

## 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 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.

Ordered List Title	Answer Choices Title
	<pre> RECONFIGURE; EXEC sp_configure 'clr enabled', '1'; EXEC sp_recomple @objname = 'TaxCalc' CREATE PROCEDURE Accounting.Amortize(@total decimal(8,2), @period int) RETURNS decimal(8,2) AS EXTERNAL NAME TaxCalc.TreyResearch.Amortize; CREATE ASSEMBLY TaxCalc FROM 'C:\temp\TreyTax.DLL' CREATE ASSEMBLY TaxCalc FROM 'C:\temp\Amortize.cs' </pre>
<p>&lt;&lt; Move Remove &gt;&gt;</p>	

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

D:\Documents and Settings\useralbo\Desktop\1.jpg  
<http://msdn.microsoft.com/en-us/library/ms131089.aspx> <http://msdn.microsoft.com/en-us/library/ms131048.aspx> <http://msdn.microsoft.com/en-us/library/ms187926.aspx>

**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 2)

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:SET TRANSACTION ISOLATION LEVEL REPEATABLE READ  
C. Move the SELECT statement at line 49 to line 57.  
D. Move the SET statement at line 46 to line 53.

**Answer: A**

#### 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 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 CONSTRAINT DF\_FullName DEFAULT (dbo.CreateFullName (FirstName, LastName)),  
B. FullName AS (FirstName + ' ' + LastName),  
C. FullName nvarchar(100) NOT NULL DEFAULT (dbo.CreateFullName (FirstName, LastName)).  
D. FullName AS (FirstName + ' ' + LastName) PERSISTED,

**Answer:** D

**Explanation:**

<http://msdn.microsoft.com/en-us/library/ms188300.aspx> <http://msdn.microsoft.com/en-us/library/ms191250.aspx>

**NEW QUESTION 10**

- (Exam Topic 2)

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

#### NEW QUESTION 10

- (Exam Topic 2)

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

#### NEW QUESTION 12

- (Exam Topic 3)

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 TYPE EmployeesInfo AS Table  
C. CREATE SCHEMA EmployeesInfo  
D. CREATE TABLE EmployeesInfo

**Answer:** B

#### Explanation:

Example Usage of Table-Valued Parameters (Database Engine)

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

(Benefits of using Table-Valued Parameters)

```
/* Create a table type. */
```

```
CREATE TYPE LocationTableType AS TABLE ( LocationName VARCHAR(50)  
, CostRate INT ); GO
```

```
/* Create a procedure to receive data for the table-valued parameter. */ CREATE PROCEDURE dbo. usp_InsertProductionLocation
```

```
@TVP LocationTableType READONLY AS
```

```
SET NOCOUNT ON
```

```
INSERT INTO AdventureWorks2012.Production.Location (Name
```

```
,CostRate
```

```
,Availability
```

```
,ModifiedDate)  
SELECT *, 0, GETDATE() FROM @TVP;  
GO
```

Also:

<http://msdn.microsoft.com/en-us/library/ms175007.aspx>(CREATE TYPE \*tabletypename\* AS TABLE)

<http://msdn.microsoft.com/en-us/library/ms175010.aspx>(table data types)

Wrong Answers:

<http://msdn.microsoft.com/en-us/library/ms174979.aspx>(CREATE TABLE) <http://msdn.microsoft.com/en-us/library/ms189462.aspx>(CREATE SCHEMA)

<http://msdn.microsoft.com/en-us/library/ms176009.aspx>(CREATE XML SCHEMA COLLECTION)

#### NEW QUESTION 15

- (Exam Topic 3)

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. A table variable
- B. The ROWNUMBER keyword
- C. An OFFSET-FETCH clause
- D. A recursive common table expression

**Answer: C**

#### Explanation:

<http://www.mssqltips.com/sqlservertip/2696/comparing-performance-for-different-sql-serverpaging-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-byclause/>

[http://www.sqlservercentral.com/blogs/juggling\\_with\\_sql/2011/11/30/using-offset-and-fetch/](http://www.sqlservercentral.com/blogs/juggling_with_sql/2011/11/30/using-offset-and-fetch/)

#### NEW QUESTION 17

- (Exam Topic 3)

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

**Answer: B**

#### Explanation:

<http://msdn.microsoft.com/en-us/library/ms181714.aspx>

#### NEW QUESTION 22

- (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 24

- (Exam Topic 4)

You need to convert the functionality of Dynamic.sql to use a stored procedure. Which Transact SQL statement should you add to the stored procedure contain?

