrderDetails.

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

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

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

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

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

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

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

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

Customer Problems Installation Issues

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

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

Index Fragmentation Issues

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

Backup Issues

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

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

Missing Data Issues

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

#### Query Performance Issues

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

#### Design Requirements

##### File Storage Requirements

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

##### Data Recovery Requirements

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

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

##### Concurrency Requirements

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

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

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

**Answer:** BE

#### Explanation:

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

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

#### NEW QUESTION 133

- (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 134

- (Exam Topic 2)

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

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

You need to upgrade the database to SQL Server 2014. The solution must minimize downtime. Which upgrade steps should you recommend? To answer, move

the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

|                    |
|--------------------|
| Fail over          |
| Fail back          |
| Upgrade SQL1       |
| Upgrade SQL2       |
| Establish a mirror |
| Break the mirror   |

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

References:<https://docs.microsoft.com/en-us/sql/database-engine/database-mirroring/upgrading-mirrored-instanc>

**NEW QUESTION 137**

- (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 141**

- (Exam Topic 2)

You are the lead database administrator (DBA) of a Microsoft SQL Server 2016 environment. All DBAs are members of the DOMAIN\JrDBAs Active Directory group.

You grant DOMAIN\JrDBAs access to the SQL Server.

You need to create a server role named SpecialDBARole that can perform the following functions:

View all databases.

View the server state.

Assign GRANT, DENY, and REVOKE permissions on logins.

You need to add DOMAIN\JrDBAs to the server role.

You also need to provide the least level of privileges necessary.

Which SQL statement or statements should you use? Choose all that apply.

- A. CREATE SERVER ROLE [SpecialDBARole] AUTHORIZATION setupadmin;
- B. ALTER SERVER ROLE [SpecialDBARole] ADD MEMBER [DOMAIN\JrDBAs];
- C. CREATE SERVER ROLE [SpecialDBARole] AUTHORIZATION securityadmin;
- D. GRANT VIEW DEFINITION TO [SpecialDBARole];
- E. CREATE SERVER ROLE [SpecialDBARole] AUTHORIZATION serveradmin;
- F. GRANT VIEW SERVER STATE, VIEW ANY DATABASE TO [SpecialDBARole];

**Answer:** BCF

**NEW QUESTION 142**

- (Exam Topic 2)

Overview

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

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

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

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

**Requirements Planned Changes**

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

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

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

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

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

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

Inventory will contain over 100 GB of data.

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

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

**Business Requirements**

You have the following requirements:

Costs for new licenses must be minimized.

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

Development effort must be minimized whenever possible.

The storage requirements for databases must be minimized.

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

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

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

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

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

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

You need to recommend a solution to allow application users to perform UPDATE operations on the database tables. The solution must meet the business requirements.

What should you recommend?

- A. Create stored procedures that use EXECUTE AS clauses.
- B. Create a user-defined database role and add users to the role.
- C. Create functions that use EXECUTE AS clauses.
- D. Create a Policy-Based Management Policy.

**Answer:** A

**Explanation:**

- EXECUTE AS Clause (Transact-SQL)

In SQL Server you can define the execution context of the following user-defined modules: functions (except inline table-valued functions), procedures, queues, and triggers.

#### NEW QUESTION 146

- (Exam Topic 2)

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

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

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

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

**Answer:** A

#### NEW QUESTION 147

- (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

EXPLICIT

ISOLATION

READ UNCOMMITTED

ROLLBACK

SERIALIZABLE

SNAPSHOT

TABLOCK

TRANSACTION

Answer Area

SET

LEVEL

BEGIN

UPDATE SalesOrder

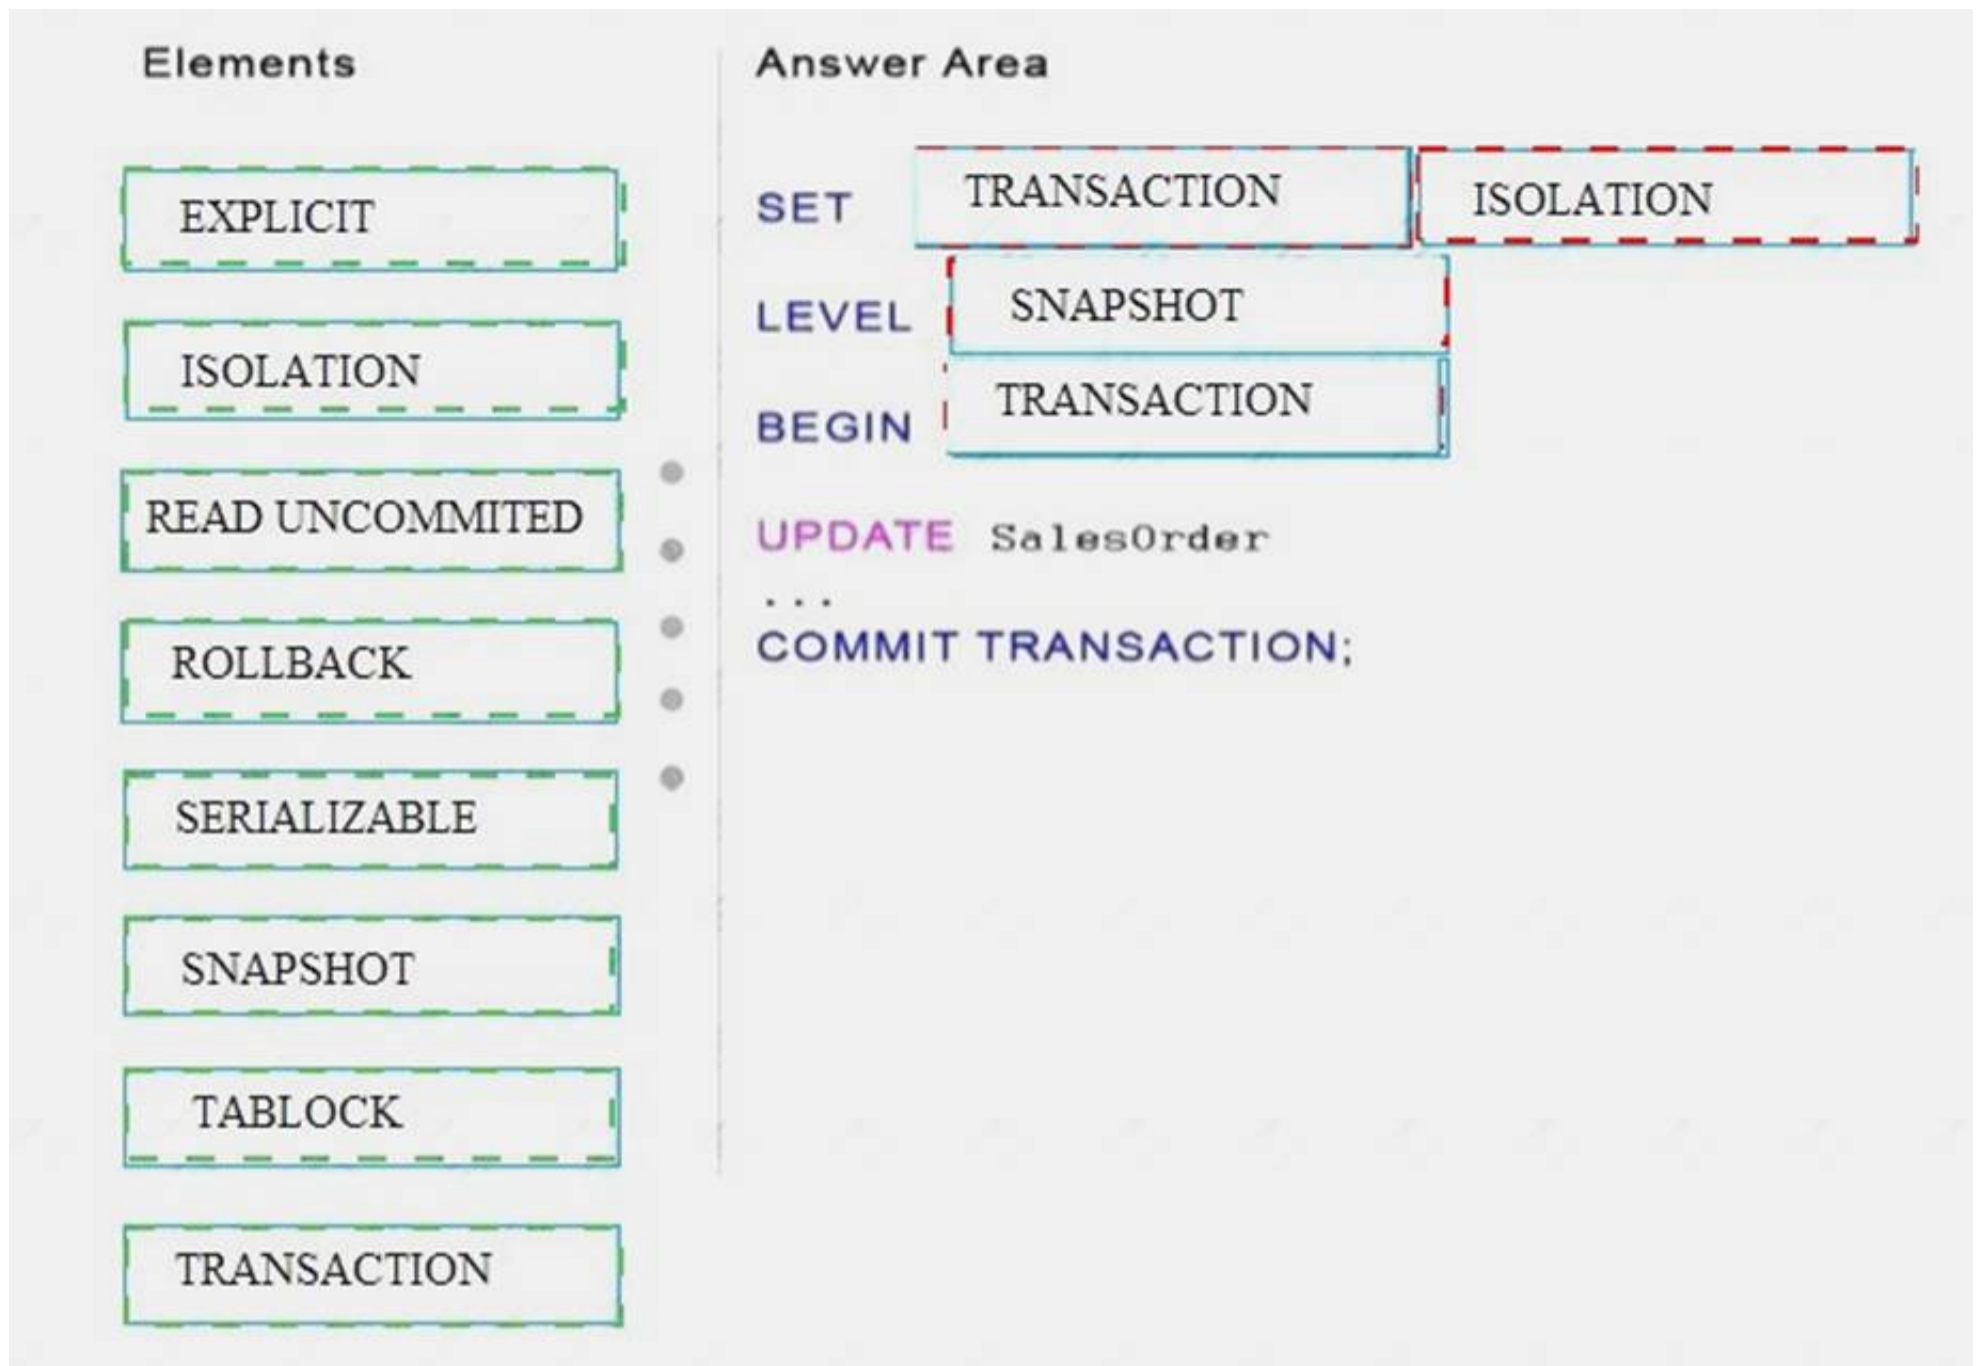
...

COMMIT TRANSACTION;

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:



#### NEW QUESTION 150

- (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.

During performance testing, you discover that database INSERT operations against the Inventory table are slow.

You need to recommend a solution to reduce the amount of time it takes to complete the INSERT operations. What should you recommend?

- A. Partition the nonclustered index.
- B. Partition the Inventory table.snapshot replication
- C. Create a column store index.Master Data Services

D. Drop the clustered index.change data capture

**Answer:** A

**Explanation:**

Scenario:

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.

**NEW QUESTION 154**

- (Exam Topic 2)

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

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

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

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

**Answer:** DE

**Explanation:**

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

References:

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

**NEW QUESTION 158**

- (Exam Topic 2)

You administer a Microsoft SQL Server 2016 database that contains a table named AccountTransaction. You discover that query performance on the table is poor due to fragmentation on the

IDX\_AccountTransaction\_AccountCode non-clustered index.

You need to defragment the index.

You also need to ensure that user queries are able to use the index during the defragmenting process. Which Transact-SQL batch should you use?

