Microsoft MCSE Certification 70-464 Exam

- **Vendor: Microsoft**
- **Exam Code: 70-464**
- **Exam Name: Developing Microsoft SQL Server 2012/2014 Databases**

QUESTION 1
Your company has a SQL Azure subscription. You implement a database named Database1. Database1 has two tables named Table1 and Table2. You create a stored procedure named sp1.Sp1 reads data from Table1 and inserts data into Table2. A user named User1 informs you that he is unable to run sp1. You verify that User1 has the SELECT permission on Table1 and Table2. You need to ensure that User1 can run sp1. The solution must minimize the number of permissions assigned to User1. What should you do?

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

Answer: D
Explanation:
http://msdn.microsoft.com/en-us/library/ms191291.aspx

QUESTION 2
You use SQL Server 2012 to maintain the data used by the applications at your company. You plan to create a table named Table1 by using the following statement. (Line numbers are included for reference only.) You need to ensure that Table1 contains a column named UserName. The UserName column will:
- Store string values in any language.
- Accept a maximum of 200 characters.
- Be case-insensitive and accent-insensitive.

Which code segment should you add at line 03?

```
01 CREATE TABLE dbo.table1(
02    ID int IDENTITY(1,1) NOT NULL,
03
04    Email varchar(100) NULL,
05    CONSTRAINT PK_table1 PRIMARY KEY CLUSTERED(ID ASC)
06 )
```

A. `UserName nvarchar(200) COLLATE Latin1_General_CS_AS NOT NULL,`
B. `UserName varchar(200) COLLATE Latin1_General_CI_AI NOT NULL,`
C. `UserName varchar(200) COLLATE Latin 1_General_CS_AS NOT NULL,`
D. `UserName nvarchar(200) COLLATE Latin1_General_CI_AI NOT NULL,`

Answer: D
Explanation:
A lot of the questions had a minor change that you need to pay attention to. Otherwise if you simply remember the answer, you will get it wrong.
Example 1. a question that had the table column specifications of "English" would normally result in the answer being "varchar", but the question now changed to "international", so you need to know that the answer is "nvarchar".
Example 2. "case sensitive, accent sensitive" would usually be "CS_AS", but now the question

changed to "case insensitive, accent sensitive" so its "CI_AS".
So while to the untrained eye, the questions appear the same. They are only 99% the same, but the specifications may have changed slightly. So pay attention.
http://msdn.microsoft.com/en-us/library/ms184391.aspx
http://msdn.microsoft.com/en-us/library/ms143726.aspx
http://msdn.microsoft.com/en-us/library/ff848763.aspx

QUESTION 3
Drag and Drop Question
You have a table named Table1 that contains 1 million rows. Table1 contains a column named Column1 that stores sensitive information. Column1 uses the nvarchar (16) data type. You have a certificate named Cert1. You need to replace Column1 with a new encrypted column named Column2 that uses one-way hashing. Which code segment should you execute before you remove Column1? To answer, move the appropriate code segments from the list of code segments to the answer area and arrange them in the correct order.

```
OPEN SYMMETRIC KEY Key1
DECRYPTION BY CERTIFICATE Cert1;
```

```
CREATE SYMMETRIC KEY Key1 WITH ALGORITHM = S
HA1
ENCRYPTION BY CERTIFICATE Cert1;
```

```
ALTER TABLE Table1
ADD Column2 nvarchar(256);
```

```
ALTER TABLE Table1
ADD Column2 varbinary(256);
```

```
CLOSE SYMMETRIC KEY;
```

```
CREATE CREDENTIAL Cred1 WITH IDENTITY = 'Use
r1', SECRET = 'P@ssw0rd';
```

```
UPDATE table1 SET Column2 = EncryptByKey
(Key_GUID('Key1'),Column1);
```

```
CREATE SYMMETRIC KEY Key1 WITH ALGORITHM = A
ES_256
ENCRYPTION BY CERTIFICATE Cert1;
```

Answer:

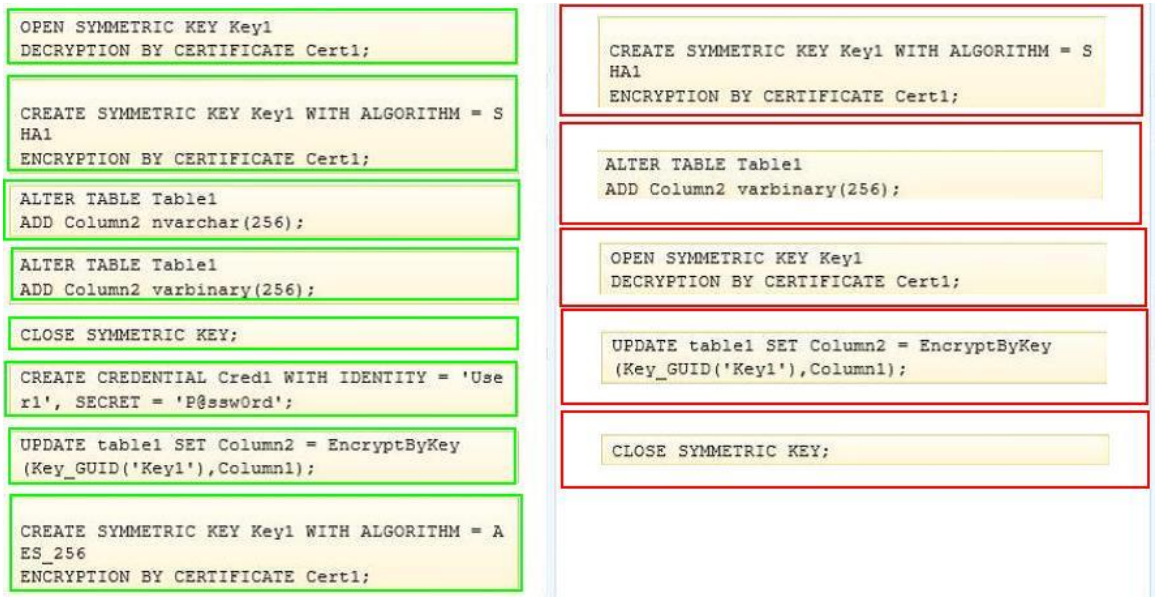

QUESTION 4
You review a query that runs slowly. The query accesses data in a table named Schema1.Table1. The following is the relevant portion of the execution plan for the query:

```
<MissingIndexes>
  <MissingIndexGroup Impact="95.8296">
    <MissingIndex Database="DB1" Schema="Schema1" Table="Table1">
      <ColumnGroup Usage="EQUALITY">
        <Column Name="Column1" ColumnId="14" />
      </ColumnGroup>
      <ColumnGroup Usage="INEQUALITY">
        <Column Name="Column2" ColumnId="17" />
        <Column Name="Column3" ColumnId="21" />
      </ColumnGroup>
      <ColumnGroup Usage="INCLUDE">
        <Column Name="Column4" ColumnId="11" />
      </ColumnGroup>
    </MissingIndex>
  </MissingIndexGroup>
</MissingIndexes>
```

You need to create the missing index. Which code segment should you execute?

A. CREATE NCNCLUSTERED INDEX 1X1 on Schema1.Table1 (Column1) INCLUDE (Column4)
B. CREATE NCNCLUSTERED INDEX 1X1 on Schema1.Table1 (Column1)
C. CREATE NONCLUSTERED INDEX 1X1 on Schema1.Table1 (Column1, Column2, Column3) INCLUDE(Column4)
D. CREATE NONCLUSTERED INDEX 1X1 on Schema1.Table1 (Column1) INCLUDE(Column4) WHERE Column2 <> Column3

Answer: C

QUESTION 5
You are creating a table to support an application that will cache data outside of SQL Server. The

application will detect whether cached values were changed before it updates the values. You need to create the table, and then verify that you can insert a row into the table. Which code segment should you use?

A.
```
CREATE TABLE Table1
(
  ID int IDENTITY(1,1),
  Name varchar(100),
  Version uniqueidentifier DEFAULT (NEWID())
)
INSERT INTO Table1 (Name, Version)
VALUES ('Smith, Ben', NEWID())
```

B.
```
CREATE TABLE Table1
(
  ID int IDENTITY(1,1),
  Name varchar(100),
  Version uniqueidentifier DEFAULT (NEWID())
)
INSERT INTO Table1 (Name)
VALUES ('Smith, Ben')
```

C.
```
CREATE TABLE Table1
(
  ID int IDENTITY(1,1),
  Name varchar(100),
  Version rowversion
)
INSERT INTO Table1 (Name)
VALUES ('Smith, Ben')
```

D.
```
CREATE TABLE Table1
(
  ID int IDENTITY(1,1),
  Name varchar(100),
  Version rowversion
)
INSERT INTO Table1 (Name, Version)
VALUES ('Smith, Ben', NEWID())
```

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

Answer: C
Explanation:

http://msdn.microsoft.com/en-us/library/ms182776.aspx
http://msdn.microsoft.com/en-us/library/ms187942.aspx
http://msdn.microsoft.com/en-us/library/ms190348.aspx

QUESTION 6
Drag and Drop Question
You run the following code segment:

```
CREATE TABLE dbo.Customers
(
  Id int CONSTRAINT Check_ID PRIMARY KEY,
  CustomerName varchar(50),
  Details xml
);
GO
CREATE PRIMARY XML INDEX PXML_Customers
  ON dbo.Customers (Details);
GO
```

After you add 10,000 rows to Customers, you discover that the index is fragmented. You need to defragment the index in the least amount of time. Which code segment should you execute? To answer, drag the appropriate value to the correct location in the code segment in the answer area. (Answer choices may be used once, more than once, or not at all.)

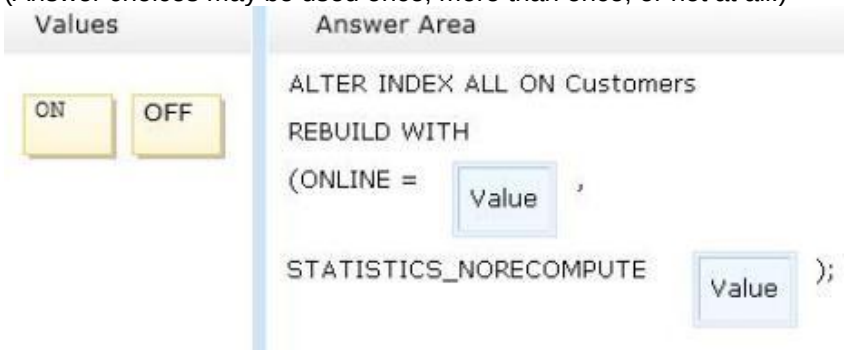

Answer:

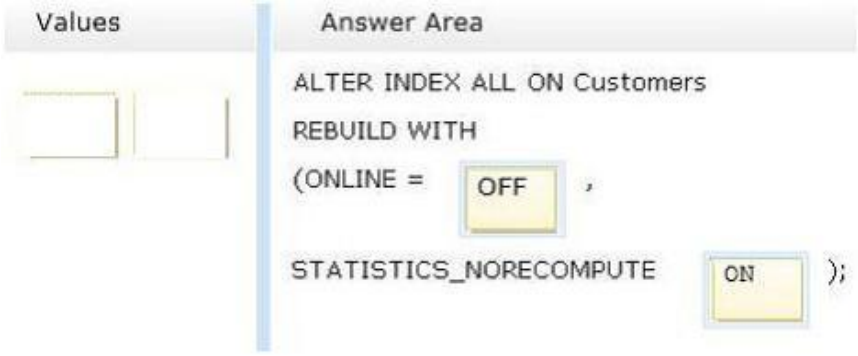

QUESTION 7
You execute the following code:

```
CREATE TABLE dbo.Projects
( Id int,
details XML);
GO

INSERT INTO Projects (Id,details)
VALUES
(1,
N'<Project Name="Project1">
<Tasks>
  <Task Name="T1"><IsFinished>true"</IsFinished></Task>
  <Task Name="T2"><IsFinished>true"</IsFinished></Task>
</Tasks>
</Project>'),
(2,
N'<Project Name="Project2">
<Tasks>
  <Task Name="T_1"><IsFinished>false</IsFinished></Task>
</Tasks>
</Project>');
```

You need to select the task that has an IsFinished value of true from the Project that has an Id value of 1. Which code segment should you use?

A.
```
SELECT Projects.details.query('Project/Tasks/Task/[@IsFinished="true"]')
FROM Projects
WHERE Projects.Id = 1;
```

B.
```
SELECT Projects.details.query('//Task/IsFinished="true"')
FROM Projects
WHERE Projects.Id = 1;
```

C.
```
SELECT Projects.details
FROM Projects
WHERE Projects.Id = 1 AND Details LIKE '%true%';
```

D.
```
SELECT Projects.details.query('//Task[@IsFinished="true"]')
FROM Projects
WHERE Projects.Id = 1;
```

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

Answer: B

QUESTION 8
You execute the following code:

```
CREATE TABLE dbo.Customers
(
  id int PRIMARY KEY,
  CustomerName char(10)
)
```

You create a nonclustered index named IX_CustomerName on the CustomerName column.

You execute the following query:

```
SELECT * FROM dbo.Customers
WHERE LEFT(CustomerName,1) = 'a'
```

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

A. Replace LEFT(CustomerName ,1) = 'a' with CustomerName LIKE 'a%'.
B. Partition the table and use the CustomerName column for the partition scheme.
C. Replace LEFT(CustomerName ,1) = 'a' with SUBSTRING(CustomerName ,1/1) = 'a'.
D. Replace IX_CustomerName with a clustered index.

Answer: A
Explanation:
http://msdn.microsoft.com/en-us/library/ms179859.aspx
http://msdn.microsoft.com/en-us/library/ms187748.aspx

QUESTION 9
You have a SQL Server 2012 database named DB1. You have a backup device named Device1. You discover that the log file for the database is full. You need to ensure that DB1 can complete transactions. The solution must not affect the chain of log sequence numbers (LSNs). Which code segment should you execute?

A. `BACKUP LCG DB1 TO Device1 WITH COPY_ONLY`
B. `BACKUP LOG DB1 TO Device1`
C. `BACKUP LOG DB1 TO Device1 WITH NCRECCVERY`
D. `BACKUP LOG D31 TO Device1 WITH TRUNCATE ONLY`

Answer: B
Explanation:
http://msdn.microsoft.com/en-us/library/ms186865.aspx
http://msdn.microsoft.com/en-us/library/ms179478.aspx
http://msdn.microsoft.com/en-us/library/ms190925.aspx

QUESTION 10
You execute the following code:

```
CREATE TABLE Department(
  DepartmentID smallint IDENTITY(1,1) NOT NULL,
  DepartmentName varchar(100) NOT NULL,
  GroupName varchar(100) NOT NULL,
  CONSTRAINT PK_Department_DepartmentID
    PRIMARY KEY CLUSTERED (DepartmentID ASC)
);
GO

CREATE UNIQUE NONCLUSTERED INDEX
  AK_Department_DepartmentName ON
  Department
(
  DepartmentName ASC
);
GO
```

You run the following query:

```
SELECT DepartmentID
      ,DepartmentName
      ,GroupName
  FROM Department
  WHERE DepartmentName = '1234';
```

The execution plan for the query is shown in the exhibit.

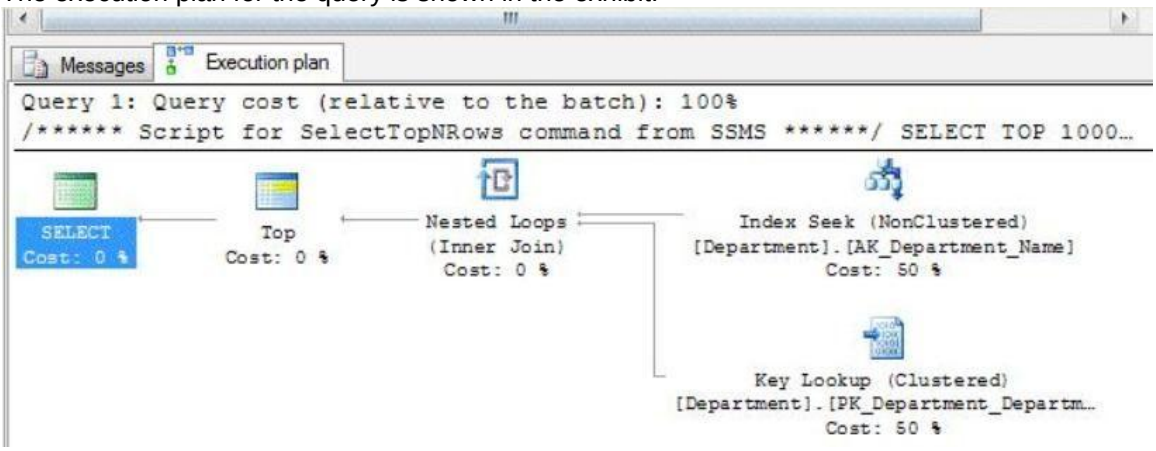

You need to prevent the key lookup.

A.
```
DROP INDEX AK_Department_DepartmentName
ON Department;
GO
CREATE INDEX AK_Department_DepartmentName
ON Department (DepartmentName, GroupName);
GO
```

B.
```
the SELECT statement to use the WITH(INDEX(AK_Department_DepartmentName)) query hint
```

C.
```
DROP INDEX AK_Department_DepartmentName
ON Department;
GO
CREATE INDEX AK_Department_DepartmentName
ON Department (DepartmentName)
INCLUDE (GroupName);
GO
```

D.
```
the SELECT statement to use the WITH(INDEX( PK_Department_DepartmentID)) query hint
```

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

Answer: C

QUESTION 11
You have a database for a mission-critical web application. The database is stored on a SQL Server 2012 instance and is the only database on the instance. The application generates all T-SQL statements dynamically and does not use stored procedures. You need to maximize the amount of memory available for data caching. Which advanced server option should you modify?

A. scan for Startup Procs
B. Allow Triggers to Fire Others
C. Enable Contained Databases
D. Optimize for Ad hoc Workloads

Answer: D

QUESTION 12
Drag and Drop Question
You have a database that contains three tables. The tables are configured as shown in the following table:

Table	Primary key index
SalesOrderHeader	PK_SalesOrderHeader_SalesOrderID
Employee	PK_Employee_EmployeeID
Contact	PK_Contact_ContactID

You have the following query:

```
SELECT soh.SalesPersonID,
  c.FirstName + ' ' + c.LastName AS FullName,
  c.EmailAddress,
  e.Title,
  soh.SubTotal,
  YEAR(soh.OrderDate) AS Year
FROM SalesOrderHeader soh
  INNER JOIN Employee e
    ON soh.SalesPersonID = e.EmployeeID
  INNER JOIN Contact c
    ON e.ContactID = c.ContactID
WHERE soh.OrderDate >= '1/1/2012'
```

The execution plan for the query is shown in the exhibit.

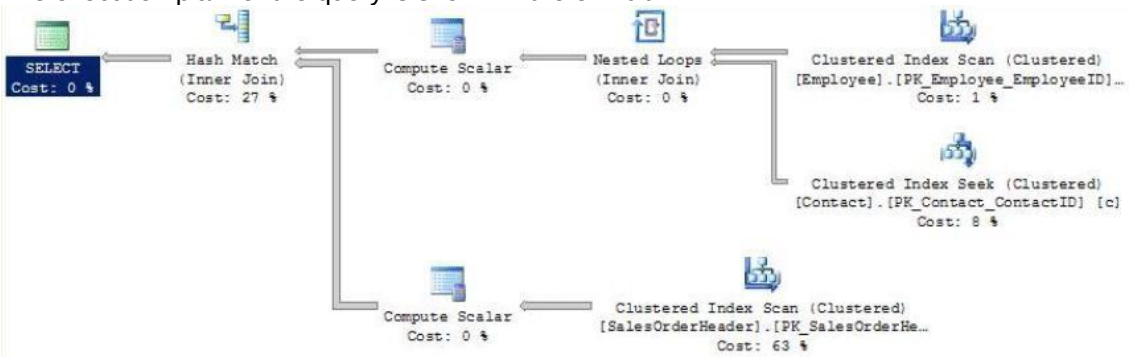

You need to create one index to minimize the amount of time it takes to execute the query. What should you do? To answer, drag the appropriate columns to the correct locations in the answer area. (Answer choices may be used once, more than once, or not at all.)

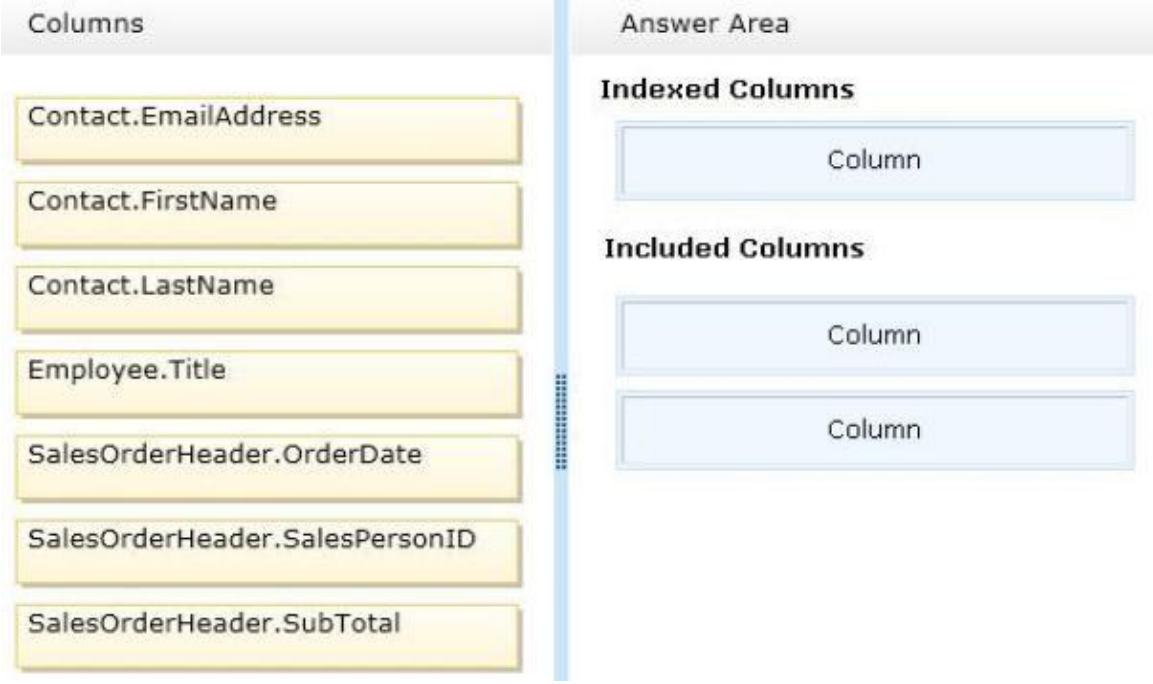

Answer:

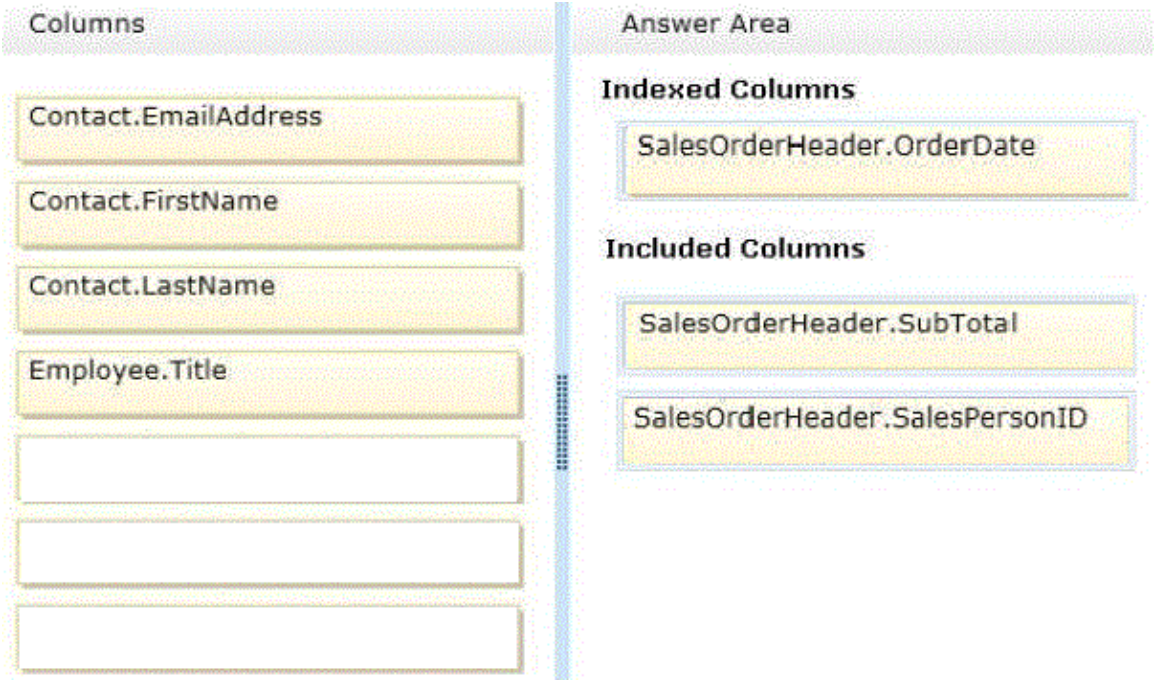

QUESTION 13
You use SQL Server 2012 to store data used by an e-commerce application. You develop a stored procedure named sp1. Sp1 is used to read the price of all the products sold on the e-commerce site. You need to ensure that sp1 can read data even while another transaction is modifying the price of a product. Sp1 must only read committed data. Which transaction isolation level should you use in sp1?

A. Serializable
B. Snapshot
C. Repeatable read
D. Read committed

Answer: B

QUESTION 14
Drag and Drop Question
You have a table named Table1. Table1 has 1 million rows. Table1 has a columnstore index for a column named Column1. You need to import data to Table1. The solution must minimize the amount of time it takes to import the data. What should you do? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

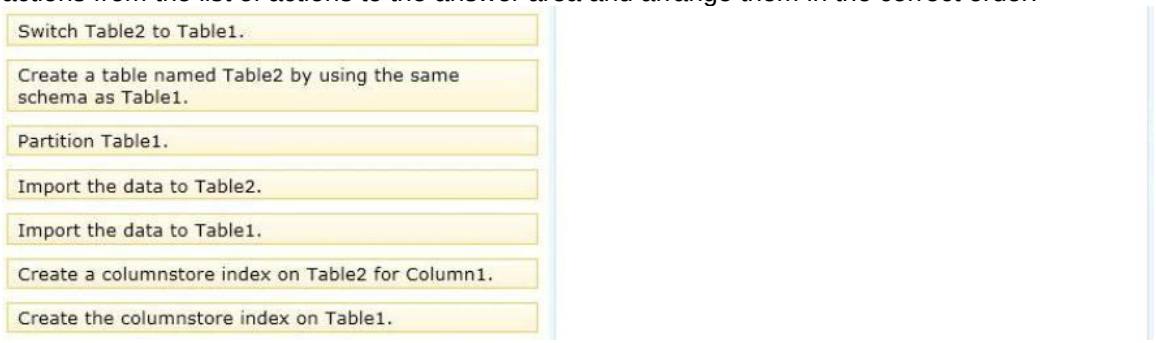

Answer:

Switch Table2 to Table1.	Create a table named Table2 by using the same schema as Table1.
Create a table named Table2 by using the same schema as Table1.	Partition Table1.
Partition Table1.	Import the data to Table2.
Import the data to Table2.	Create a columnstore index on Table2 for Column1.
Import the data to Table1.	Switch Table2 to Table1.
Create a columnstore index on Table2 for Column1.	
Create the columnstore index on Table1.	

QUESTION 15

You have a table named Rooms that contains three columns. You execute the following query:

```
SELECT [Id],
   [RoomName],
   [Position]
FROM [dbo].[Rooms]
WHERE [RoomName] = 'Room1'
```

You discover the execution plan shown in the exhibit.

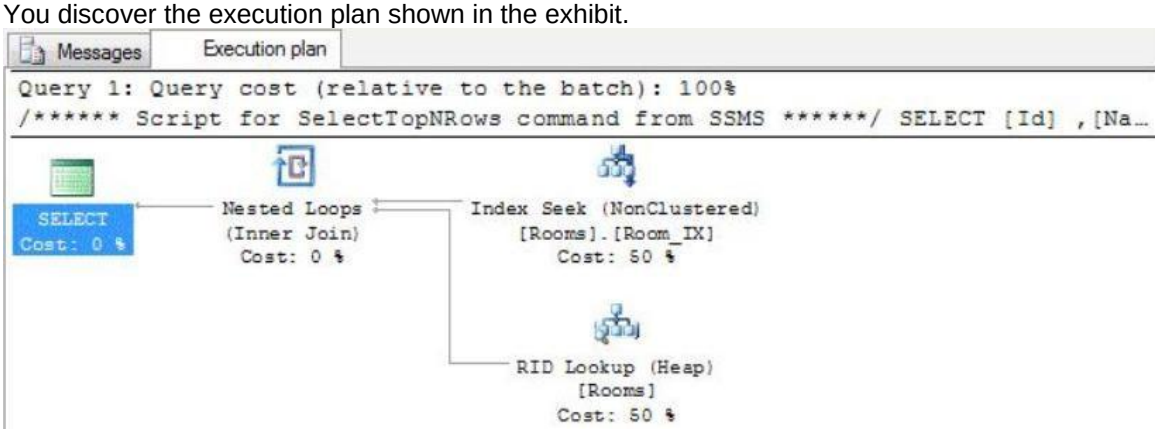

You need to recommend a solution to reduce the amount of time it takes to execute the query. What should you do? More than one answer choice may achieve the goal. Select the BEST answer.

A. use the WITH (INDEX(Room_IX),NOLOCK) query hint.
B. Create a clustered index for Id.
C. Include the RoomName column and the Position column in the Room_IX index.
D. Create a nonclustered index for RoomName, Id, and Position.

Answer: D

QUESTION 16

You have a server that has SQL Server 2012 installed. You need to identify which parallel execution plans are running in serial. Which tool should you use?

A. Data Profile Viewer
B. Database Engine Tuning Advisor
C. Performance Monitor

D. Extended Events

Answer: D
Explanation:
http://msdn.microsoft.com/en-us/library/bb677278.aspx
http://msdn.microsoft.com/en-us/library/bb630282.aspx
http://www.sql-server-performance.com/2006/query-execution-plan-analysis/
http://www.simple-talk.com/sql/learn-sql-server/understanding-and-using-parallelism-in-sql-server/
http://www.sqlservercentral.com/articles/SQL+Server+2012/At+last%2c+execution+plans+show+true+thread+reservations./92458/
http://sqlblog.com/blogs/paul_white/archive/2011/12/23/forcing-a-parallel-query-execution-plan.aspx
http://sqlblog.com/blogs/paul_white/archive/2012/05/02/parallel-row-goals-gone-rogue.aspx
http://msdn.microsoft.com/en-us/library/bb895310.aspx
http://msdn.microsoft.com/en-us/library/bb895313.aspx
http://msdn.microsoft.com/en-us/library/hh231122.aspx

QUESTION 17
You have a database named database1. Database developers report that there are many deadlocks. You need to implement a solution to monitor the deadlocks. The solution must meet the following requirements:
- Support real-time monitoring.
- Be enabled and disabled easily.
- Support querying of the monitored data.

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

A. an Extended Events session
B. a SQL Server Profiler template
C. log errors by using trace flag 1204
D. log errors by using trace flag 1222

Answer: A
Explanation:
http://www.sqlservercentral.com/blogs/james-sql-footprint/2012/08/12/monitor-deadlock-in-sql-2012/
http://blogs.technet.com/b/mspfe/archive/2012/06/28/how_2d00_to_2d00_monitor_2d00_deadlocks_2d00_in_2d00_sql_2d00_server.aspx

QUESTION 18
You have a SQL Server 2012 database named database1. Database1 contains a table named OrderDetails. For a given sales order, you need to retrieve the OrderID, Quantity, and LineTotal columns for all of the items in the OrderDetails table. The solution must ensure that the results can be joined to other tables. Which code segment should you execute?

