

Exam Questions 70-464

Developing Microsoft SQL Server 2012 Databases

<https://www.2passeasy.com/dumps/70-464/>



NEW QUESTION 1

- (Exam Topic 1)

Which data type should you use for CustomerID?

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

Answer: D

Explanation:

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.

int: -2^{31} (-2,147,483,648) to $2^{31}-1$ (2,147,483,647) (just 10 digits max)

bigint: -2^{63} (-9,223,372,036,854,775,808) to $2^{63}-1$ (9,223,372,036,854,775,807)

<http://msdn.microsoft.com/en-us/library/ms176089.aspx> <http://msdn.microsoft.com/en-us/library/ms187745.aspx>

NEW QUESTION 2

- (Exam Topic 1)

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

- A.

```
OPEN CERT1;
ALTER PROCEDURE Accounting.usp_InsertInvoice
WITH ENCRYPTION;
CLOSE CERT1;
```
- B.

```
OPEN CERT2;
ALTER PROCEDURE Accounting.usp_InsertInvoice
WITH ENCRYPTION;
CLOSE CERT2;
```
- C.

```
ADD SIGNATURE TO Accounting.usp_InsertInvoice
BY CERTIFICATE CERT1;
```
- D.

```
ADD SIGNATURE TO Accounting.usp_InsertInvoice
BY CERTIFICATE CERT2;
```

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

Answer: D

NEW QUESTION 3

- (Exam Topic 1)

You need to create a function that filters invoices by CustomerID. The SELECT statement for the function is contained in InvoicesByCustomer.sql.

Which code segment should you use to complete the function?

- A. CREATE FUNCTION Accounting.fnInvoicesByCustomertest (@CustID varchar(11))
 RETURNS @TblInvoices TABLE (CustomerID bigint, CustomerName NVARCHAR(255),
 InvoiceID bigint, InvoiceDate date, Total decimal(8,2), PONumber bigint)
 AS
- B. CREATE FUNCTION Accounting.fnInvoicesByCustomer (@CustID varchar(11))
 RETURNS @tblInvoices TABLE (CustomerID bigint, CustomerName NVARCHAR(255),
 InvoiceID bigint, InvoiceDate date, Total decimal(8,2), PONumber bigint)
 AS
 INSERT INTO @tblInvoices
- C. CREATE FUNCTION Accounting.fnInvoicesByCustomer (@CustID varchar(11))
 RETURNS xml
 AS
 RETURN
- D. CREATE FUNCTION Accounting.fnInvoicesByCustomertest (@CustID varchar(11))
 RETURNS @TblInvoices TABLE (CustomerID bigint, CustomerName NVARCHAR(255),
 InvoiceID bigint, InvoiceDate date, Total decimal(8,2), PONumber bigint)
 AS

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

Answer: A

NEW QUESTION 4

- (Exam Topic 1)

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>

NEW QUESTION 5

- (Exam Topic 1)

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

Answer: B

NEW QUESTION 6

- (Exam Topic 1)

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
 ,0xA694FBEA88C9354E5E2567C30A2A69E8FB4C44A9\

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. RESTORE CERTIFICATE CERT2
FROM FILE='CERT2.CER'
WITH PRIVATE KEY (FILE = 'CERT2.KEY',
DECRYPTION BY PASSWORD='p@ssw0rd1');
- B. CREATE CERTIFICATE CERT1
FROM FILE='CERT1.CER'
WITH PRIVATE KEY (FILE = 'CERT1.KEY',
DECRYPTION BY PASSWORD='p@ssw0rd1');
- C. CREATE CERTIFICATE CERT2
ENCRYPTION BY PASSWORD='p@ssw0rd1'
WITH SUBJECT = 'EncryptionCertificate';
- D. CREATE CERTIFICATE CERT1
ENCRYPTION BY PASSWORD='p@ssw0rd1'
WITH SUBJECT = 'EncryptionCertificate';

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

Answer: B

NEW QUESTION 7

- (Exam Topic 1)

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

Answer: C

NEW QUESTION 8

- (Exam Topic 2)

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.

NEW QUESTION 9

- (Exam Topic 2)

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.

NEW QUESTION 10

- (Exam Topic 2)

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

Answer: A

Explanation:

<http://msdn.microsoft.com/en-us/library/ms177603.aspx>

NEW QUESTION 10

- (Exam Topic 3)

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>

NEW QUESTION 13

- (Exam Topic 4)

Which data type should you use for ProductType?

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

Answer: C

NEW QUESTION 15

- (Exam Topic 5)

You need to resolve the performance issues of the usp_getOpenings stored procedure. Which three actions should you perform? Each correct answer presents part of the solution.

- A. Delete lines 05 through 08
B. Replace lines 12, 13, and 14 with the Transact-SQL segment: WHERE (CONTAINS(o.Description, 'ISABOUT(' + @keyword + ' weight (.5)))
C. Create a full text index on the Description column
D. Replace lines 12, 13, and 14 with the Transact_SQL segment: WHERE (CONTAINS(o.Description, @keyword))
E. Replace lines 12, 13, and 14 with the Transact SQL Segment: WHERE (Contains(o.Description, 'FORMSOF(INFLECTIONAL, '+@keyword+')))

Answer: ACE

Explanation:

Scenario: 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.

The FORMSOF term performs matches using other linguistic forms of the word. The following is the FORMSOF term syntax:

FORMSOF (<generation_type>,<match_words>)

The generation type specifies how Microsoft Windows Search chooses the alternative word forms. The INFLECTIONAL value chooses alternative inflection forms for the match words. If the word is a verb, alternative tenses are used. If the word is a noun, the singular, plural, and possessive forms are used to detect matches.

References:

<https://docs.microsoft.com/en-us/windows/win32/search/-search-sql-formsof>

NEW QUESTION 19

- (Exam Topic 5)

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

NEW QUESTION 22

- (Exam Topic 5)

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. ROLLBACK DISTRIBUTED TRANSACTION @transactionID
- C. BEGIN TRANSACTION WITH MARK @transactionID
- D. COMMIT TRANSACTION @transactionID

Answer: C

NEW QUESTION 26

- (Exam Topic 5)

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

Answer: D

NEW QUESTION 31

- (Exam Topic 5)

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

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

Answer: D

NEW QUESTION 33

- (Exam Topic 6)

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

NEW QUESTION 38

- (Exam Topic 6)

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:element>
    </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['+ @itemNumber+ ']/productID', int)
- D. SET @productID = @items.value'/Root/Product/@productID', int)

Answer: B

NEW QUESTION 43

- (Exam Topic 6)

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

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

Answer: C

Explanation:

A partitioned view is a view defined by a UNION ALL of member tables structured in the same way, but stored separately as multiple tables in either the same instance of SQL Server or in a group of autonomous instances of SQL Server servers, called federated database servers.