- A. ALTER INDEX IDX\_AccountTransaction\_AccountCode ONAccountTransaction.AccountCode REORGANIZE
- B. ALTER INDEX ALL ON AccountTransaction REBUILD
- C. ALTER INDEX IDX\_AccountTransaction\_AccountCode ONAccountTransaction.AccountCode REBUILD
- D. CREATE INDEX IDXAccountTransactionAccountCode ONAccountTransaction.AccountCode WITH DROP EXISTING

**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 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 relate to the certificate.

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

**NEW QUESTION 161**

- (Exam Topic 3)

Overview

General Overview

ADatum Corporation has offices in Miami and Montreal.

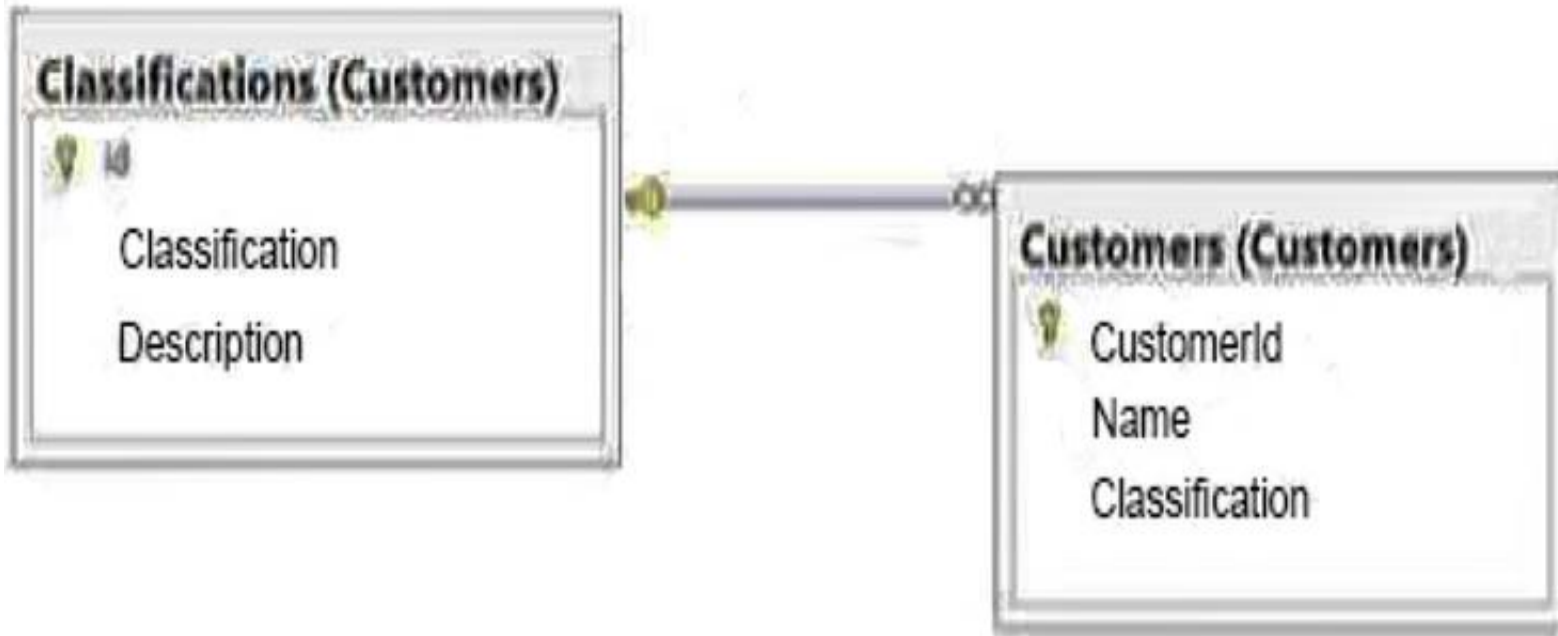
The network contains a single Active Directory forest named adatum.com. The offices connect to each other by using a WAN link that has 5-ms latency. A. Datum standardizes its database platform by using SQL Server 2014 Enterprise edition.

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. A view
- B. A temporary table
- C. A table variable
- D. A function

**Answer: A**

**Explanation:**

- 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.

#### NEW QUESTION 163

- (Exam Topic 3)

General Overview

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

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

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

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

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

**Current System**  
The current desktop application uses data stored in a SQL Server 2005 database named DesABCOppAppDB. 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 164

- (Exam Topic 3)

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

- Rebuild the indexes by using a new fill factor.
- Run a custom T-SQL command.
- Back up the databases.

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

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

**Answer:** C

#### Explanation:

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

#### NEW QUESTION 166

- (Exam Topic 3)

##### General Overview

You are the Senior Database Administrator (DBA) for a manufacturing company named Fairstone Manufacturing.

Fairstone Manufacturing is based in the New York area. The company has two offices: a main office in the city and a branch office just outside the city. The company has four factories where their products are manufactured. Two factories are in the New York area and the other two factories are in Washington.

##### Network Connectivity

The two offices are connected by a 10 Mbps dedicated WAN link. SQL Server Environment

The main office has four SQL Server 2012 Standard Edition servers named MainDB1, MainDB2, MainDB3 and MainDB4. The branch office has two SQL Server 2012 Standard Edition servers named BranchDB1 and BranchDB2. The main office has a Development department. All databases used by the Development department are hosted on MainDB3 and MainDB4. MainDB1 and MainDB2 host the following databases:

Products Manufacturing Sales

HR

Customers DailyReportsTemp

BranchDB1 and BranchDB2 host the same databases as MainDB1 and MainDB2. The DailyReportsTemp database is a temporary database that is recreated every day and used for reporting purposes.

One of the tables in the Customer database lists all the customers. Another table linked to the customers table contains a list of classifications for the customers.

The classifications are Hot, Warm and Cold based on the number of orders placed by the customers in the last year. The customers are classified according to the following criteria:

Hot - Over 100 orders placed in a year.

Warm - Between 50 and 100 orders placed in a year. Cold - Under 50 orders placed in a year.

Stored Procedures

Three tables in the Manufacturing database are modified by a stored procedure named ManProc1. A segment of code from ManProc1 is as follows:

```
CREATE PROCEDURE Manufacturing.ManProc1
AS
BEGIN TRANSACTION
UPDATE Manufacturing.Type ...
UPDATE Manufacturing.Version ...
UPDATE Manufacturing.Revision ...
COMMIT TRANSACTION
GO
```

The same three tables are also modified by a stored procedure named ManProc2. A segment of code from ManProc2 is as follows:

```
CREATE PROCEDURE Manufacturing.ManProc2
AS
BEGIN TRANSACTION
UPDATE Manufacturing.Revision ...
UPDATE Manufacturing.Type ...
UPDATE Manufacturing.Version ...
COMMIT TRANSACTION
GO
```

A product list in the Products database is updated using information from tables in the Manufacturing database by a stored procedure named ProductUpdateProc. Locks on tables in the Manufacturing database often cause ProductUpdateProc to take a long time to complete.

A list of manufacturing processes required to create each product is stored in tables in the Manufacturing database and updated by a stored procedure named ProcessUpdateProc. The ProcessUpdateProc stored procedure contains several UPDATE statements. The UPDATE statements are configured to be called in a specific order. The ProcessUpdateProc stored procedure continues to run in the event of a failure of one of the UPDATE statements.

This can cause inaccurate results in the manufacturing process list. Sales Director Statement

The Sales Director has made the following observations about the current database design:

The current customer classification system needs to be changed.

Currently the customers are classified by the number of orders placed in the last year.

This information is an unreliable guide as it does not take in to account the size of the orders.

I would suggest a trial run of a classification system based on the revenue generated by the orders placed in the last year.

We may add more than the current three classification types in future.

We should have a method of recording changes to the classifications.

IT Manager Statement

The IT Manager has listed the following requirements for the SQL Server and database environment:

We need to provide a group of users from the IT and Manufacturing departments the minimum administrative rights to view database information and server state for the Manufacturing database on MainDB1.

The Sales database takes too long to back up due to the large amount of historical sales order data in the database. We need to reduce the backup time for this database.

The DailyReportsTemp database takes four hours to back up. We need to be able to recover the DailyReportsTemp database in less than one hour if the database storage hardware fails.

We need to be able to immediately return the Manufacturing database to its previous state if the ProcessUpdateProc stored procedure fails to update the process information correctly.

I also want the ProcessUpdateProc stored procedure to stop running in the event of a failure of one of the UPDATE statements.

IT Administrators need to be able to monitor the disk space used on the SQL Servers by running real-time reports on the disk usage.

The Developers would like to install second instances of SQL Server on MainDB3 and MainDB4.

They would like to assign each instance to specific processors on the SQL Servers.

You need to enable the Developers to assign SQL Server instances on MainDB3 and MainDB4 to specific processors on the servers. What should you configure?

A. Windows System Resource Manager (WSRM)

B. Resource Governor

C. A Maintenance Plan

D. Processor Affinity

**Answer: D**

## NEW QUESTION 168

- (Exam Topic 3)

You administer a Microsoft SQL Server 2012 instance.

You need to configure a new database to support FILETABLES. What should you do? Choose all that apply.

A. Disable FILESTREAM on the Database.

B. Enable FILESTREAM on the Server Instance.

C. Configure the Database for Partial Containment.

D. Create a non-empty FILESTREAM file group.

E. Enable Contained Databases on the Server Instance.

F. Set the FILESTREAM directory name on the Database.

**Answer: BDF**

## Explanation:

References:

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

## NEW QUESTION 170

- (Exam Topic 3)

General Overview

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

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

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

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

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

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

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

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

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

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

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

The GenInfo table is used for reports.

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

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

Current System

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

SQL Servers

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

Design Requirements

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

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

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

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

Deadlocks must be analyzed with the use of Deadlock Graphs.

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

Software licensing and database storage costs must be minimized.

Development effort must be minimized.

The Tempdb databases must be monitored for insufficient free space.

Failed authentication requests must be logged.

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

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

SPUpdateSalesInfo queries data in the OrderTotals table the second time.

You need to configure a synchronization solution to copy data from the Current\_Inventory database the DesABCOppAppDB database.

What should you configure?

- A. Transactional Replication.
- B. Database Mirroring.
- C. Snapshot Replication.
- D. Incremental Backups

**Answer: A**

## NEW QUESTION 171

- (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 175

- (Exam Topic 3)

You are using dynamic management views to monitor an SQL Server server named SQL1. A database administrator named Dba1 must monitor the health of SQL1.

You need to ensure that Dba1 can access dynamic management views for SQL1.

The solution must use the principle of least privilege. Which permissions should you assign to Dba1?

- 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 177**

- (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)
WITH SAMPLE 100 PERCENT
GO
```

Does this meet the goal?

- A. Yes
- B. No

**Answer:** A

**Explanation:**

FULLSCAN: Compute statistics by scanning all rows in the table or indexed view. FULLSCAN and SAMPLE 100 PERCENT have the same results.

References: [https://technet.microsoft.com/en-us/library/ms187348\(v=sql.110\).aspx](https://technet.microsoft.com/en-us/library/ms187348(v=sql.110).aspx)

**NEW QUESTION 182**

- (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)
WITH FULLSCAN, NORECOMPUTE
GO
```

Does this meet the goal?

- A. Yes
- B. No

**Answer:** A

**Explanation:**

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

USE AdventureWorks2012;

GO

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

Note: NORECOMPUTE

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

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

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

**NEW QUESTION 186**

.....

## About ExamBible

*[Your Partner of IT Exam](#)*

## Found in 1998

ExamBible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, ExamBible has its unique advantages that other companies could not achieve.

## Our Advances

### \* 99.9% Uptime

All examinations will be up to date.

### \* 24/7 Quality Support

We will provide service round the clock.

### \* 100% Pass Rate

Our guarantee that you will pass the exam.

### \* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

### NEW QUESTION 1

- (Exam Topic 1)

You administer a single server that contains a Microsoft SQL Server 2016 default instance on which several production databases have been deployed. You plan to install a new ticketing application that requires the deployment of a database on the server. The SQL login for this application requires sysadmin permissions. You need to ensure that the login for the ticketing application cannot access other production databases. What should you do?

- A. Use the SQL Server default instance and enable Contained Databases.
- B. Use the SQL Server default instance and configure a user-defined server rol
- C. Add the login for the ticketing application to this role.
- D. Install a new named SQL Server instance on the server.
- E. Install a new default SQL Server instance on the server.

**Answer: C**

#### Explanation:

SQL Server supports multiple instances of SQL Server on a single server or processor, but only one instance can be the default instance. All others must be named instances. A computer can run multiple instances of SQL Server concurrently, and each instance runs independently of other instances.  
References: [https://msdn.microsoft.com/en-us/library/ms143531\(v=SQL.105\).aspx](https://msdn.microsoft.com/en-us/library/ms143531(v=SQL.105).aspx)

### NEW QUESTION 2

- (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 backu
- 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 backu
- G. Then, restore the latest log backup.
- H. Perform a page restore.
- 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. Restore the latest full backu
- 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 3**

- (Exam Topic 1)

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

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

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

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

Which permission should you grant?

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

**Answer:** A

**Explanation:**

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

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

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

**NEW QUESTION 4**

- (Exam Topic 1)

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

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

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

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

You need to optimize query performance.

