

Microsoft

Exam 70-465

Designing Database Solutions for Microsoft SQL Server

Version: Demo

[Total Questions: 10]

Topic break down

Topic	No. of Questions
Topic 1: Litware, Inc	1
Topic 3: Fabrikam, Inc	1
Topic 5: Mix Questions	8

Topic 1, Litware, Inc

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. 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.

Question No : 1 - (Topic 1)

You need to recommend a disk monitoring solution that meets the business requirements.

What should you include in the recommendation?

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

Answer: B

Reference:

Dynamic Management Views and Functions (Transact-SQL)

Topic 3, Fabrikam, Inc Background

Corporate Information

Fabrikam, Inc. is a retailer that sells electronics products on the Internet. The company has

a headquarters site and one satellite sales office.

You have been hired as the database administrator, and the company wants you to change the architecture of the Fabrikam ecommerce site to optimize performance and reduce downtime while keeping capital expenditures to a minimum. To help with the solution, Fabrikam has decided to use cloud resources as well as on-premise servers.

Physical Locations

All of the corporate executives, product managers, and support staff are stationed at the headquarters office. Half of the sales force works at this location.

There is also a satellite sales office. The other half of the sales force works at the satellite office in order to have sales people closer to clients in that area. Only sales people work at the satellite location.

Problem Statement

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

Configuration

Windows Logins

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

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

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

Server Configuration

The IT department has configured two physical servers with Microsoft Windows Server 2012 R2 and SQL Server 2014 Enterprise Edition and one Windows Azure Server. There

are two tiers of storage available for use by database files only a fast tier and a slower tier. Currently the data and log files are stored on the fast tier of storage only. If a possible use case exists, management would like to utilize the slower tier storage for data files.

The servers are configured as shown in the following table:

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

Database

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

```
CREATE DATABASE ProdDB  
GO
```

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

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

The columns in the Product table are defined as shown in the following table:

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

The SalesOrderDetail table holds the details about each sale. It is in the Sales schema owned by the SalesStaff Windows group.

This table is constantly being updated, inserted into, and read.

The columns in the SalesOrderDetail table are defined as shown in the following table:

Column	Data type
SalesOrderDetailID	INT
ProductID	INT
SalePrice	SMALLMONEY
SaleQuantity	INT

Database Issues

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

Requirements

Database

The database should be configured to maximize uptime and to ensure that very little data is lost in the event of a server failure. To help with performance, the database needs to be modified so that it can support in-memory data, specifically for the Product table, which the CIO has indicated should be a memory-optimized table. The auto-update statistics option is set off on this database.

Only product managers are allowed to add products or to make changes to the name, description, price, cost, and supplier. The changes are made in an internal database and pushed to the Product table in ProdDB during system maintenance time. Product managers and others working at the headquarters location also should be able to generate reports that include supplier and cost information.

Customer data access

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

Customer support data access

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

Sales force data access

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

Historical Data

The system should keep historical information about customers who access the site so that sales people can see how frequently customers log in and how long they stay on the site. The information should be stored in a table called Customer Access. Supporting this requirement should have minimal impact on production website performance.

Backups

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

Database Maintenance

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

Project milestones completed

- ✍ Revoked all existing read and write access to the database, leaving the schema ownership in place.
- ✍ Configured an Azure storage container secured with the storage account name MyStorageAccount with the primary access key StorageAccountKey on the cloud file server.
- ✍ SQL Server 2014 has been configured on the satellite server and is ready for use.
- ✍ On each database server, the fast storage has been assigned to drive letter F:, and the slow storage has been assigned to drive letter D:.

Question No : 2 DRAG DROP - (Topic 3)

The business requires a satellite office to have a local copy of the data to report against.

You want to implement a solution to support the requirements. You need to establish a new Availability Group between the two servers.