A. CREATE PROC Production.ProductsAfterDate (@ProductID AS varchar(11), @CreationDate AS date)AS...

B. CREATE PROC Production.ProductsAfterDate(@sqlstring AS nvarchar(1000)) AS...

C. CREATE PROC Production.ProductsAfterDate(@sqlstring AS nvarchar(1000), OUTPUT @ProductID AS varchar(11), OUTPUT @CreationDate AS date)AS...

D. CREATE PROC Production.ProductsAfterDate(@sqlstring AS nvarchar(1000),@ProductID AS varchar(11),@CreationDate AS date) AS...

**Answer: C**

#### Explanation:

@sqlstring, @ProductID, and @CreationData need to be declared as parameters. @ProductID, and @CreationData should be output parameters.

#### NEW QUESTION 26

- (Exam Topic 4)

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

**NEW QUESTION 27**

- (Exam Topic 4)

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

 B. EXEC sp_recompile @objname = 'Production.ProcessProduct';

 C. RECONFIGURE;

 D. Exec SP_CONFIGURE 'clr enabled', '1';

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

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

 G. CREATE TYPE Production.ProcessProduct
EXTERNAL NAME ProductionAssembly.ProcessProducts.Process;

```

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

**Answer:** ACDE

**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>

**NEW QUESTION 30**

- (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 34

- (Exam Topic 5)

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:

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.

#### NEW QUESTION 35

- (Exam Topic 5)

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\_reassign\_proxy, sp\_add\_job, sp\_add\_jobstep, and sp\_grant\_login\_to\_proxy system stored procedures.
- B. The CREATE CREDENTIAL command and the sp\_add\_proxy, sp\_add\_job, sp\_add\_jobstep, and sp\_grant\_proxy\_to\_subsystem system stored procedures.
- C. The CREATE CREDENTIAL command and the sp\_reassign\_proxy, sp\_add\_job, sp\_add\_jobstep, and sp\_grant\_login\_to\_proxy system stored procedures.
- D. 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.

**Answer:** B

#### NEW QUESTION 39

- (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 40

- (Exam Topic 5)

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 (READPAST)
- D. Change line 09 in usp\_GetOpenings to: FROM Openings o (NOLOCK)

**Answer:** D

#### NEW QUESTION 45

- (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 48

- (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 51

- (Exam Topic 6)

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

### NEW QUESTION 55

- (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 57

- (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 61**

- (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 66**

- (Exam Topic 6)

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

**Answer:** D

**NEW QUESTION 71**

- (Exam Topic 6)

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. BATCHSIZE=0.
- C. BATCHSIZE=1000.
- D. LASTROW = 1000.

**Answer:** C

**NEW QUESTION 76**

- (Exam Topic 6)

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. Execute sp\_recompile 'usp\_GetOrdersByProduct'.
- C. Add WITH RECOMPILE to line 03 in usp\_GetOrdersByProduct.
- D. Add WITH (FORCESEEK) to line 07 in usp\_GetOrdersByProduct.

**Answer:** C

**Explanation:**

Ref: [http://msdn.microsoft.com/en-us/librAry/ms190439\(v=sql.90\).aspx](http://msdn.microsoft.com/en-us/librAry/ms190439(v=sql.90).aspx)

**NEW QUESTION 80**

- (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 82**

- (Exam Topic 6)

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

- A. ALTER OrdersADD originalOrder XML (ValidateOrder);
- B. ALTER OrdersADD originalOrder XML;
- C. ALTER OrdersADD originalOrder varchar(max);
- D. ALTER OrdersADD originalOrder varbinary(max);

Answer: D

**NEW QUESTION 86**

- (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 87**

- (Exam Topic 7)

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.

| Actions                                                | Answer Area |
|--------------------------------------------------------|-------------|
| Create a nonclustered hash index.                      |             |
| Create a clustered columnstore index.                  |             |
| Create a partitioning scheme for use by the table.     |             |
| Use an ALTER INDEX REBUILD on a specific partition.    |             |
| Use an ALTER INDEX REORGANIZE on a specific partition. |             |

- A. Mastered
- B. Not Mastered

Answer: A

**Explanation:**

Box 1:

Create a clustered columnstore index.

Box 2:

Create a partitioning scheme for use by the table.

Box 3:

Use an ALTER INDEX REORGANIZE on a specific partition.

Note:

\* Creating a partitioned table or index typically happens in four parts:

- ▶ Create a filegroup or filegroups and corresponding files that will hold the partitions specified by the partition scheme.