References:

<https://docs.microsoft.com/en-us/sql/t-sql/statements/create-view-transact-sql?view=sql-server-2017>

NEW QUESTION 48

- (Exam Topic 6)

You discover that the usp_GetOrdersAndItems stored procedure 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 ALLOW_SNAPSHOT_ISOLATION OFF statement.

Answer: BD

NEW QUESTION 51

- (Exam Topic 7)

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: A

NEW QUESTION 53

- (Exam Topic 7)

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	Answer Area
<pre> , UserId int NOT NULL INDEX ix_UserId NONCLUSTERED HASH WITH (BUCKET_COUNT=2),</pre>	
<pre> , UserId int NOT NULL INDEX x_UserId NONCLUSTERED HASH WITH (BUCKET_COUNT=900000),</pre>	
<pre> POSLocation int NOT NULL, StatusID int NOT NULL, CreateDate datetime2 NOT NULL, Price money)</pre>	
<pre> POSTransactionId int NOT NULL PRIMARY KEY CLUSTERED</pre>	
<pre> POSTransactionId int NOT NULL</pre>	
<pre>ALTER DATABASE CoffeeTransactions ADD FILEGROUP [CoffeeTransactions_inmem] CONTAINS MEMORY_OPTIMIZED_DATA</pre>	
<pre>ON [CoffeeTransactions_inmem]</pre>	
<pre>WITH (MEMORY_OPTIMIZED=ON, DURABILITY=SCHEMA_ONLY)</pre>	
<pre> POSTransactionId int NOT NULL PRIMARY KEY CLUSTERED HASH WITH (BUCKET_COUNT=1000000)</pre>	
<pre> , UserId int NOT NULL NONCLUSTERED INDEX ix_UserId,</pre>	
<pre>CREATE TABLE dbo.POSTransaction (POSTransactionId int NOT NULL PRIMARY KEY NONCLUSTERED HASH WITH (BUCKET_COUNT=1)</pre>	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1:

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

Box 2:

```
CREATE TABLE dbo.POSTransaction (  
    POSTransactionId int NOT NULL  
    PRIMARY KEY NONCLUSTERED  
    HASH WITH (BUCKET_COUNT=1)
```

Box 3:

```

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

```

Box 4:

```

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

```

Box 5:

```

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

```

Box 6:

```

WITH (MEMORY_OPTIMIZED=ON,
DURABILITY=SCHEMA_ONLY)

```

Box 7:

```

ON [CoffeeTransactions_inmem]

```

Note:

* MEMORY_OPTIMIZED_DATA

First create a memory-optimized data filegroup and add a container to the filegroup. Then create a memory-optimized table.

* 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.

* Example:

```

-- create a durable (data will be persisted) memory-optimized table
-- two of the columns are indexed
CREATE TABLE dbo.ShoppingCart (
ShoppingCartId INT IDENTITY(1,1) PRIMARY KEY NONCLUSTERED,
UserId INT NOT NULL INDEX ix_UserId NONCLUSTERED HASH WITH (BUCKET_COUNT=1000000),
CreatedDate DATETIME2 NOT NULL, TotalPrice MONEY
) WITH (MEMORY_OPTIMIZED=ON) GO

```

NEW QUESTION 57

- (Exam Topic 7)

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.

Tasks	Maintenance Steps
an identity for UserActivityID starting at the next value	Convert UserActivity to use
a sequence for UserActivityID starting at the next value	Task
on-disk tables using the partitioning scheme	Copy UserActivity metadata to create UserActivity_Archive as
in-memory tables using the partitioning scheme	Task
UserActivity and UserActivity_Archive	After copying UserActivity metadata to create UserActivity_Staging, create a view on top of
UserActivity, UserActivity_Staging, and UserActivity_Archive	Task
Alter the partition function and UserActivity_Staging constraints	After switching a new partition from UserActivity_Staging into UserActivity_Archive,
Alter the partition function and UserActivity_Archive constraints	Task

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Tasks	Maintenance Steps
an identity for UserActivityID starting at the next value	Convert UserActivity to use a sequence for UserActivityID starting at the next value
a sequence for UserActivityID starting at the next value	
on-disk tables using the partitioning scheme	Copy UserActivity metadata to create UserActivity_Archive as on-disk tables using the partitioning scheme
in-memory tables using the partitioning scheme	
UserActivity and UserActivity_Archive	After copying UserActivity metadata to create UserActivity_Staging, create a view on top of UserActivity and UserActivity_Archive
UserActivity, UserActivity_Staging, and UserActivity_Archive	
Alter the partition function and UserActivity_Staging constraints	After switching a new partition from UserActivity_Staging into UserActivity_Archive, Alter the partition function and UserActivity_Archive constraints
Alter the partition function and UserActivity_Archive constraints	

NEW QUESTION 58

- (Exam Topic 7)

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.

NEW QUESTION 62

- (Exam Topic 7)

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

NEW QUESTION 67

- (Exam Topic 7)

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

Answer: A

NEW QUESTION 68

- (Exam Topic 8)

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. Server5 contains 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

NEW QUESTION 69

- (Exam Topic 8)

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)

NEW QUESTION 71

- (Exam Topic 8)

You have database objects that were created by using the following script:

```
CREATE TABLE dbo.Customer
(
    CustomerID int IDENTITY(1,1),
    FirstName nvarchar(50) NOT NULL,
    LastName nvarchar(50) NOT NULL,
    CreationDate datetime NOT NULL,
    CONSTRAINT PK_Customer
        PRIMARY KEY (CustomerID)
);
GO
CREATE NONCLUSTERED INDEX IX_Customers_CreationDate
ON dbo.Customer (CreationDate)
INCLUDE (LastName, FirstName)
WHERE CreationDate >= '1/1/2008';
GO
CREATE PROCEDURE dbo.usp_GetCustomersByDate
    @CreationDate datetime
AS
SELECT LastName,
    FirstName,
    CreationDate
FROM dbo.Customer
WHERE CreationDate > @CreationDate;
GO
```

The dbo.Customers table has 1 million rows.
 You discover that usp_GetCustomersByDate takes a long time to complete.
 The query plan used by the stored procedure is shown in the exhibit. (Click the Exhibit button.)



You need to ensure that usp_GetCustomersByDate completes as quickly as possible. What should you do?

- A. Modify the stored procedure to include the OPTIMIZE FOR UNKNOWN query hint.
- B. Execute the sp_recompile 'dbo.GetCustomersByDate' statement.
- C. Execute the ALTER INDEX IX_Customers_CreationDate WITH REBUILD statement.
- D. Modify the stored procedure to include the OPTIMIZE FOR('1/1/2008') query hint.

Answer: A

NEW QUESTION 72

- (Exam Topic 8)

You have a SQL Server 2012 database named Database1.