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

Code Blocks

```
EXEC master.dbo.sp_addlinkedserver @server
= N'SATELLITE_SERVER', @srvproduct=N'SQL Server'
```

```
CREATE ENDPOINT [Hadr_endpoint]
STATE = STARTED AS TCP (LISTENER_PORT =
5022)
FOR DATA_MIRRORING (ROLE = ALL)
```

```
ALTER AVAILABILITY GROUP
[ProdDB_AG] JOIN;
```

```
CREATE AVAILABILITY GROUP [ProdDB_AG]
FOR DATABASE [ProdDB] REPLICA ON
N'HQ_Server' WITH
(Endpoint_URL = N'TCP://
HQ_Server.OurDomain.com:5022',
FAILOVER_MODE = MANUAL,
AVAILABILITY_MODE =
ASYNCHRONOUS_COMMIT,
SECONDARY_ROLE(ALLOW_CONNECTIONS =
READ_ONLY)),
N'Satellite_Server' WITH
(Endpoint_URL = N'TCP://
Satellite_Server.OurDomain.com:5022',
FAILOVER_MODE = MANUAL,
AVAILABILITY_MODE =
ASYNCHRONOUS_COMMIT,
SECONDARY_ROLE(ALLOW_CONNECTIONS =
READ_ONLY));
```

```
CREATE ENDPOINT sql_endpoint
STATE = STARTED
AS HTTP( PATH = '/alwayson',
AUTHENTICATION = ( INTEGRATED ),
PORTS = ( CLEAR ),
SITE = 'SERVER');
```

```
CREATE AVAILABILITY GROUP [ProdDB_AG]
FOR DATABASE [ProdDB] REPLICA ON
N'HQ_Server' WITH
(Endpoint_URL = N'TCP://
HQ_Server.OurDomain.com:5022',
FAILOVER_MODE = MANUAL,
AVAILABILITY_MODE = SYNCHRONOUS_COMMIT,
SECONDARY_ROLE(ALLOW_CONNECTIONS =
READ_ONLY)),
N'Satellite_Server' WITH
(Endpoint_URL = N'TCP://
Satellite_Server.OurDomain.com:5022',
FAILOVER_MODE = MANUAL,
AVAILABILITY_MODE = SYNCHRONOUS_COMMIT,
SECONDARY_ROLE(ALLOW_CONNECTIONS =
READ_ONLY));
```

```
ALTER AVAILABILITY GROUP [ProdDB_AG]
ADD LISTENER N'ProdDB_AG_Listener' (
WITH PORT=1433);
```

```
ALTER DATABASE [ProdDB] SET HADR
AVAILABILITY GROUP = [ProdDB_AG];
```

```
RESTORE DATABASE [ProdDB] FROM DISK =
N'\\FILESERVER\SQLbackups\ProdDB.bak'
```

```
BACKUP LOG [ProdDB] TO
DISK = N'\\FILESERVER\SQLbackups
\ProdDB.trn'
WITH NORECOVERY
```

Answer Area

Answer:

Code Blocks	Answer Area
<pre>EXEC master.dbo.sp_addlinkedserver @server = N'SATELLITE_SERVER', @srvproduct=N'SQL Server'</pre>	<pre>EXEC master.dbo.sp_addlinkedserver @server = N'SATELLITE_SERVER', @srvproduct=N'SQL Server'</pre>
<pre>CREATE ENDPOINT [Hadr_endpoint] STATE = STARTED AS TCP (LISTENER_PORT = 5022) FOR DATA_MIRRORING (ROLE = ALL)</pre>	<pre>CREATE ENDPOINT [Hadr_endpoint] STATE = STARTED AS TCP (LISTENER_PORT = 5022) FOR DATA_MIRRORING (ROLE = ALL)</pre>
<pre>ALTER AVAILABILITY GROUP [ProdDB_AG] JOIN;</pre>	<pre>CREATE AVAILABILITY GROUP [ProdDB_AG] FOR DATABASE [ProdDB] REPLICA ON N'HQ_Server' WITH (ENDPOINT_URL = N'TCP:// HQ_Server.OurDomain.com:5022', FAILOVER_MODE = MANUAL, AVAILABILITY_MODE = ASYNCHRONOUS_COMMIT, SECONDARY_ROLE(ALLOW_CONNECTIONS = READ_ONLY)), N'Satellite_Server' WITH (ENDPOINT_URL = N'TCP:// Satellite_Server.OurDomain.com:5022', FAILOVER_MODE = MANUAL, AVAILABILITY_MODE = ASYNCHRONOUS_COMMIT, SECONDARY_ROLE(ALLOW_CONNECTIONS = READ_ONLY));</pre>
<pre>CREATE AVAILABILITY GROUP [ProdDB_AG] FOR DATABASE [ProdDB] REPLICA ON N'HQ_Server' WITH (ENDPOINT_URL = N'TCP:// HQ_Server.OurDomain.com:5022', FAILOVER_MODE = MANUAL, AVAILABILITY_MODE = ASYNCHRONOUS_COMMIT, SECONDARY_ROLE(ALLOW_CONNECTIONS = READ_ONLY)), N'Satellite_Server' WITH (ENDPOINT_URL = N'TCP:// Satellite_Server.OurDomain.com:5022', FAILOVER_MODE = MANUAL, AVAILABILITY_MODE = ASYNCHRONOUS_COMMIT, SECONDARY_ROLE(ALLOW_CONNECTIONS = READ_ONLY));</pre>	<pre>ALTER AVAILABILITY GROUP [ProdDB_AG] JOIN;</pre>
<pre>CREATE ENDPOINT sql_endpoint STATE = STARTED AS HTTP(PATH = '/alwayson', AUTHENTICATION = (INTEGRATED), PORTS = (CLEAR), SITE = 'SERVER');</pre>	<pre>BACKUP LOG [ProdDB] TO DISK = N'\\FILESERVER\SQLbackups \ProdDB.trn' WITH NORECOVERY</pre>
<pre>CREATE AVAILABILITY GROUP [ProdDB_AG] FOR DATABASE [ProdDB] REPLICA ON N'HQ_Server' WITH (ENDPOINT_URL = N'TCP:// HQ_Server.OurDomain.com:5022', FAILOVER_MODE = MANUAL, AVAILABILITY_MODE = SYNCHRONOUS_COMMIT, SECONDARY_ROLE(ALLOW_CONNECTIONS = READ_ONLY)), N'Satellite_Server' WITH (ENDPOINT_URL = N'TCP:// Satellite_Server.OurDomain.com:5022', FAILOVER_MODE = MANUAL, AVAILABILITY_MODE = SYNCHRONOUS_COMMIT, SECONDARY_ROLE(ALLOW_CONNECTIONS = READ_ONLY));</pre>	<pre>RESTORE DATABASE [ProdDB] FROM DISK = N'\\FILESERVER\SQLbackups\ProdDB.bak'</pre>
<pre>ALTER AVAILABILITY GROUP [ProdDB_AG] ADD LISTENER N'ProdDB_AG_Listener' (WITH PORT=1433);</pre>	<pre>ALTER DATABASE [ProdDB] SET HADR AVAILABILITY GROUP = [ProdDB_AG];</pre>
<pre>ALTER DATABASE [ProdDB] SET HADR AVAILABILITY GROUP = [ProdDB_AG];</pre>	
<pre>RESTORE DATABASE [ProdDB] FROM DISK = N'\\FILESERVER\SQLbackups\ProdDB.bak'</pre>	
<pre>BACKUP LOG [ProdDB] TO DISK = N'\\FILESERVER\SQLbackups \ProdDB.trn' WITH NORECOVERY</pre>	

Explanation:

Box 1:

```
EXEC master.dbo.sp_addlinkedserver @server  
= N'SATELLITE_SERVER', @srvproduct=N'SQL Server'
```

Box 2:

```
CREATE ENDPOINT [Hadr_endpoint]  
STATE = STARTED AS TCP (LISTENER_PORT =  
5022)  
FOR DATA_MIRRORING (ROLE = ALL)
```

