

70-464 Dumps

Developing Microsoft SQL Server 2012 Databases

<https://www.certleader.com/70-464-dumps.html>



NEW QUESTION 1

- (Exam Topic 1)

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

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

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

Answer: D**Explanation:**

<http://msdn.microsoft.com/en-us/library/ms187926.aspx> <http://msdn.microsoft.com/en-us/library/ms190782.aspx> <http://msdn.microsoft.com/en-us/library/bb669091.aspx> <http://msdn.microsoft.com/en-us/library/windows/desktop/ms709342.aspx> <http://msdn.microsoft.com/en-us/library/ms188001.aspx>

NEW QUESTION 2

- (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 3

- (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 4

- (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 5

- (Exam Topic 2)

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

- ☒ A.

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

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

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

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

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

Answer: B

Explanation:

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

NEW QUESTION 6

- (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 7

- (Exam Topic 3)

You need to recommend a solution to ensure that SQL1 supports the auditing requirements of usp_UpdateEmployeeName.

What should you include in the recommendation?

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

Answer: D

NEW QUESTION 8

- (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.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/

NEW QUESTION 9

- (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 10

- (Exam Topic 4)

You are planning the ManufacturingSteps table.

You need to define the ProductID column in the CREATE TABLE statement. Which code segment should you use?

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

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

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

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

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

D. Option D

Answer: B

Explanation:

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

NEW QUESTION 10

- (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 13

- (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 16

- (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 21

- (Exam Topic 4)

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

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

Answer: D

NEW QUESTION 24

- (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 27

- (Exam Topic 5)

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

- ☐ A. REVOKE EXEC ON usp_UpdateOpening FROM Candidates;
- ☐ B. DENY EXEC ON usp_UpdateOpening TO Candidates;
- ☐ C. ALTER PROCEDURE usp_UpdateOpening
@openingIDint,
@titlevarchar(100),
@salarydecimal(18,0),
@descriptionvarchar(8000)
WITH EXECUTE AS Administrator
AS
...
- ☐ D. ALTER PROCEDURE usp_UpdateOpening
@openingIDint,
@titlevarchar(100),
@salarydecimal(18,0),
@descriptionvarchar(8000)
WITH EXECUTE AS Company
AS
...

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

Answer: A

NEW QUESTION 32

- (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 34

- (Exam Topic 5)

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

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

Answer: C

NEW QUESTION 35

- (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 37

- (Exam Topic 6)

You need to implement a solution that addresses the performance issues of the usp_GetOrdersByProduct stored procedure.

Which statement should you execute?

- ☐ A. `CREATE INDEX IX_OrderDetails_ByProduct
ON OrderDetails (ProductID)
INCLUDE (OrderID, LineItem, UnitPrice, Total, Discount)`
- ☐ B. `CREATE INDEX IX_OrderDetails_ByProduct
ON OrderDetails (ProductID)
INCLUDE (LineItem, Quantity, UnitPrice, Total, Discount)`
- ☐ C. `CREATE INDEX IX_OrderDetails_ByProduct
ON OrderDetails (ProductID)`
- ☐ D. `CREATE INDEX IX_OrderDetails_ByProduct
ON OrderDetails (ProductID)
INCLUDE (LineItem, Quantity, UnitPrice, Discount)`

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

Answer: C

NEW QUESTION 39

- (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 44

- (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 49

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

D. Add WITH (READPAST) to the end of line 47.

Answer: B

Explanation:

The SET TRANSACTION ISOLATION LEVEL SERIALIZABLE statement specifies that no other statement can read data that has been modified but not yet committed by other transactions.