Database1 has a table named Customers. Customers contains more than 1 million rows. The database has a stored procedure that was created by using the following script:

```
CREATE PROCEDURE up_customers
    @CustomerTypeList nvarchar(400)
AS
SELECT CustomerID,
    FirstName,
    LastName
FROM dbo.customers
WHERE CustomerTypeID IN (@CustomerTypeList);
```

You need to ensure that up_customers returns rows when the following statement is executed: EXECUTE up_customers'1,2,3,4,5';
 What should you do?

- A. Update @CustomerTypelist to use the int data type.
- B. Convert @CustomerTypeList to a table variable.
- C. Convert @CustomerTypeList to an XML variable.
- D. Update @CustomerTypeList to use the XML data type.

Answer: B

NEW QUESTION 76

- (Exam Topic 8)

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. Temporary stored procedures
- C. Temporary tables
- D. Table variables

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>

NEW QUESTION 81

- (Exam Topic 8)

You discover a sudden increase in processor utilization on a server that has SQL Server installed. You need to correlate server performance and database activity for an extended time period. Which two tools should you use? Each correct answer presents part of the solution.

- A. Activity Monitor
- B. Performance Monitor
- C. SQL Server Profiler
- D. sp_who2
- E. SQL Server Extended Events

Answer: BE

Explanation:

The Performance Monitor side, we have a few SQL Server monitoring tools AKA counters that can be used when troubleshooting CPU performance. The following counters are simple and easy to use:

Processor % Processor Time == < 80% Processor % User Time == < 80% Processor % Privileged time) == < 30% References:

<https://www.sqlshack.com/sql-server-monitoring-tool-for-cpu-performance/>

NEW QUESTION 86

- (Exam Topic 8)

You use the following statement to create a table.

```
CREATE TABLE Employee
(EmployeeID INT PRIMARY KEY IDENTITY(1,1),
LastName varchar(50),
FirstName varchar(50),
DepartmentId INT,
SupervisorId INT,
OfficeId INT,
Address1 varchar(50),
Address2 varchar(50),
City varchar(50),
State char(2),
PostalCode varchar(10),
Country char(2))
```

You have the following queries.

```
SELECT FirstName, LastName, EmployeeId, OfficeID, DepartmentID
FROM Employee
WHERE FirstName = 'Ben' AND LastName = 'Smith'

SELECT FirstName, LastName, EmployeeId, OfficeID, DepartmentID
FROM Employee
WHERE LastName = 'Smith'
```

You need to create an index to minimize the execution time of the queries.