- ▶ Create a partition function that maps the rows of a table or index into partitions based on the values of a specified column.
- ▶ Create a partition scheme that maps the partitions of a partitioned table or index to the new filegroups.
- ▶ Create or modify a table or index and specify the partition scheme as the storage location.

\* Reorganizing an index uses minimal system resources.

\* From scenario:

/ The index maintenance strategy for the UserActivity table must provide the optimal structure for both maintainability and query performance.

/ 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.  
/ 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.  
\* Columnstore indexes work well for mostly read-only queries that perform analysis on large data sets. Often, these are queries for data warehousing workloads. Columnstore indexes give high performance gains for queries that use full table scans, and are not well-suited for queries that seek into the data, searching for a particular value.

**NEW QUESTION 90**

- (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 POStTransaction (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 POStTransaction 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 usp\_DetectSuspiciousActivity.

## NEW QUESTION 95

- (Exam Topic 7)

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                                                                                                                                              | Answer Area |
|----------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|
| <pre>DELAYED_DURABILITY = ON , TRANSACTION ISOLATION LEVEL = SNAPSHOT</pre>                                                                              |             |
| <pre>CREATE PROCEDURE dbo.usp_AddMobileLocation @POSTransactionId int, @Long float, @Lat float WITH</pre>                                                |             |
| <pre>NATIVE_COMPILATION ...</pre>                                                                                                                        |             |
| <pre>DELAYED_DURABILITY = OFF , TRANSACTION ISOLATION LEVEL = READ UNCOMMITTED</pre>                                                                     |             |
| <pre>DELAYED_DURABILITY = ON , TRANSACTION ISOLATION LEVEL = READ UNCOMMITTED</pre>                                                                      |             |
| <pre>Insert into dbo.MobileLocation ( POSTransactionId, Longitude, Latitude, CreateDate ) VALUES ( @POSTransactionId, @Long, @Lat, GETDATE() ) END</pre> |             |
| <pre>, LANGUAGE = N'English' )</pre>                                                                                                                     |             |
| <pre>AS BEGIN ATOMIC WITH (</pre>                                                                                                                        |             |
| <pre>DELAYED_DURABILITY = OFF , TRANSACTION ISOLATION LEVEL = SNAPSHOT</pre>                                                                             |             |

- A. Mastered
- B. Not Mastered

Answer: A

**Explanation:**

Box 1:

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

Box 2:

```
NATIVE_COMPILATION
...
```

Box 3:

```
AS
BEGIN ATOMIC WITH
(
```

Box 4:

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

Box 5:

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

Box 6:

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

Note:

\* From scenario:

The mobile application will need to meet the following requirements:

• Update the location of the user by using a stored procedure named usp\_AddMobileLocation.

\* DELAYED\_DURABILITY

SQL Server transaction commits can be either fully durable, the SQL Server default, or delayed durable (also known as lazy commit).

Fully durable transaction commits are synchronous and report a commit as successful and return control to the client only after the log records for the transaction are written to disk. Delayed durable transaction commits are asynchronous and report a commit as successful before the log records for the transaction are written to disk. Writing the transaction log entries to disk is required for a transaction to be durable. Delayed durable transactions become durable when the transaction log entries are flushed to disk.

**NEW QUESTION 99**

- (Exam Topic 7)

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.

**NEW QUESTION 100**

- (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 105**

- (Exam Topic 8)

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**

#### NEW QUESTION 109

- (Exam Topic 8)

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 10 MB. 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 Area