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

- A. Yes
- B. No

**Answer:** B

**Explanation:**

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

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

**NEW QUESTION 5**

- (Exam Topic 1)

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

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

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

You need to perform the following actions:

- identify databases where the PAGE verity option is not set.
- Configure full page protection for the identified databases. Solution: You run the following Transact-SQL Statement:

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

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

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

Does the solution meet the goal?

- A. Yes
- B. No

**Answer:** A

#### NEW QUESTION 6

- (Exam Topic 1)

You administer a Microsoft SQL Server 2016 server.

When transaction logs grow, SQL Server must send an email message to the database administrators. You need to configure SQL Server to send the email messages.

What should you configure?

- A. SQL Mail
- B. An Extended Events session
- C. Alerts and operators in SQL Server Agent
- D. Policies under Policy-Based Management

**Answer:** C

#### Explanation:

Operators are aliases for people or groups that can receive electronic notification when jobs have completed or alerts have been raised. The SQL Server Agent service supports the notification of administrators through operators. Operators enable notification and monitoring capabilities of SQL Server Agent.

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/database-mail/configure-sql-server-agent-mail-to-use-d>

#### NEW QUESTION 7

- (Exam Topic 1)

You administer a Windows Azure SQL Database database named Inventory that contains a stored procedure named p\_AddInventory.

Users need to be able to SELECT from all tables in the database and execute the stored procedure. You need to grant only the necessary permissions.

What should you do?

- A. Grant EXECUTE permission on p\_AddInventory to all user
- B. Grant VIEW DEFINITION to all users.
- C. Grant EXECUTE permission on p\_AddInventory to all user
- D. Add all users to the db\_datawriter role.
- E. Add all users to the db\_owner role.
- F. Grant EXECUTE permission on p\_AddInventory to all user
- G. Add all users to the db\_datareader role.

**Answer:** D

#### NEW QUESTION 8

- (Exam Topic 1)

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

You need to track all SELECT statements issued in the Contoso database only by users in a role named Sales. What should you create?

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

**Answer:** F

#### NEW QUESTION 9

- (Exam Topic 1)

You are the database administrator of a Microsoft SQL Server instance. Developers are writing stored procedures to send emails using sp\_send\_dbmail. Database Mail is enabled.

You need to configure each account's profile security and meet the following requirements:

Account SMTP1\_Account must only be usable by logins that have been given explicit permissions to use the SMTP1\_profile.

Account SMTP2\_Account must only be usable by logins who are a member of the [DatabaseMailUserRole] role in msdb.

In the table below, identify the profile type that must be used for each account. NOTE: Make only one selection in each column.

## Answer Area

| Profile type    | SMTP1_Account         | SMTP2_Account         |
|-----------------|-----------------------|-----------------------|
| Private Profile | <input type="radio"/> | <input type="radio"/> |
| Public Profile  | <input type="radio"/> | <input type="radio"/> |
| Default Profile | <input type="radio"/> | <input type="radio"/> |

- A. Mastered
- B. Not Mastered

**Answer:** A

### Explanation:

SMTP1\_Account1: Private Profile

When no profile\_name is specified, sp\_send\_dbmail uses the default private profile for the current user. I user does not have a default private profile, sp\_send\_dbmail uses the default public profile for the msdb database.

SMTP1\_Account2: Default Profile

Execute permissions for sp\_send\_dbmail default to all members of the DatabaseMailUser database role in the msdb database.

References:

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

### 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 master database. Does the solution meet the goal?

- A. Yes
- B. No

**Answer:** B

### Explanation:

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

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

References:

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

### NEW QUESTION 10

- (Exam Topic 1)

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

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

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

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

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

**Answer:** G

### Explanation:

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

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

**NEW QUESTION 13**

- (Exam Topic 1)

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

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

**Answer: C**

**NEW QUESTION 14**

- (Exam Topic 1)

You plan to migrate the db to azure.

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

What should you do?

- A. Use the Copy Database wizard
- B. Use the Database Transfer wizard
- C. Use the SQL Management Studio to deploy the db to azure
- D. Back up the databases from the local server and restore it to azure

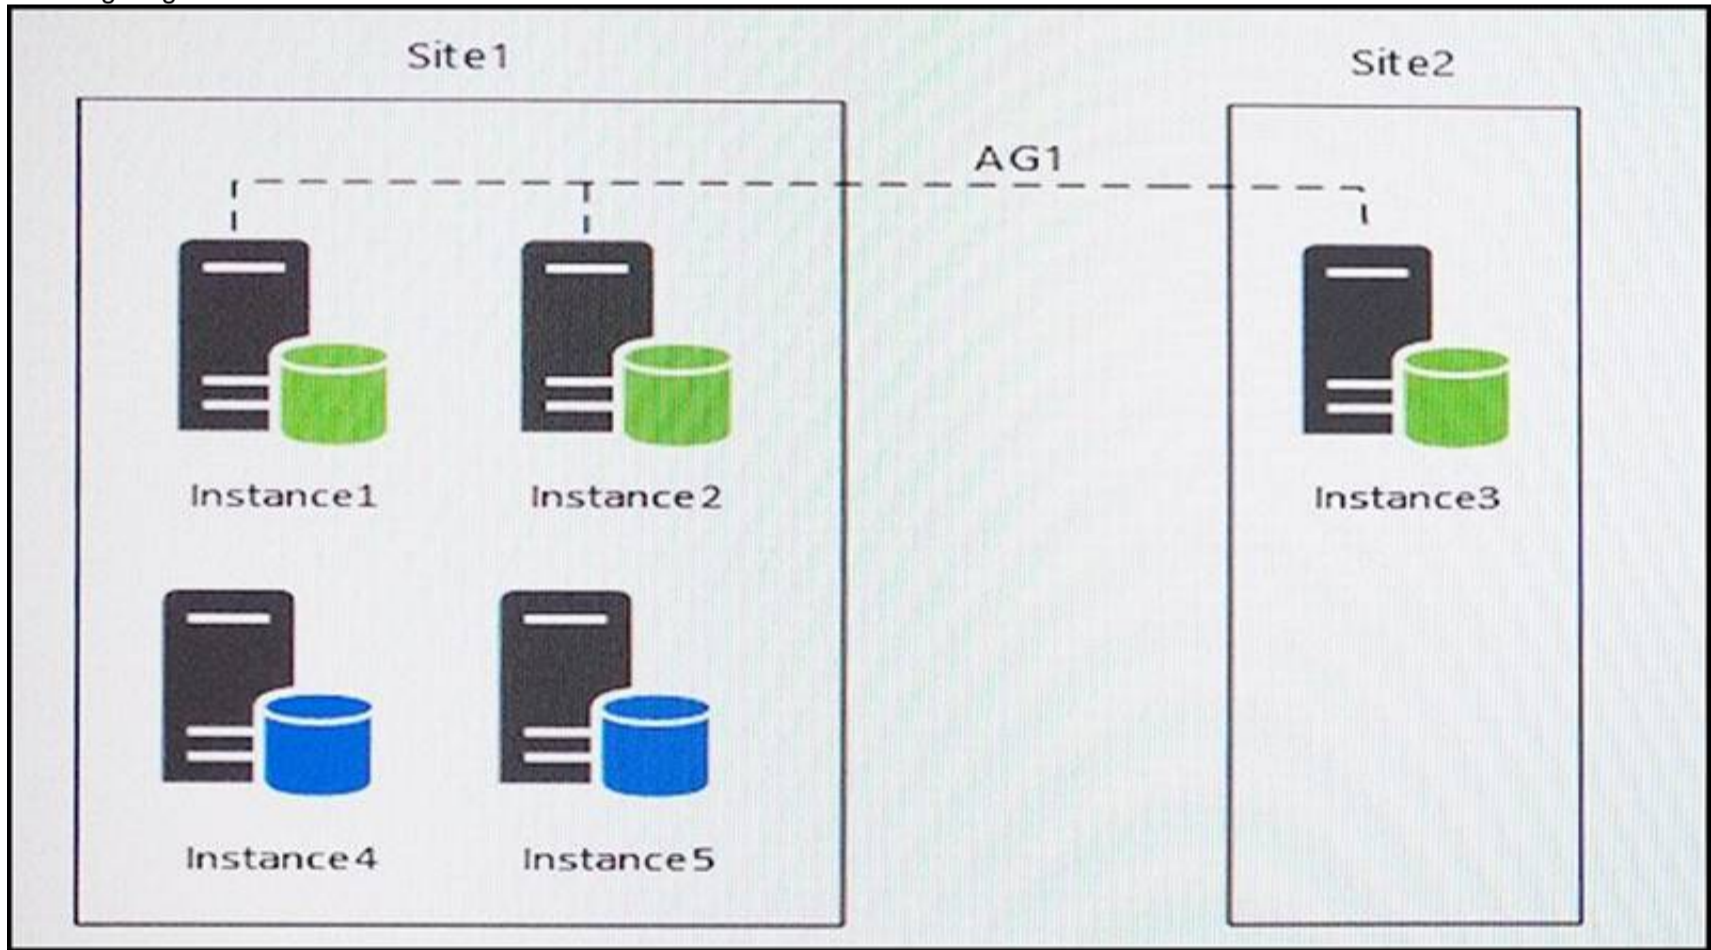
**Answer: CD**

**NEW QUESTION 17**

- (Exam Topic 1)

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

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



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

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

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

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

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

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

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

nodes in AG1.

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

You should minimize both the time required to restore

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

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

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

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

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

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

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

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

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

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

**Answer: D**

**Explanation:**

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

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

References:

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

**NEW QUESTION 22**

- (Exam Topic 1)

You administer a Microsoft SQL Server 2016 database instance.

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

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

What should you do?

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

**Answer: D**

**NEW QUESTION 24**

- (Exam Topic 1)

You are the administrator of a Microsoft SQL Server 2016 server. Some applications consume significant resources.

You need to manage the server workload by restricting resource-intensive applications. You need to dynamically limit resource consumption.

What should you do?

- A. Set up Service Broker to ensure that applications are not allowed to consume more than the specified amount of resources.
- B. Configure Resource Pools, Workload Groups, and Classifier Function, and then enable the Resource Governor.
- C. Configure Extended Events to monitor and restrict resource limits allowed by each application type.
- D. Create a new Plan Guide with a Scope Type of sql and define the resource limits for each application.

Answer: B

#### NEW QUESTION 28

- (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

Is the 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\_type is 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 each execution of a job. freq\_subday\_interval is 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 32

- (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 goals.

Some question sets might have more than one correct solution, while others might not have a correct solution.

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

You have a database named DB1 that is 640 GB and is updated frequently.

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

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

Solution: You increase the frequency of the transaction log backups to every 10 minutes. Does this meet the goal?

- A. Yes
- B. No

**Answer:** B

### NEW QUESTION 36

- (Exam Topic 1)

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

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

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

## Answer Area

| Action                   | Server instance  |
|--------------------------|--|
| Complete the backup job. | <div><div>▼</div><div>Primary server instance</div><div>Secondary server instance</div><div>Monitor server instance</div><div>Backup share file server</div></div> |
| Copy the backup job.     | <div><div>▼</div><div>Primary server instance</div><div>Secondary server instance</div><div>Monitor server instance</div><div>Backup share file server</div></div> |
| Restore the backup.      | <div><div>▼</div><div>Primary server instance</div><div>Secondary server instance</div><div>Monitor server instance</div><div>Backup share file server</div></div> |

- A. Mastered  
B. Not Mastered

**Answer:** A

**Explanation:**

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

Box 1: Primary server instance.

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

Box 2: Secondary server instance

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

Box 3: Secondary server instance.

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

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

**NEW QUESTION 39**

- (Exam Topic 1)

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

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

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

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

**Answer:** B

**Explanation:**

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

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

**NEW QUESTION 44**

- (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.  
One instance hosts a user database named HRDB. The database contains sensitive human resources data. You need to grant an auditor permission to view the SQL Server audit logs 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:** F

**Explanation:**  
Unless otherwise specified, viewing catalog views requires a principal to have one of the following:  
Membership in the sysadmin fixed server role.  
The CONTROL SERVER permission.  
The VIEW SERVER STATE permission.  
The ALTER ANY AUDIT permission.  
The VIEW AUDIT STATE permission (gives only the principal access to the sys.server\_audits catalog view).  
References: [https://technet.microsoft.com/en-us/library/cc280386\(v=sql.110\).aspx](https://technet.microsoft.com/en-us/library/cc280386(v=sql.110).aspx)

**NEW QUESTION 48**

- (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

WHERE ((sqlserver.data-base\_id>(4)) AND (severity>(15)))

(ACTION(sqlserver.client\_app\_name, sqlserver.data-base\_id,sqlserver.session\_id)

ALTER EVENT SESSION Contoso1 ON SERVER

)

GO

ADD EVENT sqlserver.error\_reported

ADD TARGET sqlserver.error\_reported

Answer Area

⏪

⏩

⏴

⏵

- A. Mastered
- B. Not Mastered

**Answer:** A

Your Partner of IT Exam

visit - <https://www.exambible.com>

#### 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.tsql_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 50

- (Exam Topic 1)

You have a Microsoft SQL Server instance that hosts a database named DB1 that contains 800 gigabyte (GB) of data. The database is used 24 hours each day.