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

Code Elements

DepartmentId

EmployeeId

FirstName

LastName

OfficeId

Answer Area

```
CREATE INDEX IX_Index3 ON dbo.Employee
(
  Code element , Code element
)
INCLUDE (
  Code element , Code element
)
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: LastName

Redesign nonclustered indexes with a large index key size so that only columns used for searching and lookups are key columns. Make all other columns that cover the query into nonkey columns. In this way, you will have all columns needed to cover the query, but the index key itself is small and efficient.

Box 2: FirstName

Box 3: DepartmentID

Non-key columns, called included columns, can be added to the leaf level of a nonclustered index to improve query performance by covering the query. That is, all columns referenced in the query are included in the index as either key or non-key columns. This allows the query optimizer to locate all the required information from an index scan; the table or clustered index data is not accessed.

Box 4: OfficeID

NEW QUESTION 90

- (Exam Topic 8)

Your network contains a server that has SQL Server 2014 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 filtered index
- C. A columnstore index
- D. A CHECK constraint

Answer: A

NEW QUESTION 93

- (Exam Topic 8)

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.)

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

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?

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

Answer: A

Explanation:

The question concerning the view that has a clause "WHERE Column1 = 'City1' is wrong. That's not what the CHECK option is made for. Actually you will be able to update ONLY the rows satisfied by that WHERE clause, that is, only the rows with the Column1 being 'City1'.

None of the answers are valid from that question. You need a trigger to achieve that.

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

NEW QUESTION 96

- (Exam Topic 8)

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.

Switch Table2 to Table1.

Create a table named Table2 by using the same schema as Table1.

Partition Table1.

Import the data to Table2.

Import the data to Table1.

Create a columnstore index on Table2 for Column1.

Create the columnstore index on Table1.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Create a table named Table2 by using the same schema as Table1. Note: Table2 is the staging table.

Box 2: Partition Table1

Box 3: Import the data to Table2.

Box 4: Create a columnstore index on Table2 for Column1. Box 5: Switch Table2 to Table1

Note:

* An xVelocity memory optimized columnstore index, groups and stores data for each column and then joins all the columns to complete the whole index.

Columnstore indexes can transform the data warehousing experience for users by enabling faster performance for common data warehousing queries such as filtering, aggregating, grouping, and star-join queries.

* Tables that have a columnstore index cannot be updated. There are three ways to work around this problem.

A) To update a table with a columnstore index, drop the columnstore index, perform any required INSERT, DELETE, UPDATE, or MERGE operations, and then rebuild the columnstore index.

B) (applies in this scenario) Partition the table and switch partitions. For a bulk insert, insert data into a staging table, build a columnstore index on the staging table, and then switch the staging table into an empty partition. For other updates, switch a partition out of the main table into a staging table, disable or drop the columnstore index on the staging table, perform the update operations, rebuild or re-create the columnstore index on the staging table, and then switch the staging table back into the main table.

C) Place static data into a main table with a columnstore index, and put new data and recent data likely to change, into a separate table with the same schema that does not have a columnstore index.

Reference: Best Practices: Updating Data in a Columnstore Index

NEW QUESTION 99

- (Exam Topic 8)

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

Answer: C

NEW QUESTION 102

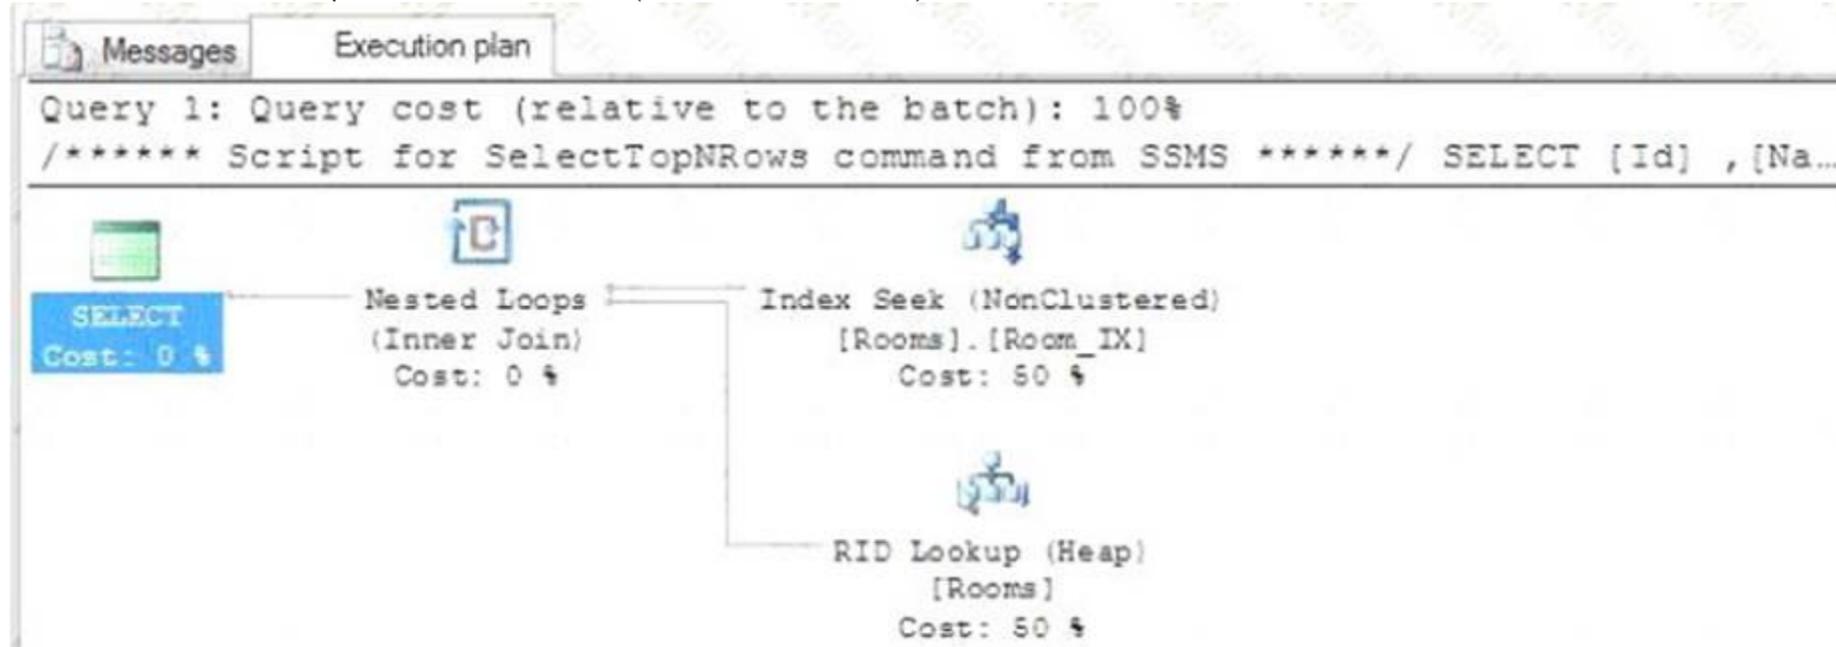
- (Exam Topic 8)

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. (Click the Exhibit button.)



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

Answer: B

NEW QUESTION 103

- (Exam Topic 8)

You have a Microsoft SQL Azure database that contains a table named Customers.

You have a table-valued function named TopCustomers that returns a list of all the customers that have purchased items during the last 12 months. The ID of the customer is passed as an argument to the TopCustomers function.

You need to create a query that returns a list of all the Customer names and the purchase dates.

The solution must return only customers that have purchased an item during the last 12 months. What should you add to the query?

- A. OUTER JOIN
- B. CROSS JOIN
- C. CROSS APPLY
- D. OUTER APPLY

Answer: C

NEW QUESTION 106

- (Exam Topic 8)

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

- ▶ Procedure1 must take a parameter that ensures that multiple rows of data can pass into the stored procedure.
- ▶ Procedure2 must use business logic that resides in a Microsoft .NET Framework assembly. You need to identify the appropriate technology for each stored procedure.

Which technologies should you identify?

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

Technologies	Answer Area
Common language runtime (CLR)	Procedure1 Technology
Extensible Markup Language (XML)	Procedure2 Technology
a table-valued parameter (TVP)	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

<http://msdn.microsoft.com/en-us/library/ms131102.aspx> <http://msdn.microsoft.com/en-us/library/bb522446.aspx> <http://msdn.microsoft.com/en-us/library/bb510489.aspx>

NEW QUESTION 109

- (Exam Topic 8)

You have two existing tables, one named COUNTRY and the other named STATES. The tables are defined as follows:

```
CREATE TABLE COUNTRY
(
Country_Abbr CHAR(3) PRIMARY KEY CLUSTERED,
Country_Description VARCHAR(30) Not Null
)
CREATE TABLE STATES
(
State_Abbr CHAR(2) PRIMARY KEY CLUSTERED,
State_Description VARCHAR(30) Not Null,
Country_Abbr CHAR(3) Not Null
)
```

You need to set up a rule that every STATE.Country_Abbr must match an existing record in the COUNTRY table. Develop the solution by selecting and arranging the required code blocks in the correct order. You may not need all of the code blocks.

Code Blocks	Answer Area
REFERENCES STATES (Country_Abbr)	
REFERENCES COUNTRY (Country_Abbr)	
ON STATES	
FOREIGN KEY (Country_Abbr)	
ON COUNTRY	
ADD CONSTRAINT FK_StateCountry	
ON COUNTRY_ABBR	
ALTER TABLE COUNTRY	
ADD FOREIGN KEY FK_StateCountry	
ALTER TABLE STATES	

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Note:

To allow naming of a FOREIGN KEY constraint, and for defining a FOREIGN KEY constraint on multiple columns, use the following SQL syntax:

MySQL / SQL Server / Oracle / MS Access:

```
ALTER TABLE Orders ADD CONSTRAINT fk_PerOrders FOREIGN KEY (P_Id) REFERENCES Persons(P_Id)
```

NEW QUESTION 110

- (Exam Topic 8)

You are a SQL Server 2014 Developer. A database that you work on contains two tables that are defined as follows:

```
CREATE TABLE Product (
ProductID int IDENTITY(1,1) PRIMARY KEY,
ProductName varchar(30) NOT NULL,
LastUpdatedDate smalldatetime,
LastUpdatedBy varchar(128))