```
DBCC SHRINKFILE (ContentDatabase_Log, 10);  
GO
```

```
ALTER DATABASE ContentDatabase  
SET RECOVERY SIMPLE;  
GO
```

```
ALTER DATABASE ContentDatabase  
SET RECOVERY FULL;  
GO
```

```
ALTER DATABASE ContentDatabase  
SET PAGE_VERIFY CHECKSUM;  
GO
```

```
BACKUP LOG ContentDatabase  
WITH TRUNCATE_ONLY
```

```
DBCC SHRINKFILE (ContentDatabase_Log, 7168);  
GO
```

- A. Mastered
- B. Not Mastered

**Answer: A**

#### Explanation:

Box 1:

```
ALTER DATABASE ContentDatabase  
SET RECOVERY SIMPLE;  
GO
```

Box 2:

```
DBCC SHRINKFILE (ContentDatabase_Log, 10);  
GO
```

Box 3:

```
ALTER DATABASE ContentDatabase  
SET RECOVERY FULL;  
GO
```

Note:

\* Shrinking a log file to a specified target size

The following example shrinks the log file in the AdventureWorks database to 1 MB. To allow the DBCC SHRINKFILE command to shrink the file, the file is first truncated by setting the database recovery model to SIMPLE.

Transact-SQL

```
USE AdventureWorks2012; GO
```

```
-- Truncate the log by changing the database recovery model to SIMPLE. ALTER DATABASE AdventureWorks2012
SET RECOVERY SIMPLE; GO
-- Shrink the truncated log file to 1 MB.
DBCC SHRINKFILE (AdventureWorks2012_Log, 1); GO
-- Reset the database recovery model. ALTER DATABASE AdventureWorks2012 SET RECOVERY FULL;
GO
```

\* If the log file does not shrink (after dbcc shrinkfile)

Typically it is the log file that appears not to shrink. This is usually the result of a log file that has not been truncated. You can truncate the log by setting the database recovery model to SIMPLE, or by backing up the log and then running the DBCC SHRINKFILE operation again.

\* DBCC SHRINKFILE shrinks the size of the specified data or log file for the current database, or empties a file by moving the data from the specified file to other files in the same filegroup, allowing the file to be removed from the database.

Arguments include: target\_size

Is the size for the file in megabytes, expressed as an integer.

#### NEW QUESTION 114

- (Exam Topic 8)

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

**Answer:** E

#### Explanation:

<http://www.techrepublic.com/blog/programming-and-development/comparing-sql-serverconstraints-and-dmltrig> <http://msdn.microsoft.com/en-us/library/ms178110.aspx>

#### NEW QUESTION 117

- (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 120

- (Exam Topic 8) 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

#### NEW QUESTION 122

- (Exam Topic 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. Partition the table and use the CustomerName column for the partition scheme.
- B. Replace IX\_CustomerName with a clustered index.
- C. Replace LEFT(CustomerName,1) = 'a' with CustomerName LIKE 'a%'.
- D. Replace LEFT(CustomerName,1) = 'a' with SUBSTRING(CustomerName,1,1) = 'a'.

Answer: C

#### Explanation:

<http://msdn.microsoft.com/en-us/library/ms179859.aspx> <http://msdn.microsoft.com/en-us/library/ms187748.aspx>

#### NEW QUESTION 125

- (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 129**

- (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 133**

- (Exam Topic 8)

You have a Microsoft SQL Azure database named DBAzure1. You create a table in DBAzure1 by using the following script:

```
CREATE TABLE dbo.Customers
(
    CustomerId int NOT NULL,
    CustomerName nvarchar(50) NULL,
    CustomerContact nvarchar(50) NULL,
    CustomerDetails nvarchar(200) NULL,
    CONSTRAINT PK_Customers PRIMARY KEY CLUSTERED (CustomerId)
)
ON [PRIMARY]
GO
```

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

- A. A CHECK constraint
- B. A filtered index
- C. A columnstore index
- D. A UNIQUE constraint

Answer: D

**NEW QUESTION 137**

- (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<br>Technology |
| Extensible Markup Language (XML) | Procedure2<br>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 138**

- (Exam Topic 8)

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.

**NEW QUESTION 139**

- (Exam Topic 8)

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

Answer: A

**NEW QUESTION 144**

- (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 145**

- (Exam Topic 8)

You create a disk-based table that contains the following script:

```
CREATE TABLE dbo.Products
(
    ProductId bigint IDENTITY(1,1),
    Name nvarchar(50) NULL,
    Description nvarchar(max) NULL,
    SKU char(10) NULL,
    CONSTRAINT PK_Products PRIMARY KEY CLUSTERED (ProductId)
) ON [PRIMARY]
GO
```

You need to prevent duplicate values in the SKU field. Which five code segments should you use?  
 To answer, move the appropriate code segments from the list of code segments to the answer area and arrange them in the correct order.

Code segments	Answer Area
(SKU)	
ALTER TABLE dbo.Products	
CHECK	
FOREIGN KEY	
UNIQUE	
ADD CONSTRAINT	
CK_SKU	

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: ALTER TABLE dbo.Products Box 2: ADD CONSTRAINT

Box 3: CK\_SKU

Box 4: UNIQUE

Box 5: (SKU)

Note: The SQL command is:

ALTER TABLE <tablename> ADD CONSTRAINT

<constraintname> UNIQUE (

<columnname>

)

**NEW QUESTION 149**

- (Exam Topic 8)

You have a SQL Server 2012 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

**NEW QUESTION 154**

- (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
ALTER TABLE	Code
dbo	Code
dbo.Customers	Code
EXEC sp_rename	
TRANSFER	
Schema2	

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

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

**NEW QUESTION 159**

- (Exam Topic 8)

You need to identify which nonclustered indexes are unused by queries.

How should you complete the statement? To answer, drag the appropriate values to the correct locations. Each value 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.

### Values

sys.dm\_db\_index\_operational\_stats

sys.dm\_db\_index\_physical\_stats

sys.dm\_db\_index\_usage\_stats

sys.indexes

sys.tables



### Answer Area