You implement indexes and set the value of the Auto Update Statistics option set to True.

Users report that queries take a long time to complete.

You need to identify statistics that have not been updated for a week for tables where more than 1,000 rows changed.

How should you complete the Transact-SQL statement? To answer, configure the appropriate Transact-SQL segments in the answer area.

## Answer Area

```
SELECT OBJECT_NAME(id), name, 

|               |   |
|---------------|---|
|               | ▼ |
| rowcnt        |   |
| stats_date    |   |
| rowmodctr     |   |
| stats_collect |   |

 (id, indid), 

|               |   |
|---------------|---|
|               | ▼ |
| rowcnt        |   |
| stats_date    |   |
| rowmodctr     |   |
| stats_collect |   |


```

```
FROM sys.sysindexes
```

```
WHERE 

|               |   |
|---------------|---|
|               | ▼ |
| rowmodctr     |   |
| stats_collect |   |
| stats_date    |   |
| rowcnt        |   |

 (id, indid) <= DATEADD(DAY, -7, GETDATE())
```

```
AND 

|               |   |
|---------------|---|
|               | ▼ |
| stats_collect |   |
| rowmodctr     |   |
| stats_date    |   |
| rowcnt        |   |

 > 1000
```

```
AND id IN (SELECT object_id FROM sys.tables)
```

- A. Mastered
- B. Not Mastered

**Answer:** A

#### Explanation:

Box 1: stats\_date See example below. Box 2: rowmodctr See examplebelow. Box 3: stats\_date

You need to identify statistics that have not been updated for a week. Box 4: rowmodctr

You need to identify that more than 1,000 rows changed.

Rowmodctr counts the total number of inserted, deleted, or updated rows since the last time statistics were updated for the table.

Example: We will query every statistics object which was not updated in the last day and has rows modified since the last update. We will use the rowmodctr field of sys.sysindexes because it shows how many rows were inserted, updated or deleted since the last update occurred. Please note that it is not always 100% accurate in SQL Server 2005 and later, but it can be used to check if any rows were modified.

--Get the list of outdated statistics

```
SELECT OBJECT_NAME(id),name,STATS_DATE(id, indid),rowmodctr FROM sys.sysindexes
```

```
WHERE STATS_DATE (id, indid)<=DATEADD(DAY,-1,GETDATE())
AND rowmodctr>0
AND id IN (SELECT object_id FROM sys.tables) GO
After collecting this information, we can decide which statistics require an update.
```

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/system-compatibility-views/sys-sysindexes-transact-sq>

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

#### NEW QUESTION 51

- (Exam Topic 1)

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

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

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

Drive D has 1 terabyte of free disk space.

Drive E has 5 GB of free disk space.

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

Five percent of the existing data is modified each day.

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

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

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

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

failover to this secondary database.

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

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

**Answer:** K

#### Explanation:

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

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

#### NEW QUESTION 55

- (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 60

- (Exam Topic 1)

You administer all the deployments of Microsoft SQL Server 2016 in your company.

You need to ensure that an OLTP database that uses a storage area network (SAN) remains available if any of the servers fail.

You also need to minimize the amount of storage used by the database. Which configuration should you use?

- A. Two servers configured in different data centers SQL Server Availability Group configured in Synchronous-Commit Availability Mode One server configured as an Active Secondary
- B. SQL Server that includes an application database configured to perform transactional replication
- C. Two servers configured in the same data center SQL Server Availability Group configured in AsynchronousCommit Availability Mode One server configured as an Active Secondary

- D. Two servers configured in different data centers SQL Server Availability Group configured in AsynchronousCommit Availability Mode
- E. Two servers configured in the same data center A primary server configured to perform log-shipping every 10 minutes A backup server configured as a warm standby
- F. Two servers configured on the same subnet SQL Server Availability Group configured in Synchronous-Commit Availability Mode
- G. SQL Server that includes an application database configured to perform snapshot replication
- H. Two servers configured in a Windows Failover Cluster in the same data center SQL Server configured as a clustered instance

**Answer:** H

**Explanation:**

A Windows Server Failover Cluster (WSFC) is a group of independent servers that work together to increase the availability of applications and services. SQL Server takes advantage of WSFC services and capabilities to support Always On availability groups and SQL Server Failover Cluster Instances.

References:

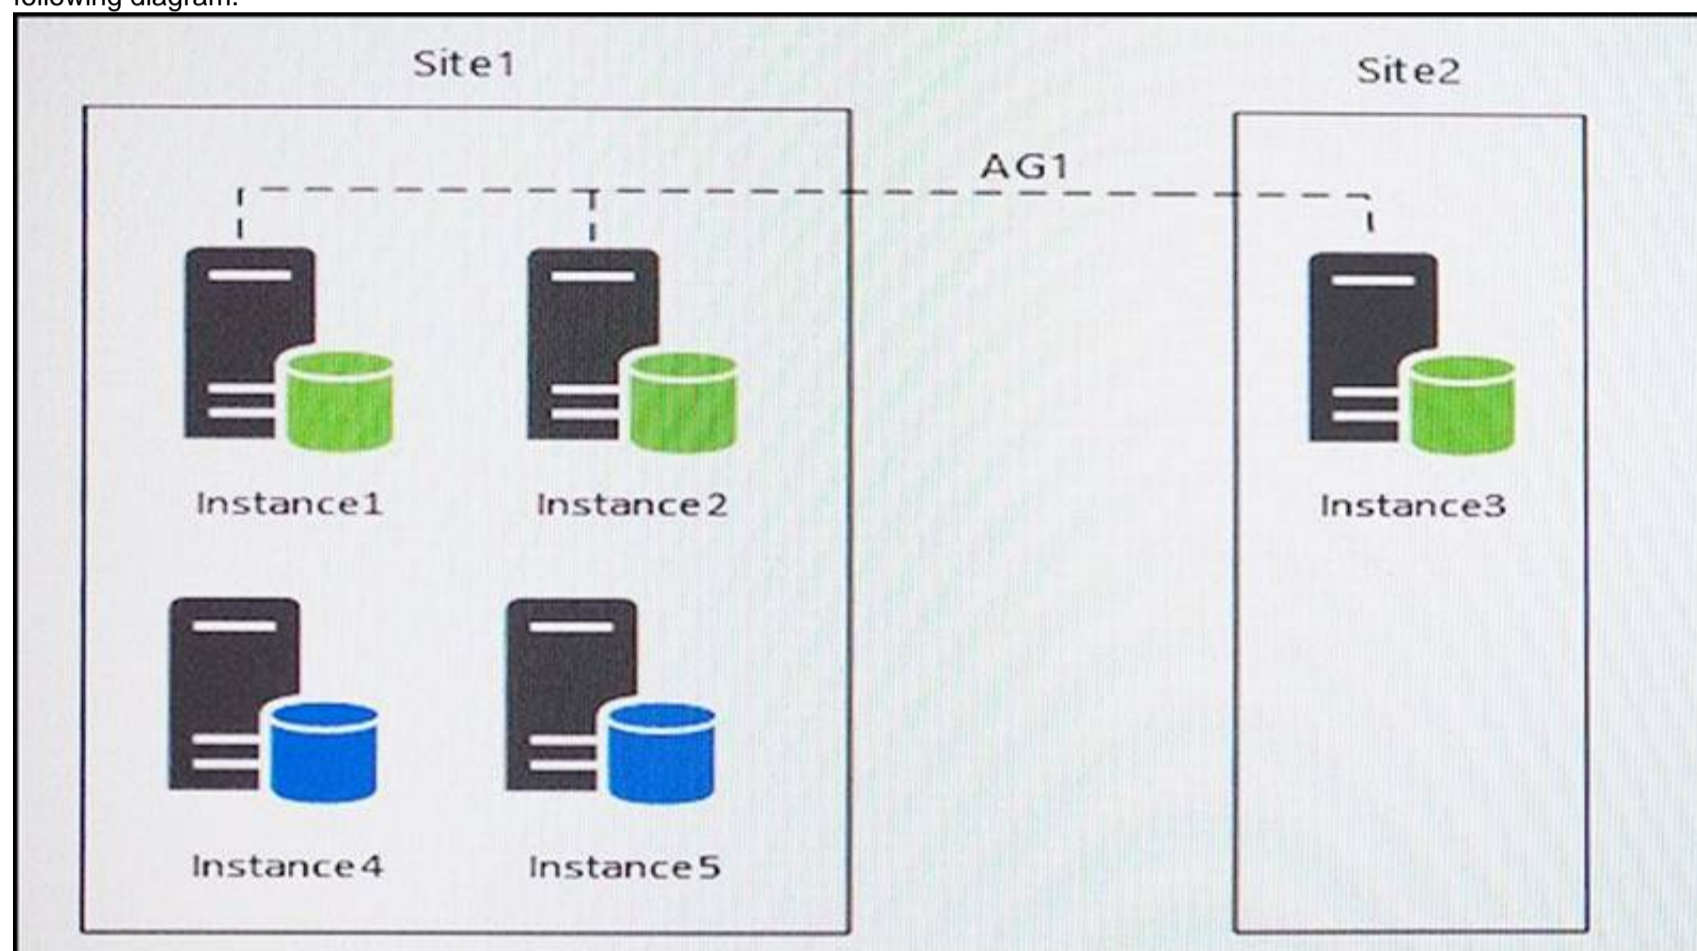
<https://docs.microsoft.com/en-us/sql/sql-server/failover-clusters/windows/windows-server-failover-clustering-ws>

**NEW QUESTION 65**

- (Exam Topic 1)

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

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



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

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

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

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

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

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

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

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

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

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

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

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

Reporting system: This solution accesses data in DB1 with a login that is mapped to a database user that is a member of the db\_datareader role. The user has EXECUTE permissions on the database. Queries make no changes to the data. The queries must be load balanced over variable read-only replicas.

Operations system: This solution accesses data in DB1 with a login that is mapped to a database user that is a member of the db\_datareader and db\_datawriter

roles. The user has EXECUTE permissions on the database. Queries from the operations system will perform both DDL and DML operations. The wait statistics monitoring requirements for the instances are described in the following table.

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

You need to analyze the wait type and statistics for specific instanced in the environment. Which object should you use to gather information about each instance? To answer, drag the appropriate objects to the correct instances. Each object may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.  
NOTE: Each correct selection is worth one point.

### Objects

Sys.dm\_os\_wait\_stats

Sys.dm\_exec\_connections

Sys.dm\_exec\_requests

Sys.dm\_exec\_procedure\_stats

Sys.dm\_exec\_sessions

Sys.dm\_exec\_query\_stats

Sys.dm\_exec\_query\_re-source\_semaphores

Sys.dm\_exec\_ses-sion\_wait\_stats

### Answer Area

| Instance  | Object            |
|-----------|-------------------|
| Instance1 | <div>Object</div> |
| Instance4 | <div>Object</div> |
| Instance5 | <div>Object</div> |

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Instance 1: sys.dm\_exec\_query\_stats  
From Scenario: Instance1 requirement: Aggregate statistics since last server restart. sys.dm\_exec\_query\_stats returns aggregate performance statistics for cachedquery plans in SQL Server.  
Instance 4: sys.dm\_os\_wait\_stats  
sys.dm\_os\_wait\_statsreturns information about all the waits encountered by threads that executed. From Scenario: Instance4 requirement: Identify the most prominent wait types.

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

Instance 5:sys.dm\_exec\_session\_wait\_stats  
From Scenario: Instance5 requirement: Identify all wait types for queries currently running on the server. sys.dm\_exec\_session\_wait\_stats returns information about all the waits encountered by threads that executed for each session.

NEW QUESTION 70

- (Exam Topic 1)  
You administer a database that is used for reporting purposes. The database has a large fact table that contains three hundred million rows. The table includes a clustered columnstore index and a nonclustered index on the ProductID column. New rows are inserted into the table every day. Performance of queries that filter the Product ID column have degraded significantly. You need to improve the performance of the queries. Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

| Actions   | Answer Area |
|---|-------------|
| Drop the clustered columnstore index.                         |             |
| Create a nonclustered index on ProductID.                     |             |
| Drop and recreate the clustered columnstore index.            |             |
| Create a nonclustered columnstore index on ProductID.         |             |
| Recreate the clustered columnstore index using DROP EXISTING. |             |
| Create a clustered rowstore index on ProductID.               |             |
| Rebuild the clustered columnstore index.                      |             |

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Step 1: Drop the clustered columnstore index

Step 2: Create a clustered rowstore index on ProductID.

Rowstore indexes perform best on queries that seek into the data, searching for a particular value, or for queries on a small range of values. Use rowstore indexes with transactional workloads since they tend to require mostly table seeks instead of table scans.

Step 3: Create a nonclustered index on ProductID

**NEW QUESTION 73**

- (Exam Topic 2)

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

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

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

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

**Answer:** B

**Explanation:**

sys.dm\_os\_wait\_stats

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

Columns include: waiting\_tasks\_count

Number of waits on this wait type.

This counter is incremented at the start of each wait.