References:

<https://docs.microsoft.com/en-us/sql/t-sql/statements/set-transaction-isolation-level-transact-sql?view=sql-server> <https://docs.microsoft.com/en-us/sql/t-sql/queries/hints-transact-sql-table?view=sql-server-2017>

NEW QUESTION 53

- (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 57

- (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 60

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

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
)
GO
```

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 61

- (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

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.

Answer Area

- 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 66

- (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 70

- (Exam Topic 7)

You need to create the usp.AssignUser stored procedure.

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

Code Blocks	Answer Area
<pre>IF @StatusID IS NULL RAISERROR (N'The transaction does not exist.',16,1)</pre>	
<pre>WITH NATIVE_COMPILATION, SCHEMABINDING, EXECUTE AS OWNER</pre>	
<pre>CREATE PROCEDURE dbo.usp_AssignUser @UserId int, @POSTransactionId int</pre>	
<pre>WITH (TRANSACTION ISOLATION LEVEL = READ COMMITTED, LANGUAGE = N'us_english')</pre>	
<pre>UPDATE dbo.POSTransaction SET UserId=@UserId WHERE POSTransactionId=@POSTransactio nId END</pre>	
<pre>AS BEGIN</pre>	
<pre>DECLARE @StatusID int SELECT @StatusID=StatusId FROM dbo.POSTransaction WHERE POSTransactionId=@POSTransactionI d</pre>	
<pre>IF @StatusID IS NULL THROW 51000, N'The transaction does not exist.', 1</pre>	
<pre>WITH (TRANSACTION ISOLATION LEVEL = REPEATABLE READ, LANGUAGE = N'us_english')</pre>	
<pre>AS BEGIN ATOMIC</pre>	

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Box 1:

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

Box 2:

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

Box 3:

```
AS
BEGIN ATOMIC
```

Box 4:

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

Box 5:

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

Box 6:

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

Box 7:

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

Note:

* From scenario: The mobile application will need to meet the following requirements:

/Communicate with web services that assign a new user to a micropayment by using a stored procedure named usp_AssignUser.

* Example:

```
create procedure dbo.OrderInsert(@OrdNo integer, @CustCode nvarchar(5))
```

```
with native_compilation, schemabinding, execute as owner as
```

```
begin atomic with
```

```
(transaction isolation level = snapshot, language = N'English')
```

```
declare @OrdDate datetime = getdate();
```

```
insert into dbo.Ord (OrdNo, CustCode, OrdDate) values (@OrdNo, @CustCode, @OrdDate); end
```