CREATE TABLE ProductAudit (
ProductAuditID int IDENTITY(1,1) PRIMARY KEY,
OldProductID int NOT NULL,
OldProductName varchar(30) NOT NULL,
UpdatedDate smalldatetime,
UpdatedBy varchar(128))
```

Product is an important table that has sensitive audit requirements. You need to create a trigger that supports the following requirements:

1. Every row that is inserted or updated in Product will reflect its actual LastUpdatedDate and LastUpdatedBy values in the Product table.
 2. Any row that is updated or deleted must write a new record reflecting the OLD values into the ProductAudit table.
 3. Any error that occurs during the course of the trigger's execution must prevent the changes from happening. Develop the solution by selecting and arranging the required code blocks in the correct order.
- You may not need all of the code blocks.

Code Blocks	Answer Area
<pre>DECLARE @OldProductId int, @OldProductName varchar (30) SELECT @OldProductId = ProductId, @OldProductName = ProductName FROM deleted INSERT ProductAudit (OldProductID, OldProductName, UpdatedDate, Update dBy) SELECT @OldProductID, @OldProductName, SUSER_NAME (), GETDATE() UPDATE Product SET LastUpdatedBy = SUSER_NAME(), LastUpdatedDate = GETDATE() FROM Product AS p INNER JOIN inserted AS i ON p.ProductID = i.ProductID UPDATE Product SET LastUpdatedBy = SUSER_NAME(), LastUpdatedDate = GETDATE() FROM Product AS p INNER JOIN inserted AS i ON p.ProductID = i.ProductID INSERT ProductAudit (OldProductID, OldProductName, UpdatedDate, Update dBy) SELECT d.ProductID, d.ProductName, SUSER_NAME (), GETDATE() FROM deleted AS d END COMMIT TRANSACTION IF @@ERROR <> 0 ROLLBACK CREATE TRIGGER ProductAuditTrigger ON Product FOR INSERT, UPDATE, DELETE AS BEGIN</pre>	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Note:

* Executing a ROLLBACK TRANSACTION or COMMIT TRANSACTION Transact-SQL statement inside a stored procedure or trigger is possible, but doing so may cause errors.

NEW QUESTION 112

- (Exam Topic 8)

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 NONCLUSTERED INDEX IX1 on Schema1.Table1 (Column1) INCLUDE (Column4) WHERE Column2 <> Column3
- B. CREATE NONCLUSTERED INDEX IX1 on Schema1.Table1 (Column1)
- C. CREATE NONCLUSTERED INDEX IX1 on Schema1.Table1 (Column1, Column2, Column3) INCLUDE (Column4)
- D. CREATE NONCLUSTERED INDEX IX1 on schema1.Table1 (Column1) INCLUDE (Column4)

Answer: C

NEW QUESTION 117

- (Exam Topic 8)

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.)

Code Segments	Answer Area			
ALTER SCHEMA	Code	Code	Code	Code
ALTER TABLE				
dbo				
dbo.Customers				
EXEC sp_rename				
TRANSFER				
Schema2				

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

<http://msdn.microsoft.com/en-us/library/ms173423.aspx>

NEW QUESTION 118

- (Exam Topic 8)

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;
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Note:

* Box 1 – Box 2:

/ You need to move the data from the first partition to a new table named SalesHistory.

/ First create the new table, then move the contents of the first partition.

*(Box 3 Box 4) Drop the partition scheme and then the partition function and the recreate them (box 5-box6). First recreate the partition function.

/You need, starting on December 31, 2012, repartition dbo.Sales to support new sales data for three months.

/ A partition function can be dropped only if there are no partition schemes currently using the partition function. If there are partition schemes using the partition function, DROP PARTITION FUNCTION returns an error.

NEW QUESTION 122

- (Exam Topic 8)

You plan to deploy two stored procedures name USP_1 and USP_2 that read data from a database. Your company identifies the following requirements for each stored procedure:

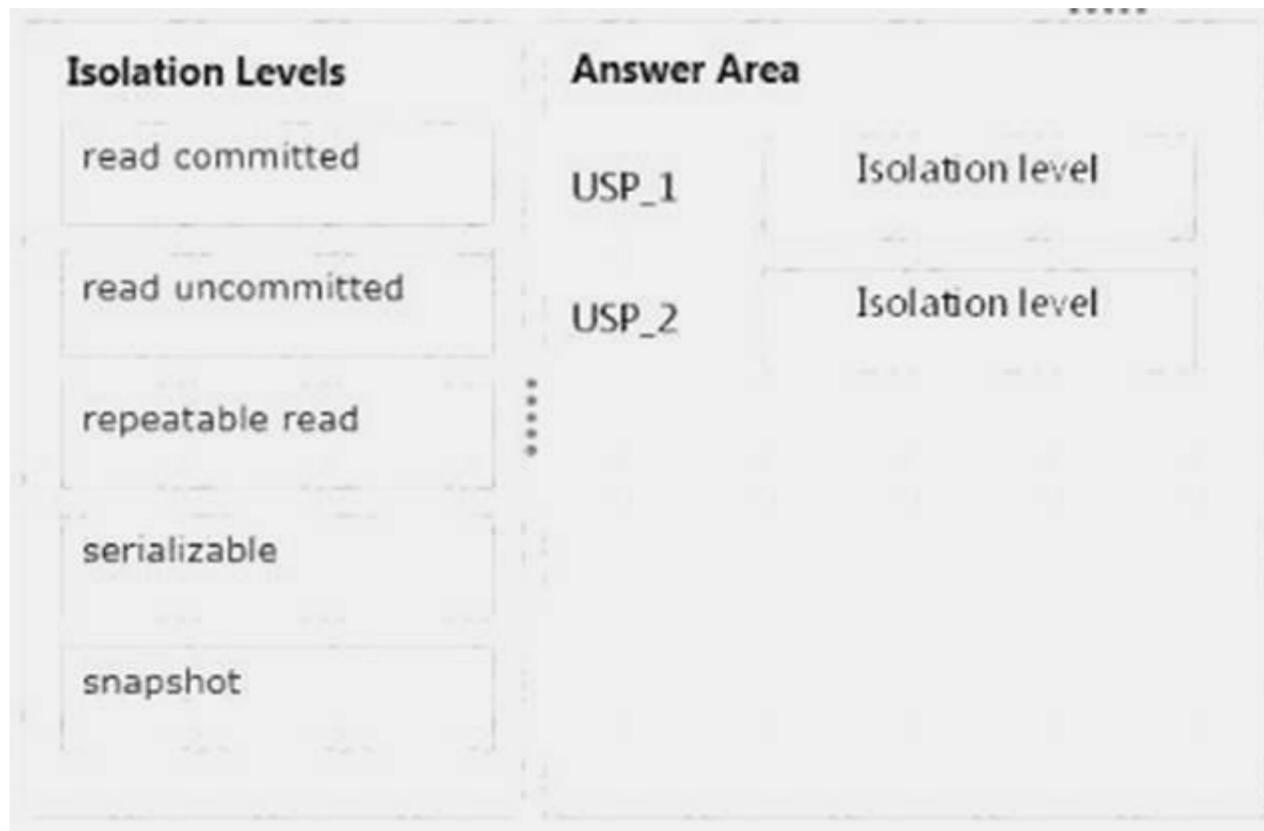
▶ USP_1 cannot allow dirty reads.

▶ USP_2 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.)



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

* read committed READ COMMITTED

Specifies that shared locks are held while the data is being read to avoid dirty reads, but the data can be changed before the end of the transaction, resulting in nonrepeatable reads or phantom data. This option is the SQL Server default.

* SERIALIZABLE

Places a range lock on the data set, preventing other users from updating or inserting rows into the data set until the transaction is complete. This is the most restrictive of the four isolation levels. Because concurrency is lower, use this option only when necessary. This option has the same effect as setting HOLDLOCK on all tables in all SELECT statements in a transaction.

NEW QUESTION 126

- (Exam Topic 8)

You have a text file that contains an XML Schema Definition (XSD).

You have a table named Schema1.Table1.

You have a stored procedure named Schema1.Proc1 that accepts an XML parameter named Param1.

You need to store validated XML data in Schema1.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>

NEW QUESTION 128

- (Exam Topic 8)

You have the following stored procedure.