**NEW QUESTION 75**

- (Exam Topic 2)

Overview

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

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

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

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

Requirements Planned Changes

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

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

**Business Requirements**

You have the following requirements:

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

The encryption strategy must minimize changes to the databases and to the applications. You need to recommend a solution to improve the performance of usp\_UpdateInventory. The solution must minimize the amount of development effort. What should you include in the recommendation?

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

**Answer:** A

**Explanation:**

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

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

**NEW QUESTION 77**

- (Exam Topic 2)

You plan to create a database.

The database will be used by a Microsoft .NET application for a special event that will last for two days. During the event, data must be highly available. After the event, the database will be deleted. You need to recommend a solution to implement the database while minimizing costs. The solution must not affect any existing applications. What should you recommend? More than one answer choice may achieve the goal. Select the BEST answer.

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

**Answer:** D

**NEW QUESTION 81**

- (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 84**

- (Exam Topic 2)

You administer a Microsoft SQL Server 2016 server.

The MSSQLSERVER service uses a domain account named CONTOSO\SQLService. You plan to configure Instant File Initialization.

You need to ensure that Data File Autogrow operations use Instant File Initialization. What should you do? Choose all that apply.

- A. Restart the SQL Server Agent Service.
- B. Disable snapshot isolation.
- C. Restart the SQL Server Service.
- D. Add the CONTOSO\SQLService account to the Perform Volume Maintenance Tasks local security policy.
- E. Add the CONTOSO\SQLService account to the Server Operators fixed server role.
- F. Enable snapshot isolation.

**Answer:** CD

**Explanation:**

How To Enable Instant File Initialization

Open Local Security Policy and go to Local Policies → User Rights Assignment.

Double click Perform Volume Maintenance Tasks and add your SQL Server database engine service account.

Restart the SQL Server service using SQL Server Configuration Manager and this setting should now be enabled.

References:

<http://msdn.microsoft.com/en-us/library/ms175935.aspx>

**NEW QUESTION 89**

- (Exam Topic 2)

Overview

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

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

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

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

Requirements Planned Changes

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

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

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

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

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

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

Inventory will contain over 100 GB of data.

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

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

Business Requirements

You have the following requirements:

Costs for new licenses must be minimized.

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

Development effort must be minimized whenever possible.

The storage requirements for databases must be minimized.

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

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

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

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

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

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

- A. a SQL Server Agent alert
- B. a dynamic management view
- C. a maintenance plan
- D. an audit

**Answer:** B

**Explanation:**

Dynamic Management Views and Functions (Transact-SQL)

### NEW QUESTION 93

- (Exam Topic 2)

Overview

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

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

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

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

Requirements Planned Changes

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

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

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

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

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

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

Inventory will contain over 100 GB of data.

The Inventory table will have two indexes: a clustered index on the primary key and a nonclustered index.

The column that is used as the primary key will use the identity property.

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

Business Requirements

You have the following requirements:

Costs for new licenses must be minimized.

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

Development effort must be minimized whenever possible.

The storage requirements for databases must be minimized.

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

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

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

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

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

The encryption strategy must minimize changes to the databases and to the applications. You need to recommend an isolation level for usp\_UpdateOrderDetails.

Which isolation level should you recommend?

- A. Read committed
- B. Repeatable read
- C. Read uncommitted
- D. Serializable

**Answer: B**

**Explanation:**

- Scenario: 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.

- REPEATABLE READ Specifies that statements cannot read data that has been modified but not yet committed by other transactions and that no other transactions can modify data that has been read by the current transaction until the current transaction completes.

### NEW QUESTION 94

- (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 98

- (Exam Topic 2)

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

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

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

**Answer:** A

**Explanation:**

If we are going to encrypt credit card number for storage, then we should have Data Encryption Key(DEK) for encrypting the credit card number.

**NEW QUESTION 102**

- (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           | uniqueidentifier |
| lastModified | datetime         |
| modifiedBy   | Varchar(200)     |

Backup Issues

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

Users report that when they search product names, the search results exclude product names that contain

accents, unless the search string includes the accent. Missing Data Issues

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

Query Performance Issues

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

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

Design Requirements

File Storage Requirements

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

Data Recovery Requirements

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

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

Concurrency Requirements

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

What should you include in the recommendation?

- A. Windows logins
- B. Server roles
- C. Contained users
- D. Database roles

**Answer: C**

**Explanation:**

- Scenario: 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.

- Creating contained users enables the user to connect directly to the contained database. This is a very significant feature in high availability and disaster recovery scenarios such as in an AlwaysOn solution. If the users are contained users, in case of failover, people would be able to connect to the secondary without creating logins on the instance hosting the secondary. This provides an immediate benefit.

**NEW QUESTION 107**

- (Exam Topic 2)

Overview

General Overview

ADatum Corporation has offices in Miami and Montreal.

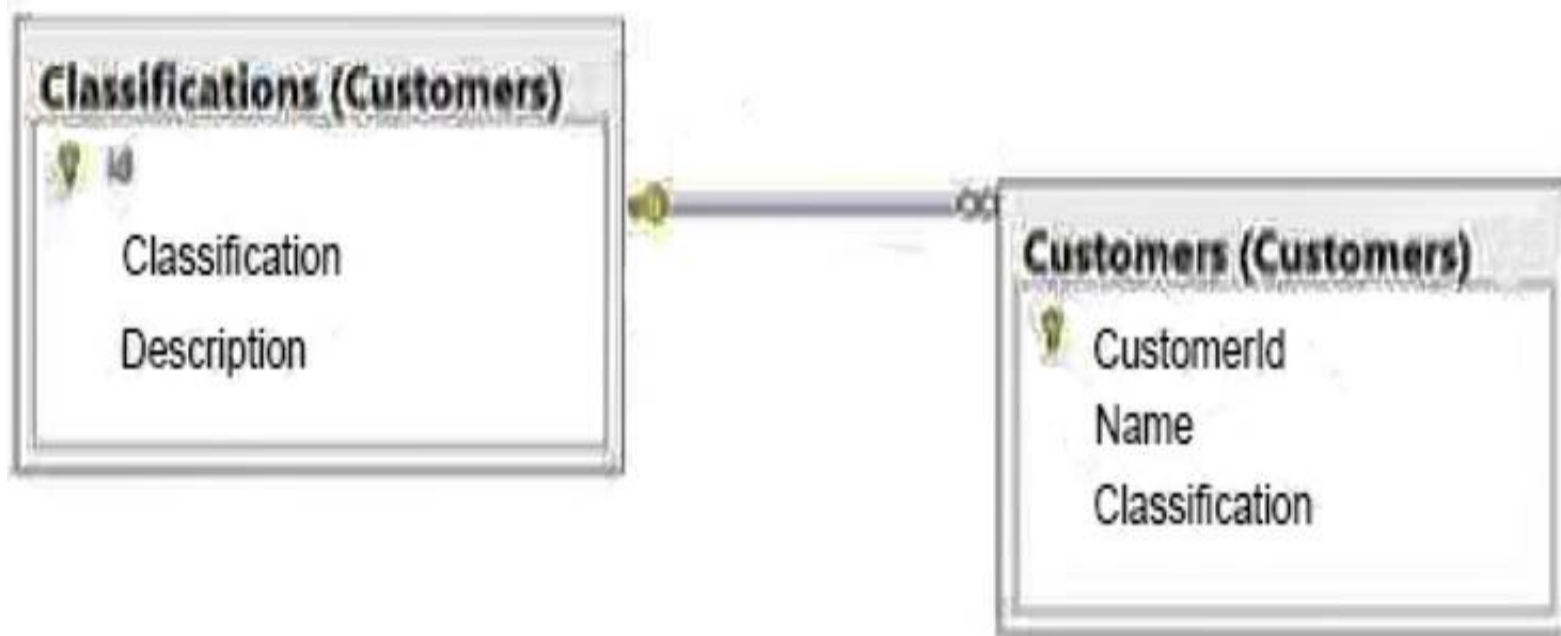
The network contains a single Active Directory forest named adatum.com. The offices connect to each other by using a WAN link that has 5-ms latency. A. Datum standardizes its database platform by using SQL Server 2014 Enterprise edition.

Databases

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

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

Classifications (Customers)



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 disaster recovery strategy for the Inventory database. What should you include in the recommendation?

- A. Log shipping
- B. SQL Server Failover Clustering
- C. AlwaysOn availability groups
- D. Peer-to-peer replication

**Answer:** A

#### Explanation:

Scenario:

- You must be able to recover data from the Inventory database if a storage failure occurs. You have a Recovery Point Objective (RPO) of one hour.
- A. Datum Corporation has offices in Miami and Montreal.
- SQL Server Log shipping allows you to automatically send transaction log backups from a primary database on a primary server instance to one or more secondary databases on separate secondary server instances. The transaction log backups are applied to each of the secondary databases individually.

#### NEW QUESTION 108

- (Exam Topic 2)

You administer a Microsoft SQL Server 2016 instance.

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

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

**Answer:** C

#### Explanation:

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

References:

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

#### NEW QUESTION 112

- (Exam Topic 2)

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

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

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

The stored procedure must be designed to maximize concurrency.

What should you include in the design? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

| Actions  | Work Area |
|--|-----------|
| Raise an error in a catch block  |           |
| Commit the transaction in a finally block                                    |           |
| Read the @@ROWCOUNT system variable  |           |
| Perform the update in a try block  |           |
| Raise an error and roll back the transaction if the row count is less than 1 |           |
| Issue a SELECT statement to count the number of rows                         |           |
| Set the isolation level to serializable                                      |           |
| Begin an explicit transaction  |           |

- A. Mastered  
 B. Not Mastered

**Answer:** A

**Explanation:**

Note:

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

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

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

- Generate a constraint violation error.

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

**NEW QUESTION 116**

- (Exam Topic 2)

Overview

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

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

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

will use the application.

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

Requirements Planned Changes

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

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

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

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

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

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

Inventory will contain over 100 GB of data.

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

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

Business Requirements

You have the following requirements:

Costs for new licenses must be minimized.

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

Development effort must be minimized whenever possible.

The storage requirements for databases must be minimized.

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

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

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

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

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

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

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

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

**Answer: A**

**Explanation:**

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

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

Transparent Data Encryption (TDE)