A.
```
CREATE FUNCTION dbo.GetOrderDetails(@OrderID int)
RETURNS TABLE
AS
RETURN
(SELECT OrderID, Quantity, LineTotal
  FROM Sales.OrderDetails
  WHERE OrderID = @OrderID);
```

B.
```
CREATE PROC dbo.GetOrderDetails(@OrderID int)
AS
SELECT OrderID, Quantity, LineTotal
  FROM Sales.OrderDetails
  WHERE OrderID = @OrderID;
```

C.
```
CREATE FUNCTION dbo.GetOrderDetails(@OrderID int)
RETURNS @retOrderDetails TABLE
(
  OrderID int NOT NULL,
  Quantity int NOT NULL,
  LineTotal decimal NULL
)
AS
BEGIN
  INSERT @retOrderDetails
  SELECT OrderID, Quantity, LineTotal
    FROM Sales.OrderDetails
    ORDER BY @OrderID;
  RETURN;
END;
```

D.
```
CREATE VIEW dbo.GetOrderDetails
AS
SELECT OrderID, Quantity, LineTotal
  FROM Sales.OrderDetails;
```

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

Answer: A

QUESTION 19
Drag and Drop Question
You are planning two stored procedures named SProc1 and SProc2. You identify the following requirements:
- SProc1 must return a table.
- SProc2 must return a status code.
You need to identify which options must be implemented to meet each stored procedure

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

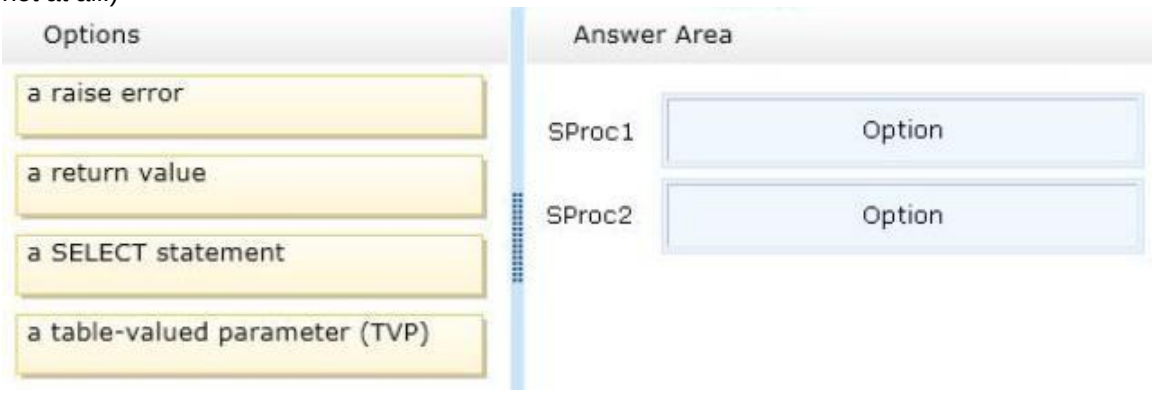

Answer:

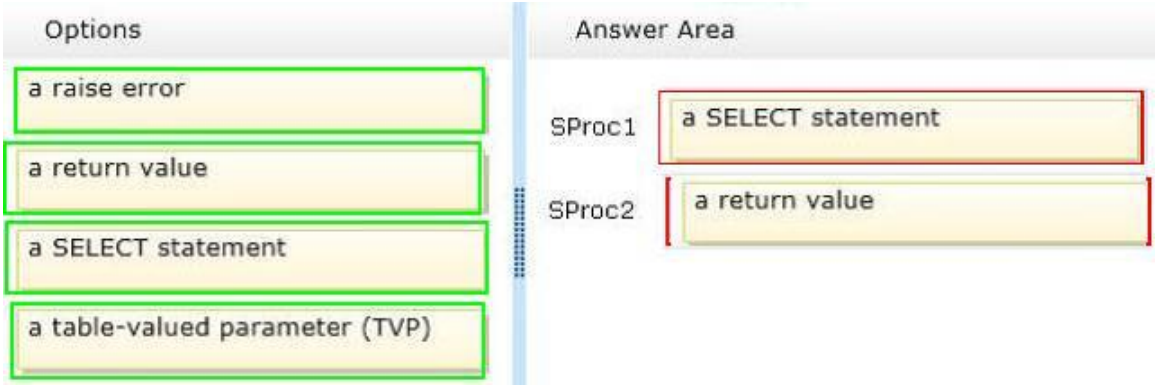

QUESTION 20

You have a database that contains a user-defined function named Schema1.Udf1 and two tables named Schema1.Table1 and Schemal.Table2. Schema1.Table1 has 1 million rows. The schema for Schema1.Table1 is configured as shown in the following table:

Column	Data type
CountryID	int
CustomerName	varchar(50)

Schema1.Udf1 was defined by using the following code:

```
CREATE FUNCTION Schema1.Udf1(@CountryID int)
RETURNS TABLE
AS
RETURN
SELECT Country
FROM Shema1.Table2
WHERE CountryID = @CountryID
```

You need to write a query that will contain the following columns:

- Country
- CountryID
- CustomerName

The solution must meet the following requirements:

- Rows must be returned only if the function returns data.

- The amount of time it takes the query to execute must be minimized.

Which query should you use?

A.
```
SELECT t.CountryID,
  u.Country,
  t.CustomerName
FROM Schema1.Table1 AS t
  LEFT JOIN Schema1.Udf1(t.CountryID) AS u;
```

B.
```
SELECT t.CountryID,
  u.Country,
  t.CustomerName
FROM Schema1.Table1 AS t
  OUTER APPLY Schema1.Udf1(t.CountryID) AS u;
```

C.
```
SELECT t.CountryID,
  u.Country,
  t.CustomerName
FROM Schema1.Table1 AS t
  CROSS APPLY Schema1.Udf1(t.CountryID) AS u;
```

D.
```
SELECT t.CountryID,
  u.Country,
  t.CustomerName
FROM Schema1.Table1 AS t
  INNER JOIN Schema1.Udf1(t.CountryID) AS u;
```

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

Answer: C

QUESTION 21

You have a database hosted on SQL Azure. You are developing a script to create a view that will be used to update the data in a table. The following is the relevant portion of the script. (Line numbers are included for reference only.) You need to ensure that the view can update the data in the table, except for the data in Column1. Which code segment should you add at line 06?

```
01 CREATE VIEW View1
02 AS
03 SELECT
04 ...
05 WHERE Column1 = 'City1'
06
```

A. WITH ENCRYPTION
B. WITH VIEW_METADATA
C. WITH CHECK OPTION
D. WITH SCHEMABINDING

Answer: C
Explanation:

http://msdn.microsoft.com/en-us/library/ms187956.aspx

QUESTION 22
You have an application that uses a view to access data from multiple tables. You need to ensure that you can insert rows into the underlying tables by using the view. What should you do?

A. Define the view by using the SCHEMABINDING option.
B. Define the view by using the CHECK option.
C. Create an INSTEAD OF trigger on the view.
D. Materialize the view.

Answer: C
Explanation:
http://msdn.microsoft.com/en-us/library/ms180800.aspx
http://msdn.microsoft.com/en-us/library/ms187956.aspx

QUESTION 23
You have a SQL Azure database. You execute the following code:

```
CREATE SCHEMA Sales;
GO

CREATE TABLE Sales.Customers
(
  CustomerID int IDENTITY(1,1) PRIMARY KEY,
  FaxNumber char(10) SPARSE NULL,
  CustomerName varchar(100) NOT NULL,
  EmailAddress varchar(100) NOT NULL
);
GO

CREATE PROCEDURE Sales.CustomersByFaxNumber
  @FaxNumber char(10)
AS
SELECT CustomerID,
  CustomerName
FROM Sales.Customers
WHERE FaxNumber = @FaxNumber
```

The Sales.Customers table will contain 100,000 rows. You expect the FaxNumber column to contain a null value for 70 percent of the rows. You need to create an index to support Sales.CustomersByFaxNumber. The solution must minimize the disk storage requirements. Which code segment should you execute?

A.
```
CREATE INDEX IX_Customers ON Customers (FaxNumber)
WHERE FaxNumber IS NOT NULL
```
B.
```
CREATE INDEX IX_Customers ON Customers (FaxNumber)
WITH FILLFACTOR=0
```
C.
```
CREATE INDEX IX_Customers ON Customers (CustomerName)
INCLUDE (FaxNumber)
```

D. `CREATE INDEX IX_Customers ON Customers (FaxNumber)`
E. `CREATE INDEX IX_Customers ON Customers (FaxNumber) WHERE FaxNumber IS NULL`

Answer: A

QUESTION 24
You run the following code:

```
CREATE TABLE dbo.Orders
(
  Id int CONSTRAINT PK_Order_Id PRIMARY KEY,
  Amount decimal,
  Details xml
);
```

You need to ensure that the root node of the XML data stored in the Details column is <Order_Details>. What should you implement? More than one answer choice may achieve the goal. Select the BEST answer.

A. an XML index
B. an XML schema collection
C. a user-defined data type
D. a Data Definition Language (DDL) trigger
E. a data manipulation language (DML) trigger

Answer: B
Explanation:
http://msdn.microsoft.com/en-us/library/ms187856.aspx

QUESTION 25
Drag and Drop Question
You have a table named Customers that has a clustered index defined on the ID column. You write a script to create a stored procedure. You need to complete the script for the stored procedure. The solution must minimize the number of locks and deadlocks. What should you do? To answer, drag the appropriate option to the correct location in the answer area. (Answer choices may be used once, more than once, or not at all.)

READ COMMITED

SERIALIZABLE

WITH(UPDLOCK)

WITH(XLOCK)

```
CREATE PROCEDURE Proc1 (@ParamID int)
AS
SET TRANSACTION ISOLATION LEVEL
BEGIN TRANSACTION
DECLARE @var as NCHAR(10)
Select @var = Name
FROM dbo.Customers
WHERE ID = @ParamID
...
UPDATE dbo.Customers
SET Name = @var
WHERE ID = @ParamID
COMMIT TRANSACTION;
GO
```

Answer:

SERIALIZABLE

WITH(XLOCK)

```
CREATE PROCEDURE Proc1 (@ParamID int)
AS
SET TRANSACTION ISOLATION LEVEL READ COMMITED
BEGIN TRANSACTION
DECLARE @var as NCHAR(10)
Select @var = Name
FROM dbo.Customers WITH(UPDLOCK)
WHERE ID = @ParamID
...
UPDATE dbo.Customers
SET Name = @var
WHERE ID = @ParamID
COMMIT TRANSACTION;
GO
```

QUESTION 26
You are creating a table named Orders. You need to ensure that every time a new row is added to the Orders table, a user-defined function is called to validate the row before the row is added to the table. What should you use? More than one answer choice may achieve the goal. Select the BEST answer.

A. a FOREIGN KEY constraint
B. a data manipulation language (DML) trigger
C. a DEFAULT constraint
D. a CHECK constraint
E. a Data Definition Language (DDL) trigger

Answer: D
Explanation:
http://www.techrepublic.com/blog/programming-and-development/comparing-sql-server-constraints-and-dmltriggers/402
http://msdn.microsoft.com/en-us/library/ms178110.aspx

QUESTION 27
You execute the following code:

```
CREATE TABLE UserInfo
(
  ID int NOT NULL IDENTITY (1, 1)
  CONSTRAINT PK_UserInfo PRIMARY KEY CLUSTERED,
  UserName varchar(100) NOT NULL,
  Manager varchar(100) NULL,
  HireDate date NOT NULL,
  PerformanceReviewScore int NULL
);
```

You have a stored procedure that includes the following SELECT statement:

```
SELECT UserName, PerformanceReviewScore
FROM UserInfo
WHERE Manager = 'Ben Smith';
```

You need to create a covering index on UserInfo. Which code segment should you execute?

A.
```
CREATE NONCLUSTERED INDEX [IX_Covering_Index] ON UserInfo
(
  [UserName] ASC,
  [PerformanceReviewScore] ASC,
);
```

B.
```
CREATE NONCLUSTERED INDEX [IX_Covering_Index] ON UserInfo
(
  [Manager] ASC
);
```

C.
```
CREATE NONCLUSTERED INDEX [IX_Covering_Index] ON UserInfo
(
  [UserName] ASC,
  [Manager] ASC
);
```

D.
```
CREATE NONCLUSTERED INDEX [IX_Covering_Index] ON UserInfo
(
  [Manager] ASC,
  [PerformanceReviewScore] ASC,
  [UserName] ASC
);
```

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

Answer: D

QUESTION 28
You plan to execute the following code:

```
CREATE TABLE dbo.Table1
(
  datavalue varchar(20)
);
GO

BEGIN TRANSACTION;

INSERT INTO Table1 VALUES('entry1');

  BEGIN TRANSACTION;
    INSERT INTO Table1 VALUES('entry2');
  COMMIT TRANSACTION;

INSERT INTO Table1 VALUES('entry3');

ROLLBACK TRANSACTION;
Go
```

You need to identify how many rows will be in dbo.Table1 after you execute the code. How many rows should you identify?

A. 3
B. 2
C. 1
D. 0

Answer: D

QUESTION 29
You have an index for a table in a SQL Azure database. The database is used for Online Transaction Processing (OLTP). You discover that the index consumes more physical disk space than necessary. You need to minimize the amount of disk space that the index consumes. What should you set from the index options?

A. `STATISTICS_NORECOMPUTE = ON`
B. `STATISTICS_NORECOMPUTE = OFF`
C. `FILLFACTOR = 0`
D. `FILLFACTOR = 80`

Answer: C
Explanation:
http://msdn.microsoft.com/en-us/library/ms177459.aspx
http://msdn.microsoft.com/en-us/library/ms188783.aspx

QUESTION 30
Drag and Drop Question
You execute the following code:

```
CREATE TABLE Customers
(
  id int primary key,
  name nchar(10)
)
GO
```

You discover that the Customers table was created in the dbo schema. You need to create a code segment to move the table to another schema named Schema2. What should you create? To answer, drag the appropriate code segments to the correct location in the answer area. (Answer choices may be used once, more than once, or not at all.)

Answer:

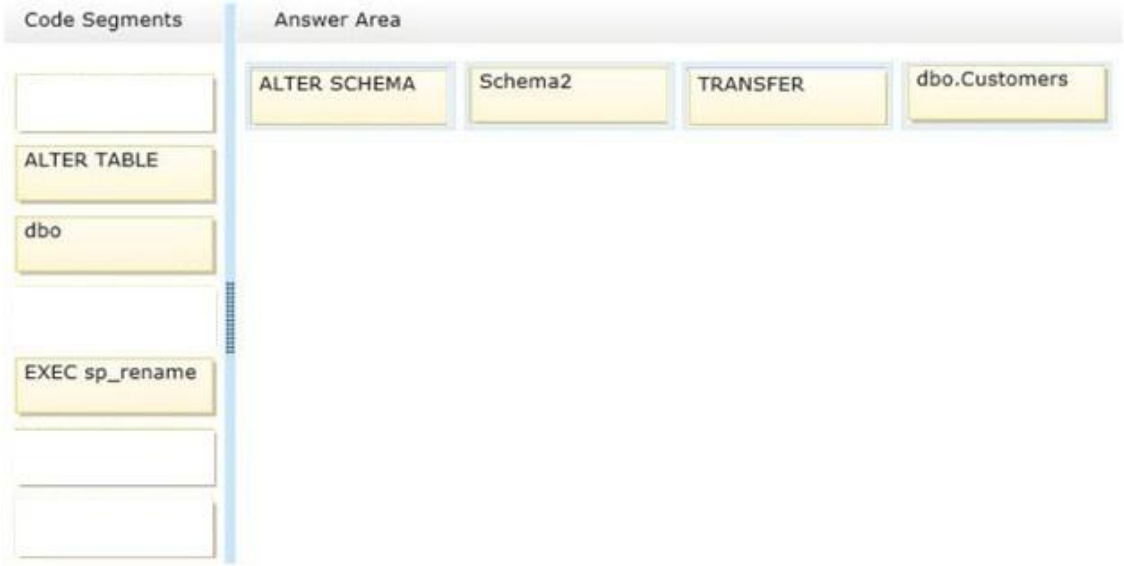

Explanation:

Note: ALTER SCHEMA newschema TRANSFER oldschema.TABLE This will transfer the table defined under "oldschema" and transfer it to "newschema".

QUESTION 31

You plan to design an application that temporarily stores data in a SQL Azure database. You need to identify which types of database objects can be used to store data for the application. The solution must ensure that the application can make changes to the schema of a temporary object during a session. Which type of objects should you identify?

A. common table expressions (CTEs)
B. table variables
C. temporary tables
D. temporary stored procedures

Answer: C
Explanation:
http://msdn.microsoft.com/en-us/library/ms175972.aspx
http://msdn.microsoft.com/en-us/library/ms189084.aspx
http://msdn.microsoft.com/en-us/library/ms175010.aspx
http://msdn.microsoft.com/en-us/library/bb510489.aspx
http://msdn.microsoft.com/en-us/library/ms187926.aspx
http://zacksfiasco.com/post/2010/01/21/SQL-Server-Temporary-Stored-Procedures.aspx

QUESTION 32
You have a text file that contains an XML Schema Definition (XSD). You have a table named Schemal.Table1. You have a stored procedure named Schemal.Proc1 that accepts an XML parameter named Param1. You need to store validated XML data in Schemal.Table1. The solution must ensure that only valid XML data is accepted by Param1. What should you do? (Each correct answer presents part of the solution. Choose all that apply.)

A. Define an XML column in Table1 by using an XML schema collection.
B. Create an XML schema collection in the database from the text file.
C. Declare Param1 var1 as type XML and associate the variable to the XML schema collection.
D. use the modify method to insert the XML schema into each row of the XML column in Table1.

Answer: ABD
Explanation:
http://msdn.microsoft.com/en-us/library/bb510420.aspx
http://msdn.microsoft.com/en-us/library/ms187856.aspx
http://msdn.microsoft.com/en-us/library/ms176009.aspx
http://msdn.microsoft.com/en-us/library/hh403385.aspx
http://msdn.microsoft.com/en-us/library/ms184277.aspx

QUESTION 33
Drag and Drop Question
You plan to deploy two stored procedures named SP1 and SP2 that read data from the database. Your company identifies the following requirements for each stored procedure:
- SP1 must allow dirty reads.
- SP2 must place range locks on the data to ensure read consistency.

You need to identify which isolation level you must set for each stored procedure. The solution must minimize the number of locks. Which isolation level should you identify? To answer, drag the appropriate isolation level to the correct stored procedure in the answer area. (Answer choices may be used once, more than once, or not at all.)

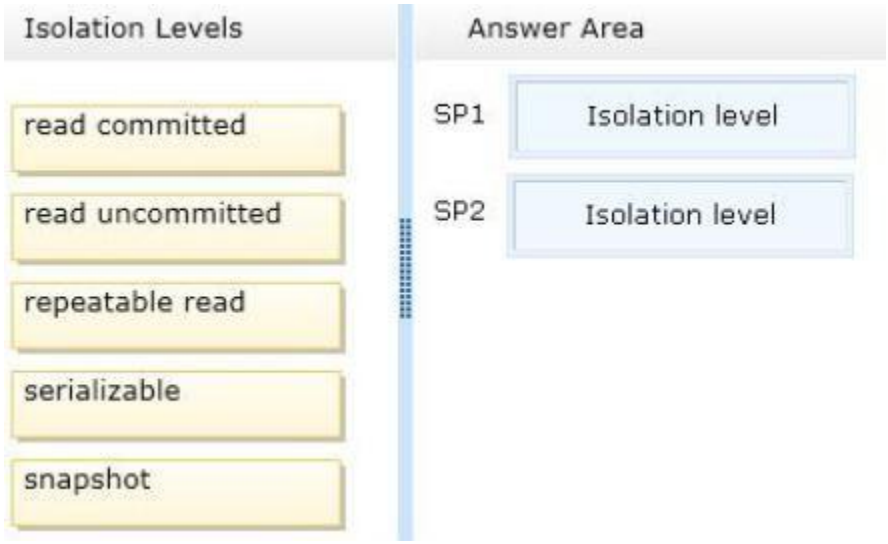

Answer:

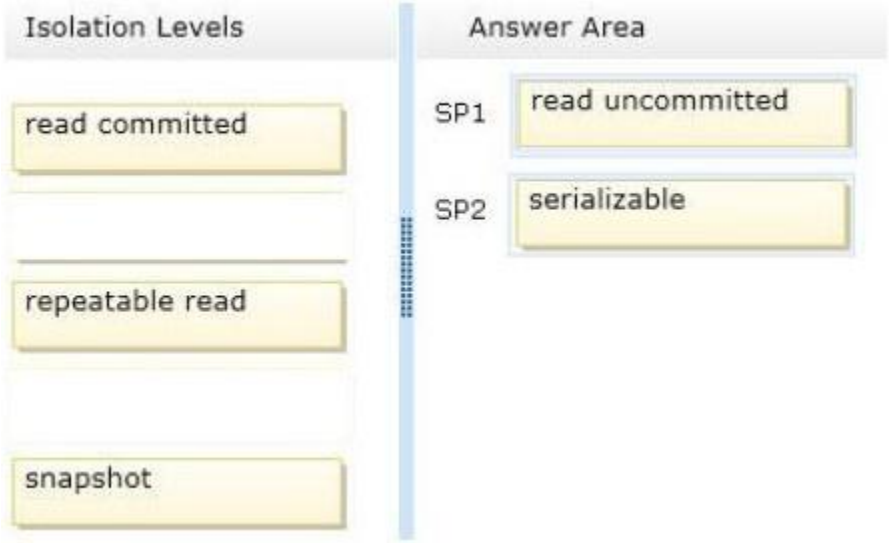

QUESTION 34

You use SQL Server 2014. The physical server is a dedicated database server that has 120GB of RAM available. There is approximately 50GB of storage space available on a slow local disk. You create a new stored procedure. You decide you need to temporarily hold approximately 300,000 rows from two tables, from which you will compute two complex business scores. The stored procedure will use temporary storage defined as follows:

```
AccountNumber char(10) NOT NULL
YearToDateSalesTotal decimal(15,2) NULL
SalesScore int NULL
FutureSalesExpectationScore int NULL
```

The code will make several passes through the data, applying complex calculations before writing the data to a permanent disk-based table in the same database from which it reads the data. For this stored procedure, you need to deal with temporary data in the most efficient way to minimize physical disk pressure. What should you do? More than one answer choice may achieve the goal. Select the BEST answer.

A.
```
CREATE TYPE dbo.AccountScoringModel as TABLE
(
     AccountNumber char(10) COLLATE Latin1_General_100_BIN2 NOT NULL ,
     YearToDateSalesTotal decimal(15,2) NULL,
     SalesScore int NULL,
     FutureSalesExpectationScore int NULL,
     INDEX AccountNumber HASH (AccountNumber) WITH (BUCKET_COUNT = 25000)
) WITH ( MEMORY_OPTIMIZED = ON )
GO
DECLARE @AccountScoring as dbo.AccountScoringModel
```

B.
```
DECLARE @AccountScoring as TABLE
(
     AccountNumber char(10) NOT NULL,
     YearToDateSalesTotal decimal(15,2) NULL,
     SalesScore int NULL,
     FutureSalesExpectationScore int NULL
)
```

C.
```
CREATE TABLE #AccountScoring
(
     AccountNumber char(10) NOT NULL,
     YearToDateSalesTotal decimal(15,2) NULL,
     SalesScore int NULL,
     FutureSalesExpectationScore int NULL
)
```

D.
```
CREATE TYPE dbo.AccountScoringModel as TABLE
(
     AccountNumber char(10) COLLATE Latin1_General_100_BIN2 NOT NULL ,
     YearToDateSalesTotal decimal(15,2) NULL,
     SalesScore int NULL,
     FutureSalesExpectationScore int NULL,
     INDEX AccountNumber HASH (AccountNumber) WITH (BUCKET_COUNT = 120)
) WITH ( MEMORY_OPTIMIZED = ON )
GO
DECLARE @AccountScoring as dbo.AccountScoringModel
```

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

Answer: A
Explanation:
- You must specify a value for the BUCKET_COUNT parameter when you create the memory-optimized table. In most cases the bucket count should be between 1 and 2 times the number of distinct values in the index key. If the index key contains a lot of duplicate values, on average there are more than 10 rows for each index key value, use a nonclustered index instead.
- You may not always be able to predict how many values a particular index key may have or will have. Performance should be acceptable if the BUCKET_COUNT value is within 5 times of the actual number of key values.

QUESTION 35
You have a SQL Server 2012 database named Database1. You execute the following code:

```
CREATE TABLE Sales
(
  ID int IDENTITY(1,1) NOT NULL PRIMARY KEY,
  OrderDate char(10) NOT NULL,
  Amount decimal
);
GO

CREATE INDEX IX_Sales_OrderDate
  ON Sales(OrderDate)
  INCLUDE (ID, Amount);
GO

CREATE PROC usp_Proc1(
  @date1 datetime,
  @date2 datetime
)
AS
SELECT ID, OrderDate, Amount
  FROM Sales
  WHERE CAST(OrderDate AS datetime)
    BETWEEN @date1 AND @date2
  ORDER BY ID;
GO
```

You insert 3 million rows into Sales. You need to reduce the amount of time it takes to execute Proc1. What should you do?

A. ProductType varchar(11) '@ProductType',
B. ProductType varchar(11) 'ProductType/ID',
C. ProductType varchar(11) 'ProductType/@ID',
D. ProductType varchar(11) 'ProductType1'.

Answer: D

QUESTION 36

You have a SQL Azure database. You execute the following script:

```
CREATE TABLE dbo.Table1
(
  Column1 int PRIMARY KEY,
  Column2 varchar(50) SPARSE NULL
)
```

You add 1 million rows to Table1. Approximately 85 percent of all the rows have a null value for Column2. You plan to deploy an application that will search Column2. You need to create an index on Table1 to support the planned deployment. The solution must minimize the storage requirements. Which code segment should you execute?

A.
```
CREATE INDEX IX_Table1 ON Table1 (Column2)
WITH FILLFACTOR-0
```
B.
```
CREATE INDEX IX_Table1 OK Table1 (Column1)
INCLUDE (Column2)
```
C.
```
CREATE INDEX IX_Table1 ON Table1 (Column2)
WHERE Column2 IS NULL
```

D. CREATE INDEX IX_Table1 ON Table1 (Column2)
WHERE Column2 IS NOT NOLL

Answer: D
Explanation:
http://msdn.microsoft.com/en-us/library/ms188783.aspx
http://msdn.microsoft.com/en-us/library/cc280372.aspx

QUESTION 37
You are creating a table named Orders. You need to ensure that every time a new row is added to the Orders table, a table that is used for auditing is updated. What should you use? More than one answer choice may achieve the goal. Select the BEST answer.

A. a CHECK constraint
B. a FOREIGN KEY constraint
C. a DEFAULT constraint
D. a data manipulation language (DML) trigger
E. a Data Definition Language (DDL) trigger

Answer: D
Explanation:
http://www.techrepublic.com/blog/programming-and-development/comparing-sql-server-constraints-and-dmltriggers/402
http://msdn.microsoft.com/en-us/library/ms178110.aspx

QUESTION 38
You have a SQL Azure database. You need to identify which keyword must be used to create a view that will be indexed. Which keyword should you identify?

A. VIEW_METADATA
B. SCHEMABINDING
C. DEFAULT
D. DISTINCT

Answer: B
Explanation:
http://msdn.microsoft.com/en-us/library/ms187956.aspx
http://msdn.microsoft.com/en-us/library/ms191432.aspx

QUESTION 39
Drag and Drop Question
You are designing two stored procedures named Procedure1 and Procedure2. You identify the following requirements:
- Procedure1 must take a parameter that ensures that multiple rows of data can pass into the stored procedure.
- Procedure2 must use business logic that resides in a Microsoft .NET Framework assembly.
You need to identify the appropriate technology for each stored procedure. Which technologies should you identify? To answer, drag the appropriate technology to the correct stored procedure in the answer area. (Answer choices may be used once, more than once, or not at all.)

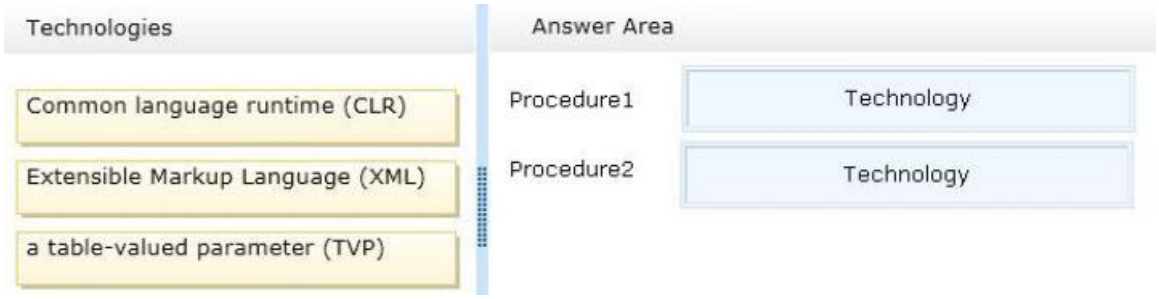

Answer:

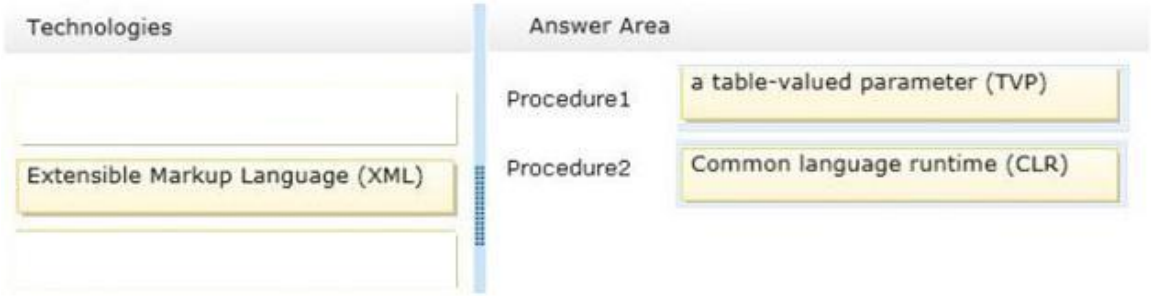

QUESTION 40

You create a view by using the following code:

```
CREATE VIEW dbo.View1
WITH VIEW_METADATA
AS
SELECT t1.col1, t1.col2, t2.*
FROM dbo.Table1 AS t1 JOIN dbo.Table2 AS t2 ON t1.col1=t2.col2;
```

Several months after you create the view, users report that the view has started to return unexpected results. You discover that the design of Table2 was modified since you created the view. You need to ensure that the view returns the correct results. Which code segment should you run?

A.
```
EXEC sp_refreshview @viewname = 'dbo.View1';
```

B.
```
ALTER dbo.View1 WITH SCHEMABINDING, VIEW_METADATA
AS
SELECT t1.col1, t1.col2, t2.*
FROM dbo.Table1 AS t1 JOIN dbo.Table2 AS t2
ON t1.col1=t2.col2;
```

C.
```
EXEC sp_refreshsqlmodule @name = 'dbo.Table2';
```

D.
```
DROP dbo.View1;
GO
CREATE dbo.View1 WITH SCHEMABINDING, VIEW_METADATA
AS
SELECT t1.col1, t1.col2, t2.*
FROM dbo.Table1 AS t1 JOIN dbo.Table2 AS t2
ON t1.col1=t2.col2;
```

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

Answer: A

QUESTION 41
Drag and Drop Question
You are planning two stored procedures named SProc1 and SProc2. You identify the following requirements:
- SProc1 must return a table.
- SProc2 must return a scalar value.

You need to identify which option must be implemented for each stored procedure to return the desired data. Which options should you identify? To answer, drag the appropriate option to the correct requirement in the answer area. (Answer choices may be used once, more than once, or not at all.)

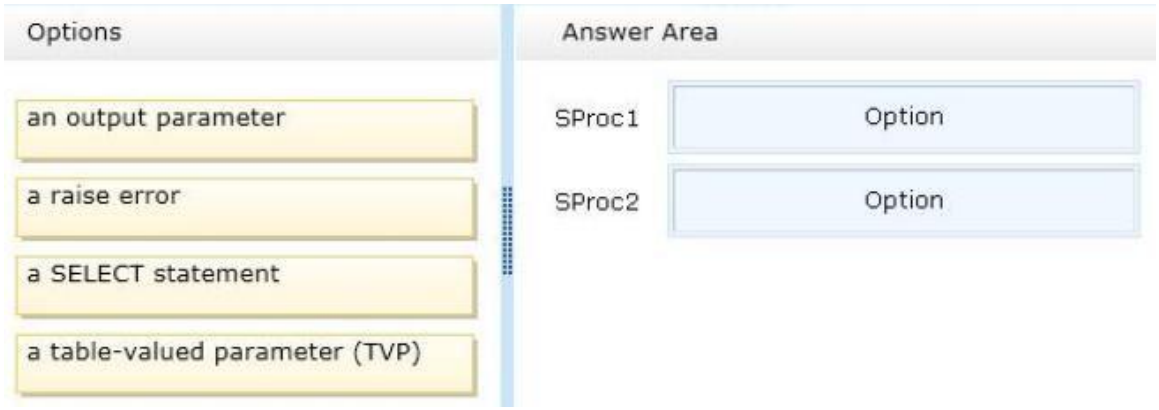

Answer:

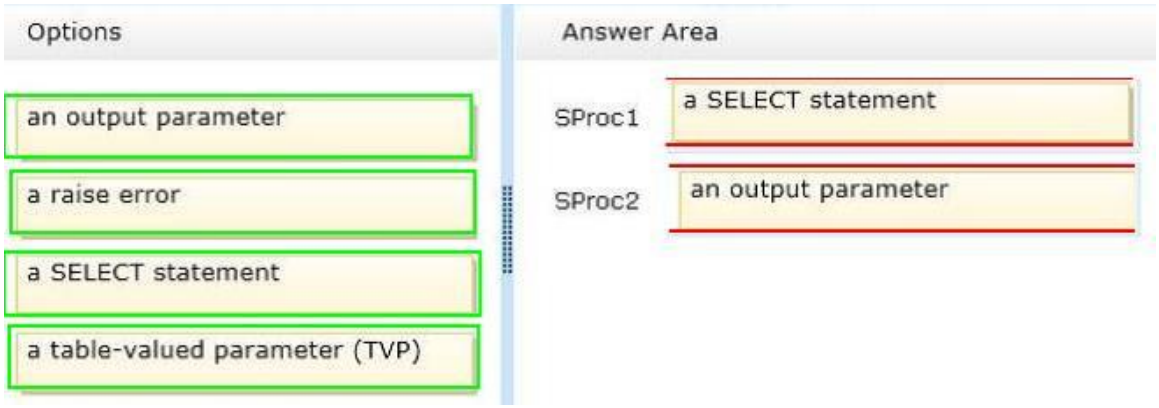

QUESTION 42
You have a SQL Server 2012 instance that hosts a single-user database. The database does not contain user-created stored procedures or user-created functions. You need to minimize the amount of memory used for query plan caching. Which advanced server option should you modify?

A. Enable Contained Databases
B. Allow Triggers to Fire Others
C. Optimize for Ad hoc Workloads
D. Scan for Startup Procs

Answer: C

QUESTION 43
Drag and Drop Question
You have a SQL Server 2012 database named Database1. Database1 has a data file named Database1_data.mdf and a transaction log named Database1jog.ldf. Database1_data.mdf is 1.5GB. Database1jog.ldf is 1.5 terabytes. A full backup of Database1 is performed every day. You need to reduce the size of the log file. The solution must ensure that you can perform transaction log backups in the future. Which code segment should you execute? To answer, move the appropriate code segments from the list of code segments to the answer area and arrange them in the correct order.