```
CREATE PROCEDURE GetFile
    @FileName nvarchar(512)
AS
SELECT *
FROM Files
WHERE FileName = @FileName
```

The stored procedure takes much longer to execute than expected.

While reviewing the execution plan of the stored procedure, you discover the following predicate for a Clustered Index Scan operator.

```
CONVERT_IMPLICIT(nvarchar(512), [SampleDatabase].[dbo].[Files].[FileName], 0)=[@1]
```

You need to resolve the performance issue. What should you do?

- A. Change the @FileName parameter from nvarchar(512) to varchar(512).
- B. Change the FileName column from varchar(512) to nvarchar(512).
- C. Add a NOLOCK query hint to the SELECT statement.
- D. Convert the table to a memory-optimized table.
- E. Add a FORCESEEK query hint to the SELECT statement.

Answer: A

Explanation:

When using a variable, make sure that the datatype matches the column's datatype. We suspect that the issue is that the variable is NVARCHAR (512) whilst the table column is VARCHAR (512). This is indicated by the CONVERT_IMPLICIT operator in the execution plan.

References:

https://sqlserverperformance.wordpress.com/2009/02/02/beware-of-convert_implicit-in-a-sql-execution-plan/

NEW QUESTION 132

- (Exam Topic 8)

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

Answer: C

NEW QUESTION 133

- (Exam Topic 8)

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.)

```
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 )
```

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?

- 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

NEW QUESTION 135

- (Exam Topic 8)

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

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

Answer: C

NEW QUESTION 140

- (Exam Topic 8)

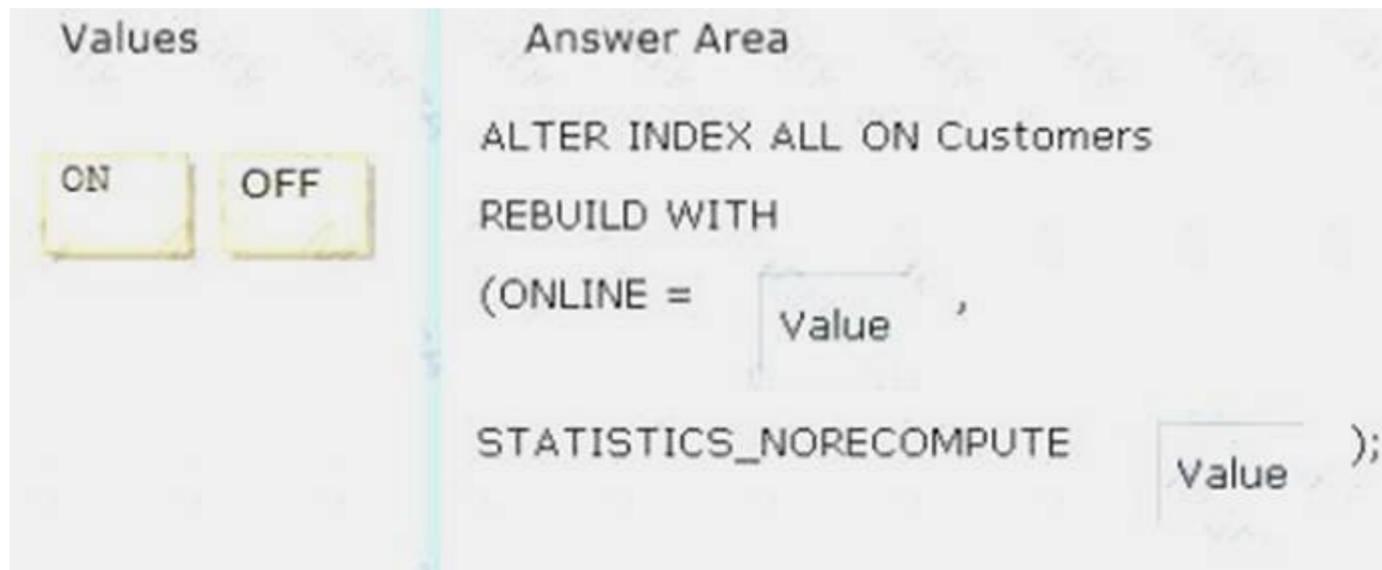
You run the following code segment:

```
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.)



- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Note:

Locking the table during the process and not recomputing statistics would be the fastest.

* Online = OFF

Table locks are applied for the duration of the index operation. An offline index operation that creates, rebuilds, or drops a clustered, spatial, or XML index, or rebuilds or drops a nonclustered index, acquires a Schema modification (Sch-M) lock on the table. This prevents all user access to the underlying table for the duration of the operation. An offline index operation that creates a nonclustered index acquires a Shared (S) lock on the table. This prevents updates to the underlying table but allows read operations, such as SELECT statements.

* STATISTICS_NORECOMPUTE = ON

Out-of-date statistics are not automatically recomputed. Reference: ALTER INDEX (Transact-SQL)

NEW QUESTION 144

- (Exam Topic 8)

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

Answer: C

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.

NEW QUESTION 147

- (Exam Topic 8)

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 identify?

To answer, drag the appropriate encryption type and data type to the correct identifier in the answer area.