## NEW QUESTION 119

- (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 120

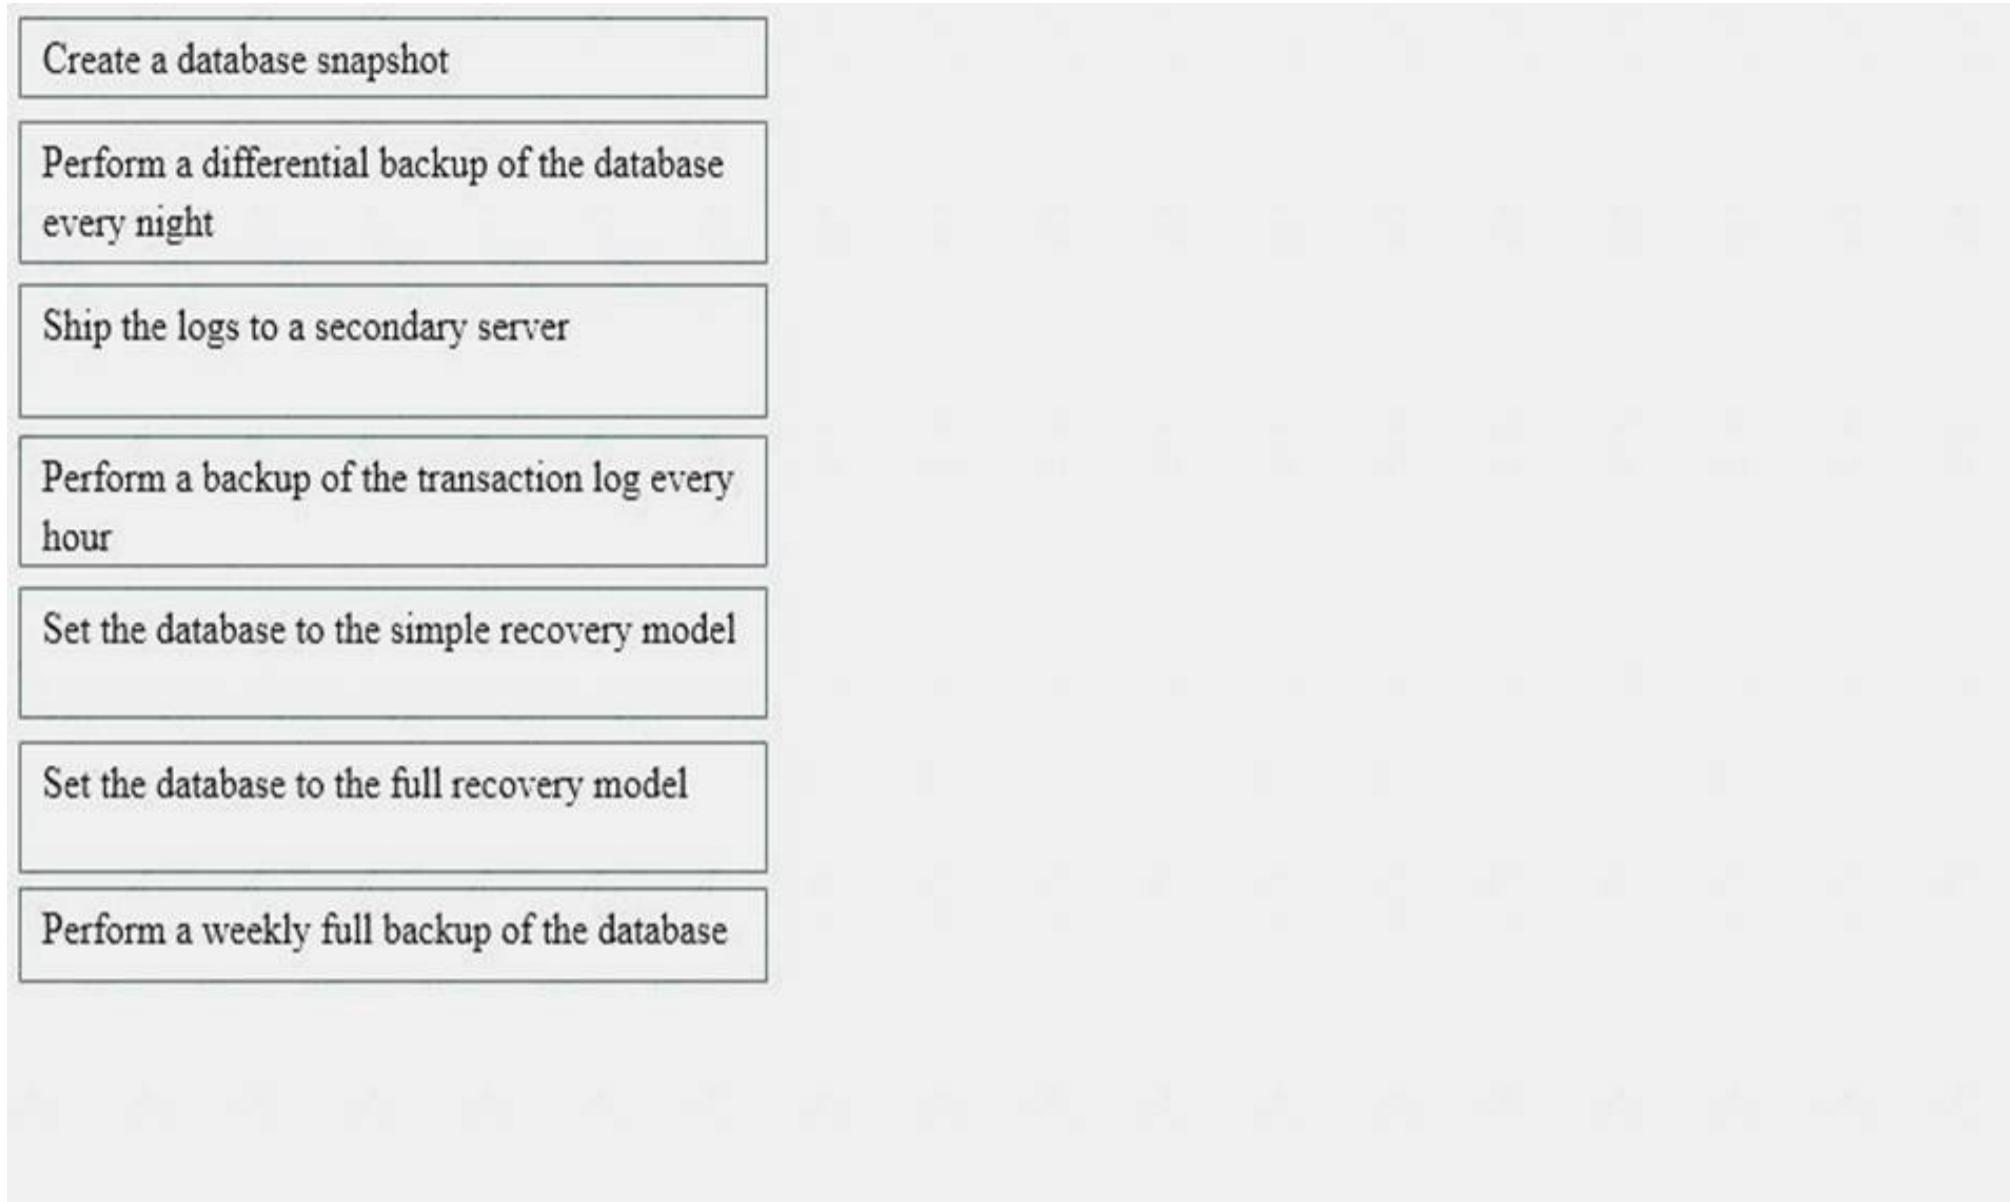
- (Exam Topic 2)

You need to recommend a backup process for an Online Transaction Processing (OLTP) database. The process must meet the following requirements:

Ensure that if a hardware failure occurs, you can bring the database online with a minimum amount of data loss.

Minimize the amount of administrative effort required to restore any lost data.

What should you include in the recommendation? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.



Create a database snapshot

Perform a differential backup of the database every night

Ship the logs to a secondary server

Perform a backup of the transaction log every hour

Set the database to the simple recovery model

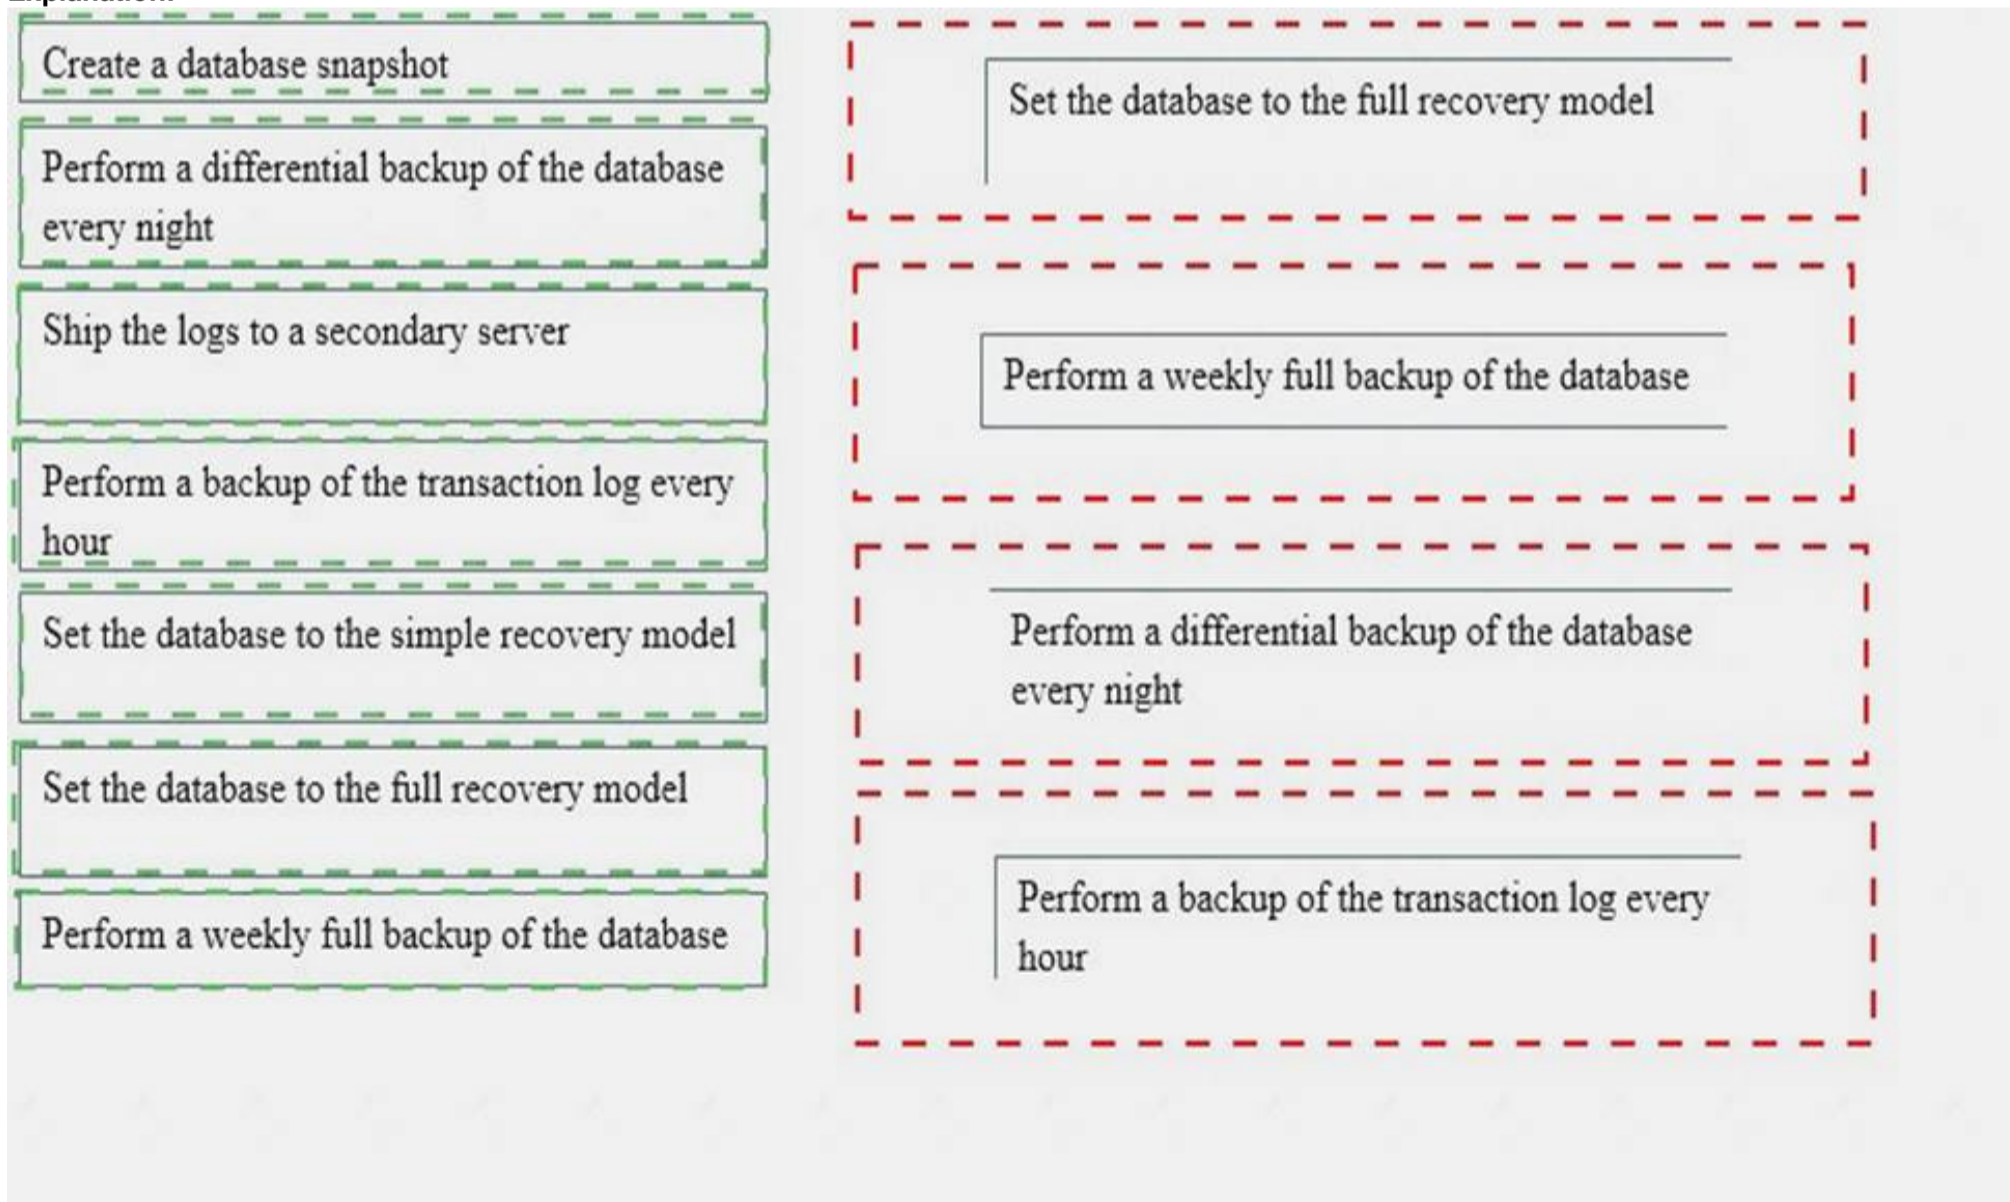
Set the database to the full recovery model

Perform a weekly full backup of the database

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**



Create a database snapshot

Perform a differential backup of the database every night

Ship the logs to a secondary server

Perform a backup of the transaction log every hour

Set the database to the simple recovery model

Set the database to the full recovery model

Perform a weekly full backup of the database

Set the database to the full recovery model

Perform a weekly full backup of the database

Perform a differential backup of the database every night

Perform a backup of the transaction log every hour

#### NEW QUESTION 124

- (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 129

- (Exam Topic 2)  
Overview  
Application Overview  
Contoso, Ltd., is the developer of an enterprise resource planning (ERP) application.  
Contoso is designing a new version of the ERP application. The previous version of the ERP application used SQL Server 2008 R2.  
The new version will use SQL Server 2014.  
The ERP application relies on an import process to load supplier data. The import process updates thousands of rows simultaneously, requires exclusive access to the database, and runs daily.  
You receive several support calls reporting unexpected behavior in the ERP application. After analyzing the calls, you conclude that users made changes directly to the tables in the database.  
Tables  
The current database schema contains a table named OrderDetails.  
The OrderDetails table contains information about the items sold for each purchase order. OrderDetails stores the product ID, quantities, and discounts applied to each product in a purchase order.  
The product price is stored in a table named Products. The Products table was defined by using the SQL\_Latin1\_General\_CP1\_CI\_AS collation.  
A column named ProductName was created by using the varchar data type. The database contains a table named Orders.  
Orders contains all of the purchase orders from the last 12 months. Purchase orders that are older than 12 months are stored in a table named OrdersOld.  
The previous version of the ERP application relied on table-level security. Stored Procedures  
The current version of the database contains stored procedures that change two tables. The following shows the relevant portions of the two stored procedures:

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

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

Customer Problems Installation Issues

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

The current version of the ERP application requires that several SQL Server logins be set up to function correctly. Most customers set up the ERP application in multiple locations and must create logins multiple times.  
Index Fragmentation Issues  
Customers discover that clustered indexes often are fragmented. To resolve this issue, the customers defragment the indexes more frequently. All of the tables affected by fragmentation have the following columns that are used as the clustered index key:  
Backup Issues  
Customers who have large amounts of historical purchase order data report that backup time is unacceptable. Search Issues  
Users report that when they search product names, the search results exclude product names that contain accents, unless the search string includes the accent.  
Missing Data Issues  
Customers report that when they make a price change in the Products table, they cannot retrieve the price that the item was sold for in previous orders.