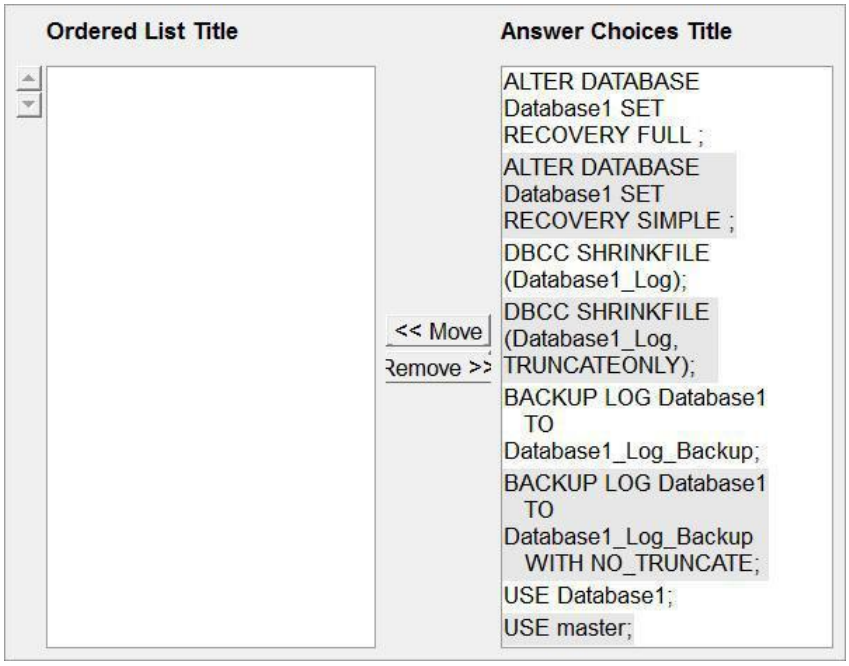

Answer:

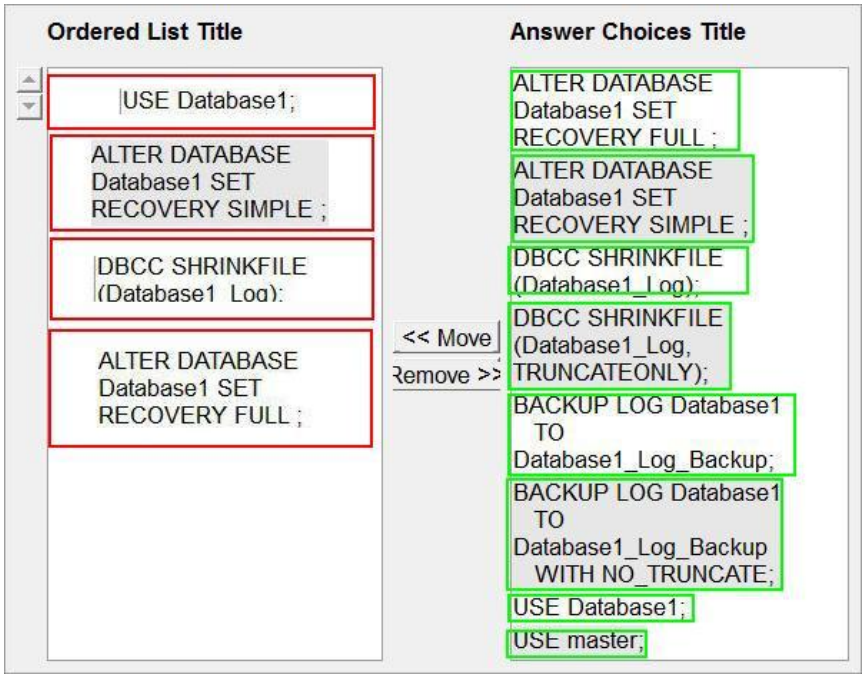

QUESTION 44

Drag and Drop Question

You have a database named database1. Each table in database1 has one index per column. Users often report that creating items takes a long time. You need to perform the following maintenance tasks:

- Identify unused indexes.
- Identify indexes that need to be defragmented.

What should you use? To answer, drag the appropriate function to the correct management task in the answer area. (Answer choices may be used once, more than once, or not at all.)

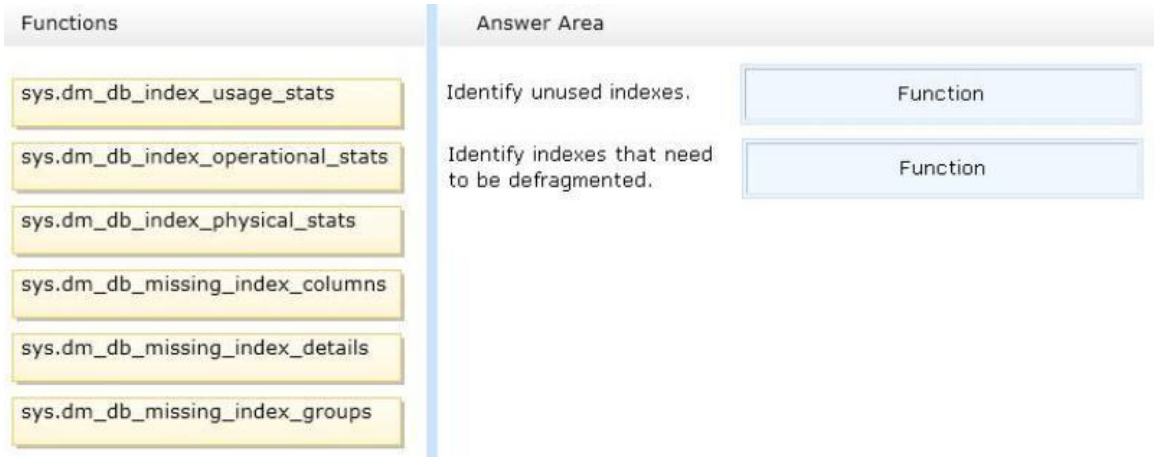

Answer:

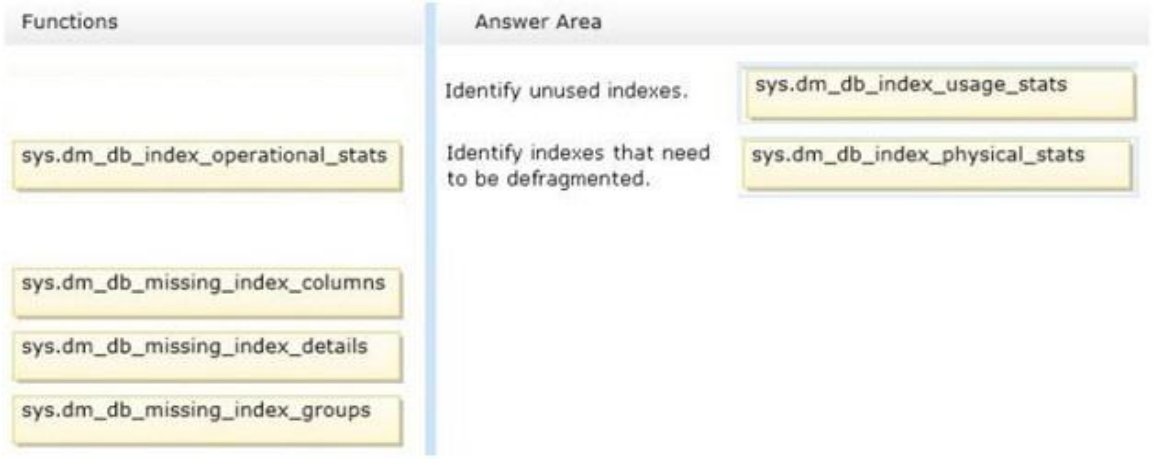

QUESTION 45

You execute the following code:

```
CREATE TABLE HumanResources.Employees
(
  EmployeeID int IDENTITY(1,1) PRIMARY KEY,
  ContactID int NOT NULL
    FOREIGN KEY REFERENCESS Person.Contact(ContactID),
  JobTitle varchar(100)
);
GO

CREATE INDEX IX_Employees
ON HumanResources.Employee(JobTitle);
GO
```

After populating the Employees table with 10,000 rows, you execute the following query:

```
SELECT EmployeeID, JobTitle
FROM HumanResources.Employee
WHERE SUBSTRING(JobTitle,1,1) = 'C'
```

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

A. Change SUBSTRING (JobTitle, l, l) = 'c' to LEFT(JobTitle,1) = 'c'
B. Change SUBSTRING(JobTitle, l, l) = 'c' to JobTitle LIKE 'c%\
C. Partition the table and use the JobTitle column for the partition scheme
D. Replace IX_Employees with a clustered index

Answer: B

Case Study 1: Scenario 1 (Question 46 ~ Question 57)

Application Information

You have two servers named SQL1 and SQL2 that have SQL Server 2012 installed. You have an application that is used to schedule and manage conferences. Users report that the application has many errors and is very slow. You are updating the application to resolve the issues. You plan to create a new database on SQL1 to support the application. A junior database administrator has created all the scripts that will be used to create the database. The script that you plan to use to create the tables for the new database is shown in Tables.sql. The script that you plan to use to create the stored procedures for the new database is shown in StoredProcedures.sql. The script that you plan to use to create the indexes for the new database is shown in Indexes.sql. (Line numbers are included for reference only.) A database named DB2 resides on SQL2. DB2 has a table named SpeakerAudit that will audit changes to a table named Speakers. A stored procedure named usp_UpdateSpeakersName will be executed only by other stored procedures. The stored procedures executing usp_UpdateSpeakersName will always handle transactions. A stored procedure named usp_SelectSpeakersByName will be used to retrieve the names of speakers. Usp_SelectSpeakersByName can read uncommitted data. A stored procedure named usp_GetFutureSessions will be used to retrieve sessions that will occur in the future.

Procedures.sql

```
01 CREATE PROCEDURE usp_UpdateSpeakerName
02   @SpeakerID int,
03   @LastName nvarchar(100)
04 AS
05
06 BEGIN TRY
07
08 UPDATE Speakers
09 SET LastName = @LastName
10 WHERE SpeakerID = @SpeakerID;
11
12 INSERT INTO SQL2.DB2.dbo.SpeakerAudit(SpeakerID, LastName)
13 VALUES (@SpeakerID, @LastName);
14
15 END TRY
16 BEGIN CATCH
17
18 END CATCH;
19
20 GO
21
22 CREATE PROCEDURE usp_SelectSpeakersByName
23   @LastName nvarchar(100)
24 AS
25 SELECT SpeakerID,
26   FirstName,
27   LastName
28 FROM Speakers
29 WHERE LastName LIKE @LastName + '%'
30
31 GO
32
33 CREATE PROCEDURE usp_InsertSessions
34   @SessionData SessionDataTable READONLY
35 AS
36 INSERT INTO Sessions
37   (SpeakerID, Title, Absract, DeliveryTime, TitleAndSpeaker)
38 SELECT SpeakerID, Title, Absract, DeliveryTime, TitleAndSpeaker
39 FROM @SessionData;
40 GO
41
42 CREATE PROCEDURE usp_UpdateSessionRoom
43   @RoomID int,
44   @SpeakerID int
45 AS
```

```
51 FROM Sessions
52 WHERE SpeakerID = @SpeakerID;
53
54 UPDATE Sessions
55 SET RoomID = @RoomID
56 WHERE SpeakerID = @SpeakerID;
57
58 COMMIT TRANSACTION;
59
60 CREATE PROCEDURE usp_AttendeesReport
61   @LastName varchar(100)
62 AS
63 SELECT FirstName + ' ' + LastName AS FullName
64 FROM Attendees
65 WHERE LastName = @LastName;
66 GO
67
68 CREATE PROCEDURE usp_GetFutureSessions
69 AS
70 SELECT SpeakerID,
71   RoomID,
72   DeliveryTime
73 FROM Sessions
74
75 GO
76
77 CREATE PROCEDURE usp_TestSpeakers
78 AS
79 EXECUTE usp_SelectSpeakersByName 'a';
80 EXECUTE usp_SelectSpeakersByName 'an';
81 EXECUTE usp_SelectSpeakersByName 'and';
82 EXECUTE usp_SelectSpeakersByName 'ander';
83 EXECUTE usp_SelectSpeakersByName 'anderson';
84 EXECUTE usp_SelectSpeakersByName 'b';
85 EXECUTE usp_SelectSpeakersByName 'bi';
86 ...
87 EXECUTE usp_SelectSpeakersByName 'zzz';
88 GO
```

Indexes.sql

```
CREATE INDEX IX_Sessions ON Sessions
(SessionID, DeliveryTime)
INCLUDE (RoomID)

GO

CREATE INDEX IX_Speakers ON Speakers
(LastName);
GO

CREATE INDEX IX_Attendees_Name ON Attendees
(FirstName, LastName);

GO

CREATE INDEX IX_Attendees_Confirmed ON Attendees
(Confirmed);
GO
```

Tables.sql

```
CREATE DATABASE Conference;
GO

ALTER DATABASE Conference
SET READ_COMMITTED_SNAPSHOT ON;
GO

CREATE TABLE Attendees
(
  AttendeeID int IDENTITY (1,1) NOT NULL,
   FirstName nvarchar(100) NOT NULL,
   LastName nvarchar(100) NOT NULL,
   EmailAddress nvarchar(100) NOT NULL,

   CONSTRAINT PK_Attendees_AttendeeID PRIMARY KEY (AttendeeID)
);
GO

CREATE TABLE Speakers
(
  SpeakerID int IDENTITY(1,1) NOT NULL,
   FirstName nvarchar(100) NOT NULL,
   LastName nvarchar(100) NOT NULL,
   Photo varbinary(max),
   CONSTRAINT PK_Speakers_SpeakerID PRIMARY KEY (SpeakerID)
);
GO

CREATE TABLE Sessions
(
  SessionID uniqueidentifier NOT NULL
    CONSTRAINT DF_SessionID DEFAULT (NEWID()),
   SpeakerID int NOT NULL,
   Title nvarchar(100) NOT NULL,
   Abstract nvarchar(max) NOT NULL,
   DeliveryTime datetime NOT NULL,
   TitleAndSpeaker nvarchar(200)

);
GO

CREATE TABLE Rooms
(
   RoomID uniqueidentifier NOT NULL CONSTRAINT DF_RoomID DEFAULT (NEWID()),
    Location varchar(100) NOT NULL
);
```

QUESTION 46
You are evaluating the index design. You need to recommend a change to Indexes.sql that will minimize the amount of time it takes for usp_AttendeesReport to execute. The solution must minimize the amount of database fragmentation. Which line of code should you use to replace line 12 of Indexes.sql?

A. `(LastName);`
B. `(FirstName) INCLUDE (LastName);`
C. `(LastName, FirstName);`
D. `(LastName) INCLUDE (FirstName);`

Answer: C

QUESTION 47
You need to recommend a solution to ensure that SQL1 supports the auditing requirements of usp_UpdateSpeakerName. What should you include in the recommendation?

A. The Distributed Transaction Coordinator (DTC)
B. Transactional replication
C. Change data capture
D. Change tracking

Answer: A

QUESTION 48
You need to modify usp_SelectSpeakersByName to support server-side paging. The solution must minimize the amount of development effort required. What should you add to usp_SelectSpeakersByName?

A. an OFFSET-FETCH clause
B. a table variable
C. the ROWNUMBER keyword
D. a recursive common table expression

Answer: A
Explanation:
http://www.mssqltips.com/sqlservertip/2696/comparing-performance-for-different-sql-server-paging-methods/
http://msdn.microsoft.com/en-us/library/ms188385.aspx
http://msdn.microsoft.com/en-us/library/ms180152.aspx
http://msdn.microsoft.com/en-us/library/ms186243.aspx
http://msdn.microsoft.com/en-us/library/ms186734.aspx
http://www.sqlserver-training.com/how-to-use-offset-fetch-option-in-sql-server-order-by-clause/
http://www.sqlservercentral.com/blogs/juggling_with_sql/2011/11/30/using-offset-and-fetch/

QUESTION 49
You are evaluating the table design. You need to recommend a change to Tables.sql that reduces the amount of time it takes for usp_AttendeesReport to execute. What should you add at line 14 of Tables.sql?

A.
```
FullName nvarchar(100) NOT NULL DEFAULT (dbo.CreateFullName(FirstName, LastName)),
```

B.
```
FullName AS (FirstName + ' ' + LastName),
```

C.
```
FullName nvarchar(100) NOT NULL CONSTRAINT DF_FullName DEFAULT (dbo.CreateFullName
(FirstName, LastName)),
```

D.
```
FullName AS (FirstName + ' ' + LastName) PERSISTED,
```

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

Answer: D
Explanation:
http://msdn.microsoft.com/en-us/library/ms188300.aspx
http://msdn.microsoft.com/en-us/library/ms191250.aspx

QUESTION 50
You need to provide referential integrity between the Sessions table and Speakers table. Which code segment should you add at line 47 of Tables.sql?

A.
```
ALTER TABLE dbo.Speakers ADD CONSTRAINT
FK_Speakers_Sessions FOREIGN KEY (SpeakerID)
REFERENCES dbo.Sessions (SessionID);
```

B.
```
ALTER TABLE dbo.Sessions ADD CONSTRAINT
FK_Sessions_Speakers FOREIGN KEY (SessionID)
REFERENCES dbo.Speakers (SpeakerID);
```

C.
```
ALTER TABLE dbo.Sessions ADD CONSTRAINT
FK_Sessions_Speakers FOREIGN KEY (SpeakerID)
REFERENCES dbo.Speakers (SpeakerID);
```

D.
```
ALTER TABLE dbo.Speakers ADD CONSTRAINT
FK_Speakers_Sessions FOREIGN KEY (SessionID)
REFERENCES dbo.Sessions (SessionID);
```

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

Answer: C

Explanation:
http://msdn.microsoft.com/en-us/library/ms189049.aspx
http://msdn.microsoft.com/en-us/library/ms179610.aspx
http://msdn.microsoft.com/en-us/library/ff878370.aspx

QUESTION 51
You need to add a new column named Confirmed to the Attendees table. The solution must meet the following requirements:
- Have a default value of false.
- Minimize the amount of disk space used.

Which code block should you use?

A.
```
ALTER TABLE Attendees
ADD Confirmed bit DEFAULT 1;
```

B.
```
ALTER TABLE Attendees
ADD Confirmed bit DEFAULT 0;
```

C.
```
ALTER TABLE Attendees
ADD Confirmed char(1) DEFAULT '1';
```

D.
```
ALTER TABLE Attendees
ADD Confirmed char(1) DEFAULT '0';
```

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

Answer: B
Explanation:
http://msdn.microsoft.com/en-us/library/ms177603.aspx

QUESTION 52
You execute usp_TestSpeakers. You discover that usp_SelectSpeakersByName uses inefficient execution plans. You need to update usp_SelectSpeakersByName to ensure that the most efficient execution plan is used. What should you add at line 30 of Procedures.sql?

A. `OPTION (FCRCESCAN)`
B. `OPTION (FCRCESEEX)`
C. `OPTION (OPTIMIZE FOR UNKNOWN)`
D. `OPTION (OPTIMIZE FOR (@lastName = 'Anderson'))`

Answer: C
Explanation:
http://msdn.microsoft.com/en-us/library/ms181714.aspx

QUESTION 53

You need to create the object used by the parameter of usp_InsertSessions. Which statement should you use?

A. CREATEXML SCHEMA COLLECTION SessionDataTable
B. CREATETYPE SessionDataTable AS Table
C. CREATESCHEMA SessionDataTable
D. CREATETABLE Sessior.EataTable

Answer: B
Explanation:
Because usp_InsertSessions is making use of the INSERT…SELECT TSQL statement.

QUESTION 54
Developers report that usp_UpdateSessionRoom periodically returns error 3960. You need to prevent the error from occurring. The solution must ensure that the stored procedure returns the original values to all of the updated rows. What should you configure in Procedures.sql?

A. Replace line 46 with the following code:
SET TRANSACTION ISOLATION LEVEL SERIALIZABLE
B. Replace line 46 with the following code:
SETTRANSACTIONISOLATIONLEVELREPEATABLEREAD
C. Move the SELECT statement at line 49 to line 57
D. Move the SET statement at line 46 to line 53

Answer: A

QUESTION 55
You discover that usp.SelectSpeakersByName executes slowly if usp_UpdateSpeakerName executes simultaneously. You need to minimize the execution time of usp.SelectSpeakersByName. The solution must not affect the performance of the other stored procedures. What should you update?

A. Usp_UpdateSpeakerName to use the NOLOCK query hint
B. Usp_UpdateSpeakerName to use snapshot isolation
C. Usp_SelectSpeakersByName to use the NOLOCK query hint
D. Usp_SelectSpeakersByName to use snapshot isolation

Answer: C
Explanation:
NOLOCK: is equivalent to READUNCOMMITTED.
READUNCOMMITTED: specifies that dirty reads are allowed.

QUESTION 56
While testing usp.GetFutureSessions, you discover that IX_Sessions is accessed by a scan rather than a seek. You need to minimize the amount of time it takes to execute usp_GetFutureSessions. What should you do? (Each correct answer presents part of the solution. Choose all that apply.)

☐ A. Change line 02 of Indexes.sql to:

```
(DeliveryTime, SessionID)
```

☐ B. At line 04 of Indexes.sql, add:

```
WHERE GETDATE() < DeliveryTime;
```

☐ C. Change line 02 of Indexes.sql to:

```
(SpeakerID, RoomID, DeliveryTime)
```

☐ D. Change line 74 of Procedures.sql to:

```
WHERE GETDATE() > DeliveryTime;
```

☐ E. Change line 74 of Procedures.sql to:

```
WHERE GETDATE() < DeliveryTime;
```

☐ F. At line 04 of Indexes.sql, add:

```
WHERE GETDATE() > DeliveryTime;
```

A. Option A
B. Option B
C. Option C
D. Option D
E. Option E
F. Option F

Answer: BE
Explanation:
Future delivery dates.

QUESTION 57
You need to ensure that if any of the statements in usp_UpdateSpeakerName return an error message, all of the changes executed by usp_UpdateSpeakerName are not committed to the database. What should you do in Procedures.sql? (Each correct answer presents part of the solution. Choose all that apply.)

A. Add the following at line 17:

```
ROLLBACK TRANSACTION
```

B. Add the following at line 05:

```
BEGIN TRANSACTION SpeakerUpdate
```

C. Add the following at line 05:

```
SAVE TRANSACTION SpeakerUpdate
```

D. Add the following at line 17:

```
ROLLBACK TRANSACTION SpeakerUpdate
```

E. Add the following at line 07:

```
BEGIN TRANSACTION
```

A. Option A
B. Option B
C. Option C
D. Option D
E. Option E

Answer: BD

Case Study 2: Scenario 2 (Question 58 ~ Question 68)

All Information

Your company receives invoices in XML format from customers. Currently, the invoices are stored as files and processed by a desktop application. The application has several performance and security issues. The application is being migrated to a SQL Server-based solution. A schema named InvoiceSchema has been created for the invoices xml. The data in the invoices is sometimes incomplete. The incomplete data must be stored and processed as-is. Users cannot filter the data provided through views. You are designing a SQL Server database named DB1 that will be used to receive, process, and securely store the invoice data. A third-party Microsoft .NET Framework component will be purchased to perform tax calculations. The third-party tax component will be provided as a DLL file named Treytax.dll and a source code file named Amortize.cs. The component will expose a class named TreyResearch and a method named Amortize(). The files are located in c:\temp\. The following graphic shows the planned tables:

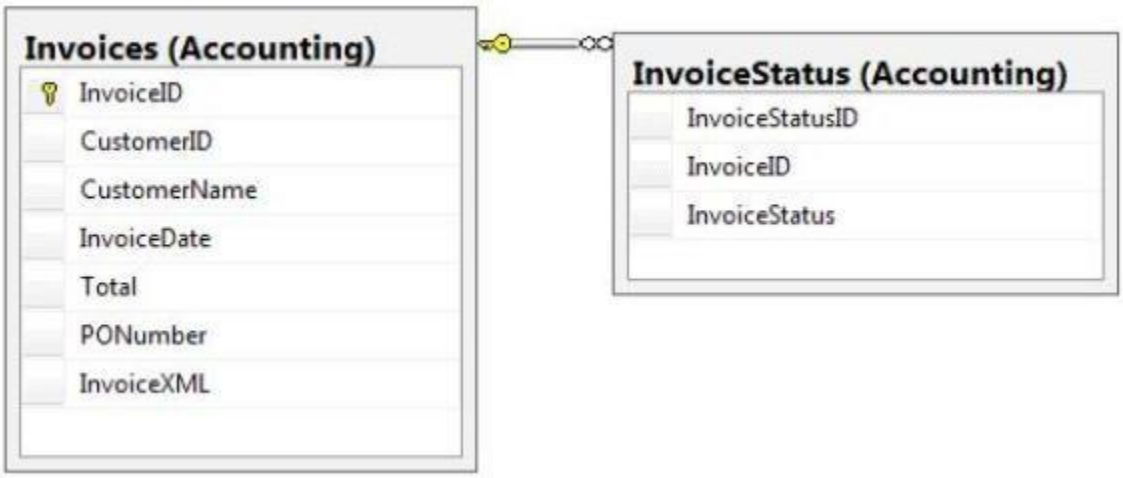

You have a sequence named Accounting.InvoiceID_Seq. You plan to create two certificates named CERT1 and CERT2. You will create CERT1 in master. You will create CERT2 in DB1. You have a legacy application that requires the ability to generate dynamic T-SQL statements against DB1. A sample of the queries generated by the legacy application appears in Legacy.sql.

Application Requirements

The planned database has the following requirements:

- All stored procedures must be signed.
- The original XML invoices must be stored in the database.
- An XML schema must be used to validate the invoice data.
- Dynamic T-SQL statements must be converted to stored procedures.
- Access to the .NET Framework tax components must be available to T-SQL objects.
- Columns must be defined by using data types that minimize the amount of space used by each table.
- Invoices stored in the InvoiceStatus table must refer to an invoice by the same identifier used by the Invoice table.
- To protect against the theft of backup disks, invoice data must be protected by using the highest level of encryption.
- The solution must provide a table-valued function that provides users with the ability to filter invoices by customer.
- Indexes must be optimized periodically based on their fragmentation by using the minimum amount of administrative effort.

Usp_InsertInvoices.sql

```
01 CREATE PROCEDURE InsertInvoice @XML nvarchar(1000)
02 AS
03 DECLARE @XmlDocumentHandle INT;
04 DECLARE @XmlDocument nvarchar(1000);
05 SET @XmlDocument = @XML;
06
07 EXEC sp_xml_preparedocument @XmlDocumentHandle OUTPUT, @XmlDocument;
08
09 INSERT INTO DB1.Accounting.Invoices (
10   InvoiceID,
11   InvoiceXML,
12   CustomerID,
13   CustomerName,
14   InvoiceDate,
15   Total,
16   PONumber
17 )
18 SELECT (NEXT VALUE FOR Accounting.InvoiceID_Seq),
19   @XML, * FROM OPENXML (@XmlDocumentHandle, '/Invoice',2)
20   WITH (
21     CustomerID nvarchar(11) 'Customer/@ID',
22     CustomerName nvarchar(50) 'Customer/@Name',
23     InvoiceDate date 'InvoiceDate',
24     Total decimal(8, 2) 'Total',
25     PONumber bigint 'PONumber'
26   );
27
28 EXEC sp_xml_removedocument @XmlDocumentHandle;
```

Invoices.xml

All customer IDs are 11 digits. The first three digits of a customer ID represent the customer's country. The remaining eight digits are the customer's account number. The following is a sample of a customer invoice in XML format:

```
01 <?xml version="1.0"?>
29 <Invoice InvoiceDate="2012-02-20">
30   <Customer ID="00156590099" Name="Litware" />
31   <Total>125</Total>
32   <PONumber>1666</PONumber>
33 </Invoice>
```

InvoiceByCustomer.sql

```
01 (SELECT CustomerID,
34   CustomerName,
35   InvoiceID,
36   InvoiceDate,
37   Total,
38   PONumber
39   FROM Accounting.Invoices
40   WHERE CustomerID=@CustID);
```

Legacy.sql

Legacy.sql

```
01 DECLARE @sqlstring AS nvarchar(1000);
41 DECLARE @CustomerID AS varchar(11), @Total AS decimal(8,2);
42
43 SET @sqlstring=N'SELECT CustomerID, InvoiceID, Total
44   FROM Accounting.Invoices
45   WHERE CustomerID=@CustomerID AND Total > @Total;';
46
47 EXEC sys.sp_executesql
48   @statement=@sqlstring,
49   @params=N'@CustomerID AS varchar(11), @Total AS decimal(8,2)',
50   @CustomerID=999, @Total=500;
```

CountryFromID.sql

```
01 CREATE FUNCTION CountryFromID (@CustomerID varchar(11)) RETURNS varchar(20)
51 AS
52 BEGIN
53   DECLARE @Country varchar(20);
54   SET @CustomerID = LEFT(@CustomerID,3);
55   SELECT @Country = CASE @CustomerID
56     WHEN '001'
57       THEN 'United States'
58     WHEN '002'
59       THEN 'Spain'
60     WHEN '003'
61       THEN 'Japan'
62     WHEN '004'
63       THEN 'China'
64     WHEN '005'
65       THEN 'Brazil'
66     ELSE 'Other'
67   END;
68   RETURN @CustomerID;
69 END;
```

IndexManagement.sql

```
01 DECLARE @IndexTable TABLE (
70   TableName varchar(100), IndexName varchar(100), Fragmentation int, RowNumber int
71   );
72 DECLARE @TableName sysname, @IndexName sysname, @Fragmentation int,
73   @RowNumber int, @sqlcommand varchar(1000);
74
75 INSERT INTO @IndexTable (TableName, IndexName, Fragmentation, Rownumber)
76   SELECT OBJECT_NAME(i.Object_id),
77     i.name AS IndexName,
78     indexstats.avg_fragmentation_in_percent,
79     ROW_NUMBER() OVER(ORDER BY i.name DESC) AS 'RowNumber'
80   FROM sys.dm_db_index_physical_stats(DB_ID(), NULL, NULL, NULL, 'DETAILED')
81     AS indexstats INNER JOIN sys.indexes AS i
82     ON i.OBJECT_ID = indexstats.OBJECT_ID AND i.index_id = indexstats.index_id;
83
84 DECLARE @counter int = 0;
85
86 WHILE @counter < (SELECT RowNumber FROM @indextable)
87   BEGIN
88     SET @counter = @counter + 1;
89     WITH t AS (
90       SELECT TableName, IndexName, Fragmentation
91       FROM @IndexTable WHERE RowNumber = @counter
92       )
93     SELECT
94       @TableName= TableName,
95       @IndexName = IndexName,
96       @Fragmentation = Fragmentation
97     FROM t;
98
99     IF @Fragmentation <= 30
100       BEGIN
101         SET @sqlCommand =
102           N'ALTER INDEX '+@indexName+N' ON '+@TableName+N' REORGANIZE';
103         EXEC sp_executesql @sqlCommand;
104       END;
105     ELSE
106       BEGIN
107         SET @sqlCommand=N'ALTER INDEX '+@indexName+N' ON '+@TableName+N' REBUILD';
108         EXEC sp_executesql @sqlCommand;
109       END;
110     END;
```

QUESTION 58

You need to modify InsertInvoice to comply with the application requirements. Which code segment should you execute?

A.
```
OPEN CERT1;
ALTER PROCEDURE Accounting.usp_AuthPayment
WITH ENCRYPTION;
CLOSE CERT1;
```

B.
```
OPEN CERT2;
ALTER PROCEDURE Accounting,usp_AuthPayment
WITH ENCRYPTION;
CLOSE CERT2;
```

C.
```
ADD SIGNATURE TO Accounting.usp_AuthPayment
BY CERTIFICATE CERT1;
```

D.
```
ADD SIGNATURE TO Accounting.usp_AuthPayment
BY CERTIFICATE CERT2;
```

Answer: D

Explanation:

http://msdn.microsoft.com/en-us/library/bb669102.aspx

QUESTION 59
Which data type should you use for CustomerID?

A. varchar(11)
B. bigint
C. nvarchar(11)
D. char(11)

Answer: D
Explanation:
http://msdn.microsoft.com/en-us/library/ms176089.aspx
http://msdn.microsoft.com/en-us/library/ms187745.aspx

QUESTION 60
You need to modify the function in CountryFromID.sql to ensure that the country name is returned instead of the country ID. Which line of code should you modify in CountryFromID.sql?

A. 06
B. 04
C. 19
D. 05

Answer: C
Explanation:
http://msdn.microsoft.com/en-us/library/ms186755.aspx
http://msdn.microsoft.com/en-us/library/ms191320.aspx

QUESTION 61
Drag and Drop Question
You have a SQL Server 2012 database named database1. Databaset has a data file named database1_data.mdf and a transaction log file named database1_log.ldf. Database1_data.mdf is 1.5GB. Database1_log.ldf is 1.5 terabytes. A full backup of Database1 is performed every day. You need to reduce the size of the log file. The solution must ensure that you can perform transaction log backups in the future. Which code segment should you execute? To answer, move the appropriate code segments from the list of code segments to the answer area and arrange them in the correct order.

A.
```
CREATE ASSEMBLY TaxCalc FROM 'c:\temp\TreyTax.DLL'
```

B.
```
EXEC SP_CONFIGURE 'clr enabled', '1';
```

C.
```
CREATE FUNCTION Accounting.Amortize(
  @total decimal(8,2),@period int
  )RETURNS decimal(8,2)
AS EXTERNAL NAME TaxCalc.TreyResearch.Amortize;
```

D.
```
EXEC sp_recompile @objname = 'TaxCalc'
```

E.
```
CREATE ASSEMBLY TaxCalc FROM 'C:\temp\Amortize.cs';
```

F.
```
RECONFIGURE;
```

A. Option A
B. Option B
C. Option C
D. Option D
E. Option E
F. Option F

Answer: ACDE

QUESTION 62
You need to create the InvoiceStatus table in DB1. How should you define the InvoiceID column in the CREATE TABLE statement?