Encryption Types

- Transparent data encryption (TDE)
- Encrypting File System (EFS)
- Cell-level encryption
- BitLocker Drive Encryption (BitLocker)

Data Types

- decimal
- varchar
- varbinary
- money

Answer Area

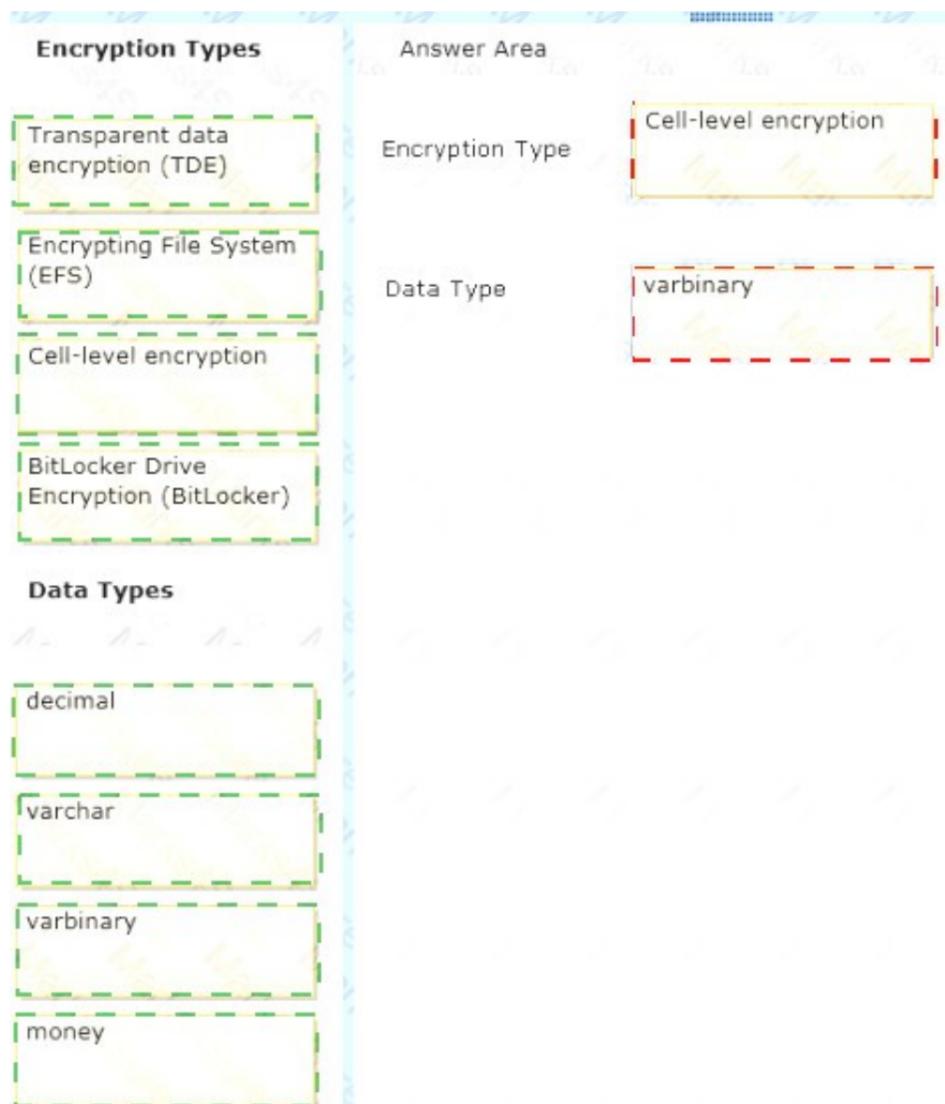
Encryption Type

Data Type

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:



NEW QUESTION 148

- (Exam Topic 8)

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

Answer: A

NEW QUESTION 150

- (Exam Topic 8)

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

Answer: C

NEW QUESTION 151

- (Exam Topic 8)

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. ROLLBACK DISTRIBUTED TRANSACTION @transactionID
- C. BEGIN TRANSACTION WITH MARK @transactionID
- D. COMMIT TRANSACTION @transactionID

Answer: C

NEW QUESTION 152

- (Exam Topic 8)

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.)

Functions	Answer Area
sys.dm_db_index_usage_stats	Identify unused indexes.
sys.dm_db_index_operational_stats	Identify indexes that need to be defragmented.
sys.dm_db_index_physical_stats	
sys.dm_db_missing_index_columns	
sys.dm_db_missing_index_details	
sys.dm_db_missing_index_groups	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Note:

* sys.dm_db_index_usage_stats

Returns counts of different types of index operations and the time each type of operation was last performed.

* sys.dm_db_index_physical_stats

Returns size and fragmentation information for the data and indexes of the specified table or view.

NEW QUESTION 157

- (Exam Topic 8)

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. Create an INSTEAD OF trigger on the view.
- B. Define the view by using the SCHEMABINDING option.
- C. Define the view by using the CHECK option.
- 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>

NEW QUESTION 162

- (Exam Topic 8)

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

NEW QUESTION 163

- (Exam Topic 8)

You plan to create a new column in a disk-based 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. FileStream
B. nvarchar
C. image
D. varbinary

Answer: A

NEW QUESTION 167

- (Exam Topic 8)

You plan to create a custom aggregation function named Median.

You plan to deploy Median to a SQL Server 2014 server named Server1.

You need to ensure that Median can access a web service named WebApp1. The solution must minimize the number of changes made to the database.

You create a Microsoft .NET Framework class that contains the definition of Median. You upload a certificate to Server1.

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.

Actions	Answer Area
Execute the CREATE AGGREGATE statement.	
Modify the TRUSTWORTHY property of the database.	
Execute the CREATE ASSEMBLY statement.	
Execute the CREATE FUNCTION statement.	
Use the certificate to add a digital signature to the assembly.	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Execute the CREATE ASSEMBLY statement.

Box 2: Use the certificate to add a digital signature to the assembly. Box 3: Execute the CREATE AGGREGATE statement.

Note:

* CREATE AGGREGATE

Creates a user-defined aggregate function whose implementation is defined in a class of an assembly in the .NET Framework. For the Database Engine to bind the aggregate function to its implementation, the .NET Framework assembly that contains the implementation must first be uploaded into an instance of SQL Server by using a CREATE ASSEMBLY statement.

NEW QUESTION 172

- (Exam Topic 8)

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

Answer: A

Explanation:

<http://msdn.microsoft.com/en-us/library/ms187956.aspx> <http://msdn.microsoft.com/en-us/library/ms191432.aspx>

NEW QUESTION 176

- (Exam Topic 8)