Box 3:

```
CREATE AVAILABILITY GROUP [ProdDB_AG]  
FOR DATABASE [ProdDB] REPLICA ON  
N'HQ_Server' WITH  
(ENDPOINT_URL = N'TCP://  
HQ_Server.OurDomain.com:5022',  
FAILOVER_MODE = MANUAL,  
AVAILABILITY_MODE =  
ASYNCHRONOUS_COMMIT,  
SECONDARY_ROLE(ALLOW_CONNECTIONS =  
READ_ONLY)),  
N'Satellite_Server' WITH  
(ENDPOINT_URL = N'TCP://  
Satellite_Server.OurDomain.com:5022',  
FAILOVER_MODE = MANUAL,  
AVAILABILITY_MODE =  
ASYNCHRONOUS_COMMIT,  
SECONDARY_ROLE(ALLOW_CONNECTIONS =  
READ_ONLY));
```

Box 4:

```
ALTER AVAILABILITY GROUP  
[ProdDB_AG] JOIN;
```

Box 5:

```
BACKUP LOG [ProdDB] TO  
DISK = N'\\FILESERVER\SQLbackups  
\ProdDB.trn'  
WITH NORECOVERY
```

Box 6:

```
RESTORE DATABASE [ProdDB] FROM DISK =  
N'\\FILESERVER\SQLbackups\ProdDB.bak'
```

Box 7:

```
ALTER DATABASE [ProdDB] SET HADR  
AVAILABILITY GROUP = [ProdDB_AG];
```

Note:

* The following table lists the basic tasks involved in creating and configuring an availability group and indicates which Transact-SQL statements to use for these tasks. The AlwaysOn Availability Groups tasks must be performed in the sequence in which they are presented in the table.

(step 2) Create database mirroring endpoint (once per SQL Server instance)

CREATE ENDPOINT endpointName ... FOR DATABASE_MIRRORING

(step 3) Create availability group

CREATE AVAILABILITY GROUP

(step 4) Join secondary replica to availability group

ALTER AVAILABILITY GROUP group_name JOIN

(step 5-6)

Prepare the secondary database

BACKUP and RESTORE.

Create backups on the server instance that hosts the primary replica.

Restore backups on each server instance that hosts a secondary replica, using RESTORE WITH NORECOVERY.

(step 7)

Start data synchronization by joining each secondary database to availability group




ALTER DATABASE database_name SET HADR AVAILABILITY GROUP = group_name

Topic 5, Mix Questions

Question No : 3 - (Topic 5)

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.

Reference: Maintenance Plans

Question No : 4 - (Topic 5)

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. SQL Server 2014 Enterprise
- B. SQL Server 2014 Standard
- C. SQL Azure
- D. SQL Server 2014 Express with Advanced Services

Answer: B

Explanation:

Programmability (AMO, ADOMD.Net, OLEDB, XML/A, ASSL) supported by Standard and Enterprise editions only.

Reference: Features Supported by the Editions of SQL Server 2014

Question No : 5 - (Topic 5)

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.

Reference:

sys.dm_db_wait_stats (Windows Azure SQL Database)

Question No : 6 - (Topic 5)

You have two SQL Server 2014 instances named SQLDev and SQLProd.

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

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

What should you use?

More than one answer choice may achieve the goal. Select the BEST answer.

A. Triggers

B. Contained database

C. SQL Server Analysis Services (SSAS) scripts

D. Extended Events

E. SQL Server Integration Services (SSIS) scripts

Answer: B

Explanation:

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





Reference: Contained Databases

Question No : 7 - (Topic 5)

You administer a SQL Server 2014 instance.

Users report that the SQL Server has seemed slow today. A large database was being restored for much of the day, which could be causing issues.

You want to write a query of the system views that will report the following:

-  Number of users that have a connection to the server
-  Whether a user's connection is active
-  Whether any connections are blocked
-  What queries are being executed

Whether the database restore is still executing and, if it is, what percentage of the restore is complete.