A.
```
InvoiceID bigint
DEFAULT (NEXT VALUE FOR Accounting.InvoiceID_Seq) NOT NULL,
```

B.
```
InvoiceID bigint DEFAULT ((NEXT VALUE
FOR Accounting.InvoiceID_Seq OVER
(ORDER BY InvoiceStatusID))) NOT NULL FOREIGN
KEY REFERENCES Accounting.Invoices(InvoiceID),
```

C.
```
InvoiceID bigint FOREIGN KEY REFERENCES
Accounting.Invoices(InvoiceID) NOT NULL,
```

D.
```
InvoiceID bigint DEFAULT ((NEXT VALUE
FOR Accounting.InvoiceID_Seq
OVER (ORDER BY InvoiceStatusID))) NOT NULL,
```

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

Answer: C

QUESTION 63
You execute IndexManagement.sql and you receive the following error message:
"Msg 512, Level 16, State 1, Line 12 Subquery returned more than 1 value. This is not permitted when the subquery follows =, ! = , <, <= , >, > = or when the subquery is used as an expression."
You need to ensure that IndexManagement.sql executes properly. Which WHILE statement should you use at line 18?

A. `WHILE @counter < (SELECT COUNT(RowNumber) FROM @indextable)`
B. `WHILE @counter < (SELECT SUM(RowNumber) FROM @indextable)`
C. `WHILE SUM(@RowNumber) < (SELECT @counter FROM @indextable)`
D. `WHILE COUNT(@RowNumber) < (SELECT @counter FROM @indextable)`

Answer: A

QUESTION 64
You attempt to process an invoice by using usp_InsertInvoice.sql and you receive the following error message:
"Msg 515, Level 16, State 2, Procedure usp_InsertInvoice, Line 10 Cannot insert the value NULL into column 'InvoiceDate', table 'DB1.Accounting.Invoices'; column does not allow nulls. INSERT fails."
You need to modify usp_InsertInvoice.sql to resolve the error. How should you modify the INSERT statement?

A. InvoiceDate date `Customer/@InvoiceDate1',
B. InvoiceDate varchar(100) 'InvoiceDate',
C. InvoiceDate varchar(100) 'Customer/InvoiceDate',
D. InvoiceDate date `@InvoiceDate',

Answer: D
Explanation:
http://msdn.microsoft.com/en-us/library/ms187926.aspx
http://msdn.microsoft.com/en-us/library/ms190782.aspx
http://msdn.microsoft.com/en-us/library/bb669091.aspx
http://msdn.microsoft.com/en-us/library/windows/desktop/ms709342.aspx
http://msdn.microsoft.com/en-us/library/ms188001.aspx

QUESTION 65
You are testing disaster recovery procedures. You attempt to restore DB1 to a different server and you receive the following error message:
"Msg 33111, Level 16, State 3, Line 1 Cannot find server certificate with thumbprint `OxA694FBEA88C9354E5E2567C30A2A69E8FB4C44A9'.
Msg 3013, Level 16, State 1, Line 1 RESTORE DATABASE is terminating abnormally."
You need to ensure that you can restore DB1 to a different server. Which code segment should you execute?

A.
```
CREATE CERTIFICATE CERT1
  ENCRYPTION BY PASSWORD='p@ssw0rd1'
  WITH SUBJECT = 'EncryptionCertificate';
```

B.
```
CREATE CERTIFICATE CERT1
  FROM FILE='CERT1.CER'
  WITH PRIVATE KEY (FILE = 'CERT1.KEY',
  DECRYPTION BY PASSWORD='p@ssw0rd1');
```

C.
```
RESTORE CERTIFICATE CERT2
  FROM FILE='CERT2.CER'
  WITH PRIVATE KEY (FILE = 'CERT2.KEY',
  DECRYPTION BY PASSWORD='p@ssw0rd1');
```

D.
```
CREATE CERTIFICATE CERT2
  ENCRYPTION BY PASSWORD='p@ssw0rd1'
  WITH SUBJECT = 'EncryptionCertificate';
```

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

Answer: B

QUESTION 66
You need to modify the function in CountryFromID.sql to ensure that the country name is returned instead of the country ID. Which line of code should you modify in CountryFromID.sql?

A. 04
B. 05
C. 06
D. 19

Answer: D
Explanation:
http://msdn.microsoft.com/en-us/library/ms186755.aspx
http://msdn.microsoft.com/en-us/library/ms191320.aspx

QUESTION 67
You need to convert the functionality of Legacy.sql to use a stored procedure. Which code segment should the stored procedure contain?

A.
```
CREATE PROC usp_InvoicesByCustomerAboveTotal(
  @sqlstring AS nvarchar(1000),
  @CustomerID AS char(11),
  @Total AS decimal(8,2))
AS
...
```

B.
```
CREATE PROC usp_InvoicesByCustomerAboveTotal(
  @sqlstring AS nvarchar(1000))
AS
...
```

C.
```
CREATE PROC usp_InvoicesByCustomerAboveTotal(
  @sqlstring AS nvarchar(1000),
  OUTPUT @CustomerID AS char(11),
  OUTPUT @Total AS decimal(8,2))
AS
...
```

D.
```
CREATE PROC usp_InvoicesByCustomerAboveTotal (
  @CustomerID AS char(11), @Total AS decimal(8,2))
AS
...
```

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

Answer: D
Explanation:
http://msdn.microsoft.com/en-us/library/ms187926.aspx
http://msdn.microsoft.com/en-us/library/ms190782.aspx
http://msdn.microsoft.com/en-us/library/bb669091.aspx
http://msdn.microsoft.com/en-us/library/windows/desktop/ms709342.aspx
http://msdn.microsoft.com/en-us/library/ms188001.aspx

QUESTION 68
Drag and Drop Question
You need to build a stored procedure that amortizes the invoice amount. Which code segment should you use to create the stored procedure? To answer, move the appropriate code segments from the list of code segments to the answer area and arrange them in the correct order.

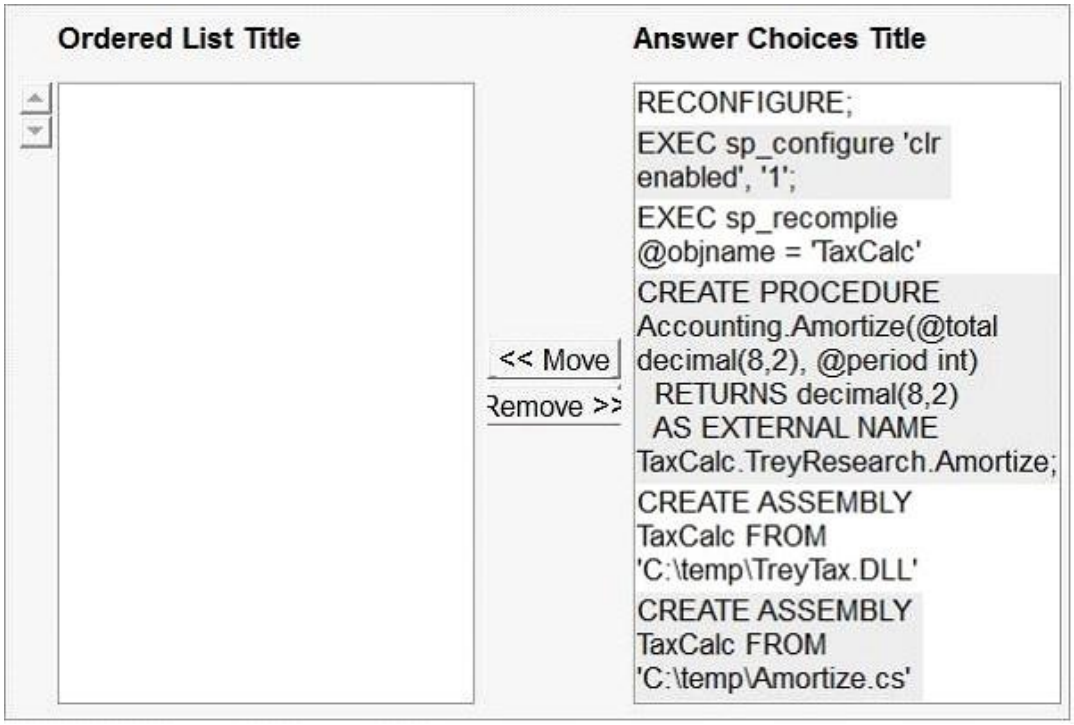

Answer:

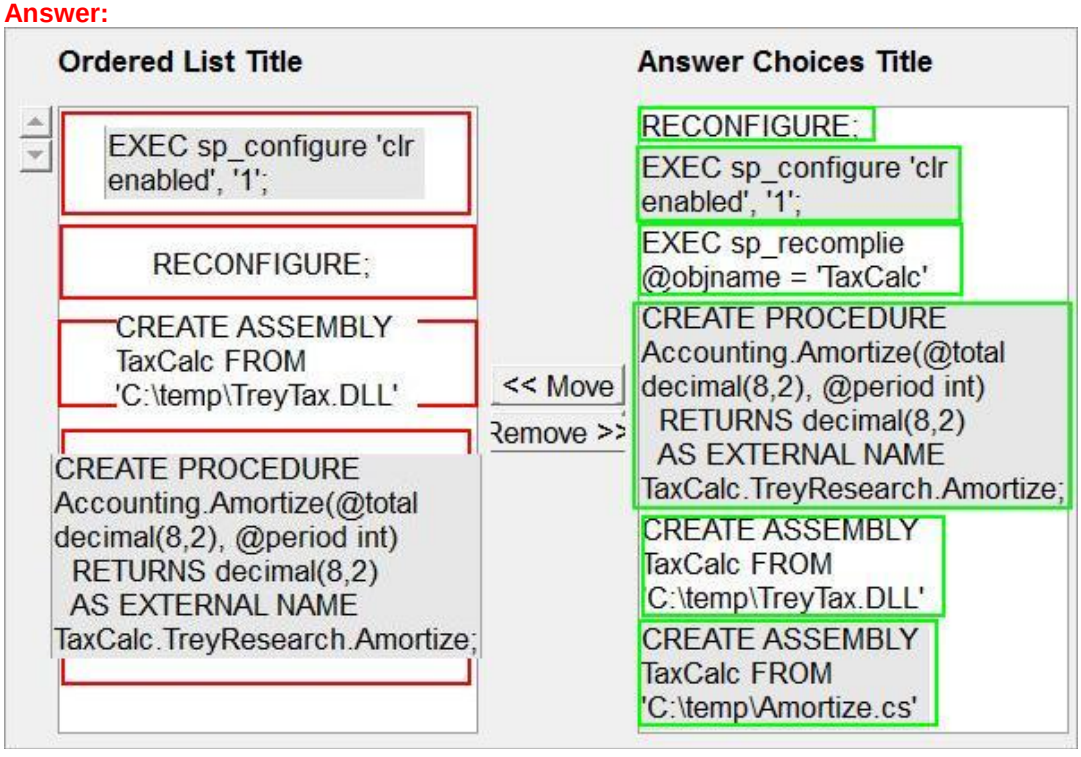

Case Study 3: Scenario 3 (Question 69 ~ Question 79)

Application Information

You are a database administrator for a manufacturing company. You have an application that stores product data. The data will be converted to technical diagrams for the manufacturing process.

The product details are stored in XML format. Each XML must contain only one product that has a root element named Product. A schema named Production.ProductSchema has been created for the products xml. You develop a Microsoft .NET Framework assembly named ProcessProducts.dll that will be used to convert the XML files to diagrams. The diagrams will be stored in the database as images. ProcessProducts.dll contains one class named ProcessProduct that has a method name of Convert(). ProcessProducts.dll was created by using a source code file named ProcessProduct.es. All of the files are located in C:\Products\. The application has several performance and security issues. You will create a new database named ProductsDB on a new server that has SQL Server 2012 installed. ProductsDB will support the application. The following graphic shows the planned tables for ProductsDB:

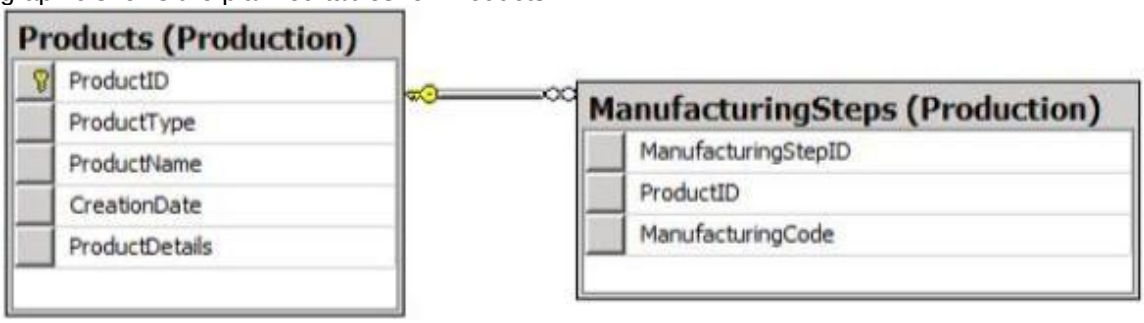

You will also add a sequence named Production.ProductID_Seq. You plan to create two certificates named DBCert and ProductsCert. You will create ProductsCert in master. You will create DBCert in ProductsDB. You have an application that executes dynamic T-SQL statements against ProductsDB. A sample of the queries generated by the application appears in Dynamic.sql.

Application Requirements

The planned database has the following requirements:

- All stored procedures must be signed.
- The amount of disk space must be minimized.
- Administrative effort must be minimized at all times.
- The original product details must be stored in the database.
- An XML schema must be used to validate the product details.
- The assembly must be accessible by using T-SQL commands.
- A table-valued function will be created to search products by type.
- Backups must be protected by using the highest level of encryption.
- Dynamic T-SQL statements must be converted to stored procedures.
- Indexes must be optimized periodically based on their fragmentation.
- Manufacturing steps stored in the ManufacturingSteps table must refer to a product by the same identifier used by the Products table.

ProductDetails_Insert.sql

```
01 CREATE PROCEDURE Production.ProductDetails_Insert @XML nvarchar(1000)
02 AS
03 DECLARE @handle INT;
04 DECLARE @document nvarchar(1000);
05 SET @document = @XML;
06
07 EXEC sp_xml_preparedocument @handle OUTPUT, @document;
08
09 INSERT INTO PRODUCTSDB.Production.Invoices (
10   ProductID,
11   ProductDetails,
12   ProductType,
13   ProductName,
14   CreationDate
15 )
16 SELECT (NEXT VALUE FOR Production.ProductID_Seq),
17   @XML, * FROM OPENXML (@handle, '/Invoice',2)
18   WITH (
19     ProductType nvarchar(11) 'ProductType/ID',
20     ProductName nvarchar(50) '@ProductName',
21     CreationDate date 'CreationDate'
22   );
23
24 EXEC sp_xml_removedocument @handle;
```

Product. xml

All product types are 11 digits. The first five digits of the product id reference the category of the product and the remaining six digits are the subcategory of the product. The following is a sample customer invoice in XML format:

```
01 <?xml version="1.0"?>
25 <Product ProductName="Widget">
26   <ProductType ID="00156590099" />
27   <CreationDate>2011-08-05</CreationDate>
28   </Invoice>
```

ProductsByProductType.sql

```
01 (SELECT ProductID,
29   ProductType,
30   CreationDate
31   FROM Production.Products
32     WHERE ProductType=@ProductType);
```

Dynamic.sql

```
01 DECLARE @tsql AS nvarchar(500);
33 DECLARE @ProductType AS varchar(11), @CreationDate AS date;
34
35 SET @sqlstring=N'SELECT ProductID, ProductType, CreationDate
36   FROM Production.Product
37   WHERE ProductID=@ProductID AND CreationDate > @CreationDate;';
38
39 EXEC sys.sp_executesql
40   @statement=@sqlstring,
41   @params=N'@ ProductType AS varchar(11), @CreationDate AS date',
42 @ProductType=00125061246, @Total='2012-05-10';
```

CategoryFromType.sql

```
01 CREATE FUNCTION CategoryFromType (@Type varchar(11)) RETURNS nvarchar(20)
43 AS
44 BEGIN
45   DECLARE @Category AS varchar(20);
46   SET @Category = LEFT(@Category,5);
47   SELECT @Category = CASE @Type
48     WHEN '00001'
49       THEN 'Bikes'
50     WHEN '00002'
51       THEN 'Wheels'
52     ...
53     ELSE 'Other'
54   END;
55 RETURN @Category;
56 END;
```

IndexManagement.sql

```
01 DECLARE @IndexTable TABLE (
57   TableName varchar(100), IndexName varchar(100), Fragmentation int, RowNumber int
58   );
59 DECLARE @TableName sysname, @IndexName sysname, @Fragmentation int,
60   @RowNumber int, @sqlcommand varchar(1000);
61
62 INSERT INTO @IndexTable (TableName, IndexName, Fragmentation, Rownumber)
63   SELECT OBJECT_NAME(i.Object_id),
64     i.name AS IndexName,
65     indexstats.avg_fragmentation_in_percent,
66     ROW_NUMBER() OVER(ORDER BY i.name DESC) AS 'RowNumber'
67   FROM sys.dm_db_index_physical_stats(DB_ID(), NULL, NULL, NULL, 'DETAILED')
68     AS indexstats INNER JOIN sys.indexes AS i
69     ON i.OBJECT_ID = indexstats.OBJECT_ID AND i.index_id = indexstats.index_id;
70
71 DECLARE @counter int = 0;
72
73 WHILE @counter < (SELECT RowNumber FROM @indextable)
74   BEGIN
75     SET @counter = @counter + 1;
76     WITH t AS (
77       SELECT TableName, IndexName, Fragmentation
78       FROM @IndexTable WHERE RowNumber = @counter
79       )
80     SELECT
81       @TableName= TableName,
82       @IndexName = IndexName,
83       @Fragmentation = Fragmentation
84     FROM t;
85
86     IF @Fragmentation <= 30
87       BEGIN
88         SET @sqlCommand =
89           N'ALTER INDEX '+@indexName+N' ON '+@TableName+N' REORGANIZE';
90         EXEC sp_executesql @sqlCommand;
91       END;
92     ELSE
93       BEGIN
94         SET @sqlCommand=N'ALTER INDEX '+@indexName+N' ON '+@TableName+N' REBUILD';
95         EXEC sp_executesql @sqlCommand;
96       END;
97     END;
```

QUESTION 69

Which code segment should you use to define the ProductDetails column?

A. `ProductDetails varchar(MAX) NULL`
B. `ProductDetails xml NULL`
C. `ProductDetails xml (CONTENT Production.ProductDetailsSchema) NULL`
D. `ProductDetails xml (DOCUMENT Production.ProductDetailsSchema) NULL`

Answer: A

QUESTION 70
You need to prepare the database to use the .NET Framework ProcessProducts component. Which code segments should you execute? (Each correct answer presents part of the solution. Choose all that apply.)

A.
```
CREATE ASSEMBLY ProductionAssembly FROM 'C:\Products\ProcessProducts.DLL'
```

B.
```
RECONFIGURE;
```

C.
```
EXEC sp_recompile @objname = 'Production.ProcessProduct';
```

D.
```
CREATE TYPE Production.ProcessProduct
EXTERNAL NAME ProductionAssembly.ProcessProductss.Process;
```

E.
```
Exec SP_CONFIGURE 'clr enabled', '1';
```

F.
```
CREATE PROCEDURE Production.ProcessProduct(
  @ProductID int, @ProductType varchar(11)
  )
AS EXTERNAL NAME ProductionAssembly.ProcessProducts.Process;
```

G.
```
CREATE ASSEMBLY ProductionAssembly FROM 'C:\Products\ProcessProducts.cs';
```

A. Option A
B. Option B
C. Option C
D. Option D
E. Option E
F. Option F
G. Option G

Answer: ABDE
Explanation:
http://msdn.microsoft.com/en-us/library/ms131048.aspx
http://msdn.microsoft.com/en-us/library/ms131052.aspx
http://msdn.microsoft.com/en-us/library/ms189524.aspx
http://msdn.microsoft.com/en-us/library/ms345106.aspx
http://msdn.microsoft.com/en-us/library/ms131107.aspx

QUESTION 71
You are planning the ManufacturingSteps table. You need to define the ProductID column in the CREATE TABLE statement. Which code segment should you use?

A.
```
ProductID bigint
DEFAULT (NEXT VALUE FOR Production.ProductID_Seq) NOT NULL,
```

B.
```
ProductID bigint FOREIGN KEY REFERENCES
Production.Product(ProductID) NOT NULL,
```

C.
```
ProductID bigint DEFAULT
((NEXT VALUE FOR Production.ProductID_Seq OVER
(ORDER BY ManufacturingStepID))) NOT NULL,
```

D.
```
ProductID bigint DEFAULT
((NEXT VALUE FOR Production.ProductID_Seq OVER
(ORDER BY ManufacturingStepID)))
NOT NULL FOREIGN KEY REFERENCES
Production.Product(ProductID),
```

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

Answer: B
Explanation:
http://msdn.microsoft.com/en-us/library/ms189049.aspx
http://msdn.microsoft.com/en-us/library/ms179610.aspx
http://msdn.microsoft.com/en-us/library/ff878370.aspx

QUESTION 72
You execute IndexManagement.sql and you receive the following error message:
"Msg 512, Level 16, State 1, Line 12 Subquery returned more than 1 value. This is not permitted when the subquery follows =, ! = , <, <= , >, >= or when the subquery is used as an expression."
You need to ensure that IndexManagement.sql executes properly. Which WHILE statement should you use at line 18?

A. `WHILE @counter < (SELECT SUM(RowNumber) FROM @indextable)`
B. `WHILE COUNT(@RowNumber) < (SELECT @counter FROM @indextable)`
C. `WHILE SUM(@RowNumber) < (SELECT @counter FROM @indextable)`
D. `WHILE @counter < (SELECT COUNT(RowNumber) FROM @indextable)`

Answer: D

QUESTION 73
An administrator provides a digital certificate named ServerCert. You need to implement Transparent Data Encryption (TDE) on ProductsDB. Which code segment should you use?

A.
```
USE PRODUCTSDB;
GO
CREATE DATABASE ENCRYPTION KEY WITH ALGORITHM = TRIPLE_DES_3KEY
ENCRYPTION BY SERVER CERTIFICATE PRODUCTSCERT;
GO
ALTER DATABASE PRODUCTSDB SET ENCRYPTION ON;
GO
```

B.
```
USE PRODUCTSDB;
GO
CREATE DATABASE ENCRYPTION KEY WITH ALGORITHM = TRIPLE_DES_3KEY
ENCRYPTION BY SERVER CERTIFICATE DBCERT;
GO
ALTER DATABASE PRODUCTSDB SET ENCRYPTION ON;
GO
```

C.
```
USE PRODUCTSDB;
GO
CREATE DATABASE ENCRYPTION KEY WITH ALGORITHM = AES_256
ENCRYPTION BY SERVER CERTIFICATE PRODUCTSCERT;
GO
ALTER DATABASE PRODUCTSDB SET ENCRYPTION ON;
GO
```

D.
```
USE PRODUCTSDB;
GO
CREATE DATABASE ENCRYPTION KEY WITH ALGORITHM = AES_256
ENCRYPTION BY SERVER CERTIFICATE DBCERT;
GO
ALTER DATABASE PRODUCTSDB SET ENCRYPTION ON;
GO
```

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

Answer: C
Explanation:
http://msdn.microsoft.com/en-us/library/bb934049.aspx

QUESTION 74
You need to modify Production.ProductDetails_Insert to comply with the application requirements.
Which code segment should you execute?

A.
```
OPEN PRODUCTSCERT;
ALTER PROCEDURE Production.ProductDetails_Insert
  WITH ENCRYPTION;
CLOSE PRODUCTSCERT;
```

B.
```
OPEN DBCERT;
ALTER PROCEDURE Production.ProductDetails_Insert
  WITH ENCRYPTION;
CLOSE DBCERT;
```

C.
```
ADD SIGNATURE TO Production.ProductDetails_Insert
  BY CERTIFICATE DBCERT;
```

D.
```
ADD SIGNATURE TO Production.ProductDetails_Insert
  BY CERTIFICATE PRODUCTSCERT;
```

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

Answer: C
Explanation:
http://msdn.microsoft.com/en-us/library/bb669102.aspx

QUESTION 75
You need to create a function that will use a SELECT statement in ProductsByProductType.sql. Which code segment should you use to complete the function?

A.
```
CREATE FUNCTION Production.fnProductsByProductType (@ProductType varchar(11))
RETURNS @tblInvoices TABLE (ProductID bigint, ProductType varchar(11), CreationDate
date)
AS
INSERT INTO @tblInvoices
```

B.
```
CREATE FUNCTION Production.fnProductsByProductType (@ProductType varchar(11))
RETURNS TABLE
AS
RETURN
```

C.
```
CREATE FUNCTION Production.fnProductsByProductType (@ProductType varchar(11))
RETURNS @TblInvoices TABLE (ProductID bigint, ProductType varchar(11), CreationDate
date)
AS
```

D.
```
CREATE FUNCTION Production.fnProductsByProductType (@ProductType varchar(11))
RETURNS xml
AS
RETURN
```

A. Option A
B. Option B
C. Option C

D. Option D

Answer: B
Explanation:
http://msdn.microsoft.com/en-us/library/ms191320.aspx
http://msdn.microsoft.com/en-us/library/ms186755.aspx

QUESTION 76
You execute IndexManagement.sql and you receive the following error message:
"Msg 512, Level 16, State 1, Line 80 Subquery returned more than 1 value. This is not permitted when the subquery follows =, ! = , <, <= , >, > = or when the subquery is used as an expression."
You need to ensure that IndexManagement.sql executes properly. Which WHILE statement should you use at line 86?

A. `WHILE @counter < (SELECT COUNT(RowNumber) FROM @indextable)`
B. `WHILE @counter < (SELECT SUM(RowNumber) FROM @indextable)`
C. `WHILE SUM(@RowNumber) < (SELECT @counter FROM @indextable)`
D. `WHILE COUNT(@RowNumber) < (SELECT @counter FROM @indextable)`

Answer: A

QUESTION 77
While testing the CategoryFromType function, you discover that the function is returning 'Other'. You need to update CategoryFromType to return the category name. Which line of code should you modify in CategoryFromType.sql?

A. 04
B. 05
C. 12
D. 14

Answer: B

QUESTION 78
You are testing disaster recovery procedures. When you attempt to restore ProductsDB to another server, you receive the following error message:
"Msg 33111, Level 16, State 3, Line 5 Cannot find server certificate with thumbprint ' 0x9D876A3468B911E1BA4CFCBF4724019B\
Msg 3013, Level 16, State 1, Line 5 RESTORE DATABASE is terminating abnormally."
You need to ensure that you can restore ProductsDB to another server. Which code segment should you execute on the other server?

A.
```
RESTORE CERTIFICATE DBCERT
  FROM FILE='DBCERT.CER'
  WITH PRIVATE KEY (FILE = 'c:\DBCERT.KEY',
  DECRYPTION BY PASSWORD = 'SecretP@ss');
```

B.
```
CREATE CERTIFICATE PRODUCTSCERT
  ENCRYPTION BY PASSWORD = 'SecretP@ss'
  WITH SUBJECT = 'SecurityCertificate';
```

C.
```
CREATE CERTIFICATE DBCERT
  ENCRYPTION BY PASSWORD = 'SecretP@ss'
  WITH SUBJECT = 'SecurityCertificate';
```

D.
```
CREATE CERTIFICATE PRODUCTSCERT
  FROM FILE='PRODUCTSCERT.CER'
  WITH PRIVATE KEY (FILE = 'c:\PRODUCTSCERT.KEY',
  DECRYPTION BY PASSWORD = 'SecretP@ss');
```

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

Answer: D

QUESTION 79
Which data type should you use for ProductType?

A. varchar(11)
B. nvarchar(11)
C. char(11)
D. bigint

Answer: C

Case Study 4: Scenario 4 (Question 80 ~ Question 85)

Application Information
You have two servers named SQL1 and SQL2. SQL1 has SQL Server 2012 Enterprise installed. SQL2 has SQL Server 2008 Standard installed. You have an application that is used to manage employees and office space. Users report that the application has many errors and is very slow. You are updating the application to resolve the issues. You plan to create a new database on SQL1 to support the application. The script that you plan to use to create the tables for the new database is shown in Tables.sql. The script that you plan to use to create the stored procedures for the new database is shown in StoredProcedures.sql. The script that you plan to use to create the indexes for the new database is shown in Indexes.sql. A database named DB2 resides on SQL2. DB2 has a table named EmployeeAudit that will audit changes to a table named Employees. A stored procedure named usp_UpdateEmployeeName will be executed only by other stored procedures. The stored procedures executing usp_UpdateEmployeeName will always handle transactions. A stored procedure named usp_SelectEmployeesByName will be used to retrieve the names of

employees. Usp_SelectEmployeesByName can read uncommitted data. A stored procedure named usp_GetFutureOfficeAssignments will be used to retrieve office assignments that will occur in the future.

StoredProcedures.sql

```
01 CREATE PROCEDURE usp_UpdateEmployeeName
02   @EmployeesInfo EmployeesInfo READONLY
03 AS
04
05 BEGIN TRY
06
07 UPDATE Employees
08 SET LastName = ei.LastName
09 FROM Employees e
10   INNER JOIN @ EmployeesInfo ei ON e.EmployeeID = ei.EmployeeID;
11
12 INSERT INTO SQL2.DB2.dbo.EmployeeAudit(EmployeeID, LastName)
13 SELECT EmployeeID, LastName
14 FROM @EmployeesInfo;
15
16 END TRY
17 BEGIN CATCH
18
19 END CATCH;
20
21 GO
22
23 CREATE PROCEDURE usp_SelectEmployeesByName
24   @LastName nvarchar(100)
25 AS
26 SELECT EmployeeID,
27   FirstName,
28   LastName
29 FROM Employees
30 WHERE LastName LIKE @LastName + '%'
```

```
31
32 GO
33
34 CREATE PROCEDURE usp_UpdateOffice
35   @OfficeID int,
36   @EmployeeID int
37 AS
38 SET TRANSACTION ISOLATION LEVEL SNAPSHOT
39 BEGIN TRANSACTION;
40
41 SELECT OfficeID,
42   OfficeName
43 FROM Offices
44 WHERE EmployeeID = @EmployeeID;
45
46 UPDATE Offices
47 SET EmployeeID = @EmployeeID,
48   StartDate = GETDATE()
49 WHERE OfficeID = @OfficeID;
50
51 COMMIT TRANSACTION;
52
53 CREATE PROCEDURE usp_GetFutureOfficeAssignments
54 AS
55 SELECT EmployeeID,
56   OfficeID,
57   StartDate
58 FROM Offices
59 WHERE StartDate > GETDATE();
60 GO
61
```

Indexes.sql

```
01 CREATE INDEX IX_Offices ON Offices
02 (EmployeeID, StartDate)
03 INCLUDE (OfficeID)
04
05 GO
06
07 CREATE INDEX IX_Employees ON Employees
08 (LastName);
09 GO
10
```

Tables.sql

```
01 CREATE DATABASE HumanResources;
02 GO
03
04 ALTER DATABASE HumanResources
05 SET ALLOW_SNAPSHOT_ISOLATION ON;
06 GO
07
08 USE HumanResources
09 GO
10
11 CREATE TABLE Employees
12 (
13   EmployeeID int IDENTITY(1,1) NOT NULL,
14    FirstName nvarchar(100) NOT NULL,
15    LastName nvarchar(100) NOT NULL,
16
17 );
18 GO
19
20 CREATE TABLE Offices
21 (
22   OfficeID int IDENTITY(1,1) NOT NULL,
23    EmployeeID int NOT NULL,
24    OfficeName nvarchar(100) NOT NULL,
25    StartDate datetime NOT NULL
26 );
27 GO
```

QUESTION 80

You need to modify usp_SelectEmployeesByName to support server-side paging. The solution must minimize the amount of development effort required. What should you add to usp_SelectEmployeesByName?

A. an OFFSET-FETCH clause
B. a table variable
C. the ROWNUMBER keyword
D. a recursive common table expression

Answer: A

Explanation:

http://www.mssqltips.com/sqlservertip/2696/comparing-performance-for-different-sql-server-paging-methods/
http://msdn.microsoft.com/en-us/library/ms188385.aspx
http://msdn.microsoft.com/en-us/library/ms180152.aspx
http://msdn.microsoft.com/en-us/library/ms186243.aspx
http://msdn.microsoft.com/en-us/library/ms186734.aspx
http://www.sqlserver-training.com/how-to-use-offset-fetch-option-in-sql-server-order-by-clause/

http://www.sqlservercentral.com/blogs/juggling_with_sql/2011/11/30/using-offset-and-fetch/

QUESTION 81
You need to provide referential integrity between the Offices table and Employees table. Which code segment or segments should you add at line 27 of Tables.sql? (Each correct answer presents part of the solution. Choose all that apply.)

A.
```
ALTER TABLE dbo.Offices ADD CONSTRAINT
PK_Offices_EmployeeID PRIMARY KEY (EmployeeID);
```

B.
```
ALTER TABLE dbo.Employees ADD CONSTRAINT
FK_Employees_Offices FOREIGN KEY (OfficeID)
REFERENCES dbo.Offices (OfficeID);
```

C.
```
ALTER TABLE dbo.Employees ADD CONSTRAINT
PK_Employees_EmployeeID PRIMARY KEY (EmployeeID);
```

D.
```
ALTER TABLE dbo.Offices ADD CONSTRAINT
FK_Offices_Employees FOREIGN KEY (EmployeeID)
REFERENCES dbo.Employees (EmployeeID);
```

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

Answer: CD
Explanation:
http://msdn.microsoft.com/en-us/library/ms189049.aspx

QUESTION 82
You execute usp_SelectEmployeesByName multiple times, passing strings of varying lengths to @>LastName. You discover that usp_SelectEmployeesByName uses inefficient execution plans. You need to update usp_SelectEmployeesByName to ensure that the most efficient execution plan is used. What should you add at line 31 of StoredProcedures.sql?

A. `OPTION (KEEPFIXED PLAN)`
B. `OPTION (KEEP PLAN)`
C. `OPTION (ROBUST PLAN)`
D. `OPTION (OPTIMIZE FOR UNKNOWN)`

Answer: D
Explanation:
http://msdn.microsoft.com/en-us/library/ms181714.aspx

QUESTION 83
You need to add a new column named Confirmed to the Employees table. The solution must meet

the following requirements:
- Have a default value of TRUE.
- Minimize the amount of disk space used.

Which code segment should you use?