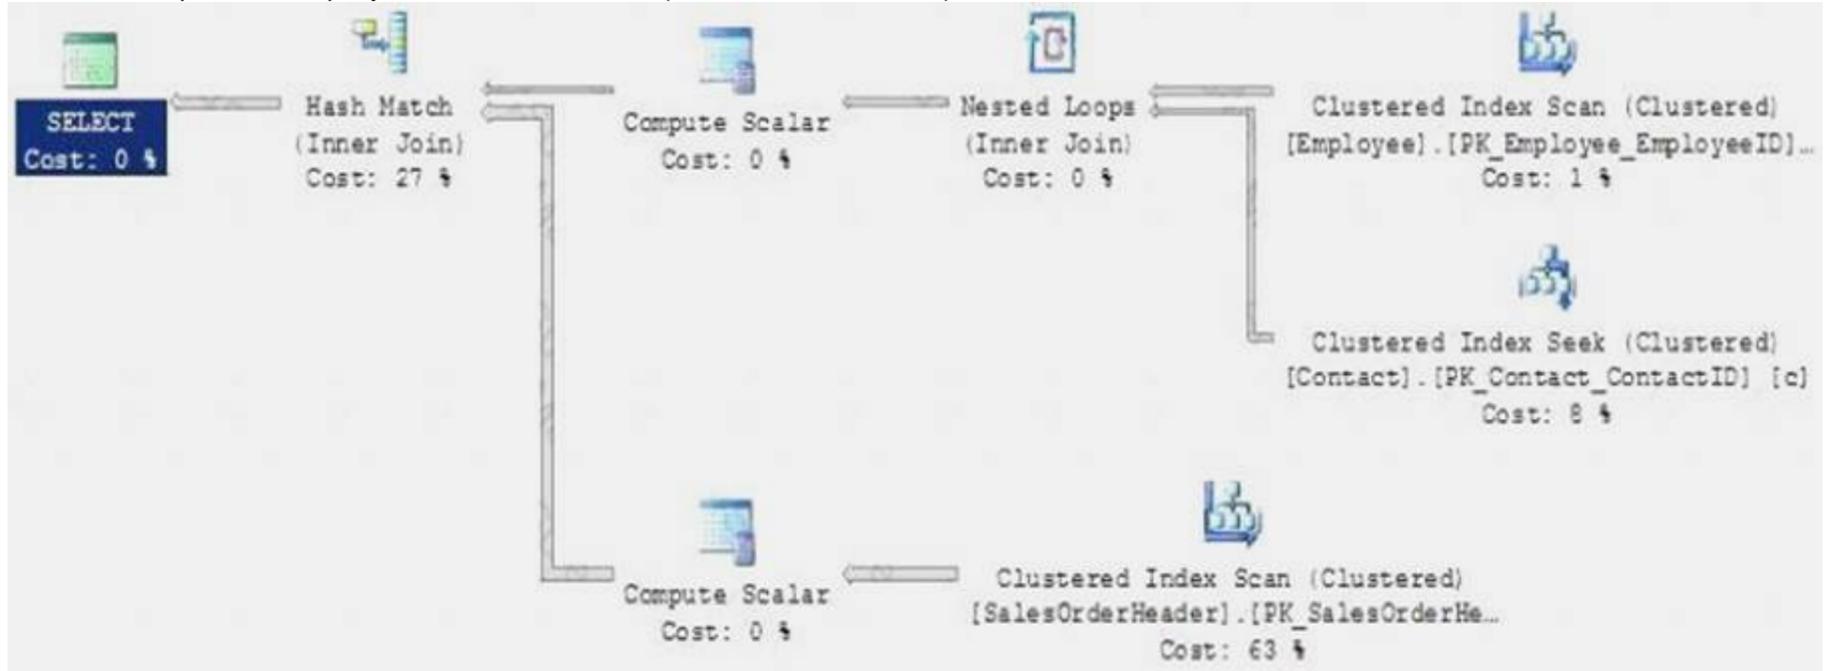
You have a database that contains three tables. The tables are configured as shown in the following table.

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. (Click the Exhibit button.)



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.)

Columns	Answer Area
Contact.EmailAddress	Indexed Columns Column
Contact.FirstName	Included Columns Column
Contact.LastName	Column
Employee.Title	
SalesOrderHeader.OrderDate	
SalesOrderHeader.SalesPersonID	
SalesOrderHeader.SubTotal	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Note:

Covering index: A type of index that includes all the columns that are needed to process a particular query. For example, your query might retrieve the FirstName and LastName columns from a table, based on a value in the ContactID column. You can create a covering index that includes all three columns.

NEW QUESTION 177

- (Exam Topic 8)

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	<ul style="list-style-type: none"> Table data will not be updated. The table will contain historical calculations. The table will contain 10 million records.
Table2	<ul style="list-style-type: none"> 20% of the table data will be updated weekly. The table will contain 25 million records.
Table3	<ul style="list-style-type: none"> 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.

Index Types	Answer Area
Columnstore Index	Table1
Nonclustered Index	Table2
	Table3

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Index Types	Answer Area
Columnstore Index	Table1
Nonclustered Index	Table2
	Table3

NEW QUESTION 180

- (Exam Topic 8)

The database contains a disk-based table named ContentTable that has 1 million rows and a column named Fax. Fax allows null values. You need to update Fax to meet the following requirements:

- Prevent null values from being used.
- Always use an empty string instead of a null value.

Which statement or statements should you execute? (Each correct answer presents part of the solution. Choose all that apply.)

- A. ALTER TABLE dbo.ContentTable
ADD CONSTRAINT
DF_ContentTable_Fax
DEFAULT '' FOR Fax
- B. ALTER TABLE ContentTable
ALTER COLUMN
Fax char(10) NOT NULL
- C. ALTER TABLE ContentTable
ADD COLUMN
Fax char(10) NOT NULL
- D. ALTER TABLE ContentTable
DROP COLUMN Fax
- E. UPDATE ContentTable
SET Fax = ''
WHERE Fax IS NULL

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

Answer: ABE

Explanation:

E: First change the NULLs to ''.

A: Then set the default to the column to ''.

B: Finally add the NOT NULL constraint to the column.

NEW QUESTION 181

- (Exam Topic 8)

You are designing a new database table that will be used for reporting. You define the table by using the following statement.

```
CREATE TABLE Customers
(CustomerId BIGINT,
LastName VARCHAR(255),
FirstName VARCHAR(255),
SalesPersonId INT,
Address1 VARCHAR(255),
Address2 VARCHAR(255),
City VARCHAR(255),
State CHAR(2),
PostalCode VARCHAR(10),
Country CHAR(2))
```

You need to store the data in the table by using the least amount of storage space possible. Which storage option should you use?

- A. a clustered index
- B. a clustered columnstore index
- C. a nonclustered index
- D. In-Memory OLTP

Answer: B

Explanation:

Columnstore indexes work well for mostly read-only queries that perform analysis on large data sets. This would fit this scenario as the table will be used for reporting.

Columnstore Index benefits include Columnstore Index benefits high compression rates, which improve query performance by using a smaller in-memory footprint.

In turn, query performance can improve because SQL Server can perform more query and data operations in-memory.

Use the columnstore index to achieve up to 10x query performance gains over traditional row-oriented storage, and up to 7x data compression over the uncompressed data size.

References: [https://msdn.microsoft.com/en-us/library/gg492088\(v=sql.120\).aspx](https://msdn.microsoft.com/en-us/library/gg492088(v=sql.120).aspx)

NEW QUESTION 183

- (Exam Topic 8)

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

NEW QUESTION 187

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual 70-464 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the 70-464 Product From:

<https://www.2passeasy.com/dumps/70-464/>

Money Back Guarantee

70-464 Practice Exam Features:

- * 70-464 Questions and Answers Updated Frequently
- * 70-464 Practice Questions Verified by Expert Senior Certified Staff
- * 70-464 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * 70-464 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year