#### Query Performance Issues

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

#### Design Requirements

##### File Storage Requirements

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

##### Data Recovery Requirements

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

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

##### Concurrency Requirements

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

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

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

**Answer:** BE

#### Explanation:

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

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

#### NEW QUESTION 133

- (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 134

- (Exam Topic 2)

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

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

You need to upgrade the database to SQL Server 2014. The solution must minimize downtime. Which upgrade steps should you recommend? To answer, move

the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

|                    |
|--------------------|
| Fail over          |
| Fail back          |
| Upgrade SQL1       |
| Upgrade SQL2       |
| Establish a mirror |
| Break the mirror   |

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

References:<https://docs.microsoft.com/en-us/sql/database-engine/database-mirroring/upgrading-mirrored-instanc>

**NEW QUESTION 137**

- (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 141**

- (Exam Topic 2)

You are the lead database administrator (DBA) of a Microsoft SQL Server 2016 environment. All DBAs are members of the DOMAIN\JrDBAs Active Directory group.

You grant DOMAIN\JrDBAs access to the SQL Server.

You need to create a server role named SpecialDBARole that can perform the following functions:

View all databases.

View the server state.

Assign GRANT, DENY, and REVOKE permissions on logins.

You need to add DOMAIN\JrDBAs to the server role.

You also need to provide the least level of privileges necessary.

Which SQL statement or statements should you use? Choose all that apply.

- A. CREATE SERVER ROLE [SpecialDBARole] AUTHORIZATION setupadmin;
- B. ALTER SERVER ROLE [SpecialDBARole] ADD MEMBER [DOMAIN\JrDBAs];
- C. CREATE SERVER ROLE [SpecialDBARole] AUTHORIZATION securityadmin;
- D. GRANT VIEW DEFINITION TO [SpecialDBARole];
- E. CREATE SERVER ROLE [SpecialDBARole] AUTHORIZATION serveradmin;
- F. GRANT VIEW SERVER STATE, VIEW ANY DATABASE TO [SpecialDBARole];

**Answer:** BCF

**NEW QUESTION 142**

- (Exam Topic 2)

Overview

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

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

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

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

**Requirements Planned Changes**

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

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

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

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

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

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

Inventory will contain over 100 GB of data.

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

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

**Business Requirements**

You have the following requirements:

Costs for new licenses must be minimized.

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

Development effort must be minimized whenever possible.

The storage requirements for databases must be minimized.

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

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

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

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

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

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

You need to recommend a solution to allow application users to perform UPDATE operations on the database tables. The solution must meet the business requirements.

What should you recommend?

- A. Create stored procedures that use EXECUTE AS clauses.
- B. Create a user-defined database role and add users to the role.
- C. Create functions that use EXECUTE AS clauses.
- D. Create a Policy-Based Management Policy.

**Answer:** A

**Explanation:**

- EXECUTE AS Clause (Transact-SQL)

In SQL Server you can define the execution context of the following user-defined modules: functions (except inline table-valued functions), procedures, queues, and triggers.

#### NEW QUESTION 146

- (Exam Topic 2)

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

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

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

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

**Answer:** A

#### NEW QUESTION 147

- (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

EXPLICIT

ISOLATION

READ UNCOMMITTED

ROLLBACK

SERIALIZABLE

SNAPSHOT

TABLOCK

TRANSACTION

Answer Area

SET

LEVEL

BEGIN

UPDATE SalesOrder

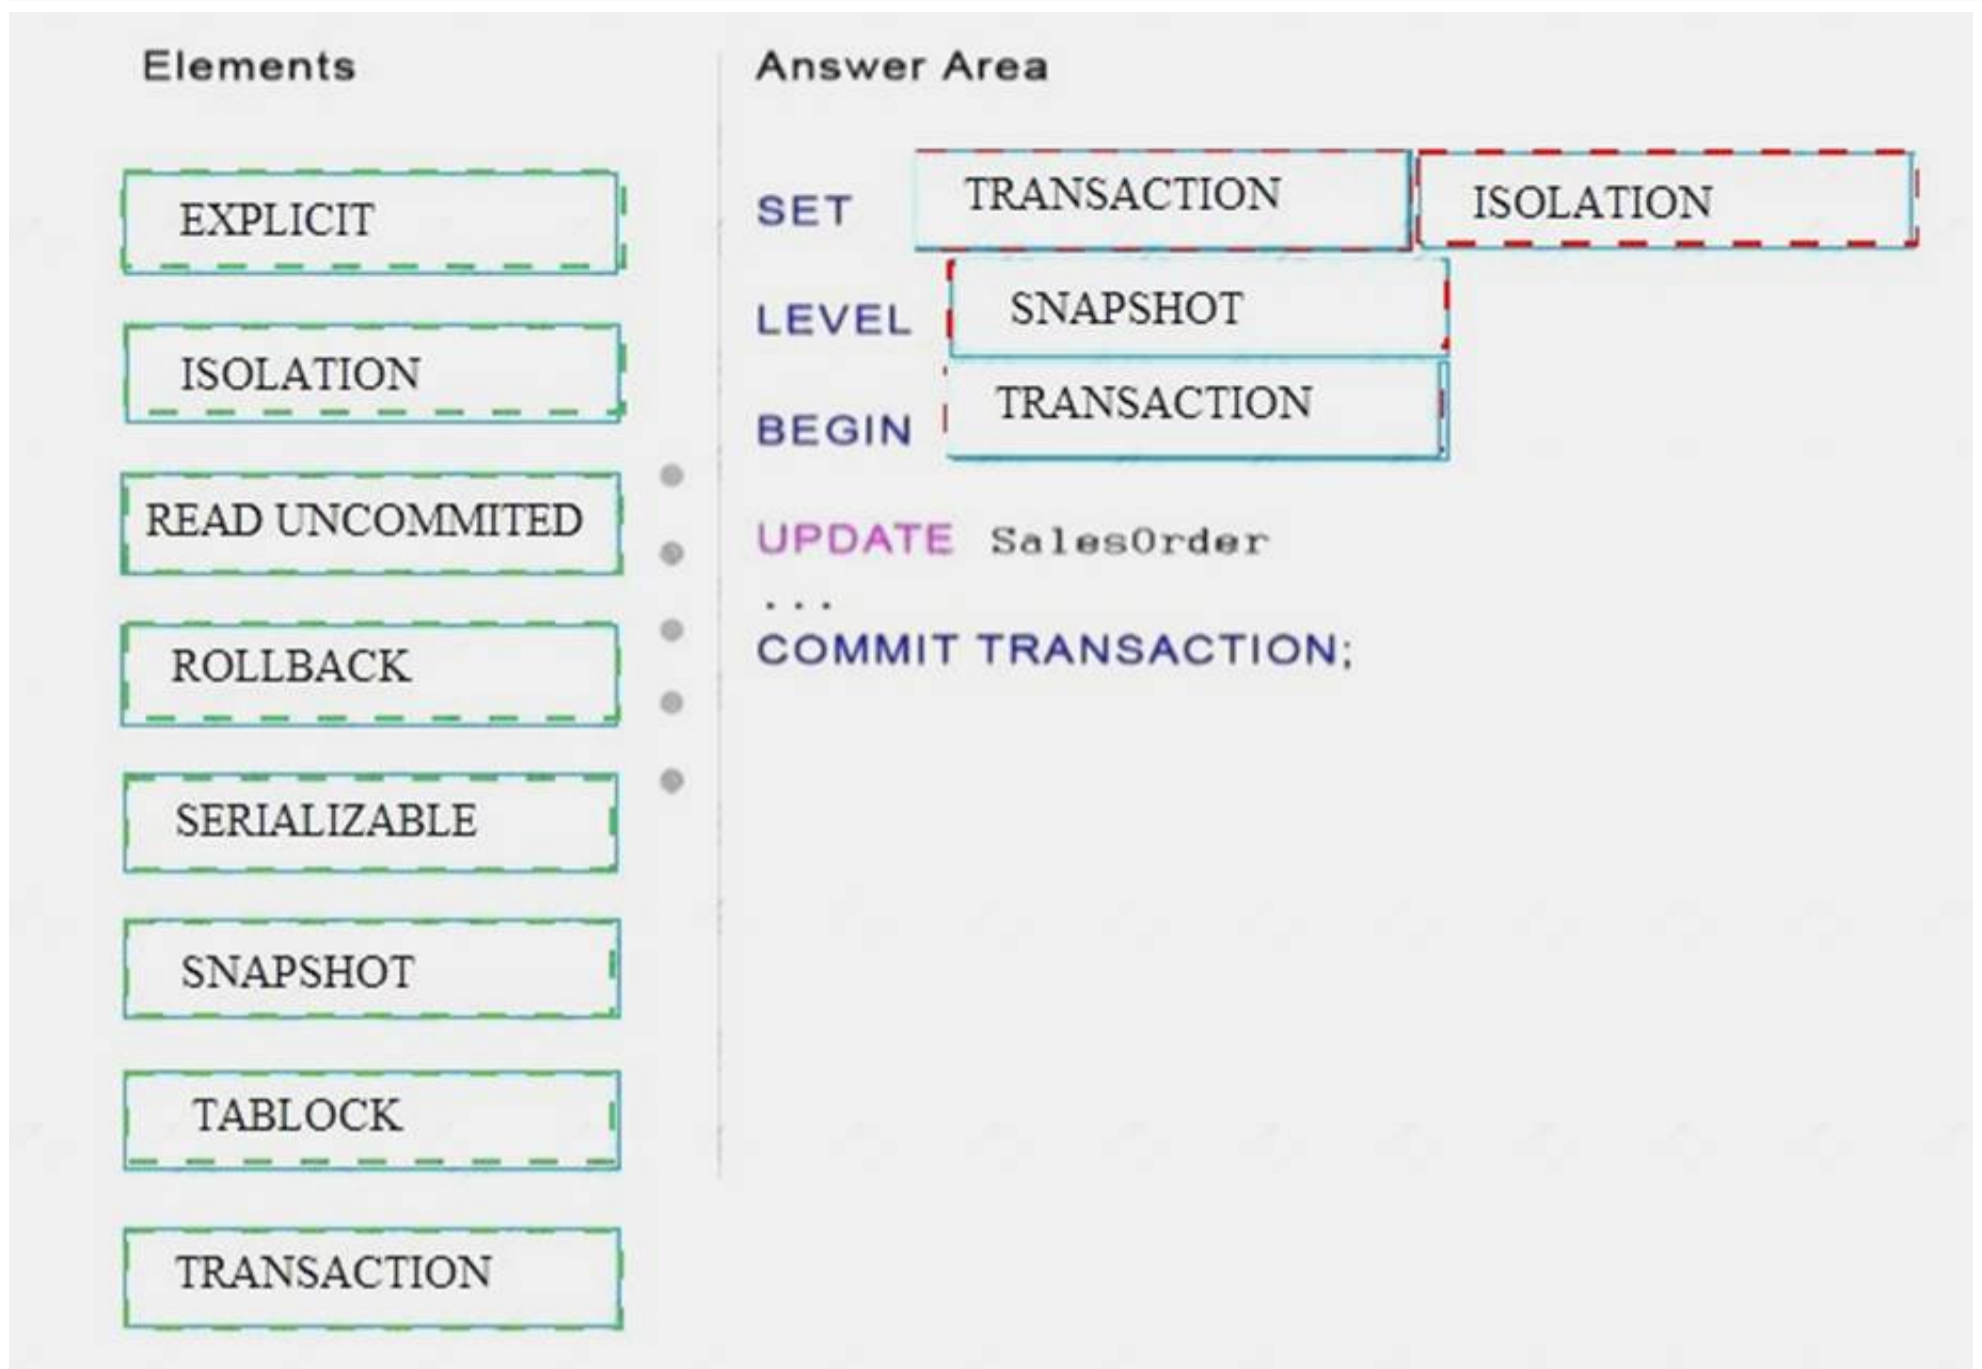
...