A.
```
ALTER TABLE Employees
ADD Confirmed char(1) DEFAULT '1';
```

B.
```
ALTER TABLE Employees
ADD Confirmed char(1) DEFAULT '0';
```

C.
```
ALTER TABLE Employees
ADD Confirmed bit DEFAULT 0;
```

D.
```
ALTER TABLE Employees
ADD Confirmed bit DEFAULT 1;
```

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

Answer: D

QUESTION 84
You need to recommend a solution to ensure that SQL1 supports the auditing requirements of usp_UpdateEmployeeName. What should you include in the recommendation?

A. change data capture
B. transactional replication
C. change tracking
D. the Distributed Transaction Coordinator (DTC)

Answer: D

QUESTION 85
You need to create the object used by the parameter of usp_UpdateEmployeeName. Which code segment should you use?

A. `CREATE XML SCHEMA COLLECTION EmployeesInfo`
B. `CREATE TABLE EmployeesInfo`
C. `CREATE TYPE EmployeesInfo AS Table`
D. `CREATE SCHEMA EmployeesInfo`

Answer: C
Explanation:
http://msdn.microsoft.com/en-us/library/bb510489.aspx

http://msdn.microsoft.com/en-us/library/ms175007.aspx
http://msdn.microsoft.com/en-us/library/ms175010.aspx
http://msdn.microsoft.com/en-us/library/ms174979.aspx
http://msdn.microsoft.com/en-us/library/ms189462.aspx
http://msdn.microsoft.com/en-us/library/ms176009.aspx

Case Study 5: Litware Inc. (Question 86 ~ Question 96)

General Overview

You are a database developer for a company named Litware, Inc. Litware has a main office in Miami. Litware has a job posting web application named WebApp1. WebApp1 uses a database named DB1. DB1 is hosted on a server named Server1. The database design of DB1 is shown in the exhibit.

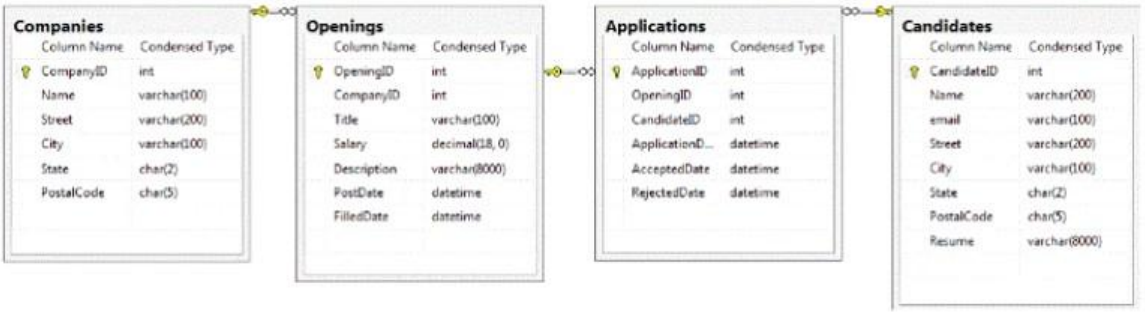

WebApp1 allows a user to log on as a job poster or a job seeker. Candidates can search for job openings based on keywords, apply to an opening, view their application, and load their resume in Microsoft Word format. Companies can add a job opening, view the list of candidates who applied to an opening, and mark an application as denied.

Users and Roles

DB1 has five database users named Company, CompanyWeb, Candidate, CandidateWeb, and Administrator. DB1 has three user-defined database roles. The roles are configured as shown in the following table:

Role name	Role member
Companies	Company Administrator CompanyWeb
Candidates	Candidate Administrator CandidateWeb
Administrators	Administrator

Keyword Search

The keyword searches for the job openings are performed by using the following stored procedure named usp_GetOpenings:

```
01  CREATE PROCEDURE usp_GetOpenings
02    @keyword varchar(max),
03    @minsalary decimal(18,0) = 0
04  AS
05  DECLARE @plural varchar(max);
06  DECLARE @ing varchar(max);
07  SET @plural = @keyword + 's';
08  SET @ing = @keyword + 'ing';
09  SELECT o.Title, o.Salary, c.Name, o.Description
10  FROM Openings o
11  INNER JOIN Companies c ON c.CompanyID = o.CompanyID
12  WHERE (o.Description LIKE '%'+@keyword+'%'
13    OR o.Description LIKE '%'+@plural+'%'
14    OR o.Description LIKE '%'+@ing+'%')
15    AND o.Salary >= @minsalary;
```

Opening Update

Updates to the Openings table are performed by using the following stored procedure named usp_UpdateOpening:

```
01  CREATE PROCEDURE usp_UpdateOpening
02    @openingID int,
03    @title varchar(100),
04    @salary decimal(18,0),
05    @description varchar(8000)
06  AS
07  UPDATE Openings
08  SET Title = @title,
09    Salary = @salary,
10    Description = @description
11  WHERE OpeningID = @openingID;
```

The Candidates role and the Companies role are assigned the GRANT EXECUTE permission for usp_UpdateOpening.

Problems and Reported Issues

Concurrency Problems

You discover that deadlocks frequently occur. You identify that a stored procedure named usp_AcceptCandidate and a stored procedure named usp_UpdateCandidate generate deadlocks. The following is the code for usp_AcceptCandidate:

```
01  CREATE PROCEDURE usp_AcceptCandidate
02    @applicationID int
03  AS
04    DECLARE @date datetime;
05    SET @date = GETDATE();
06    UPDATE Applications
07      SET AcceptedDate = @date
08      WHERE ApplicationID = @applicationID;
09    SELECT Name, email
10      FROM Candidates c
11      INNER JOIN Applications a
12        ON a.CandidateID = c.CandidateID
13      WHERE a.AcceptedDate IS NOT NULL;
```

Salary Query Issues

Users report that when they perform a search for job openings without specifying a minimum salary, only job openings that specify a minimum salary are displayed.

Log File Growth Issues

The current log file for DB1 grows constantly. The log file fails to shrink even when the daily SQL Server Agent Shrink Database task runs.

Performance Issues

You discover that a stored procedure named usp_ExportOpenings takes a long time to run and executes a table scan when it runs. You also discover that the usp_GetOpenings stored procedure takes a long time to run and that the non-clustered index on the Description column is not being used.

Page Split Issues

On DB1, many page splits per second spike every few minutes.

Requirements

Security and Application Requirements

Litware identifies the following security and application requirements:

- Only the Administrator, Company, and CompanyWeb database users must be able to execute the usp_UpdateOpening stored procedure.
- Changes made to the database must not affect WebApp1.

Locking Requirements

Litware identifies the following locking requirements:

- The usp_GetOpenings stored procedure must not be blocked by the usp_UpdateOpening stored procedure.
- If a row is locked in the Openings table, usp_GetOpenings must retrieve the latest version of the row, even if the row was not committed yet.

Integration Requirements

Litware exports its job openings to an external company as XML data. The XML data uses the following format:

```
<Opening title="web programmer" salary="75000">
  This is the description of the opening
</Opening>
```

A stored procedure named usp_ExportOpenings will be used to generate the XML data. The following is the code for usp_ExportOpenings:

```
01  CREATE PROCEDURE usp_ExportOpenings
02    @lastPost datetime
03  AS
04  SELECT Description
05    , Title
06    , Salary
07  FROM Openings
08  WHERE PostDate > @lastPost
09    AND FilledDate IS NULL
```

The stored procedure will be executed by a SQL Server Integration Services (SSIS) package named Package1. The XML data will be written to a secured folder named Folder1. Only a dedicated Active Directory account named Account1 is assigned the permissions to read from or write to Folder1.

Refactoring Requirements

Litware identifies the following refactoring requirements:

- New code must be written by reusing the following query:

```
01  SELECT Title, Salary, Description
02  FROM Openings
03  WHERE Salary >= @minsalary
04    AND FilledData IS NULL
```

- The results from the query must be able to be joined to other queries.

Upload Requirements

Litware requires users to upload their job experience in a Word file by using WebApp1. WebApp1 will send the Word file to DB1 as a stream of bytes. DB1 will then convert the Word file to text

before the contents of the Word file is saved to the Candidates table. A database developer creates an assembly named Conversions that contains the following:
- A class named Convert in the SqlConversions namespace.
- A method named ConvertToText in the Convert class that converts Word files to text.

The ConvertToText method accepts a stream of bytes and returns text. The method is used in the following stored procedure:

```
01  CREATE PROCEDURE usp_UpdateCandidate
02    @candidateID int,
03    @wordResume varbinary(max)
04  AS
05  DECLARE @textResume varchar(8000);
06  SET @textResume = ConvertToText(@wordResume);
07  UPDATE Candidates SET Resume = @textResume
08    WHERE CandidateID = @candidateID;
09  SELECT OpeningID, ApplicationDate
10  FROM Applications
11  WHERE CandidateID = @candidateID;
```

Job Application Requirements
A candidate can only apply to each job opening once.
Data Recovery Requirements
All changes to the database are performed by using stored procedures. WebAppl generates a unique transaction ID for every stored procedure call that the application makes to the database. If a server fails, you must be able to restore data to a specified transaction.

QUESTION 86
You need to create a script that automates the export of the XML data. The script must meet the integration requirements. What should you include in the script?

A. The CREATE SERVER ROLE command and the sp_add_proxy, sp_add_job, sp_add_jobstep/ and sp_grant_proxy_to_subsystem system stored procedures.
B. The CREATE CREDENTIAL command and the sp_reassign_proxy, sp_add_job, sp_add_jobstep, and sp_grant_login_to_proxy system stored procedures.
C. The CREATE CREDENTIAL command and the sp_add_proxy, sp_add_job, sp_add_jobstep, and sp_grant_proxy_to_subsystem system stored procedures.
D. The CREATE SERVER ROLE command and the sp_reassign_proxy, sp_add_job, sp_add_jobstep, and sp_grant_login_to_proxy system stored procedures.

Answer: C

QUESTION 87
You need to resolve the performance issues of the usp_ExportOpenings stored procedure. The solution must minimize the amount of hard disk space used. Which statement should you execute on DB1?

A. `EXEC sp_dboption 'DB1', 'auto create statistics', 'TRUE';`
B.
```
CREATE INDEX IX_Exp_Openings ON Openings(PostDate, FilledDate)
INCLUDE (Description, Title, Salary);
```
C.
```
CREATE INDEX IX_Exp_Openings ON Openings(PostDate)
INCLUDE (Description, Title, Salary)
WHERE FilledDate IS NULL;
```
D. `EXEC sp_recompile 'usp_ExportOpenings';`

Answer: C

QUESTION 88
You need to implement a solution that meets the job application requirements. What should you do?

A. Create a one-to-one relationship between the Candidates table and the Applications table.
B. Create a one-to-one relationship between the Openings table and the Applications table.
C. Add a UNIQUE constraint to the Applications table on the OpeningID column and the CandidateID column.
D. Add a UNIQUE constraint to the Applications table on the ApplicationID column and CandidateID column.

Answer: C

QUESTION 89
You need to identify the cause of the page split issues. Which SQL Server feature should you use?

A. DBCC TRACEOFF
B. DBCC REINDEX
C. Extended Events
D. SQL Server Profiler

Answer: C

QUESTION 90
You need to design a solution that meets the refactoring requirements. Which type of object should you include in the solution?

A. A table-valued function
B. A distributed view
C. An aggregate function
D. An indexed view

Answer: A

QUESTION 91
You need to recommend a solution that meets the concurrency problems. What should you include in the recommendation?

A. Modify the stored procedures to use the SERIALIZABLE isolation level.
B. Modify the order in which usp_UpdateCandidate accesses the Applications table and the Candidates table.
C. Modify the stored procedures to use the REPEATABLE READ isolation level.
D. Modify the order in which usp_AcceptCandidate accesses the Applications table and the Candidates table.

Answer: B

QUESTION 92
You need to implement a change to usp_ExportOpenings that meets the integration requirements. What should you modify in usp_ExportOpenings? (Each correct answer presents part of the solution. Choose all that apply?)

A. To the end of line 04, add [Opening].
B. To the end of line 05, add [Opening! title].
C. To line 10, add FOR XML RAW.

D. To line 10, add FOR XMLEXPLICIT.
E. To line 10, add FOR XML AUTO.
F. To the end of line 04, add [Opening!ELEMENT].
G. To the end of line 06, add [Opening!salary!ELEMENT].
H. To the end of line 05, add [Opening!title!ELEMENT].
I. To the end of line 06, add [Opening! salary].

Answer: ABEI
Explanation:
E: The AUTO mode generates nesting in the resulting XML by using heuristics based on the way the SELECT statement is specified. You have minimal control over the shape of the XML generated. The nested FOR XML queries can be written to generate XML hierarchy beyond the XML shape that is generated by AUTO mode heuristics.

QUESTION 93
You need to implement a solution that meets the locking requirements. Which line of code should you modify?

A.
```
Change line 07 in usp_UpdateOpening to:
UPDATE Openings WITH (UPDLOCK)
```
B.
```
Change line 09 in usp_GetOpenings to:
FROM Openings o (ROWLOCK)
```
C.
```
Change line 07 in usp_UpdateOpening to:
UPDATE Openings WITH (READPA5T)
```
D.
```
Change line 09 in usp_GetOpenings to:
FROM Openings o (NOLOCK)
```

Answer: D

QUESTION 94
You need to implement a solution that addresses the upload requirements. Which code segment should you use to implement the Conversions assembly?

A.
```
CREATE FUNCTION ConvertToText (@wordResume varbinary(max))
RETURNS varchar(8000)
AS EXTERNAL NAME Conversions.SqlConversions.ConvertToText;
```

B.
```
CREATE FUNCTION ConvertToText (@wordResume varchar(8000))
RETURNS varbinary(max)
AS EXTERNAL NAME SqlConversions.Conversions.ConvertToText;
```

C.
```
CREATE PROCEDURE ConvertToText (@wordResume varbinary(max))
AS EXTERNAL NAME Conversions.SqlConversions.ConvertToText;
```

D.
```
CREATE PROCEDURE ConvertToText (@wordResume varchar(8000))
AS EXTERNAL NAME SqlConversions.Conversions.ConvertToText;
```

A. Option A
B. Option B
C. Option C

D. Option D

Answer: A

QUESTION 95
You need to implement a solution that meets the security requirements. Which statement should you execute?

A.
```
REVOKE EXEC ON usp_UpdateOpening FROM Candidates;
```

B.
```
DENY EXEC ON usp_UpdateOpening TO Candidates;
```

C.
```
ALTER PROCEDURE usp_UpdateOpening
@openingIDint,
@titlevarchar(100),
@salarydecimal(18,0),
@descriptionvarchar(8000)
WITH EXECUTE AS Administrator
AS
...
```

D.
```
ALTER PROCEDURE usp_UpdateOpening
@openingIDint,
@titlevarchar(100),
@salarydecimal(18,0),
@descriptionvarchar(8000)
WITH EXECUTE AS Company
AS
...
```

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

Answer: A

QUESTION 96
You need to implement a solution that resolves the salary query issue. Which statement should you execute on DB1?

A.
```
UPDATE Openings SET Salary=0 WHERE Salary IS NULL;
GO
ALTER TABLE Openings
WITH NOCHECK
MODIFY COLUMN Salary NOT NULL;
GO
ALTER TABLE Openings
WITH NOCHECK
ADD CONSTRAINT DF_SALARY
DEFAULT 0 FOR Salary;
GO
```

B.
```
ALTER TABLE Openings
WITH NOCHECK
ADD CONSTRAINT DF_SALARY
DEFAULT 0 FOR Salary;
GO
ALTER TABLE Openings
WITH NOCHECK
MODIFY COLUMN Salary NULL;
GO
UPDATE Openings SET Salary=0 WHERE Salary IS NULL;
GO
```

C.
```
UPDATE Openings SET Salary=0 WHERE Salary IS NULL;
GO
ALTER TABLE Openings
WITH NOCHECK
ADD CONSTRAINT CT_SALARY
CHECK (Salary>=0);
GO
ALTER TABLE Openings
WITH NOCHECK
MODIFY COLUMN Salary NOT NULL;
GO
```

D.
```
ALTER TABLE Openings
WITH NOCHECK
ADD CONSTRAINT CT_SALARY
CHECK (Salary>=0);
GO
ALTER TABLE Openings
WITH NOCHECK
MODIFY COLUMN Salary NOT NULL;
GO
UPDATE Openings SET Salary=0 WHERE Salary IS NULL;
GO
```

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

Answer: A

Case Study 6: Coho Winery (Question 97 ~ Question 110)

Overview

You are a database developer for a company named Coho Winery. Coho Winery has an office in

London. Coho Winery has an application that is used to process purchase orders from customers and retailers in 10 different countries. The application uses a web front end to process orders from the Internet. The web front end adds orders to a database named Sales. The Sales database is managed by a server named Server1. An empty copy of the Sales database is created on a server named Server2 in the London office. The database will store sales data for customers in Europe. A new version of the application is being developed. In the new version, orders will be placed either by using the existing web front end or by loading an XML file. Once a week, you receive two files that contain the purchase orders and the order details of orders from offshore facilities. You run the usp_ImportOders stored procedure and the usp_ImportOrderDetails stored procedure to copy the offshore facility orders to the Sales database. The Sales database contains a table named Orders that has more than 20 million rows.

Database

Definitions

Database and Tables

The following scripts are used to create the database and its tables:

```
01 CREATE DATABASE Sales;
02 GO
03 USE Sales;
04 GO
05 CREATE TABLE Products
06 (
07   ProductID int IDENTITY(1,1) NOT NULL,
08   Name nvarchar(100) NOT NULL,
09   UnitPrice decimal(18,2) NOT NULL,
10   Discontinued bit NOT NULL DEFAULT 0,
11   CONSTRAINT PK_Products PRIMARY KEY (ProductID)
12 );
13 GO
14
15 CREATE TABLE Customers
16 (
17   CustomerID int IDENTITY(1,1) NOT NULL,
18   Name nvarchar(200) NOT NULL,
19   Email nvarchar(200) NOT NULL,
20   Phone nvarchar(10) NOT NULL,
21   Address1 nvarchar(200) NOT NULL,
22   Address2 nvarchar(200) NULL,
23   City nvarchar(200) NOT NULL,
24   State char(2) NOT NULL,
25   ZIP char(5) NOT NULL,
26   CONSTRAINT PK_Customers PRIMARY KEY (CustomerID)
27 );
28 GO
29
30 CREATE TABLE Orders
31 (
32   OrderID int IDENTITY(1,1) NOT NULL,
33   CustomerID int NOT NULL,
34   OrderDate datetime NOT NULL DEFAULT GETDATE(),
35   DeliveryDate datetime NOT NULL,
36   ShipDate datetime NULL,
37   Amount decimal(18,2) NOT NULL,
38   CONSTRAINT PK_Orders PRIMARY KEY(OrderID)
39 );
40 GO
41
42 ALTER TABLE Orders
43   ADD CONSTRAINT FK_Orders_Customers
44   FOREIGN KEY(CustomerID)
45   REFERENCES Customers(CustomerID);
46 GO
47
48 CREATE TABLE OrderDetails
49 (
50   OrderID int NOT NULL,
51   LineItem int NOT NULL,
52   ProductID int NOT NULL,
53   Quantity int NOT NULL,
54   UnitPrice decimal(18,2) NOT NULL,
55   Total decimal(18,2) NOT NULL,
56   Discount decimal(18,2) NULL,
57   CONSTRAINT PK_OrderDetails PRIMARY KEY(OrderID,
58 );
59 GO
60
61 ALTER TABLE OrderDetails
62   ADD CONSTRAINT FK_OrderDetails_Orders
63   FOREIGN KEY(OrderID)
64   REFERENCES Orders(OrderID);
65 GO
66
67 ALTER TABLE OrderDetails
68   ADD CONSTRAINT FK_OrderDetails_Products
69   FOREIGN KEY(ProductID)
70   REFERENCES Products(ProductID);
71 GO
```

Stored Procedures

The following are the definitions of the stored procedures used in the database:

```
51 AS
52   SELECT OrderID, DeliveryDate, Amount
53   FROM Orders
54   WHERE ShipDate IS NULL
55   ORDER BY DeliveryDate;
56 GO
57
58 CREATE PROCEDURE usp_GetOrdersByProduct
59   @productID int
60
61 AS
62 SELECT OrderID, LineItem, Quantity,
63   UnitPrice, Total, Discount
64 FROM OrderDetails
65
66 WHERE ProductID = @productID;
67 GO
68
69 CREATE PROCEDURE usp_ImportOrders
70 AS
71 BULK INSERT Orders
72   FROM 'f:\orders\orders.tbl'
73   WITH
74     (
75       FIELDTERMINATOR =' |',
76       ROWTERMINATOR =' |\n'
77     );
78 GO
79 CREATE PROCEDURE usp_ImportOrderDetails
80   @firstRow int
81 AS
82 BULK INSERT OrderDetails
83   FROM 'f:\orders\details.tbl'
84   WITH
85     (
86
87       FIRSTROW = @firstRow,
88       FIELDTERMINATOR =' |',
89       ROWTERMINATOR =' |\n'
90     );
91 GO
```

Indexes

The following indexes are part of the Sales database:

```
01 CREATE INDEX IX_Orders_ShipDate
02   ON Orders(Shipdate)
03
04   INCLUDE (CustomerID, OrderDate, Amount);
05 GO
```

Data Import

The XML files will contain the list of items in each order. Each retailer will have its own XML schema and will be able to use different types of encoding. Each XML schema will use a default namespace.

The default namespaces are not guaranteed to be unique. For testing purposes, you receive an XSD file from a customer. For testing purposes, you also create an XML schema collection named ValidateOrder. ValidateOrder contains schemas for all of the retailers. The new version of the application must validate the XML file, parse the data, and store the parsed data along with the original XML file in the database. The original XML file must be stored without losing any data.

Reported Issues

Performance Issues

You notice the following for the usp_GetOrdersAndItems stored procedure:

- The stored procedure takes a long time to complete.
- Less than two percent of the rows in the Orders table are retrieved by usp_GetOrdersAndItems.
- A full table scan runs when the stored procedure executes.
- The amount of disk space used and the amount of time required to insert data are very high.

You notice that the usp_GetOrdersByProduct stored procedure uses a table scan when the stored procedure is executed.

Page Split Issues

Updates to the Orders table cause excessive page splits on the IX_Orders_ShipDate index.

Requirements

Site Requirements

Users located in North America must be able to view sales data for customers in North America and Europe in a single report. The solution must minimize the amount of traffic over the WAN link between the offices.

Bulk Insert Requirements

The usp_ImportOrderDetails stored procedure takes more than 10 minutes to complete. The stored procedure runs daily. If the stored procedure fails, you must ensure that the stored procedure restarts from the last successful set of rows.

Index Monitoring Requirements

The usage of indexes in the Sales database must be monitored continuously. Monitored data must be maintained if a server restarts. The monitoring solution must minimize the usage of memory resources and processing resources.

QUESTION 97

You need to implement a solution that meets the site requirements. What should you implement?

A. A non-indexed view on Server2
B. A distributed view on Server2
C. A distributed view on Server1
D. A non-indexed view on Server1

Answer: C

QUESTION 98

You need to ensure that usp_AddXMLOrder can be used to validate the XML input from the retailers. Which parameters should you add to usp_AddXMLOrder on line 04 and line 05? (Each correct answer presents part of the solution. Choose all that apply.)

A. @schema varbinary(100).
B. @items varchar(max).
C. @schema sysname.
D. @items varbinary(max).
E. @items xml.
F. @schema xml.

Answer: CE

QUESTION 99
You need to implement a solution that addresses the performance issues of the usp_GetOrdersByProduct stored procedure. Which statement should you execute?

A.
```
CREATE INDEX IX_OrderDetails_ByProduct
ON OrderDetails (ProductID)
INCLUDE (LineItem, Quantity, UnitPrice, Discount)
```

B.
```
CREATE INDEX IX_OrderDetails_ByProduct
ON OrderDetails (ProductID)
INCLUDE (OrderID, LineItem, UnitPrice, Total, Discount)
```

C.
```
CREATE INDEX IX_OrderDetails_ByProduct
ON OrderDetails (ProductID)
```

D.
```
CREATE INDEX IX_OrderDetails_ByProduct
ON OrderDetails (ProductID)
INCLUDE (LineItem, Quantity, UnitPrice, Total, Discount)
```

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

Answer: C

QUESTION 100
You plan to create a stored procedure that inserts data from an XML file to the OrderDetails table. The following is the signature of the stored procedure:

```
CREATE PROCEDURE usp_InsertItems
  @items XML (ValidateOrder)
```

The following is the XSD file used to create the ValidateOrder schema collection:

```
<?xml version="1.0" encoding="UTF-16"?>
<xsd:schema
   xmlns:xsd="http://www.w3.org/2001/XMLSchema" >
<xsd:element name="root">
  <xsd:complexType mixed="true">
    <xsd:sequence>
      <xsd:element name="Product"
                   minOccurs="1" maxOccurs="unbounded">
        <xsd:complexType mixed="true">
          <xsd:sequence>
            <xsd:element name="UnitPrice" type="xsd:decimal"
                         minOccurs="1" maxOccurs="1" />
            <xsd:element name="Quantity" type="xsd:integer"
                         minOccurs="1" maxOccurs="1" />
          </xsd:sequence>
          <xsd:attribute name="lineItem"
                         type="xsd:integer" use="required"/>
          <xsd:attribute name="productID"
                         type="xsd:integer" use="required"/>
        </xsd:complexType>
      </xsd:sequence>
      <xsd:attribute name="numberItems"
                     type="xsd:integer" use="required"/>
    </xsd:complexType>
  </xsd:element>
  </xsd:schema>
```

You develop a code segment that retrieves the number of items and loops through each item. Each time the loop runs, a variable named @itemNumber is incremented. You need to develop a code segment that retrieves the product ID of each item number in the loop. Which code segment should you develop?

A. `SET @productID = @items.value('/Root/Product/productID', int)`
B. `SET @productID = @items.value('/Root/Product['+ @itemNumber+ ']/@productID', int)`
C. `SET @productID = @items.value('/Root/Product/@productID', int)`
D. `SET @productID = @items.value('/Root/Product['+ @itemNumber+ ']/productID', int)`

Answer: D

QUESTION 101
You need to ensure that a new execution plan is used by usp_GetOrdersByProduct each time the stored procedure runs. What should you do?

A. Add WITH (FORCESEEK) to line 07 in usp_GetOrdersByProduct.
B. Execute sp_recompile 'usp_GetOrdersByProduct'.
C. Execute sp_help 'usp_GetOrdersByProduct'.
D. Add WITH RECOMPILE to line 03 in usp_GetOrdersByProduct.

Answer: D
Explanation:
http://msdn.microsoft.com/en-us/librAry/ms190439(v=sql.90).aspx

QUESTION 102

You need to implement a solution that addresses the index monitoring requirements. What should you do?

A. Schedule a SQL Server Agent job that saves data from the dynamic management views to a table in the database.
B. Create a SQL Server Audit that saves data to a log file, and then create a SQL Server Audit Specification that gathers data from the DATABASE_OPERATION group.
C. Create a performance monitor Data Collector Set (DCS) that monitors the SQL Server counters.
D. Schedule a SQL Server Profiler trace, and then save the trace data to a table in the database.

Answer: A

QUESTION 103
You need to implement a solution that addresses the page split issues. Which statement should you execute?

A.
```
ALTER INDEX IX_Orders_ShipDate ON Orders
REBUILD WITH (PAD_INDEX = OFF, DROP_EXISTING = ON);
```
B.
```
ALTER INDEX IX_Orders_ShipDate ON Orders
REBUILD WITH (FILLFACTOR=50, DROP_EXISTING = ON);
```
C.
```
ALTER INDEX IX_Orders_ShipDate ON Orders
REBUILD WITH (FILLFACTOR=0, DROP_EXISTING = ON);
```
D.
```
ALTER INDEX IX_Orders_ShipDate ON Orders
REBUILD WITH (PAD_INDEX=ON/ DROP_EXISTING = ON);
```

Answer: B

QUESTION 104
You need to implement a solution that solves the performance issues of usp_GetOrdersAndItems. Which statements should you execute?

A.
```
CREATE INDEX IX_Orders_Active
ON Orders(ShipDate, DeliveryDate, Amount)
```
B.
```
CREATE INDEX IX_Orders_Active
ON Orders(DeliveryDate)
INCLUDE(Amount) WHERE ShipDate IS NXXL
```
C.
```
CREATE INDEX IX_Orders_Active
ON Orders(DeliveryDate, Amount)
WHERE ShipDate IS NULL
```
D.
```
CREATE INDEX IX_Orders_Active
ON Orders(ShipDate, DeliveryDate)
INCLUDE( Amount)
```

Answer: B

QUESTION 105
You need to modify usp_GetOrdersAndItems to ensure that an order is NOT retrieved by usp_GetOrdersAndItems while the order is being updated. What should you add to usp_GetOrdersAndItems?

A. Add (READPAST) to the end of line 06.
B. Add SET TRANSACTION ISOLATION LEVEL SNAPSHOT to line 03.
C. Add SET TRANSACTION ISOLATION LEVEL SERIALIZABLE to line 03.

D. Add (UPDLOCK) to the end of line 06.

Answer: A

QUESTION 106
You need to implement a solution that addresses the bulk insert requirements. What should you add to line 08 in usp_ImportOrderDetails?

A. LASTROW=0.
B. BATCHSIE=0.
C. BATCHSIZE=1000.
D. LASTROW=1000.

Answer: C

QUESTION 107
You discover that the usp_GetOrdersAndItems stored procedures takes a long time to complete while usp_AddOrder or usp_AddXMLOrder run. You need to ensure that usp_GetOrdersAndItems completes as quickly as possible. What should you do? (Each correct answer presents part of the solution. Choose all that apply.)

A. Set the isolation level of the usp_GetOrdersAndItems stored procedure to SERIALIZABLE.
B. Execute the ALTER DATABASE Sales SET ALLOW_SNAPSHOT_ISOLATION ON statement.
C. Set the isolation level of the usp_AddOrder stored procedure to SERIALIZABLE.
D. Set the isolation level of the usp_GetOrdersAndItems stored procedure to SNAPSHOT.
E. Set the isolation level of the usp_AddOrder stored procedure to SNAPSHOT.
F. Execute the ALTER DATABASE Sales SET ALLOWSNAPSHOTISOLATION OFF statement.

Answer: BD

QUESTION 108
You need to modify the Orders table to store the XML data used by the retailers. Which statement should you execute?

A.
```
ALTER Orders
ADD originalOrder XML (ValidateOrder);
```
B.
```
ALTER Orders
ADD originalOrder XML;
```
C.
```
ALTER Orders
ADD originalOrdervarchar(max);
```
D.
```
ALTER Orders
ADD originalOrdervarbinary(max);
```

Answer: D

QUESTION 109
You need to modify usp.GetOrdersAndItems to ensure that an order is NOT retrieved by usp_GetOrdersAndItems while the order is being updated. What should you add to usp.GetOrdersAndItems?

A. Add WITH (NOLOCK) to the end of line 47.
B. Add SET TRANSACTION ISOLATION LEVEL READ COMMITTED to line 44.

C. Add SET TRANSACTION ISOLATION LEVEL READ UNCOMMITTED to line 44.
D. Add WITH (READPAST) to the end of line 47.

Answer: B

QUESTION 110
You need to ensure that a new execution plan is used by usp_GetOrdersByProduct each time the stored procedure runs. What should you do?

A. Execute sp_help usp_GetOrdersByProduct.
B. Add WITH (FORCESEEK) to line 69 in usp.GetOrdersByProduct.
C. Add WITH RECOMPILE to line 64 in usp.GetOrdersByProduct.
D. Execute sp_recompile usp.GetOrdersByProduct.

Answer: B

QUESTION 111
Drag and Drop Question
You have a SQL Server 2012 database named Database1. Database1 has a data file named Database1_data.mdf and a transaction log named Database1jog.ldf. Database1_data.mdf is 1.5GB. Database1jog.ldf is 1.5 terabytes. A full backup of Database1 is performed every day. You need to reduce the size of the log file. The solution must ensure that you can perform transaction log backups in the future. Which code segment should you execute? To answer, move the appropriate code segments from the list of code segments to the answer area and arrange them in the correct order.

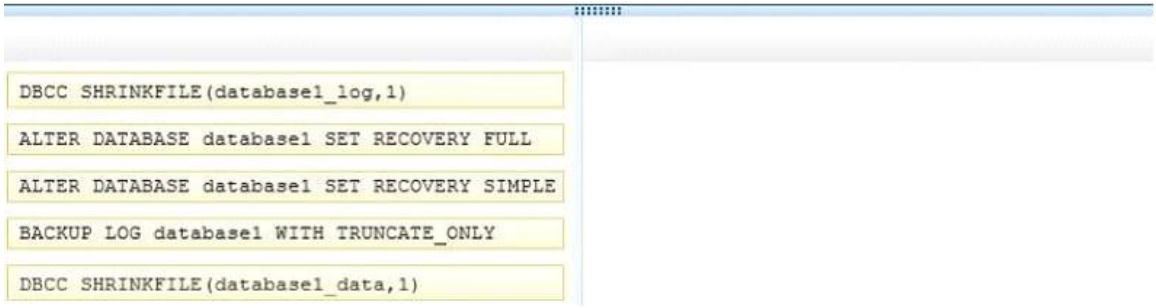

Answer:

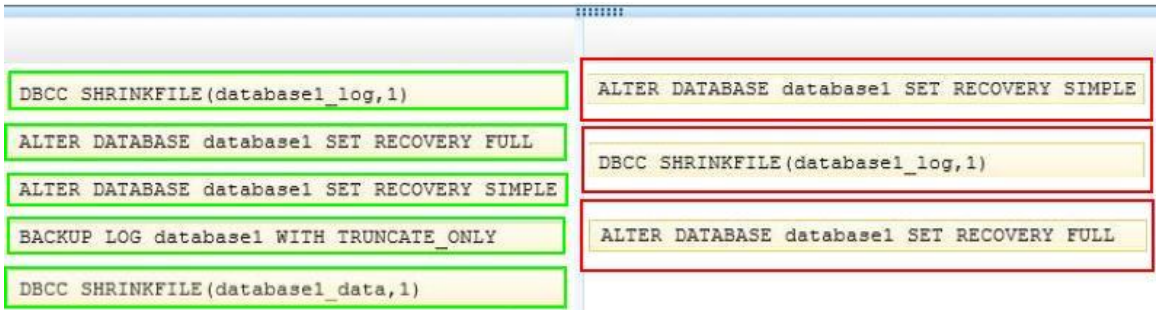

QUESTION 112
Drag and Drop Question
You plan to deploy SQL Server 2012. You must create two tables named Table1 and Table2 that will have the following specifications:
- Table1 will contain a date column named Column1 that will contain a null value approximately 80 percent of the time.