Which system objects should you use in your query to best achieve this task?

- A.** sys.dm_exec_requests, sys.dm_exec_sessions, sys.objects
- B.** sys.dm_exec_sessions, sys.dm_exec_query_stats, sys.dm_exec_query_text, sys.objects
- C.** sys.sysprocesses, sys.dm_exec_query_text, sys.objects
- D.** sys.dm_exec_requests, sys.dm_exec_sessions, sys.dm_exec_query_text

Answer: D

Explanation:

* sys.dm_exec_requests

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

* sys.dm_exec_sessions

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

* sys.dm_exec_query_text

Returns the text of the SQL batch that is identified by the specified sql_handle.

References:

sys.dm_exec_requests (Transact-SQL)

sys.dm_exec_sessions (Transact-SQL)

Incorrect:

* sys.dm_exec_query_stats

Returns aggregate performance statistics for cached query plans in SQL Server. The view contains one row per query statement within the cached plan, and the lifetime of the rows are tied to the plan itself.

* sys.objects

Contains a row for each user-defined, schema-scoped object that is created within a database.

Question No : 8 DRAG DROP - (Topic 5)

You plan to deploy three highly available SQL Server environments that will use SQL Server 2014.

You identify the following specifications for each environment as shown following table.

Environment	Number of nodes	SQL Server edition	Automatic failover required
ENV1	3	Standard	Yes
ENV2	3	Enterprise	Yes
ENV3	4	Enterprise	Yes

You need to recommend which high-availability feature is required for each environment.

Which features should you identify?

To answer, drag the appropriate feature to the correct requirement in the answer area.

Features		Answer Area
AlwaysOn availability groups	ENV1	Feature
database mirroring	ENV2	Feature
log shipping	ENV3	Feature
peer-to-peer replication		

Answer:

Features		Answer Area
AlwaysOn availability groups	ENV1	database mirroring
database mirroring	ENV2	AlwaysOn availability groups
log shipping	ENV3	AlwaysOn availability groups
peer-to-peer replication		

Explanation:

Features		Answer Area
AlwaysOn availability groups	ENV1	database mirroring
database mirroring	ENV2	AlwaysOn availability groups
log shipping	ENV3	AlwaysOn availability groups
peer-to-peer replication		

* Always on availability groups

The AlwaysOn Availability Groups feature is a high-availability and disaster-recovery solution that provides an enterprise-level alternative to database mirroring. Introduced in SQL Server 2012, AlwaysOn Availability Groups maximizes the availability of a set of user databases for an enterprise.

Deploying AlwaysOn Availability Groups requires a Windows Server Failover Clustering (WSFC) cluster.

References:



AlwaysOn Availability Groups (SQL Server)

Features Supported by the Editions of SQL Server 2014

Question No : 9 DRAG DROP - (Topic 5)

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		

Answer:

Isolation Levels		Answer Area
read committed	SP1	repeatable read
read uncommitted	SP2	read uncommitted
repeatable read		
serializable		

Explanation:

Isolation Levels		Answer Area
read committed	SP1	repeatable read
	SP2	read uncommitted
serializable		

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

Question No : 10 - (Topic 5)

You have a SQL Server 2014 database named DB1.

You plan to import a large number of records from a SQL Azure database to DB1.

You need to recommend a solution to minimize the amount of space used in the transaction log during the import operation.

What should you include in the recommendation?

- A. The bulk-logged recovery model
- B. The full recovery model
- C. A new partitioned table
- D. A new log file
- E. A new file group

Answer: A

Explanation:

Compared to the full recovery model, which fully logs all transactions, the bulk-logged recovery model minimally logs bulk operations, although fully logging other transactions. The bulk-logged recovery model protects against media failure and, for bulk operations, provides the best performance and least log space usage.

Note:

The bulk-logged recovery model is a special-purpose recovery model that should be used only intermittently to improve the performance of certain large-scale bulk operations, such as bulk imports of large amounts of data.

Reference: Recovery Models (SQL Server)