COMMIT TRANSACTION;

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:



#### NEW QUESTION 150

- (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.

During performance testing, you discover that database INSERT operations against the Inventory table are slow.

You need to recommend a solution to reduce the amount of time it takes to complete the INSERT operations. What should you recommend?

- A. Partition the nonclustered index.
- B. Partition the Inventory table.snapshot replication
- C. Create a column store index.Master Data Services

D. Drop the clustered index.change data capture

**Answer:** A

**Explanation:**

Scenario:

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.

**NEW QUESTION 154**

- (Exam Topic 2)

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

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

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

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

**Answer:** DE

**Explanation:**

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

References:

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

**NEW QUESTION 158**

- (Exam Topic 2)

You administer a Microsoft SQL Server 2016 database that contains a table named AccountTransaction. You discover that query performance on the table is poor due to fragmentation on the

IDX\_AccountTransaction\_AccountCode non-clustered index.

You need to defragment the index.

You also need to ensure that user queries are able to use the index during the defragmenting process. Which Transact-SQL batch should you use?

- A. ALTER INDEX IDX\_AccountTransaction\_AccountCode ONAccountTransaction.AccountCode REORGANIZE
- B. ALTER INDEX ALL ON AccountTransaction REBUILD
- C. ALTER INDEX IDX\_AccountTransaction\_AccountCode ONAccountTransaction.AccountCode REBUILD
- D. CREATE INDEX IDXAccountTransactionAccountCode ONAccountTransaction.AccountCode WITH DROP EXISTING

**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 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 relate to the certificate.

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

**NEW QUESTION 161**

- (Exam Topic 3)

Overview

General Overview

ADatum Corporation has offices in Miami and Montreal.

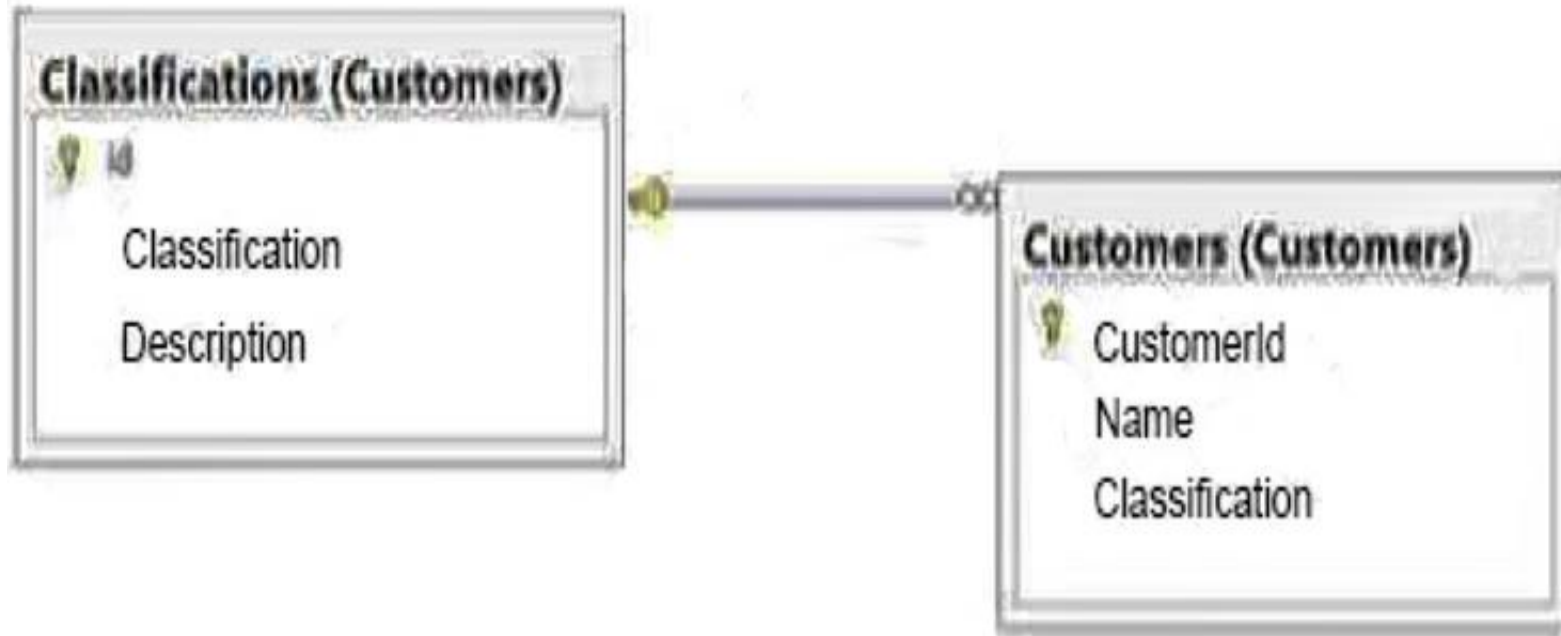
The network contains a single Active Directory forest named adatum.com. The offices connect to each other by using a WAN link that has 5-ms latency. A. Datum standardizes its database platform by using SQL Server 2014 Enterprise edition.

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. A view
- B. A temporary table
- C. A table variable
- D. A function

**Answer: A**

**Explanation:**

- 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.

#### NEW QUESTION 163

- (Exam Topic 3)

General Overview

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

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

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

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

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

**Current System**  
The current desktop application uses data stored in a SQL Server 2005 database named DesABCOppAppDB. 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 164

- (Exam Topic 3)

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

- Rebuild the indexes by using a new fill factor.
- Run a custom T-SQL command.
- Back up the databases.

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

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

**Answer:** C

#### Explanation:

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

#### NEW QUESTION 166

- (Exam Topic 3)

##### General Overview

You are the Senior Database Administrator (DBA) for a manufacturing company named Fairstone Manufacturing.

Fairstone Manufacturing is based in the New York area. The company has two offices: a main office in the city and a branch office just outside the city. The company has four factories where their products are manufactured. Two factories are in the New York area and the other two factories are in Washington.

##### Network Connectivity

The two offices are connected by a 10 Mbps dedicated WAN link. SQL Server Environment

The main office has four SQL Server 2012 Standard Edition servers named MainDB1, MainDB2, MainDB3 and MainDB4. The branch office has two SQL Server 2012 Standard Edition servers named BranchDB1 and BranchDB2. The main office has a Development department. All databases used by the Development department are hosted on MainDB3 and MainDB4. MainDB1 and MainDB2 host the following databases:

Products Manufacturing Sales

HR

Customers DailyReportsTemp

BranchDB1 and BranchDB2 host the same databases as MainDB1 and MainDB2. The DailyReportsTemp database is a temporary database that is recreated every day and used for reporting purposes.

One of the tables in the Customer database lists all the customers. Another table linked to the customers table contains a list of classifications for the customers.

The classifications are Hot, Warm and Cold based on the number of orders placed by the customers in the last year. The customers are classified according to the following criteria:

Hot - Over 100 orders placed in a year.

Warm - Between 50 and 100 orders placed in a year. Cold - Under 50 orders placed in a year.

Stored Procedures

Three tables in the Manufacturing database are modified by a stored procedure named ManProc1. A segment of code from ManProc1 is as follows:

```
CREATE PROCEDURE Manufacturing.ManProc1
AS
BEGIN TRANSACTION
UPDATE Manufacturing.Type ...
UPDATE Manufacturing.Version ...
UPDATE Manufacturing.Revision ...
COMMIT TRANSACTION
GO
```

The same three tables are also modified by a stored procedure named ManProc2. A segment of code from ManProc2 is as follows:

```
CREATE PROCEDURE Manufacturing.ManProc2
AS
BEGIN TRANSACTION
UPDATE Manufacturing.Revision ...
UPDATE Manufacturing.Type ...
UPDATE Manufacturing.Version ...
COMMIT TRANSACTION
GO
```

A product list in the Products database is updated using information from tables in the Manufacturing database by a stored procedure named ProductUpdateProc. Locks on tables in the Manufacturing database often cause ProductUpdateProc to take a long time to complete.

A list of manufacturing processes required to create each product is stored in tables in the Manufacturing database and updated by a stored procedure named ProcessUpdateProc. The ProcessUpdateProc stored procedure contains several UPDATE statements. The UPDATE statements are configured to be called in a specific order. The ProcessUpdateProc stored procedure continues to run in the event of a failure of one of the UPDATE statements.

This can cause inaccurate results in the manufacturing process list. Sales Director Statement

The Sales Director has made the following observations about the current database design:

The current customer classification system needs to be changed.

Currently the customers are classified by the number of orders placed in the last year.

This information is an unreliable guide as it does not take in to account the size of the orders.

I would suggest a trial run of a classification system based on the revenue generated by the orders placed in the last year.

We may add more than the current three classification types in future.

We should have a method of recording changes to the classifications.

IT Manager Statement

The IT Manager has listed the following requirements for the SQL Server and database environment:

We need to provide a group of users from the IT and Manufacturing departments the minimum administrative rights to view database information and server state for the Manufacturing database on MainDB1.

The Sales database takes too long to back up due to the large amount of historical sales order data in the database. We need to reduce the backup time for this database.

The DailyReportsTemp database takes four hours to back up. We need to be able to recover the DailyReportsTemp database in less than one hour if the database storage hardware fails.

We need to be able to immediately return the Manufacturing database to its previous state if the ProcessUpdateProc stored procedure fails to update the process information correctly.

I also want the ProcessUpdateProc stored procedure to stop running in the event of a failure of one of the UPDATE statements.

IT Administrators need to be able to monitor the disk space used on the SQL Servers by running real-time reports on the disk usage.

The Developers would like to install second instances of SQL Server on MainDB3 and MainDB4.

They would like to assign each instance to specific processors on the SQL Servers.

You need to enable the Developers to assign SQL Server instances on MainDB3 and MainDB4 to specific processors on the servers. What should you configure?

A. Windows System Resource Manager (WSRM)

B. Resource Governor

C. A Maintenance Plan

D. Processor Affinity

**Answer: D**

## NEW QUESTION 168

- (Exam Topic 3)

You administer a Microsoft SQL Server 2012 instance.

You need to configure a new database to support FILETABLES. What should you do? Choose all that apply.

A. Disable FILESTREAM on the Database.

B. Enable FILESTREAM on the Server Instance.

C. Configure the Database for Partial Containment.

D. Create a non-empty FILESTREAM file group.

E. Enable Contained Databases on the Server Instance.

F. Set the FILESTREAM directory name on the Database.

**Answer: BDF**

## Explanation:

References:

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

#### NEW QUESTION 170

- (Exam Topic 3)

General Overview

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

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

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

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

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

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

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

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

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

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

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

The GenInfo table is used for reports.

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

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

Current System

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

SQL Servers

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

Design Requirements

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

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

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

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

Deadlocks must be analyzed with the use of Deadlock Graphs.

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

Software licensing and database storage costs must be minimized.

Development effort must be minimized.

The Tempdb databases must be monitored for insufficient free space.

Failed authentication requests must be logged.

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

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

SPUpdateSalesInfo queries data in the OrderTotals table the second time.

You need to configure a synchronization solution to copy data from the Current\_Inventory database the DesABCOppAppDB database.

What should you configure?

- A. Transactional Replication.
- B. Database Mirroring.
- C. Snapshot Replication.
- D. Incremental Backups

**Answer: A**

#### NEW QUESTION 171

- (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 175

- (Exam Topic 3)

You are using dynamic management views to monitor an SQL Server server named SQL1. A database administrator named Dba1 must monitor the health of SQL1.

You need to ensure that Dba1 can access dynamic management views for SQL1.

The solution must use the principle of least privilege. Which permissions should you assign to Dba1?

- 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 177**

- (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)
WITH SAMPLE 100 PERCENT
GO
```

Does this meet the goal?

- A. Yes
- B. No

**Answer:** A

**Explanation:**

FULLSCAN: Compute statistics by scanning all rows in the table or indexed view. FULLSCAN and SAMPLE 100 PERCENT have the same results.

References: [https://technet.microsoft.com/en-us/library/ms187348\(v=sql.110\).aspx](https://technet.microsoft.com/en-us/library/ms187348(v=sql.110).aspx)

**NEW QUESTION 182**

- (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)
WITH FULLSCAN, NORECOMPUTE
GO
```

Does this meet the goal?

- A. Yes
- B. No

**Answer:** A

**Explanation:**

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

USE AdventureWorks2012;

GO

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

Note: NORECOMPUTE

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

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

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

**NEW QUESTION 186**

.....

## Relate Links

**100% Pass Your 70-764 Exam with ExamBible Prep Materials**

<https://www.exambible.com/70-764-exam/>

## Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - <https://www.exambible.com/>