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

You need to recommend which options must be defined for the columns. The solution must minimize the storage requirements for the tables. Which options should you recommend? To answer, drag the appropriate options to the correct column in the answer area.

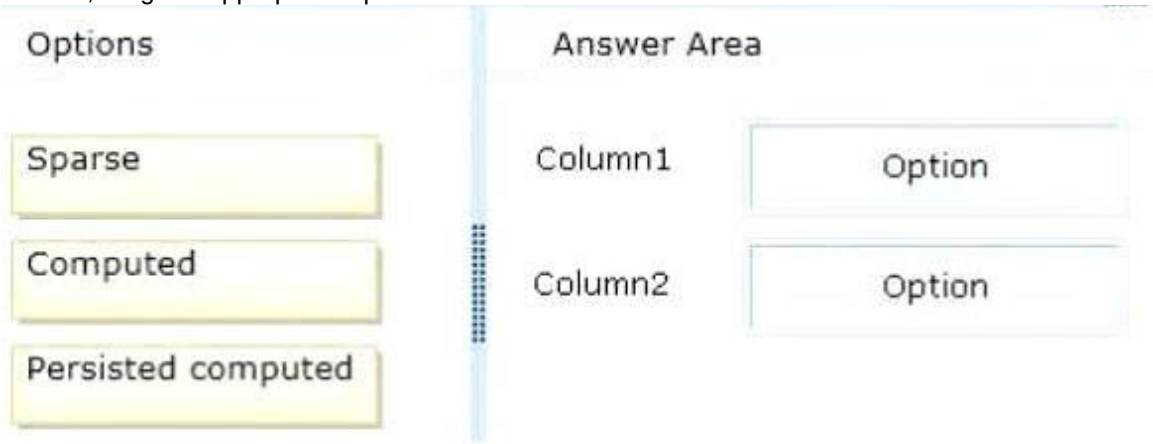

Answer:

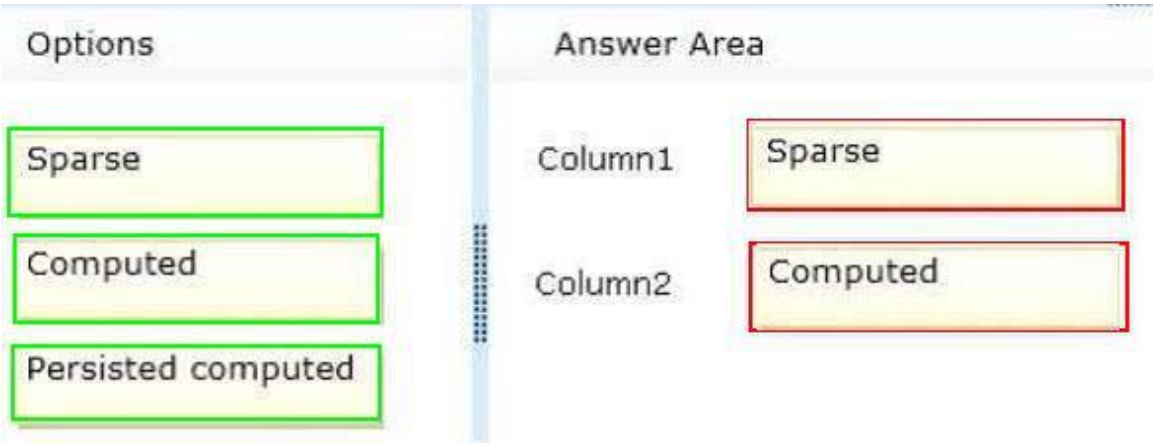

QUESTION 113

Drag and Drop Question

You are designing a database for a university. The database will contain two tables named Classes and StudentGrades that have the following specifications:

- Classes will store brochures in the XPS format.
- The brochures must be structured in folders and must be accessible by using UNC paths.
- StudentGrades must be backed up on a separate schedule than the rest of the database.

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

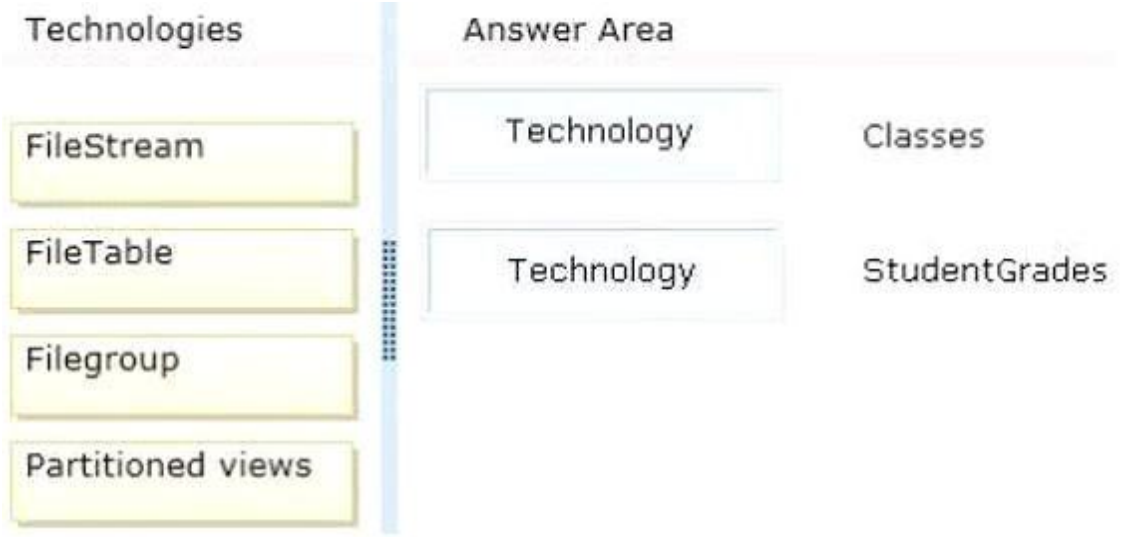

Answer:

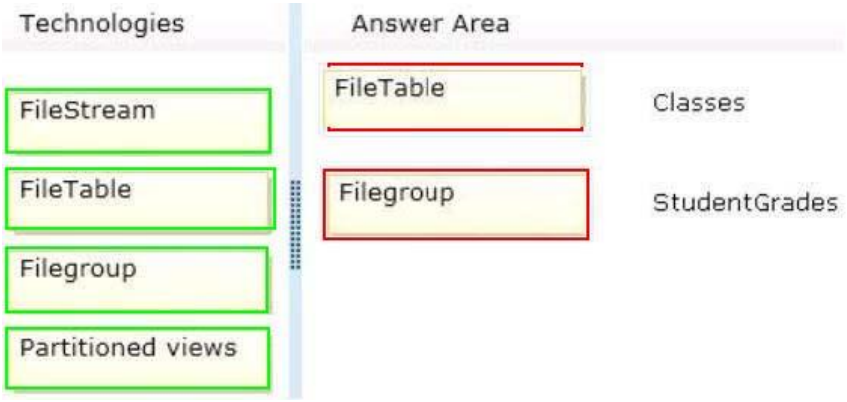

QUESTION 114

Drag and Drop Question

You have a SQL Azure database named Database1. You need to design the schema for a table named table1. Table1 will have less than one million rows. Table1 will contain the following information for each row:

Column	Description
ID	An incremental numeric value used to identify the row
Name	A string in English
Code	An alphanumeric code that has five characters
ModifiedDate	The date of the last modification

The solution must minimize the amount of space used to store each row. Which data types should you recommend for each column? To answer, drag the appropriate data type to the correct column in the answer area.

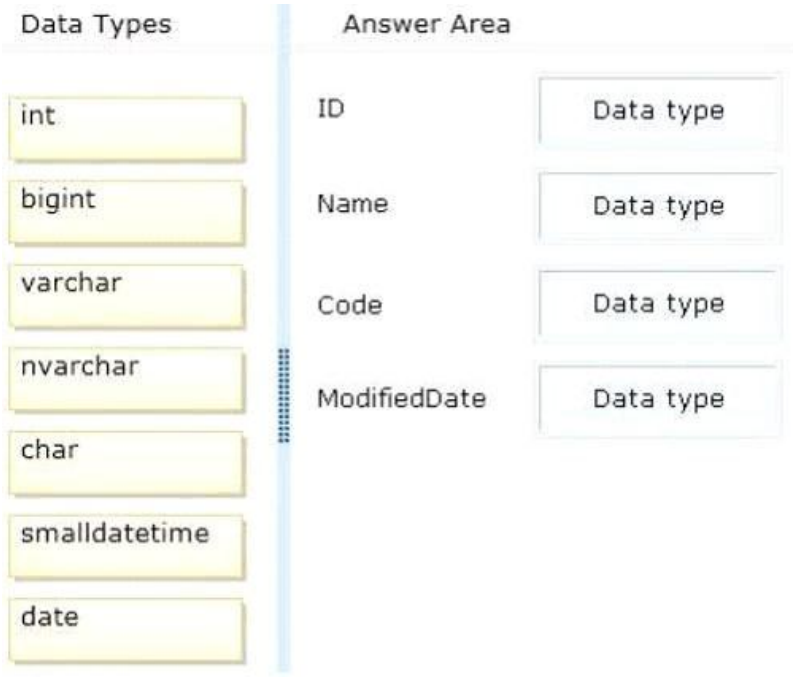

Answer:

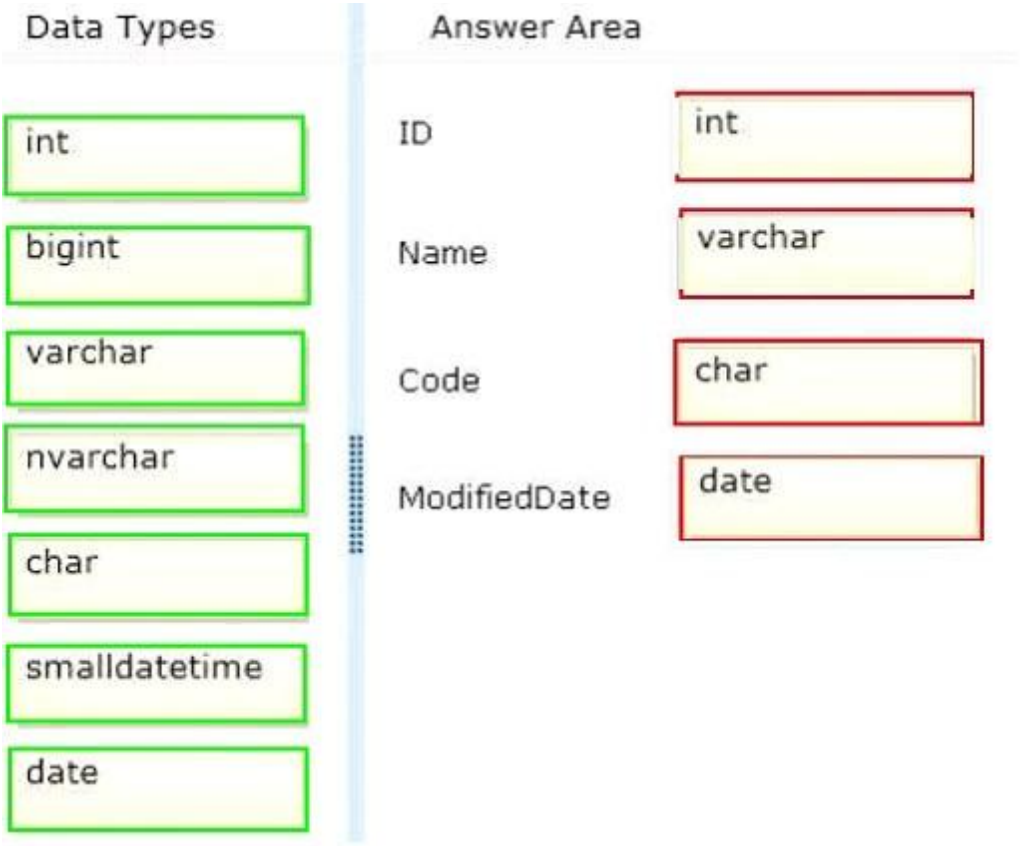

QUESTION 115

You have a database named database1. Database1 has two stored procedures named Proc1 and

Proc2 and a table named Table1. Table1 has millions of rows. Proc1 updates data in Table1. Proc2 reads data from Table1. You discover that when Proc1 is executed to update more than 4,000 rows, Proc2 is blocked. The block affects all rows, including those that are not being updated by Proc1. You need to ensure that when Proc1 is executing, Proc2 can access the data in Table1 that Proc1 is not updating. What should you change Proc1 to do? More than one answer choice may achieve the goal. Select the BEST answer.

A. Use the ROWLOCK table hint.
B. Wait for Proc2 to complete.
C. Update less than 4,000 rows simultaneously.
D. Use the PAGLOCK table hint.

Answer: C

QUESTION 116
You need to implement a solution that meets the data recovery requirements. You update each stored procedure to accept a parameter named @transactionID. What should you add next to the beginning of each stored procedure?

A. `SAVE TRANSACTION WITH MARK @transactionID`
B. `COMMIT TRANSACTION @transaction`
C. `BEGIN TRANSACTION WITH MARK @transactionID`
D. `ROLLBACK DISTRIBUTED TRANSACTION @transactionID`

Answer: C

QUESTION 117
You have a Microsoft SQL Azure database. You have the following stored procedure:

```
01 CREATE PROC up_employees
02   @ID int,
03   @Name nvarchar(50)
04 AS
05
06 SELECT Name AS OriginalName
07 FROM HR.Employees
08
09 WHERE ID = @ID;
10
11 UPDATE HR.Employees
12 SET Name = @Name
13
14 WHERE ID = @ID;
```

You discover that the stored procedure periodically fails to update HR.Employees. You need to ensure that HR.Employees is always updated when up_employees executes. The solution must minimize the amount of time required for the stored procedure to execute and the number of locks held. What should you do?

A. Add the following line of code to line 05:
`SET TRANSACTION ISOLATION LEVEL SERIALIZABLE`
B. Add the following line of code to line 05:
`SET TRANSACTION ISOLATION LEVEL SNAPSHOT`
C. Add the following line of code to line 13:

```
WITH (UPDLOCK)
```

D. Add the following line of code to line 08:
```
WITH (UPDLOCK)
```

Answer: D

QUESTION 118
You need to encapsulate a T-SQL script into a reusable user-defined object. The object must meet the following requirements:
- Permit insertions into a table variable.
- Support structured exception handling.
- Prevent changes to the definition of referenced objects.
- Support the use of the APPLY operator on the output of the object.

Which type of object should you use?

A. An inline table-valued function
B. A multi-statement table-valued function
C. A stored procedure
D. A scalar user-defined function

Answer: D

QUESTION 119
Your network contains a server named SQL1 that has SQL Server 2012 installed. SQL1 contains a database name DB1 and a table named Customers. You add an additional server named SQL2 that runs SQL Server 2012. You need to create a distributed partitioned view. The solution must minimize the amount of network traffic. What should you do? (Each correct answer presents part of the solution. Choose all that apply.)

A. Add SQL2 as a Distributor.
B. Add the Customers table to SQL2.
C. Add SQL2 as a linked server.
D. Create the view on SQL1.
E. Remove the Customers table from SQL1.
F. Create the view on SQL2.

Answer: BCDF

QUESTION 120
You plan to migrate an instance of SQL Server 2008 to a new installation of SQL Server 2012. You need to migrate alerts and e-mail notifications. Which system stored procedures should you use? (Each correct answer presents part of the solution. Choose all that apply.)

A. sp_syspolicy_create_job
B. sp_add_operator
C. sp_audit_write
D. sp_add_alert

Answer: BC
Explanation:
B: sp_add_operator
Creates an operator (notification recipient) for use with alerts and jobs.

C: sp_audit_write
Adds a user-defined audit event to the USER_DEFINED_AUDIT_GROUP. If USER_DEFINED_AUDIT_GROUP is not enabled, sp_audit_write is ignored.
Incorrect:
Not A: sp_syspolicy_create_job
No such standard stored procedure.
Not D: sp_add_alert
Creates an alert.

QUESTION 121
You have a table named Table1 that stores customer data. Each customer has a credit limit that can only be discovered by querying multiple tables. You need to ensure that the value of the credit limit is returned by executing a query on Table1. What should you create?

A. A trigger that uses a ranking function
B. A trigger that uses a table-valued function
C. A calculated column that uses a table-valued function
D. A calculated column that uses a scalar function

Answer: C
Explanation:
A table-valued parameter is scoped to the stored procedure, function, or dynamic Transact-SQL text, exactly like other parameters. Similarly, a variable of table type has scope like any other local variable that is created by using a DECLARE statement. You can declare table-valued variables within dynamic Transact-SQL statements and pass these variables as table-valued parameters to stored procedures and functions. Table-valued parameters offer more flexibility and in some cases better performance than temporary tables or other ways to pass a list of parameters.
Incorrect:
Not A: A scalar function would only be able to use other columns from the same table.

QUESTION 122
You plan to create a new table that has the following requirements:
- Uses a GUID data type as the primary key.
- Uses a clustered index as the primary key.
- Minimizes fragmentation.

You need to recommend which option to include in the CREATE statement. Which option should you include? More than one answer choice may achieve the goal. Select the BEST answer.

A. ROWGUID
B. SEQUENCE
C. NEWID
D. NEWSEQUENTIALID

Answer: D

QUESTION 123
You are creating a stored procedure named usp1. Usp1 will create a table that will be used during the execution of usp1. Only usp1 will be allowed to access the table. You need to write the code required to create the table for usp1. The solution must minimize the need to recompile the stored procedure. Which code segment should you use to create the table?

A. `CREATE TABLE oneTable`
B. `CREATE TAB1E ##oneTable`

C. `CREATE TABLE #oneTable`
D. `DECLARE goneTable TABLE`

Answer: D

QUESTION 124
You plan to modify a stored procedure to use temporary data. The stored procedure must meet the following requirements:
- Favor physical memory when physical memory is available.
- Be able to roll back changes to the temporary data.

You need to recommend which object to add to the stored procedure. Which T-SQL command should you recommend?

A. `CREATE TABLE ##Table...`
B. `CREATE TABLE Table...`
C. `CREATE VIEW Table...`
D. `CREATE PARTITION SCHEME Table...`
E. `DECLARE TABLE @ Table...`

Answer: A
Explanation:
Temporary Tables
You can create local and global temporary tables. Local temporary tables are visible only in the current session, and global temporary tables are visible to all sessions. Temporary tables cannot be partitioned. Prefix local temporary table names with single number sign (#table_name), and prefix global temporary table names with a double number sign (##table_name).

QUESTION 125
You have a database named database1. You execute the following code:

```
CREATE TABLE dbo.table1
(
  ID int IDENTITY(1,1) NOT NULL PRIMARY KEY,
  FirstName varchar(50) NOT NULL,
  LastName varchar(50) NOT NULL,
  EmailAddress varchar(200) NULL,
  Notes nvarchar(MAX) NULL,
  LastContactDate datetime NULL
)
```

You have the following query. (Line numbers are included for reference only.)

```
01 SELECT FirstName + ' ' + LastName AS Name
02 FROM dbo.table1
03 WHERE Notes LIKE '% call%' AND
04 LastContactDate >= '1/1/2010'
```

Users report that the query takes a long time to complete. You create a full-text index for the Notes column. You need to recommend changes to the query to reduce the amount of time it takes for the query to complete. Which code segment should you use to replace line 03?

A. `INNER JOIN FREETEXTTABLE(dbo.table1, notes, 'call') AS t2 ON dbo.table1.ID = t2.key WHERE`
B. `WHERE CONTAINS(notes, 'call*') AND`
C. `WHERE FREETEXT(notes, '%call%') AND`
D. `WHERE CONTAINS(notes, 'cal%') AND`

Answer: C

QUESTION 126
You have a database that is accessed by 300 concurrent users. You need to log all of the queries that become deadlocked. The solution must meet the following requirements:
- Provide a representation of the deadlock in XML format.
- Minimize the impact on the server.

What should you create?

A. A SQL Server Profiler trace
B. A script that enables trace flags
C. A SQL Server Agent job that retrieves information from the sys.dm_tran_session_transactions dynamic management views
D. A SQL Server Agent job that retrieves information from the sys.dm_tran_active_transactions dynamic management views

Answer: A
Explanation:
Analyze Deadlocks with SQL Server Profiler Use SQL Server Profiler to identify the cause of a deadlock. A deadlock occurs when there is a cyclic dependency between two or more threads, or processes, for some set of resources within SQL Server. Using SQL Server Profiler, you can create a trace that records, replays, and displays deadlock events for analysis. To trace deadlock events, add the Deadlock graph event class to a trace. This event class populates the TextData data column in the trace with XML data about the process and objects that are involved in the deadlock. SQL Server Profiler can extract the XML document to a deadlock XML (.xdl) file which you can view later in SQL Server Management Studio.

QUESTION 127
Drag and Drop Question
You have a SQL Server 2012 database named database1. Users report that queries that usually take less than one second to execute, take more than 30 seconds to execute. You need to view the server resource consumption when the queries are executed. What should you do? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Answer:

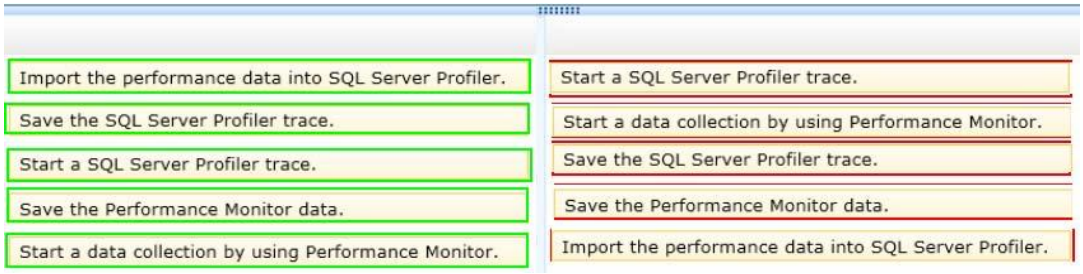

QUESTION 128

Drag and Drop Question

Your network contains a server named Server1 that runs SQL Server 2012. Server1 contains an instance named Instance1. Instance1 contains a database named ContentDatabase. ContentDatabase uses transaction log backups. The recovery model of ContentDatabase is set to FULL. You need to shrink the ContentDatabase_Log log file to 10MB. The solution must ensure that you can continue to back up the transaction log. Which three code segments should you execute? To answer, move the appropriate code segments from the list of code segments to the answer area and arrange them in the correct order.

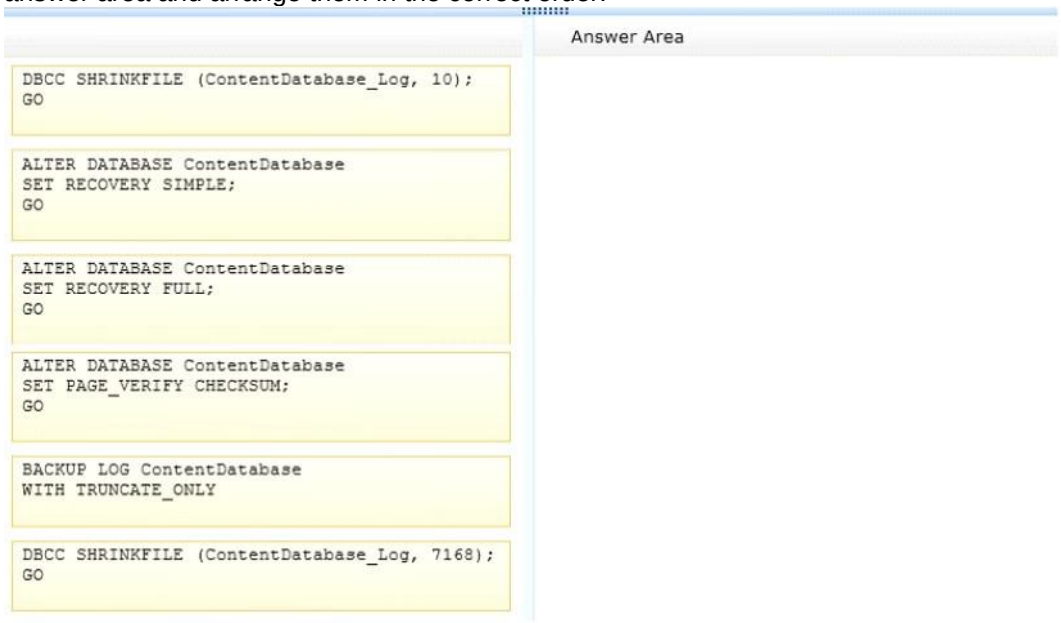

Answer:

QUESTION 129
You have a SQL Server 2012 database named DB1 that is accessed by 650 concurrent users. You need to log all of the queries to DB1 that become deadlocked. The solution must minimize the impact on the server. What should you create?

A. A SQL Server Agent job that retrieves information from the sys.dm_tran_active_transactions dynamic management views
B. A script that enables trace flags
C. A SQL Server Agent job that retrieves information from the sys.dm_tran_session_transactions dynamic management views
D. A SQL Server Profiler trace

Answer: B

QUESTION 130
You have a Microsoft SQL Azure database that contains a table named Employees.

```
CREATE TABLE HR.Employees
(
  id int primary key,
  name varchar(50)
)
```

You create a non-clustered index named EmployeeName on the name column. You write the following query to retrieve all of the employees that have a name that starts with the letters JOH:

```
SELECT * FROM HR.Employees
WHERE 'JOH' = LEFT(name,3)
```

You discover that the query performs a table scan. You need to ensure that the query uses EmployeeName. What should you do?

A. Recreate EmployeeName as a unique index
B. Recreate EmployeeName as a clustered index
C. Replace LEFT(name,3) = 'JOH' by using name like 'JOH%'
D. Replace LEFT(name,3) = 'JOH' by using substring(name, 1, 3) = 'JOH'

Answer: C

QUESTION 131
You have a database named DB1. You plan to configure change data capture on the existing tables in DB1. The database file settings for the DB1 are shown in the exhibit. You need to minimize disk contention caused by change data capture. What should you do? More than one answer choice may achieve the goal. Select the BEST answer.

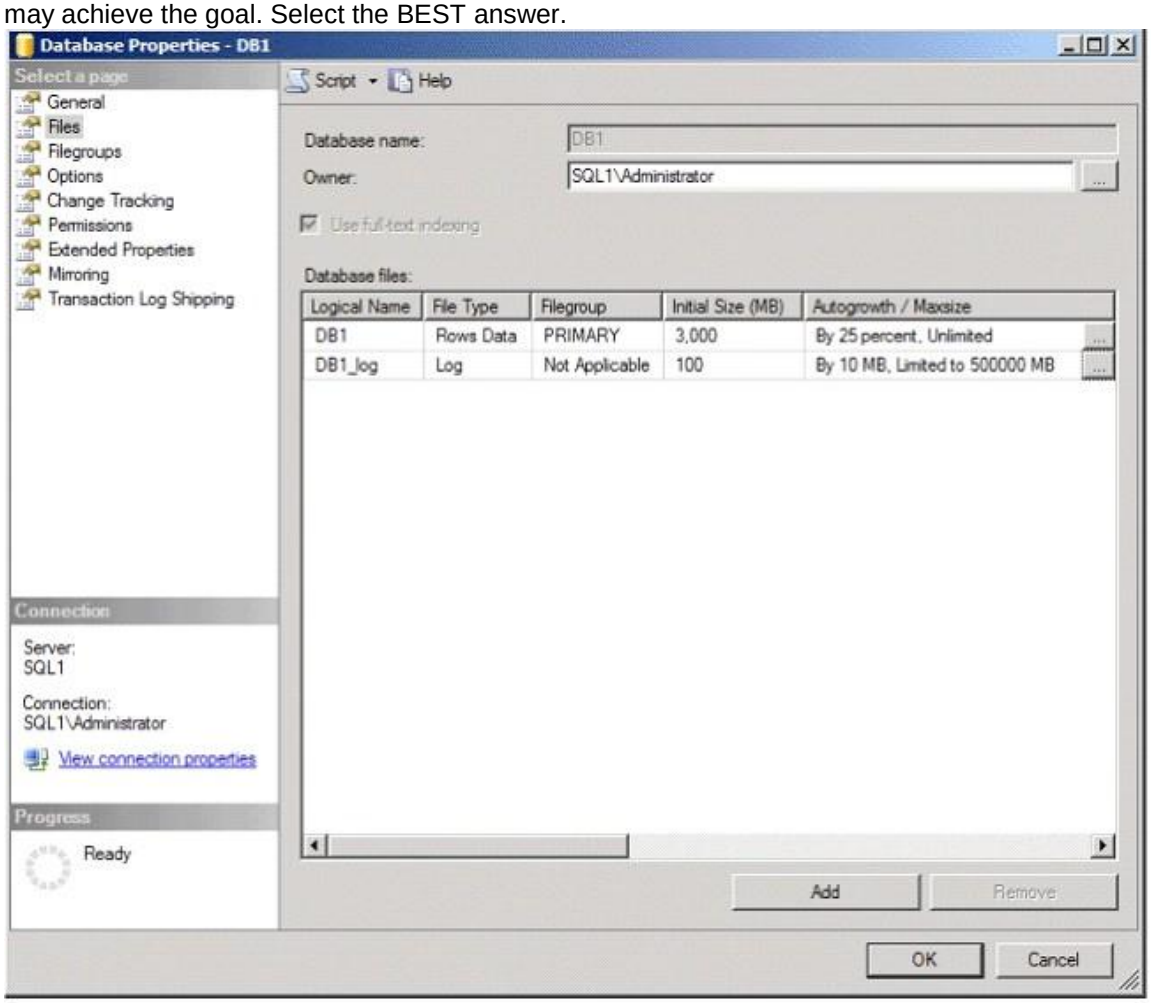

A. Increase the autogrowth value of the log file.
B. Configure change data capture to use to a secondary filegroup.
C. Increase the autogrowth value of the database file.
D. Set the database recovery model to simple.

Answer: B

QUESTION 132
You have a SQL Server 2012 database that contains a table named Users. The Users table contains usernames and passwords. You need to ensure that all new records have a password. Which code segment should you use? More than one answer choice may achieve the goal. Select the BEST answer.

A.
```
ALTER TABLE dbo.Users
  DROP Password;
GO
ALTER TABLE dbo.Users
  ADD Password varchar(30) NOT NULL;
GO
```

B.
```
ALTER TABLE dbo.Users
    ADD CONSTRAINT CK_Users_Password
    CHECK (Password IS NULL);
GO
```

C.
```
DROP TABLE dbo.Users;
GO
CREATE TABLE dbo.Users (
  CustID int PRIMARY KEY,
  Name varchar(30),
  Password varchar(30),
  CONSTRAINT CK_Users_Password
    CHECK (Password IS NOT NULL));
GO
```

D.
```
ALTER TABLE dbo.Users
  ADD CONSTRAINT CK_Users_Password
  CHECK (Password IS NOT NULL);
GO
```

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

Answer: D

QUESTION 133
Drag and Drop Question
You have a SQL Server 2012 database named DB1. DB1 contains four filegroups named FG1, FG2, FG3, and FG4. You execute the following code:

```
CREATE PARTITION FUNCTION PF1 (int)
AS RANGE LEFT FOR VALUES (20120331, 20120630, 20120930);
GO
CREATE PARTITION SCHEME PS1
AS PARTITION PF1
TO (FG1, FG2, FG3, FG4);
GO

CREATE TABLE dbo.Sales
(
  Date_key int NOT NULL,
  Customer_key int,
  Amount money
) ON PS1(Date_key);
GO
```

Two million rows are added to dbo.Sales. You need to move the data from the first partition to a new table named SalesHistory and, starting on December 31, 2012, repartition dbo.Sales to support new sales data for three months. Which code segment should you execute? To answer, move the appropriate code segments from the list of code segments to the answer area and arrange them in the correct order.

```
ALTER PARTITION FUNCTION PF1 MERGE RANGE
(20120331);
```

```
CREATE PARTITION SCHEME PS1
AS PARTITION PF1
TO (FG1, FG2, FG3, FG4);
```

```
DROP PARTITION SCHEME PS1;
```

```
CREATE PARTITION FUNCTION PF1 (int)
AS RANGE LEFT FOR VALUES
(20120630, 20120930, 20121231);
```

```
CREATE TABLE SalesHistory
(
  Date_key int NOT NULL,
  Customer_key int,
  Amount money
) ON PS1(Date_key);
```

```
ALTER TABLE SalesHistory SWITCH 1 TO Sales;
```

```
DROP PARTITION FUNCTION PF1
```

```
ALTER PARTITION FUNCTION PF1 SPLIT RANGE
(20121231);
```

```
CREATE TABLE SalesHistory
(
  Date_key int NOT NULL,
  Customer_key int,
  Amount money
) ON FG1;
```

```
ALTER TABLE Sales SWITCH 1 TO SalesHistory;
```

Answer:

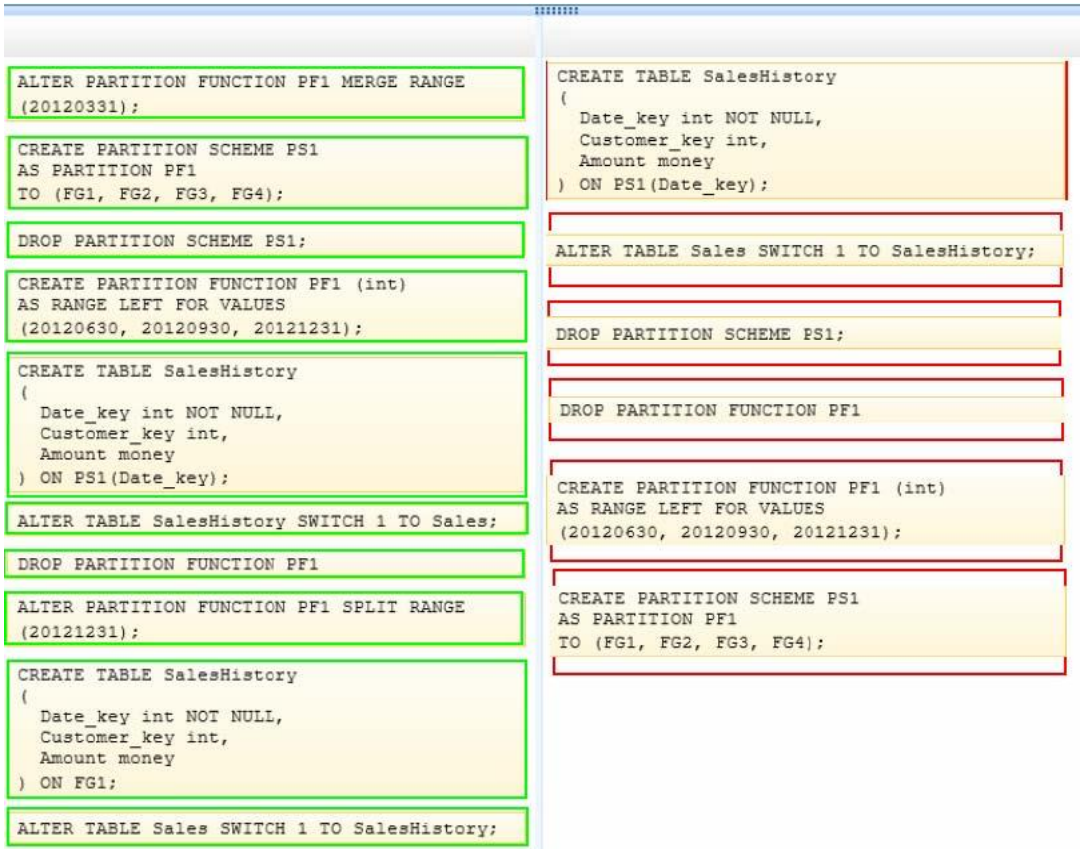

QUESTION 134

You have a SQL Server 2012 environment that contains two servers. The servers are configured as shown in the following table:

Server name	Database	Type
Server1	DB1	Principal
Server2	DB1	Mirror

After the failover is complete, a user receives the following error message when connecting to DB1 on Server2:

`"Msg 916, Level 14, State 1, Line 1 The server principal "Account1" is not able to access the database "DB1" under the current security context."`

You verify that there is a server login for Account1 on Server2. You need to ensure that Account1 can connect to DB1. What should you do? More than one answer choice may achieve the goal. Select the BEST answer.