```
SELECT database_id, a.object_id, a.index_id, b.name, a.user_seeks, a.user_scans, a.user_updates
FROM [Value] a
join [Value] b on a.object_id = b.object_id
and a.index_id = b.index_id
where a.index_id >= 2
and a.database_id = db_id()
and (a.user_seeks = 0
or a.user_scans = 0)
```

- A. Mastered
- B. Not Mastered

Answer: A

**Explanation:**

Box 1: sys.dm\_db\_index\_usage\_stats

sys.dm\_db\_index\_usage\_stats shows you how many times the index was used for user queries. It returns counts of different types of index operations and the time each type of operation was last performed in SQL Server.

Box 2: sys.indexes

sys.indexes contains a row per index or heap of a tabular object, such as a table, view, or table-valued function.

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/system-dynamic-management-views/sys-dm-db-index->

<https://docs.microsoft.com/en-us/sql/relational-databases/system-catalog-views/sys-indexes-transact-sql>

**NEW QUESTION 160**

- (Exam Topic 8) 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 ON 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 NULL

Answer: D

**Explanation:**

<http://msdn.microsoft.com/en-us/library/ms188783.aspx> <http://msdn.microsoft.com/en-us/library/cc280372.aspx>

**NEW QUESTION 161**

- (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 162**

- (Exam Topic 8)

You have a database that contains a user-defined function named Schema1.Udf1 and two tables named Schema1.Table1 and Schema1.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 Shemal.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
INNER JOIN Schema1.Udf1(t.CountryID) AS u;

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

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

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

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

Answer: B

**NEW QUESTION 164**

- (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:

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.)

Isolation Levels	Answer Area
read committed	USP_1 Isolation level
read uncommitted	USP_2 Isolation level
repeatable read	
serializable	
snapshot	

- A. Mastered
- B. Not Mastered

Answer: A

**Explanation:**

Box 1: read uncommitted

READ UNCOMMITTED is the least restrictive isolation level because it ignores locks placed by other transactions. Transactions executing under READ UNCOMMITTED can read modified data values that have not yet been committed by other transactions; these are called "dirty" reads.

Box 2: 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.

References: [https://msdn.microsoft.com/en-us/library/tcbchxcb\(v=vs.110\).aspx](https://msdn.microsoft.com/en-us/library/tcbchxcb(v=vs.110).aspx)

**NEW QUESTION 168**

- (Exam Topic 8)

You use SQL Server 2014 to maintain the data used by applications at your company. You need to run two separate SQL statements.

You must guarantee that the following three things happen:

1. Either BOTH statements succeed or BOTH statements fail as a batch.
2. If an error occurs on the first statement, SQL should not attempt to run the second statement.
3. Error information should be returned to the client. What should you do?

<p><input type="radio"/> A</p> <pre>SET XACT_ABORT ON BEGIN TRY BEGIN TRANSACTION ...Statement 1 ...Statement 2 COMMIT TRANSACTION END TRY BEGIN CATCH ROLLBACK TRANSACTION END CATCH</pre>	<p><input type="radio"/> C</p> <pre>SET XACT_ABORT ON BEGIN TRANSACTION ...Statement 1 ...Statement 2 If @@ERROR &lt;&gt; 0 ROLLBACK ELSE COMMIT TRANSACTION</pre>
<p><input type="radio"/> B</p> <pre>SET XACT_ABORT OFF BEGIN TRY ...Statement 1 END TRY BEGIN TRY ...Statement 2 END TRY BEGIN CATCH THROW END CATCH</pre>	<p><input type="radio"/> D</p> <pre>SET XACT_ABORT ON BEGIN TRY ...Statement 1 if @@ERROR &lt;&gt; 0 GOTO CATCH ...Statement 2 if @@ERROR &lt;&gt; 0 GOTO CATCH END TRY BEGIN CATCH THROW END CATCH</pre>

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

**Answer:** A

**Explanation:**

\* SET XACT\_ABORT

When SET XACT\_ABORT is ON, if a Transact-SQL statement raises a run-time error, the entire transaction is terminated and rolled back.

When SET XACT\_ABORT is OFF, in some cases only the Transact-SQL statement that raised the error is rolled back and the transaction continues processing.

**NEW QUESTION 172**

- (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 Latin1\_General\_CS\_AS NOT NULL,  
D. UserName nvarchar(200) COLLATE Latin1\_General\_CI\_AI NOT NULL,

**Answer:** D

#### NEW QUESTION 177

- (Exam Topic 8)

You have a SQL Server database.

The recovery model is set to full. The transaction log is backed up every night. You discover that the transaction log never decreases in size.

You execute the DBCC SHRINKFILE statement for the transaction log and you discover that the transaction log is unaffected.

You need to ensure that you can reduce the size of the transaction log. What should you do first?

- A. Truncate the transaction log  
B. Kill long-running transactions  
C. Change the recovery model to bulk-logged  
D. Perform a full backup

**Answer:** A

#### Explanation:

The transaction log must be truncated before running the DBCC SHRINKFILE operation.

#### NEW QUESTION 182

- (Exam Topic 8)

You have a Microsoft Azure SQL Database instance. The database contains a table named Table1 that has one million rows and a column named Column1.

Column1 allows null values.

You need to update Column1 to meet the following requirements: Prevent null values from being used

Always use a value of zero instead of a null value

Which three Transact-SQL statements should you run? Each correct answer presents part of the solution.

- A. ALTER TABLE dbo.Table1 ADD CONSTRAINTDF\_Table1\_Column1 DEFAULT 0 FOR Column1  
B. ALTER TABLE Table1 DROP COLUMN Column1  
C. ALTER TABLE Table1 ALTER COLUMNColumn int NOT NULL  
D. ALTER TABLE Table1 ADD COLUMNColumn1 int NOT NULL  
E. UPDATE Table1SET Column1 = 0WHERE Column1 IS NULL

**Answer:** ACE

#### NEW QUESTION 186

- (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 189

- (Exam Topic 8)

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 = OFF  
B. FILLFACTOR = 80  
C. FILLFACTOR = 0  
D. STATISTICS\_NORECOMPUTE = ON

**Answer:** C

#### Explanation:

<http://msdn.microsoft.com/en-us/library/ms177459.aspx> <http://msdn.microsoft.com/en-us/library/ms188783.aspx>

#### NEW QUESTION 194

- (Exam Topic 8)

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. Change the query inside Proc1 to:

```
SELECT ID, OrderDate, Amount
  FROM Sales
 WHERE OrderDate BETWEEN CONVERT(char(10),@date1,112)
 AND CONVERT(char(10),@date2,112)
 ORDER BY ID;
```

B. Change the definition of Proc1 to:

```
CREATE PROC usp_Proc1(
  @date1 int, @date2 int
)
```

C. Change the query inside Proc1 to:

```
SELECT ID, OrderDate, Amount
  FROM Sales
 WHERE CAST(OrderDate AS datetime) < @date1
 AND CAST(OrderDate AS datetime) > @date2
 ORDER BY ID;
```

D. Change the definition of Proc1 to:

```
CREATE PROC usp_Proc1(
  @date1 date, @date2 date
)
```

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

**Answer: A**

**NEW QUESTION 199**

- (Exam Topic 8)

You administer a SQL Server 2014 instance.

You have been assigned to determine the cause of frequent long-running transactions that have been tracked to the dbo.Account table, where there are many cases of blocking and deadlocks. The dbo.Account table contains more than one million rows.

Users and processes frequently search for and update data by using the AccountId column, and less frequently the AccountNumber and GovernmentId columns, all of which contain only unique values. Users frequently get lists of AccountNumber values by searching on Last Name and then First Name.

You need to modify the structure of the dbo.Account table to alleviate the issues.

How should you complete the table definition to reduce contention on the table structure? To answer, drag the appropriate code snippets to the correct locations in the CREATE TABLE statement. Each code snippet 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.

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Note:

Users and processes frequently search for and update data by using the AccountId column (Primary Key Clustered) , and less frequently the AccountNumber (Unique Clustered) and GovernmentId(Unique Clustered) columns, all of which contain only unique values. Users frequently get lists of AccountNumber values by searching on Last Name and then First Name (LastName, Firstname) INCLUDE (AccountNumber).

**NEW QUESTION 202**

- (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 204**

- (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 205**

- (Exam Topic 8)

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 SNAPSHOT
- B. Add the following line of code to line 13: 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 08: WITH (UPDLOCK)

**Answer: D**

**NEW QUESTION 210**

- (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**

- sys.dm\_db\_index\_usage\_stats
- sys.dm\_db\_index\_operational\_stats
- sys.dm\_db\_index\_physical\_stats
- sys.dm\_db\_missing\_index\_columns
- sys.dm\_db\_missing\_index\_details
- sys.dm\_db\_missing\_index\_groups

**Answer Area**

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

- 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 213**