A. Update the SID for Account1 on DB1.
B. Add Account1 to the db_datareader role.
C. Create a new database user on DB1.
D. Implement Windows authentication.

Answer: B

QUESTION 135

Drag and Drop Question

You plan to create a new table that will contain a column named Salary. Salary will contain highly

sensitive data. Salary must meet the following requirements:
- Contain numeric data.
- Contain only encrypted data that remains encrypted in memory.

You need to identify which encryption type and data type must be used for Salary. Which encryption type and data type should you identity? To answer, drag the appropriate encryption type and data type to the correct identifier in the answer area.

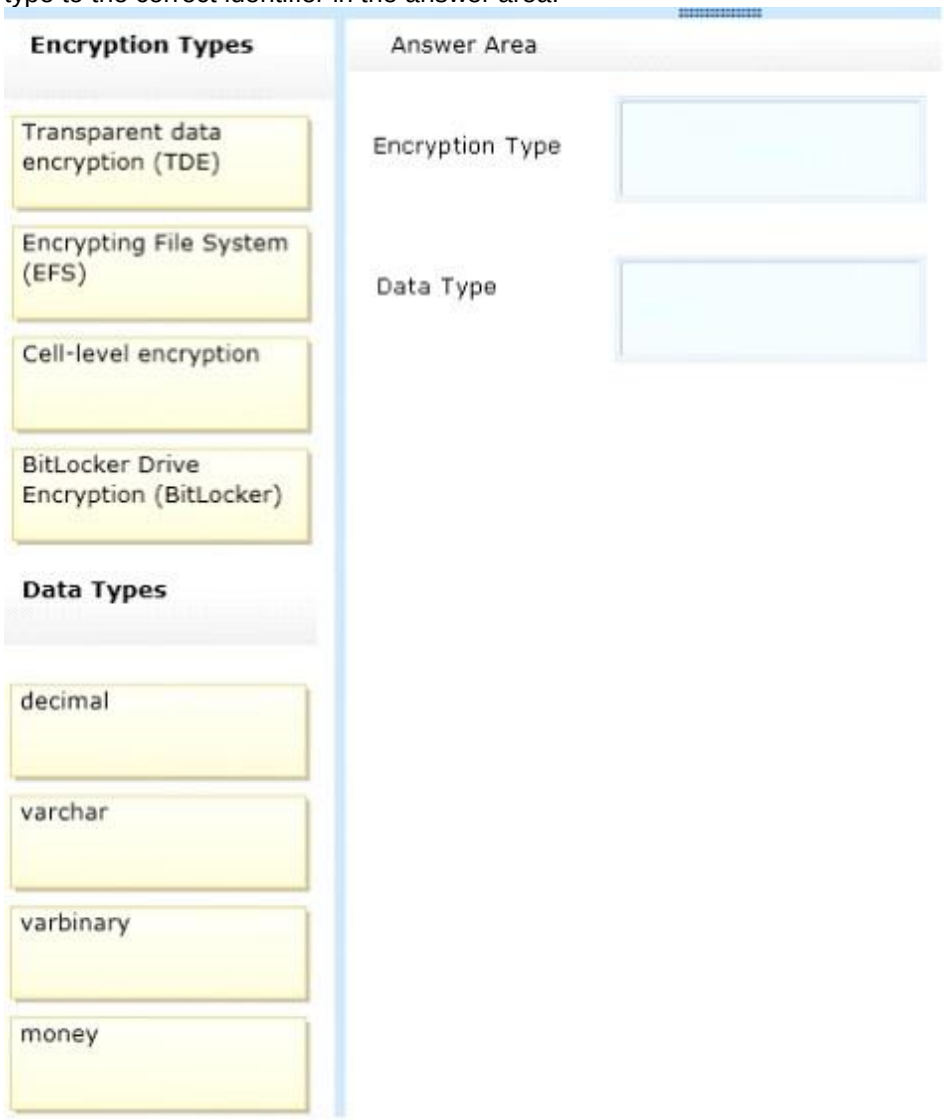

Answer:

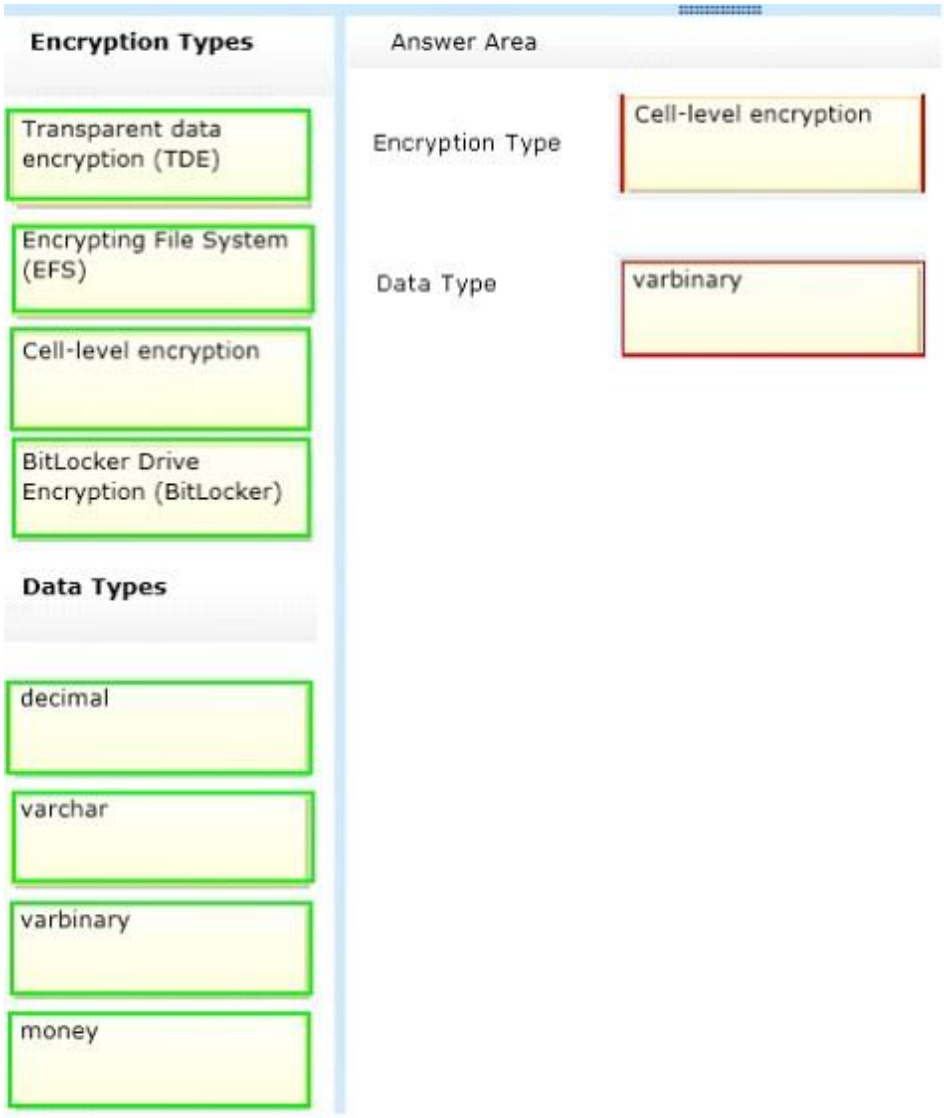

QUESTION 136
You plan to create a database that has multiple tables. The tables will contain product information. Each product has a stock-keeping unit (SKU). You need to recommend a solution to ensure that each SKU starts with the letters "ADV" and is followed by 10 digits. The solution must minimize the amount of development effort required. What should you include in the recommendation?

A. A CHECK constraint
B. A FOREIGN KEY constraint
C. A trigger
D. A user-defined data type

Answer: D

QUESTION 137
You have a Microsoft SQL Azure database. You have the following stored procedure:

```
01 CREATE PROCEDURE UpdateContact
02   @ContactID int,
03   @LastName nvarchar(50)
04 AS
05
06 SELECT LastName AS OriginalName
07 FROM Person.Contact
08
09 WHERE ContactID = @ContactID;
10 UPDATE Person.Contact
11 SET LastName = @LastName
12
13 WHERE ContactID = @ContactID;
```

You discover that the stored procedure periodically fails to update Person.Contact. You need to ensure that Person.Contact is always updated when UpdateContact executes. The solution must minimize the amount of time required for the stored procedure to execute and the number of locks held. What should you do?

A. Add the following line of code to line 08:
`WITH (UPDLOCK)`
B. Add the following line of code to line 12:
`WITH (UPDLOCK)`
C. Add the following line of code to line 05:
`SET TRANSACTION ISOLATION LEVEL SERIALIZABLE`
D. Add the following line of code to line 05:
`SET TRANSACTION ISOLATION LEVEL SNAPSHOT`

Answer: A
Explanation:
- Overall, you should use UPDLOCK when you read a value that you plan to update later in the same transaction to prevent the value from changing.
- UPDLOCK. Specifies that update locks are to be taken and held until the transaction completes. UPDLOCK takes update locks for read operations only at the row-level or page-level. If UPDLOCK is combined with TABLOCK, or a table-level lock is taken for some other reason, an exclusive (X) lock will be taken instead. When UPDLOCK is specified, the READCOMMITTED and READCOMMITTEDLOCK isolation level hints are ignored. For example, if the isolation level of the session is set to SERIALIZABLE and a query specifies (UPDLOCK, READCOMMITTED), the READCOMMITTED hint is ignored and the transaction is run using the SERIALIZABLE isolation level.

QUESTION 138
Drag and Drop Question
Your network contains a SQL Server 2012 instance named SQL1. SQL1 contains a database named DB1. DB1 contains three tables. The tables are configured as shown in the following table:

Table name	Configuration
Table1	• Table data will not be updated. • The table will contain historical calculations. • The table will contain 10 million records.
Table2	• 20% of the table data will be updated weekly. • The table will contain 25 million records.
Table3	• 40% of the table data will be updated weekly. • The table will contain 1 million records.

You plan to create indexes for the tables. You need to identify which type of index must be created for each table. The solution must minimize the amount of time required to return information from the tables. Which type of index should you create for each table? To answer, drag the appropriate index type to the correct table in the answer area.

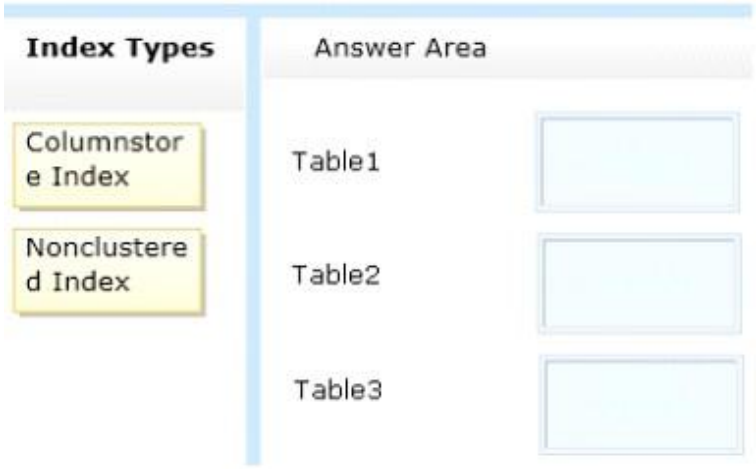

Answer:

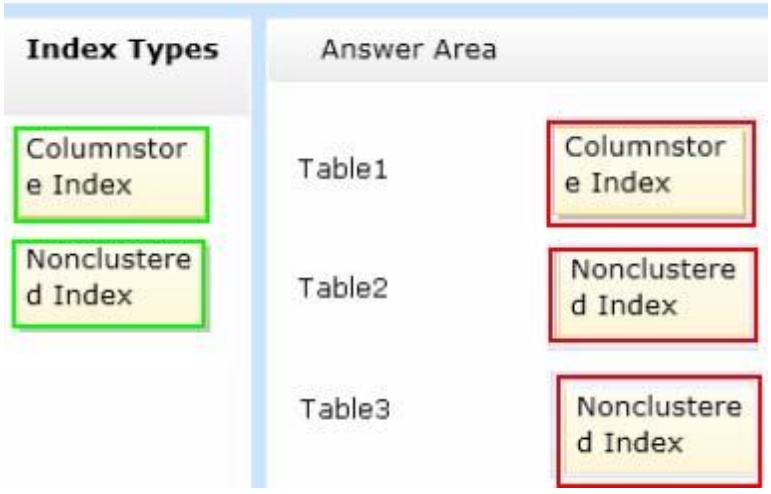

QUESTION 139

Drag and Drop Question

You plan to create a custom aggregation function named Function1. You plan to deploy Function1 to SQL Server 2012. You need to ensure that Function1 can access a web service. The solution must minimize the number of changes made to the database. You create a Microsoft .NET

Framework class that contains the definition of Function1. You upload a certificate to SQL Server. What three tasks should you perform next? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Answer:

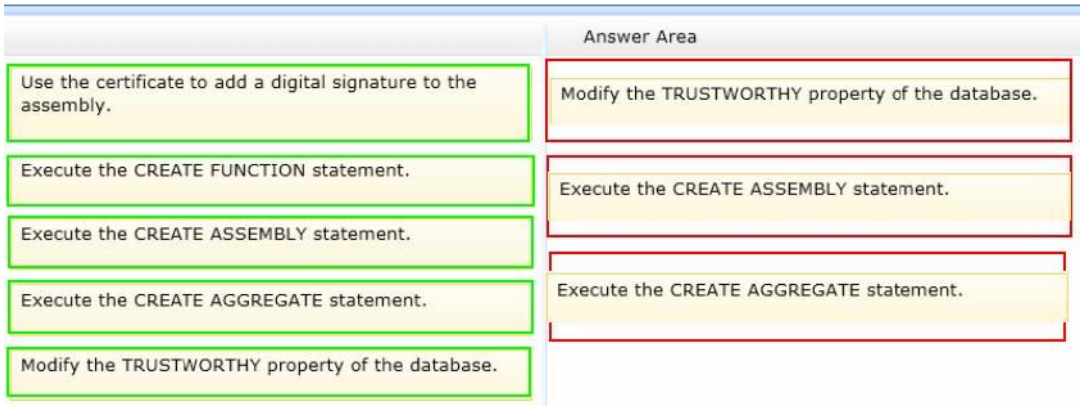

QUESTION 140

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

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

Answer: C

QUESTION 141

You administer a Microsoft SQL Server 2012 database that has multiple tables in the Sales schema. Some users must be prevented from deleting records in any of the tables in the Sales schema. You need to manage users who are prevented from deleting records in the Sales schema. You need to achieve this goal by using the minimum amount of administrative effort. What should you do?

A. Create a custom database role that includes the users.
Deny Delete permissions on the Sales schema for the custom database role.

B. Include the Sales schema as an owned schema for the db_denydatawriter role.
Add the users to the db_denydatawriter role.
C. Deny Delete permissions on each table in the Sales schema for each user.
D. Create a custom database role that includes the users.
Deny Delete permissions on each table in the Sales schema for the custom database role.

Answer: A

QUESTION 142
You are the lead database administrator (DBA) of a Microsoft SQL Server 2012 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

QUESTION 143
You administer a Microsoft SQL Server 2012 database. You provide temporary securityadmin access to User1 to the database server. You need to know if User1 adds logins to securityadmin. Which server-level audit action group should you use?

A. `SERVER_STATE_CHANGE_GROUP`
B. `SERVER_PRINCIPAL_IMPERSONATION_GROUP`
C. `SUCCESSFUL_LOGIN_GROUP`
D. `SERVER_ROLE_MEMBER_CHANGE_GROUP`

Answer: D

QUESTION 144
You administer a Microsoft SQL Server 2012 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

QUESTION 145
You administer a Microsoft SQL Server 2012 database. You need to ensure that the size of the transaction log file does not exceed 2GB. What should you do?

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

Answer: D

QUESTION 146
You administer a Microsoft SQL Server 2012 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

QUESTION 147
You administer a Microsoft SQL Server 2012 database named Contoso on a server named Server01. You need to write messages to the Application Log when users are added to or removed from a fixed server role in Server01. What should you create?

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

Answer: G

QUESTION 148
You administer a Microsoft SQL Server 2012 database named Contoso on a server named Server01. You need to be notified immediately when fatal errors occur on Server01. What should you create?

A. an Alert
B. a Server Audit Specification
C. an Extended Event session
D. a Resource Pool
E. a Policy

F. a SQL Profiler Trace
G. a Database Audit Specification

Answer: A

QUESTION 149
You administer a Microsoft SQL Server 2012 database named Contoso on a server named Server01. You need to diagnose deadlocks that happen when executing a specific set of stored procedures by recording events and playing them back on a different test server. What should you create?

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

Answer: F

QUESTION 150
You create an availability group that has replicas named HA/Server01 and HA/Server02. Currently, HA/Server01 is the primary replica. You have multiple queries that read data and produce reports from the database. You need to offload the reporting workload to the secondary replica when HA/Server01 is the primary replica. What should you do?

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

Answer: B

QUESTION 151
You administer two Microsoft SQL Server 2012 servers. Each server resides in a different, untrusted domain. You plan to configure database mirroring. You need to be able to create database mirroring endpoints on both servers. What should you do?

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

Answer: B

QUESTION 152
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

QUESTION 153
You administer two instances of Microsoft SQL Server 2012. You deploy an application that uses a database on the named instance. The application is unable to connect to the database on the named instance. You need to ensure that the application can connect to the named instance. What should you do?

A. Configure the application as data-tiered.
B. Open port 1433 on the Windows firewall on the server.
C. Configure the named SQL Server instance to use an account that is a member of the Domain Admins group.
D. Start the SQL Server Browser Service.

Answer: D

QUESTION 154
You have a SQL Server 2012 database named Database1. You execute the following code:

```
CREATE TABLE Sales
(
  ID int IDENTITY(1,1) NOT NULL PRIMARY KEY,
  OrderDate char(10) NOT NULL,
  Amount decimal
);
GO

CREATE INDEX IX_Sales_OrderDate
  ON Sales(OrderDate)
  INCLUDE (ID, Amount);
GO

CREATE PROC usp_Proc1(
  @date1 datetime,
  @date2 datetime
)
AS
SELECT ID, OrderDate, Amount
  FROM Sales
  WHERE CAST(OrderDate AS datetime)
    BETWEEN @date1 AND @date2
  ORDER BY ID;
GO
```

You insert 3 million rows into Sales. You need to reduce the amount of time it takes to execute

Proc1. What should you do?

A. Run the following:
```
ALTER TABLE Sales ALTER COLUMN OrderDate datetime NOT NULL;
```
B. Change the WHERE clause to the following:
```
WHERE OrderDate BETWEEN CAST(@date1,char(10))
AND CAST(@date2,char(10))
```
C. Remove the ORDER BY clause from the stored procedure.
D. Run the following:
```
DROP INDEX IX_Sales_OrderDate;
GO
CREATE INDEX IX_Sales_OrderDate ON Sales(OrderDate);
GO
```

Answer: C
Explanation:
http://www.c-sharpcorner.com/UploadFile/skumaar_mca/good-practices-to-write-the-stored-procedures-in-sqlserver/

QUESTION 155
Your network contains a server that has SQL Server 2012 installed. You create a table by using the following script:

```
CREATE TABLE dbo.Products
(
  id int NOT NULL,
  ProductName nvarchar(50) NULL,
  ProductManufacturer nvarchar(50) NULL,
  ProductDescription nvarchar(200) NULL,
    CONSTRAINT PK_Products PRIMARY KEY CLUSTERED (id)
)
  ON [PRIMARY]
GO
```

You need to recommend a solution to ensure that each combination of ProductName and ProductManufacturer is not duplicated. What should you recommend creating?

A. A UNIQUE constraint
B. A CHECK constraint
C. A filtered index
D. A column store index

Answer: A

Case Study 7: Fourth Coffee (Question 156 ~ Question 166)

Background
Corporate Information
Fourth Coffee is global restaurant chain. There are more than 5,000 locations worldwide.
Physical Locations
Currently a server at each location hosts a SQL Server 2012 instance. Each instance contains a database called StoreTransactions that stores all transactions from point of sale and uploads summary batches nightly. Each server belongs to the COFFECORP domain. Local computer accounts access the StoreTransactions database at each store using sysadmin and datareaderwriter roles.

Planned changes

Fourth Coffee has three major initiatives:

- The FT department must consolidate the point of sales database infrastructure.
- The marketing department plans to launch a mobile application for micropayments.
- The finance department wants to deploy an internal tool that will help detect fraud.

Initially, the mobile application will allow customers to make micropayments to buy coffee and other items on the company web site. These micropayments may be sent as gifts to other users and redeemed within an hour of ownership transfer. Later versions will generate profiles based on customer activity that will push texts and ads generated by an analytics application. When the consolidation is finished and the mobile application is in production, the micropayments and point of sale transactions will use the same database.

Existing Environment

Existing Application Environment

Some stores have been using several pilot versions of the micropayment application. Each version currently is in a database that is independent from the point of sales systems. Some versions have been used in field tests at local stores, and others are hosted at corporate servers. All pilot versions were developed by using SQL Server 2012.

Existing Support Infrastructure

The proposed database for consolidating micropayments and transactions is called CoffeeTransactions. The database is hosted on a SQL Server 2014 Enterprise Edition instance and has the following file structures:

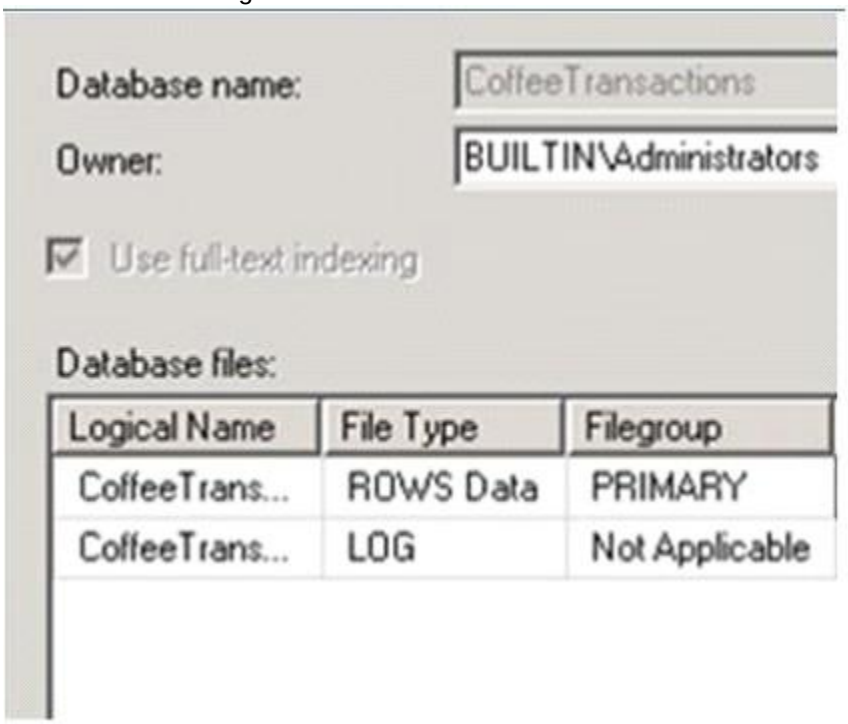

Business Requirements

General Application Solution Requirements

The database infrastructure must support a phased global rollout of the micropayment application and consolidation. The consolidated micropayment and point of sales database will be into a CoffeeTransactions database. The infrastructure also will include a new CoffeeAnalytics database for reporting on content from CoffeeTransactions. Mobile applications will interact most frequently with the micropayment database for the following activities:

- Retrieving the current status of a micropayment.
- Modifying the status of the current micropayment.
- Canceling the micropayment.

The mobile application will need to meet the following requirements:

- Communicate with web services that assign a new user to a micropayment by using a stored procedure named usp_AssignUser.
- Update the location of the user by using a stored procedure named usp_AddMobileLocation.

The fraud detection service will need to meet the following requirements:

- Query the current open micropayments for users who own multiple micropayments by using a stored procedure named usp.LookupConcurrentUsers.
- Persist the current user locations by using a stored procedure named usp_Mobilel_ocationSnapshot.
- Look at the status of micropayments and mark micropayments for internal investigations.
- Move micropayments to dbo.POSException table by using a stored procedure named ups_DetectSuspiciousActivity.
- Detect micropayments that are flagged with a StatusId value that is greater than 3 and that occurred within the last minute.

The CoffeeAnalytics database will combine imports of the POSTransaction and MobileLocation tables to create a UserActivity table for reports on the trends in activity. Queries against the UserActivity table will include aggregated calculations on all columns that are not used in filters or groupings. Micropayments need to be updated and queried for only a week after their creation by the mobile application or fraud detection services.

Performance

The most critical performance requirement is keeping the response time for any queries of the POSTransaction table predictable and fast. Web service queries will take a higher priority in performance tuning decisions over the fraud detection agent queries.

Scalability

Queries of the user of a micropayment cannot return while the micropayment is being updated, but can show different users during different stages of the transaction. The fraud detection service frequently will run queries over the micropayments that occur over different time periods that range between 30 seconds and ten minutes. The POSTransaction table must have its structure optimized for hundreds of thousands of active micropayments that are updated frequently. All changes to the POSTransaction table will require testing in order to confirm the expected throughput that will support the first year's performance requirements. Updates of a user's location can tolerate some data loss. Initial testing has determined that the POSTransaction and POSException tables will be migrated to an in-memory optimized table.

Availability

In order to minimize disruption at local stores during consolidation, nightly processes will restore the databases to a staging server at corporate headquarters.

Technical Requirements

Security

The sensitive nature of financial transactions in the store databases requires certification of the COFFECORP\Auditors group at corporate that will perform audits of the data. Members of the COFFECORP\Auditors group cannot have sysadmin or datawriter access to the database. Compliance requires that the data stewards have access to any restored StoreTransactions database without changing any security settings at a database level. Nightly batch processes are run by the services account in the COFFECORP\StoreAgent group and need to be able to restore and verify the schema of the store databases match. No Windows group should have more access to store databases than is necessary.

Maintainability

You need to anticipate when POSTransaction table will need index maintenance. When the daily maintenance finishes, micropayments that are one week old must be available for queries in UserActivity table but will be queried most frequently within their first week and will require support for in-memory queries for data within first week. The maintenance of the UserActivity table must allow frequent maintenance on the day's most recent activities with minimal impact on the use of disk space and the resources available to queries. The processes that add data to the UserActivity

table must be able to update data from any time period, even while maintenance is running. The index maintenance strategy for the UserActivity table must provide the optimal structure for both maintainability and query performance. All micropayments queries must include the most permissive isolation level available for the maximum throughput. In the event of unexpected results, all stored procedures must provide error messages in text message to the calling web service. Any modifications to stored procedures will require the minimal amount of schema changes necessary to increase the performance.

Performance

Stress testing of the mobile application on the proposed CoffeeTransactions database uncovered performance bottlenecks. The sys.dm_os_wait_stats Dynamic Management View (DMV) shows high wait_time values for WRTTELOG and PAGEIOLATCHJJP wait types when updating the MobileLocation table. Updates to the MobileLocation table must have minimal impact on physical resources.

Supporting Infrastructure

The stored procedure usp_LookupConcurrentUsers has the current implementation:

```
CREATE PROCEDURE usp_LookupConcurrentUsers
AS BEGIN
  --summary table
  CREATE TABLE #POSTransactionTemp (
  POSTransactionId int NOT NULL,
  UserId int NOT NULL,
  StatusID int NOT NULL,
  POSLocation int NOT NULL,
  CreateDate datetime2 NOT NULL,
  Price money
  )
  DECLARE @timewindow datetime2
  SET @timewindow = GETDATE();

  WITH concurrentusers
  AS
  (
  SELECT UserId, COUNT(*) concurrentsessions
  FROM dbo.POSTransaction
  WHERE CreateDate >= dateadd(second,-60, @timewindow )
  GROUP BY UserId
  HAVING COUNT(*) > 1
  )
  INSERT INTO #POSTransactionTemp
  (
POSTransactionId, UserId,
StatusID, POSLocation,
CreateDate, Price
  )

  SELECT d.*
  FROM dbo.POSTransaction d
  JOIN concurrentusers c
  on d.UserID = c.UserId
  WHERE d.CreateDate >= dateadd(second,-60, @timewindow )
  ...
  SELECT * FROM #POSTransactionTemp

  END
```

The current stored procedure for persisting a user location is defined in the following code:

```
CREATE PROCEDURE dbo.usp_MobileLocationSnapshot
AS
BEGIN

INSERT INTO CoffeeAnalytics.dbo.MobileLocationLog
SELECT * FROM CoffeeTransactions.dbo.MobileLocation

END
```

The current stored procedure for managing micropayments needing investigation is defined in the following code:

```
01  CREATE PROCEDURE dbo.usp_DetectSuspiciousActivity
02  WITH NATIVE_COMPILATION, SCHEMABINDING, EXECUTE AS OWNER
03  AS
04  BEGIN ATOMIC
05  WITH (TRANSACTION ISOLATION LEVEL = SNAPSHOT,
06  LANGUAGE = 'us_english')
07  IF EXISTS(SELECT POSTransactionId FROM dbo.POSTransaction
08  WHERE StatusID >= 4 and CreateDate >= dateadd(second,-60,
09  GETDATE() ))
10  MERGE dbo.POSException AS target
11  USING (SELECT POSTransactionId, StatusID, UserId,
12  POSLocation, CreateDate, Price FROM dbo.POSTransaction
13  WHERE StatusID >= 4 and
14  CreateDate >= dateadd(second,-60, GETDATE() ))
15  AS source (POSTransactionId, StatusID, UserId,
16  POSLocation, CreateDate, Price)
17  ON (target.POSTransactionId = source.POSTransactionId)
18  WHEN MATCHED THEN
19  UPDATE SET StatusID = source.StatusID
20  WHEN NOT MATCHED THEN
21  INSERT (POSTransactionId, StatusID, UserId,
22  POSLocation, CreateDate, Price)
23  VALUES (source.POSTransactionId, source.StatusID,
24  source.UserId, source.POSLocation,
25  source.CreateDate, source.Price);
26  END
```

The current table, before implementing any performance enhancements, is defined as follows:

```
CREATE TABLE dbo.POSTransaction (
  POSTransactionId int NOT NULL PRIMARY KEY,
  UserId int NOT NULL,
  POSLocation int NOT NULL,
  StatusID int NOT NULL,
  CreateDate datetime2 NOT NULL,
  Price money
  )
  CREATE INDEX ix_UserID on dbo.POSTransaction(UserId)
```

QUESTION 156

You need to monitor the health of your tables and indexes in order to implement the required index maintenance strategy. What should you do?

A. Query system DMVs to monitor avg_chain_length and max_chain_length.
Create alerts to notify you when these values converge.

B. Create a SQL Agent alert when the File Table:
Avg time per file I/O request value is increasing.

C. Query system DMVs to monitor total_bucket_count.
Create alerts to notify you when this value increases.

D. Query system DMVs to monitor total_bucket_count.
Create alerts to notify you when this value decreases.

Answer: A
Explanation:
From scenario:
- You need to anticipate when POSTransaction table will need index maintenance.
- The index maintenance strategy for the UserActivity table must provide the optimal structure for both maintainability and query performance.

QUESTION 157
Drag and Drop Question
You need to design the UserActivity table. Which three steps should you perform in sequence? To answer, move the appropriate three actions from the list of actions to the answer area and arrange them in the correct order.

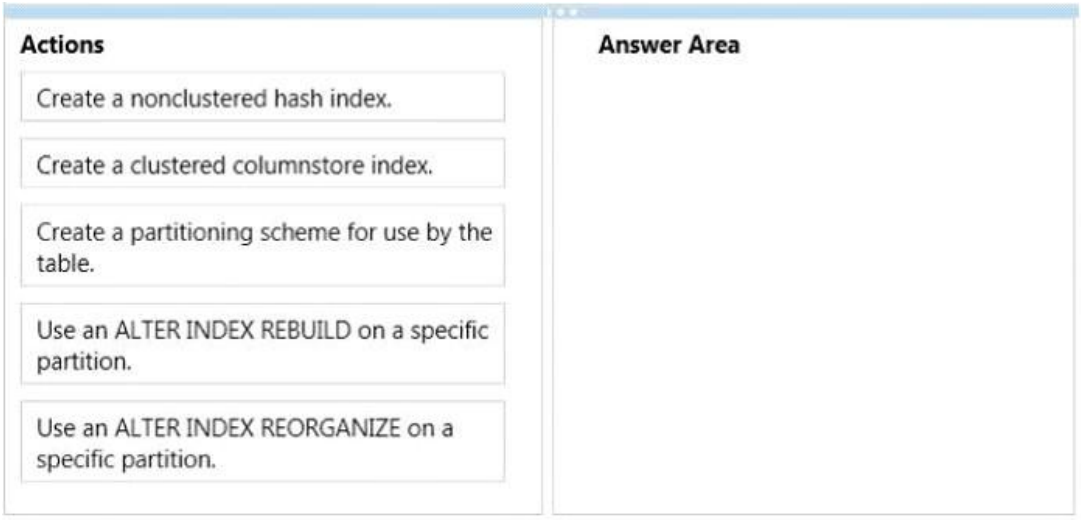

Answer:

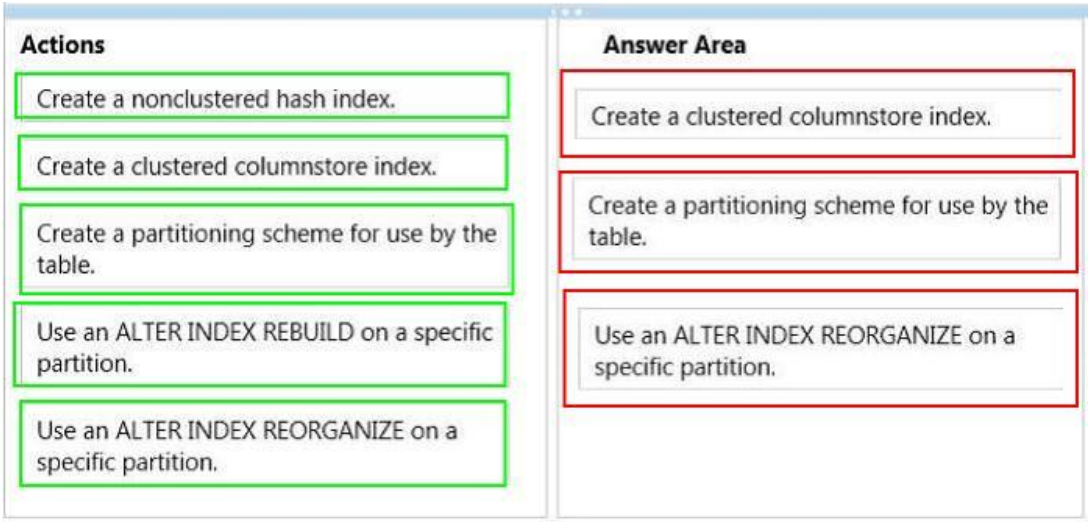

QUESTION 158
You need to implement security for the restore and audit process. What should you do?

A. Grant the COFFECORP\Auditors group ALTER ANY CONNECTION and SELECT ALL USER SECURABLES permissions.
Grant the COFFECORP\StoreAgent group ALTER ANY CONNECTION and IMPERSONATE ANY LOGIN

permissions.

B. Grant the COFFECORP\Auditors group CONNECT ANY DATABASE and IMPERSONATE ANY LOGIN permissions.
Grant the COFFECORP\StoreAgent group CONNECT ANY DATABASE and SELECT ALL USER SECURABLES permissions.
C. Grant the COFFECORP\Auditors group ALTER ANY CONNECTION and IMPERSONATE ANY LOGIN permissions.
Grant the COFFECORP\StoreAgent group ALTER ANY CONNECTION and SELECT ALL USER SECURABLES permissions.
D. Grant the COFFECORP\Auditors group CONNECT ANY DATABASE and SELECT ALL USER SECURABLES permissions.
Grant the COFFECORP\StoreAgent group CONNECT ANY DATABASE and IMPERSONATE ANY LOGIN permissions.

Answer: A

QUESTION 159
You need to modify the stored procedure usp_LookupConcurrentUsers. What should you do?

A. Add a clustered index to the summary table.
B. Add a nonclustered index to the summary table.
C. Add a clustered columnstore index to the summary table.
D. Use a table variable instead of the summary table.

Answer: A
Explanation:
Scenario: Query the current open micropayments for users who own multiple micropayments by using a stored procedure named usp.LookupConcurrentUsers.

QUESTION 160
Drag and Drop Question
You need to create the usp.AssignUser stored procedure. 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

```
    IF @StatusID IS NULL
    RAISERROR (N'The transaction does
not exist.',16,1)
```

```
WITH
NATIVE_COMPILATION, SCHEMABINDING,
EXECUTE AS OWNER
```

```
CREATE PROCEDURE dbo.usp_AssignUser
@UserId int, @POSTransactionId int
```

```
WITH (TRANSACTION ISOLATION LEVEL =
READ COMMITTED, LANGUAGE
= N'us_english')
```

```
  UPDATE dbo.POSTransaction
SET UserId=@UserId
  WHERE POSTransactionId=@POSTransactio
nId
END
```

```
AS
BEGIN
```

```
DECLARE @StatusID int
  SELECT @StatusID=StatusId
  FROM dbo.POSTransaction
WHERE POSTransactionId=@POSTransactionI
d
```

```
    IF @StatusID IS NULL
    THROW 51000, N'The transaction
does not exist.', 1
```

```
WITH (TRANSACTION ISOLATION LEVEL =
REPEATABLE READ, LANGUAGE
= N'us_english')
```

```
AS
BEGIN ATOMIC
```

Answer Area

Answer:

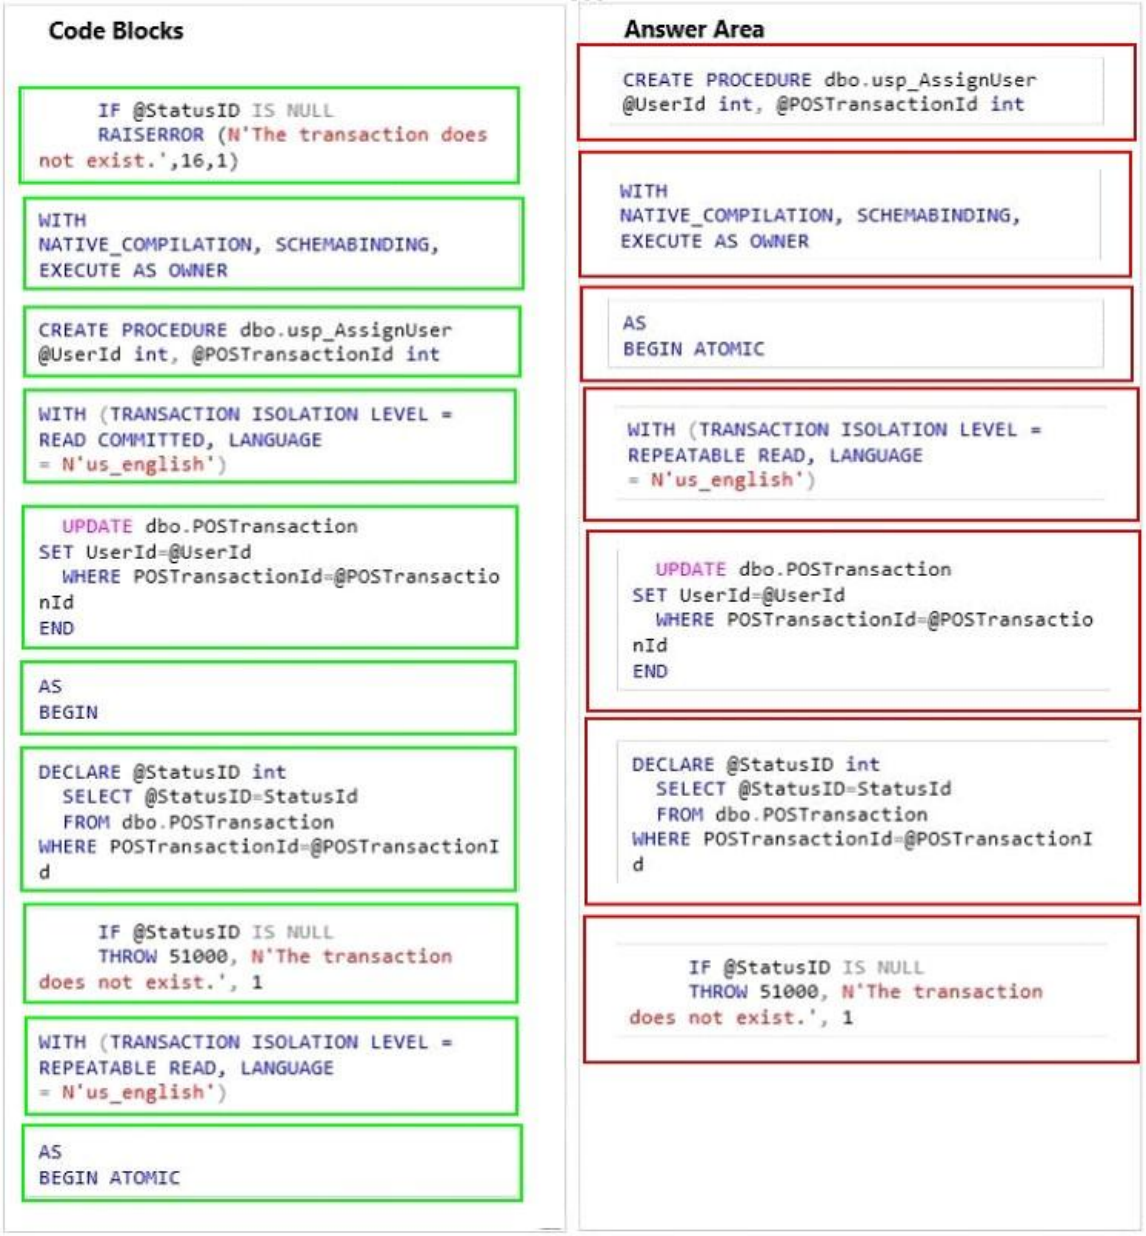

QUESTION 161
Drag and Drop Question
You need to implement a new version of usp_AddMobileLocation. Develop the solution by selecting and arranging the required code blocks in the correct order. You may not need all of the code blocks.

Code Blocks

```
DELAYED_DURABILITY = ON
,TRANSACTION ISOLATION LEVEL = SNAPSHOT
```

```
CREATE PROCEDURE dbo.usp_AddMobileLocat
ion @POSTransactionId int, @Long
float, @Lat float
WITH
```

```
NATIVE_COMPILATION
...
```

```
DELAYED_DURABILITY = OFF
,TRANSACTION ISOLATION LEVEL
= READ UNCOMMITTED
```

```
DELAYED_DURABILITY = ON
,TRANSACTION ISOLATION LEVEL
= READ UNCOMMITTED
```

```
Insert into dbo.MobileLocation
(
POSTransactionId,
Longitude,
Latitude,
CreateDate
)
VALUES
(
@POSTransactionId,
@Long,
@Lat,
GETDATE()
)
END
```

```
,LANGUAGE = N'English'
)
```

```
AS
BEGIN ATOMIC WITH
(
```

```
DELAYED_DURABILITY = OFF
,TRANSACTION ISOLATION LEVEL = SNAPSHOT
```

Answer Area

Answer:

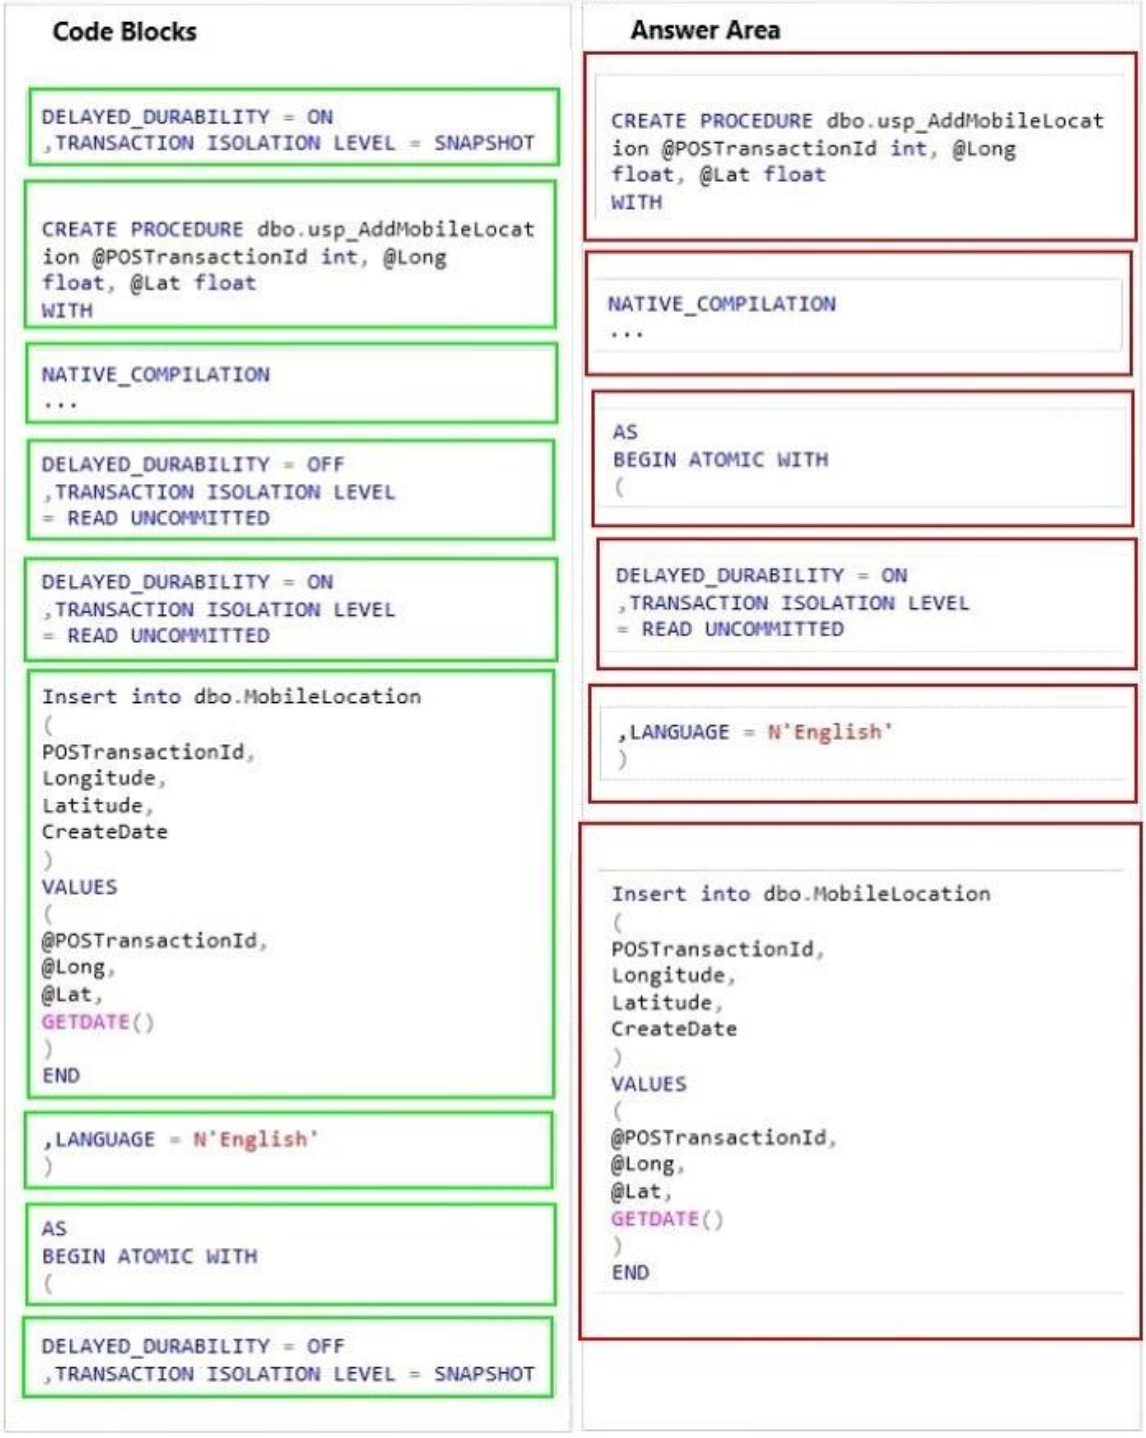

QUESTION 162

You need to modify the usp_DetectSuspiciousActivity stored procedure. Which two actions should you perform? Each correct answer presents part of the solution. (Choose two.)

☐A. Replace lines 04-06 with the following code:

```
BEGIN ATOMIC WITH
     (
          DELAYED_DURABILITY = ON,
          TRANSACTION ISOLATION LEVEL = READ UNCOMMITTED,
          LANGUAGE = N'English'
)
```

☐B. Replace lines 04-06 with the following code:

```
BEGIN ATOMIC WITH
     (
          DELAYED_DURABILITY = ON,
          TRANSACTION ISOLATION LEVEL = REPEATABLE READ
)
```

☐C. Change the logic of the stored procedure to use separate UPDATE and INSERT statements.

☐D. Replace lines 07-09 with the following code:

```
DECLARE @exists BIT = 0
IF EXISTS ( SELECT TOP 1 * FROM POSTransaction (NOLOCK) WHERE StatusID = 4 and CreateDate
>= dateadd(second,-60, GETDATE() ))
```

☐E. Replace lines 04-06 with the following code:

```
BEGIN ATOMIC WITH
     (
          TRANSACTION ISOLATION LEVEL = READ UNCOMMITTED,
          LANGUAGE = N'English'
)
```

☐F. Replace lines 07-09 with the following code:

```
DECLARE @exists BIT = 0
SELECT TOP 1 @exists = 1 FROM POSTransaction WHERE StatusID >= 4 and CreateDate >= dateadd
(second,-60, GETDATE() )
IF @exists = 1
```

A. Option A
B. Option B
C. Option C
D. Option D
E. Option E
F. Option F

Answer: DE
Explanation:
Note: Move micropayments to dbo.POSException table by using a stored procedure named ups_DetectSuspiciousActivity.

QUESTION 163
Drag and Drop Question
You need to redesign the system to meet the scalability requirements of the application. 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

```
,
UserId int NOT NULL
INDEX ix_UserId NONCLUSTERED
HASH WITH (BUCKET_COUNT=2),
```

```
,
UserId int NOT NULL
INDEX x_UserId NONCLUSTERED
HASH WITH (BUCKET_COUNT=900000),
```

```
POSLocation int NOT NULL,
StatusID int NOT NULL,
CreateDate datetime2 NOT NULL,
Price money
)
```

```
POSTransactionId int NOT NULL
PRIMARY KEY CLUSTERED
```

```
POSTransactionId int NOT NULL
```

```
ALTER DATABASE CoffeeTransactions
ADD FILEGROUP [CoffeeTransactions_inmem
] CONTAINS MEMORY_OPTIMIZED_DATA
```

```
ON [CoffeeTransactions_inmem]
```

```
WITH (MEMORY_OPTIMIZED=ON,
DURABILITY=SCHEMA_ONLY)
```

```
POSTransactionId int NOT NULL
PRIMARY KEY CLUSTERED
HASH WITH (BUCKET_COUNT=1000000)
```

```
,
UserId int NOT NULL
NONCLUSTERED INDEX ix_UserId,
```

```
CREATE TABLE dbo.POSTransaction (
```

```
POSTransactionId int NOT NULL
PRIMARY KEY NONCLUSTERED
HASH WITH (BUCKET_COUNT=1)
```

Answer Area

Answer:

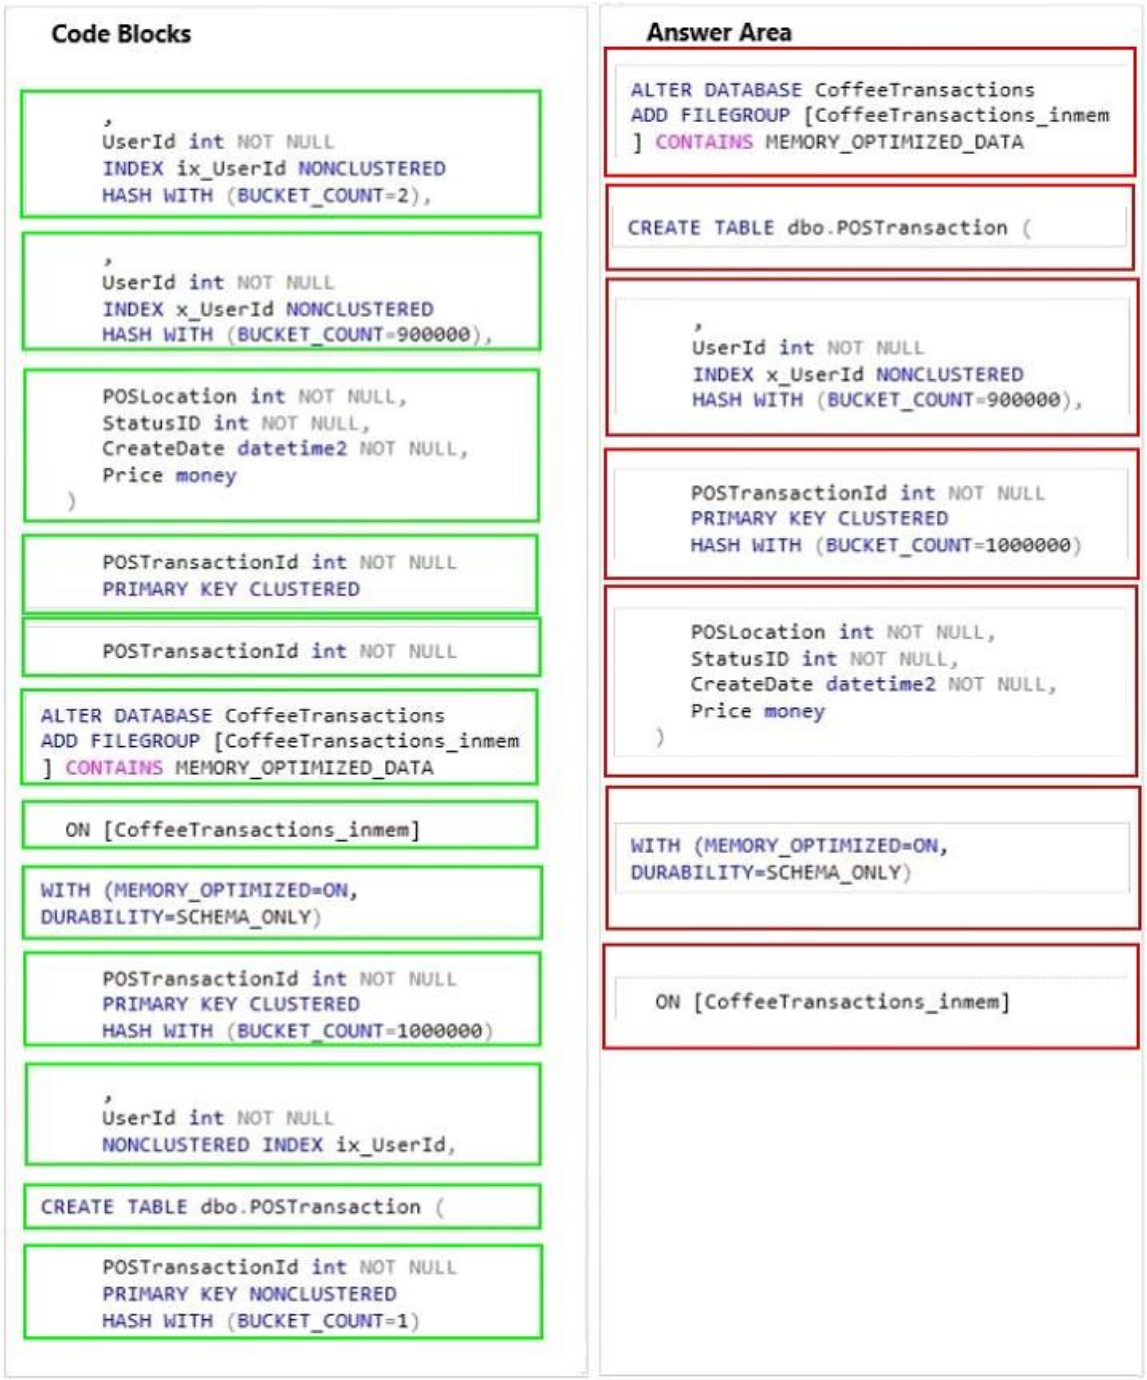

QUESTION 164

Drag and Drop Question

You need to optimize the index and table structures for POSTransaction. Which task should you use with each maintenance step? To answer, drag the appropriate tasks to the correct maintenance steps. Each task 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.

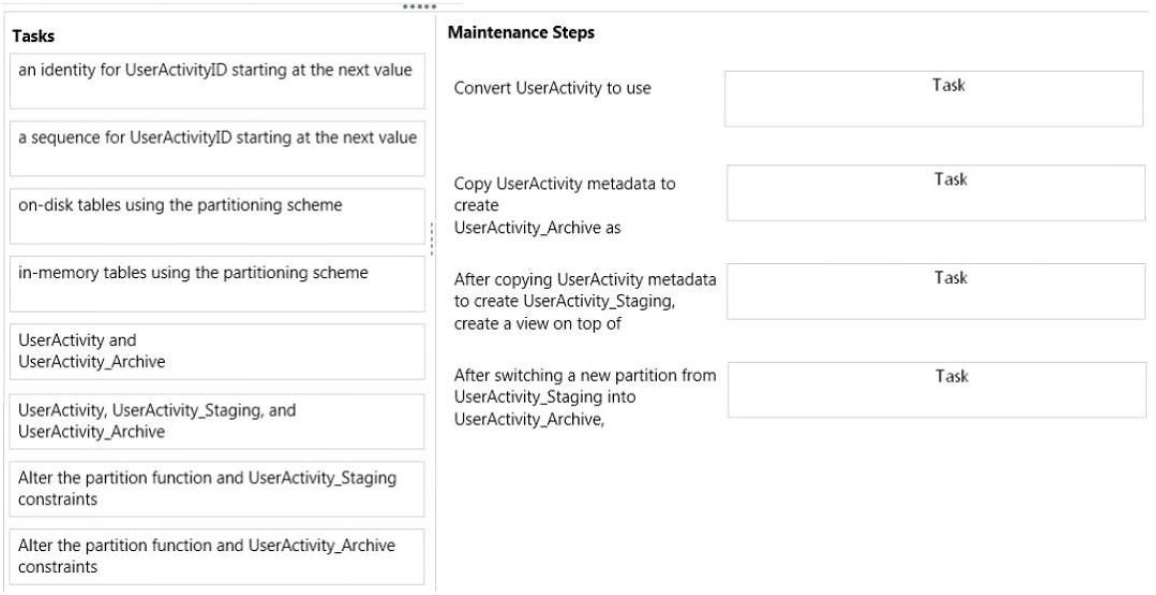

Answer:

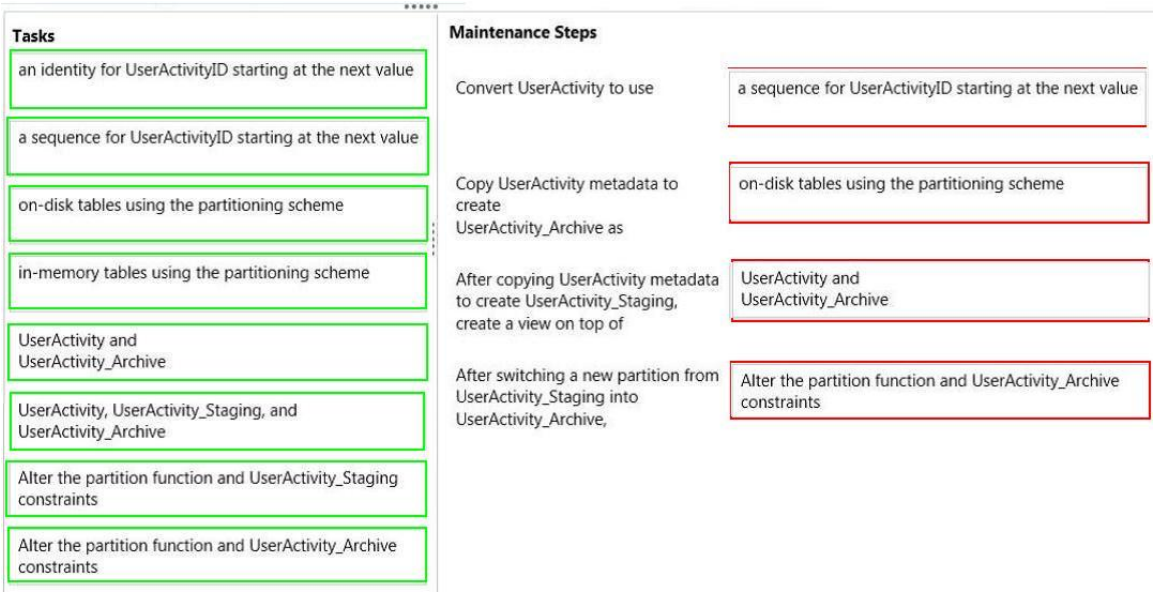

QUESTION 165
You need to optimize the index structure that is used by the tables that support the fraud detection services. What should you do?

A. Add a hashed nonclustered index to CreateDate.
B. Add a not hash nonclustered index to CreateDate.
C. Add a not hash clustered index on POSTransactionId and CreateDate.
D. Add a hashed clustered index on POSTransactionId and CreateDate.

Answer: A
Explanation:
The fraud detection service will need to meet the following requirement (among others): detect micropayments that are flagged with a StatusId value that is greater than 3 and that occurred within

the last minute.

QUESTION 166
You need to modify the stored procedure usp_LookupConcurrentUsers. What should you do?

A. Use the summary table as an in-memory optimized table with a non-hash clustered index.
B. Use the summary table as an in-memory optimized table with a non-hash nonclustered index.
C. Use a type variable instead of the summary table.
D. Add a clustered index to the summary table.

Answer: D
Explanation:
Scenario: Query the current open micropayments for users who own multiple micropayments by using a stored procedure named usp.LookupConcurrentUsers.

QUESTION 167
You need to identify which long running transactions use an index. Which dynamic management view should you use?

A. sys.dm_exec_query_optimizer_info
B. sys.dm_exec_connections
C. sys.dm_exec_query_stats
D. sys.dm_exec_sessions

Answer: B

QUESTION 168
You create a table named Customers by using the following code segment:

```
CREATE TABLE dbo.Customers
(
  id int primary key,
  name char(10)
)
```

You create a non-clustered index named IX_Name on the name column. You write the following query to retrieve all of the customers that have a name that starts with the letters SMI:

```
SELECT * FROM dbo.Customers
WHERE 'smi' = LEFT(name,3)
```

You discover that the query performs a table scan. You need to ensure that the query uses the index. What should you do?

A. Replace LEFT(name,3) = 'smi' by using name like 'smi%'
B. Replace LEFT(name,3) = 'smi' by using substring(name,l,3) = 'smi'
C. Recreate IX_Name as a unique index
D. Recreate IX Name as a clustered index

Answer: A

QUESTION 169
You plan to create a new column in a table. The column must meet the following requirements:
- Be able to store images that are larger than 1 MB each.
- Be able to access the images from Microsoft .NET Framework applications.
You need to recommend which data type must be used in the column. Which data type should you

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

A. nvarchar
B. varbinary
C. image
D. FileStream

Answer: D

QUESTION 170
You plan to modify a procedure that contains hundreds of lines of code. The modification must support the following guidelines:
- Use only tables that are not persistent in the database.
- Minimize the amount of time required to execute and recompile procedures.
You need to identify which type of table must be used to support the planned modification. Which type of table should you identify?

A. A system table
B. A partitioned table
C. A table variable
D. A temporary table

Answer: C

QUESTION 171
You have a SQL Server 2014 instance named SQL\Instance1. Instance1 contains a database named Database1. You need to recommend an index defragmentation solution for an index named ContentIndex. ContentIndex must meet the following requirements:
- Remain online during the defragmentation.
- Update distribution statistics.
- Perform defragmentation as quickly as possible.
Which type of index defragmentation solution should you include in the recommendation? More than one answer choice may achieve the goal. Select the BEST answer.

A. DBCC DBREINDEX
B. REORGANIZE
C. REBUILD
D. DBCC INDEXDEFRAG

Answer: B

QUESTION 172
Your company has a main office in London and a branch office in New York. Your network contains a server named Server5 that has SQL Server 2012 installed. Serverscontains a database name ContentDB and a table named ContentTable. You add an additional server named Server9 that runs SQL Server 2012. You need to create a distributed partitioned view. The solution must minimize the amount of network traffic. What should you do? (Each correct answer presents part of the solution. Choose all that apply.)

A. Create the view on Server5.
B. Add Server9 as a linked server.
C. Create the view on Server9.

D. Add the Customers table to Server9.
E. Add Server9 as a Distributor.
F. Remove the Customers table from Server5.

Answer: ABCD

QUESTION 173

.....