- (Exam Topic 8)

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 LOG DB1 TO Device1 WITH COPY\_ONLY
- B. BACKUP LOG DB1 TO Device1 WITH NORECOVERY
- C. BACKUP LOG DB1 TO Device1 WITH TRUNCATE\_ONLY
- D. BACKUP LOG DB1 TO Device1

**Answer:** D

**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>

**NEW QUESTION 215**

- (Exam Topic 8)

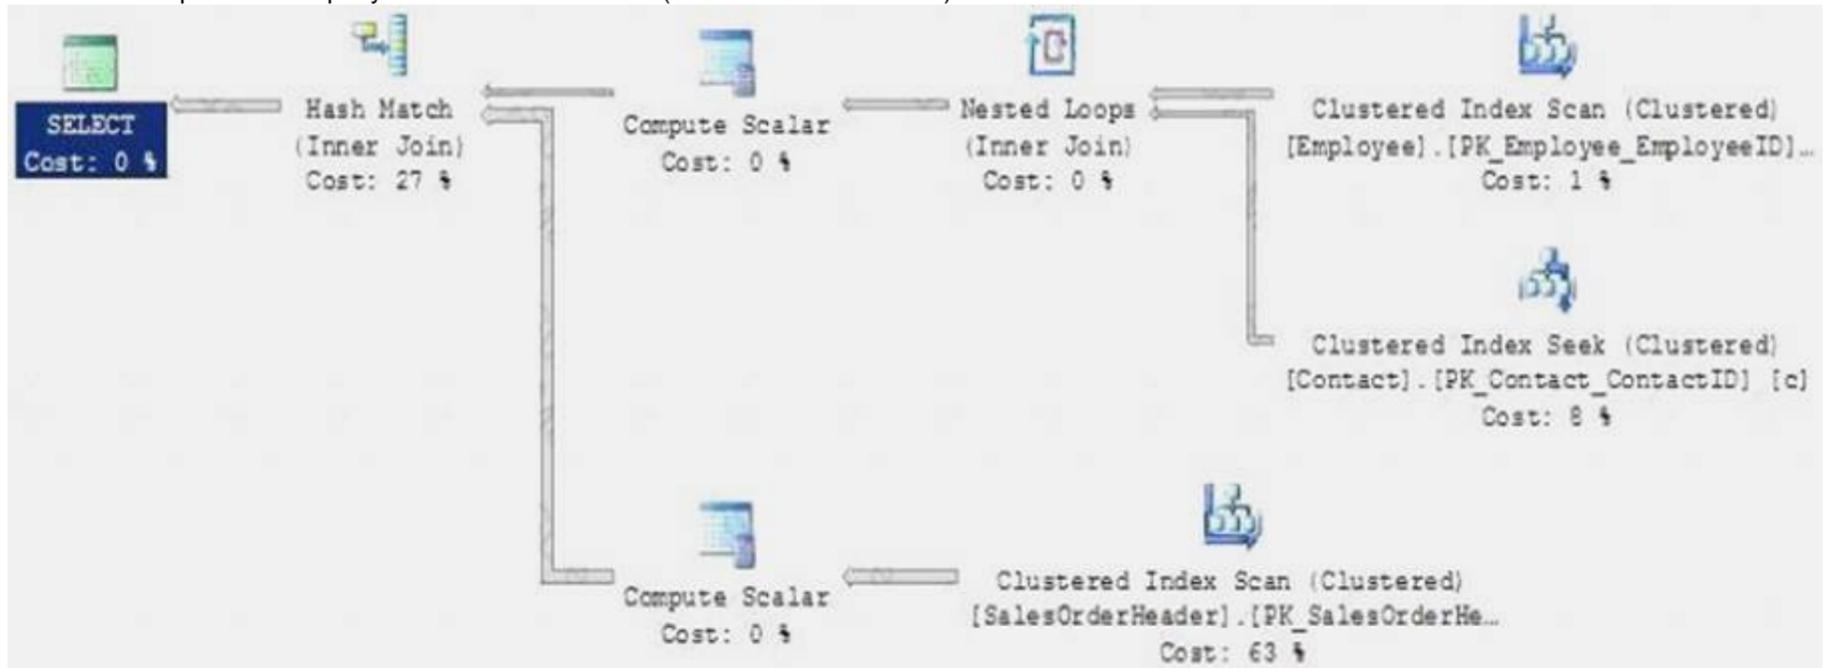
You have a database that contains three tables. The tables are configured as shown in the following table.

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.)

- 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 216**

- (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 219**

- (Exam Topic 8)

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.

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

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

**NEW QUESTION 224**

- (Exam Topic 8)

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**

**NEW QUESTION 228**

- (Exam Topic 8)

You use SQL Azure to store data used by an e-commerce application.

You develop a stored procedure named sp1. Sp1 is used to read and change the price of all the products sold on the e-commerce site.

You need to ensure that other transactions are blocked from updating product data while sp1 is executing. Which transaction isolation level should you use in sp1?

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

**Answer: C**

**NEW QUESTION 229**

- (Exam Topic 8)

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. Performance Monitor
- B. Database Engine Tuning Advisor
- C. Data Profile Viewer
- 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-sqlserver/>

<http://www.sqlservercentral.com/articles/SQL+Server+2012/At+last%2c+execution+plans+show+true+thread+r>

[http://sqlblog.com/blogs/paul\\_white/archive/2011/12/23/forcing-a-parallel-query-executionplan.aspx](http://sqlblog.com/blogs/paul_white/archive/2011/12/23/forcing-a-parallel-query-executionplan.aspx) [http://sqlblog.com/blogs/paul\\_white/archive/2012/05/02/parallel-row-goals-gone-rogue.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>

**NEW QUESTION 234**

- (Exam Topic 8)

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\_active\_transactions dynamic management views
- D. A SQL Server Agent job that retrieves information from the sys.dm\_tran\_session\_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.

**NEW QUESTION 238**

- (Exam Topic 8)

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:** A

**NEW QUESTION 242**

.....

## 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