```
go
```

* Natively compiled stored procedures are Transact-SQL stored procedures compiled to native code that access memory-optimized tables. Natively compiled stored procedures allow for efficient execution of the queries and business logic in the stored procedure.

* READ COMMITTED versus REPEATABLE READ

Read committed is an isolation level that guarantees that any data read was committed at the moment is read. It simply restricts the reader from seeing any intermediate, uncommitted, 'dirty' read. IT makes no promise whatsoever that if the transaction re-issues the read, will find the Same data, data is free to change after it was read.

Repeatable read is a higher isolation level, that in addition to the guarantees of the read committed level, it also guarantees that any data read cannot change, if the transaction reads the same data again, it will find the previously read data in place, unchanged, and available to read.

* Both RAISERROR and THROW statements are used to raise an error in Sql Server.

The journey of RAISERROR started from Sql Server 7.0, where as the journey of THROW statement has just began with Sql Server 2012. obviously, Microsoft suggesting us to start using THROW statement instead of RAISERROR. THROW statement seems to be simple and easy to use than RAISERROR.

* Explicit transactions. The user starts the transaction through an explicit BEGIN TRAN or BEGIN ATOMIC. The transaction is completed following the corresponding COMMIT and ROLLBACK or END (in the case of an atomic block).

NEW QUESTION 75

- (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 77

- (Exam Topic 8)

You have a SQL Server 2012 instance.

You plan to create an application that uses spatial data.

You need to create an object that will support the representation of the surface area of all the oceans. Which code segment should you use?

☐ A. DECLARE @g GEOGRAPHY =
GEOGRAPHY::STGeomFromText(
'FULLGLOBE',4326
);

☐ B. DECLARE @g GEOGRAPHY =
GEOGRAPHY::STGeomFromText(
'POLYGON(0 0, 0 10, 10 10, 10 0, 0 0)',4326
);

☐ C. DECLARE @g GEOGRAPHY =
GEOGRAPHY::STGeomFromText('
COMPOUNDCURVE(
CIRCULARSTRING(0 -50, 90 0, 0 50),
CIRCULARSTRING(0 50, 45 50, -90 50),
CIRCULARSTRING(-90 50, 0 0, -90 -50),
CIRCULARSTRING(-90 -50, 45 -50, 0 -50),4326
)'
);

☐ D. DECLARE @g GEOGRAPHY =
GEOGRAPHY::STGeomFromText(
'CIRCULARSTRING(0 50, 90 50, 180 50)',4326
);

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

Answer: A

NEW QUESTION 79

- (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 82

- (Exam Topic 8)

You are building a new index for an application.

You need to reduce the storage space used by the index and to minimize logical reads. What should you configure?

- A. SORT_IN_TEMPDB = ON
- B. MAXDOP = 4
- C. PAD_INDEX = OFF
- D. DATA_COMPRESSION = PAGE
- E. DATA_COMPRESSION = ROW

Answer: D

Explanation:

The data compression feature to help reduce the size of the rowstore tables and indexes. It can also help improve performance of I/O intensive workloads because the data is stored in fewer pages and queries need to read fewer pages from disk.

NEW QUESTION 83

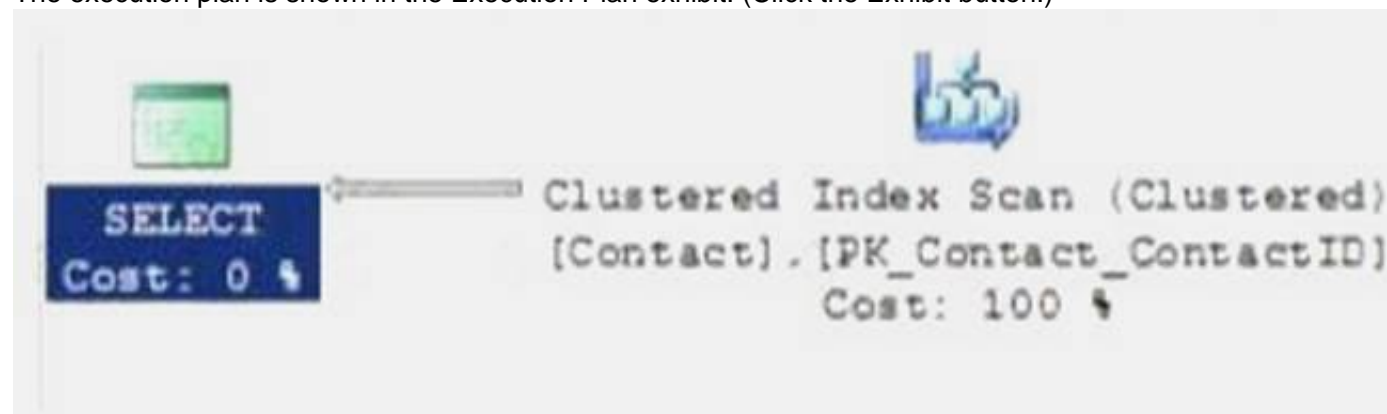
- (Exam Topic 8)

You have the following query on a disk-based table:

```
SELECT ContactID,  
       EmailAddress,  
       LastName  
FROM Person.Contact  
WHERE LastName = N'Johnson'
```

You discover that the query takes a long time to complete.

The execution plan is shown in the Execution Plan exhibit. (Click the Exhibit button.)



The index usage is show in the Index Usage exhibit. (Click the Exhibit button.)

Clustered Index Scan (Clustered)	
Scanning a clustered index, entirely or only a range.	
Physical Operation	Clustered Index Scan
Logical Operation	Clustered Index Scan
Actual Execution Mode	Row
Estimated Execution Mode	Row
Actual Number of Rows	730
Actual Number of Batches	0
Estimated I/O Cost	2.04016
Estimated Operator Cost	2.06229 (100%)
Estimated CPU Cost	0.0221262
Estimated Subtree Cost	2.06229
Number of Executions	1
Estimated Number of Executions	1
Estimated Number of Rows	82.1249
Estimated Row Size	78 B
Actual Rebinds	0
Actual Rewinds	0
Ordered	False
Node ID	0
Predicate	
[DB1].[Person].[Contact].[LastName]=CONVERT_IMPLICIT (nvarchar(4000),[@1],0)	
Object	
[DB1].[Person].[Contact].[PK_Contact_ContactID]	
Output List	
[DB1].[Person].[Contact].ContactID, [DB1].[Person]. [Contact].EmailAddress, [DB1].[Person].[Contact].LastName	

You need to reduce the amount of time it takes to complete the query. You must achieve this goal as quickly as possible. What should you do?

- A. Reorganize the index.
- B. Update statistics.
- C. Create an index on LastName.
- D. Rebuild the index.

Answer: C

NEW QUESTION 87

- (Exam Topic 8)

Your company has a SQL Azure subscription.

You implement a database named Database1. Database1 has two tables named Table1 and Table2. You create a stored procedure named sp1. Sp1 reads data from Table1 and inserts data into Table2. A user named User1 informs you that he is unable to run sp1.

You verify that User1 has the SELECT permission on Table1 and Table2.

You need to ensure that User1 can run sp1. The solution must minimize the number of permissions assigned to User1.

What should you do?

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

Answer: B

Explanation:

<http://msdn.microsoft.com/en-us/library/ms191291.aspx>

NEW QUESTION 89

- (Exam Topic 8)

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

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

You create a non-clustered index named EmployeeName on the name column.

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

You write the following query to retrieve all of the employees that have a name that starts with the letters JOH:

You discover that the query performs a table scan.

You need to ensure that the query uses EmployeeName. What should you do?

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

Answer: C

NEW QUESTION 92

- (Exam Topic 8)

You have a table named Table1 that contains one million rows. Table1 contains a column named Column1 that stores sensitive information. Column1 uses the nvarchar(16) data type. You have a certificate named Cert1.

You must add a column named Column2 that contains an encrypted version of the data from Column1. You must use two-way encryption. You plan to remove Column1 after you create Column2.

Which five Transact-SQL statements should you run in sequence before you remove Column1? To answer, move the appropriate Transact-SQL statements from the list of Transact-SQL statements to the answer area and arrange them in the correct order.

NOTE: More than one order of answer choices is correct. You will receive credit for any correct orders you select.

Code segments	Answer Area
CREATE SYMMETRIC KEY Key1 WITH ALGORITHM = SHA1 ENCRYPTION BY CERTIFICATE Cert1;	
CLOSE SYMMETRIC KEY;	
ALTER TABLE TABLE1 ADD Column2 varbinary(256);	
UPDATE table1 SET Column2 = EncryptByKey(Key_GUID (Key1),Column1);	
ALTER TABLE Table1 ADD Column2 nvarchar(256);	
CREATE SYMMETRIC KEY Key1 WITH ALGORITHM = AES_256 ENCRYPTION BY CERTIFICATE Cert1;	
CREATE CREDENTIAL Cred1 WITH IDENTITY = 'User1', SECRET = 'P@ssword';	
OPEN SYMMETRIC KEY Key1 DECRYPTION BY CERTIFICATE Cert1;	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/encrypt-a-column-of-data?view=sq> <https://docs.microsoft.com/en-us/sql/t-sql/statements/close-symmetric-key-transact-sql?view=sql-server-2017>

NEW QUESTION 96

- (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 97

- (Exam Topic 8)

You are planning two stored procedures named SProc1 and SProc2. You identify the following requirements:

- ▶ SProc1 must return a table.
- ▶ SProc2 must return a scalar value.

You need to identify which option must be implemented for each stored procedure to return the desired data. Which options should you identify?

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

Options		Answer Area
an output parameter		SProc1
a raise error		
a SELECT statement		
a table-valued parameter (TVP)		SProc2

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Note:

- * a table (a set of rows) can be returned through a SELECT statement
- * a scalar can be returned through an output parameter.
- * incorrect: TVP is used for input not output.

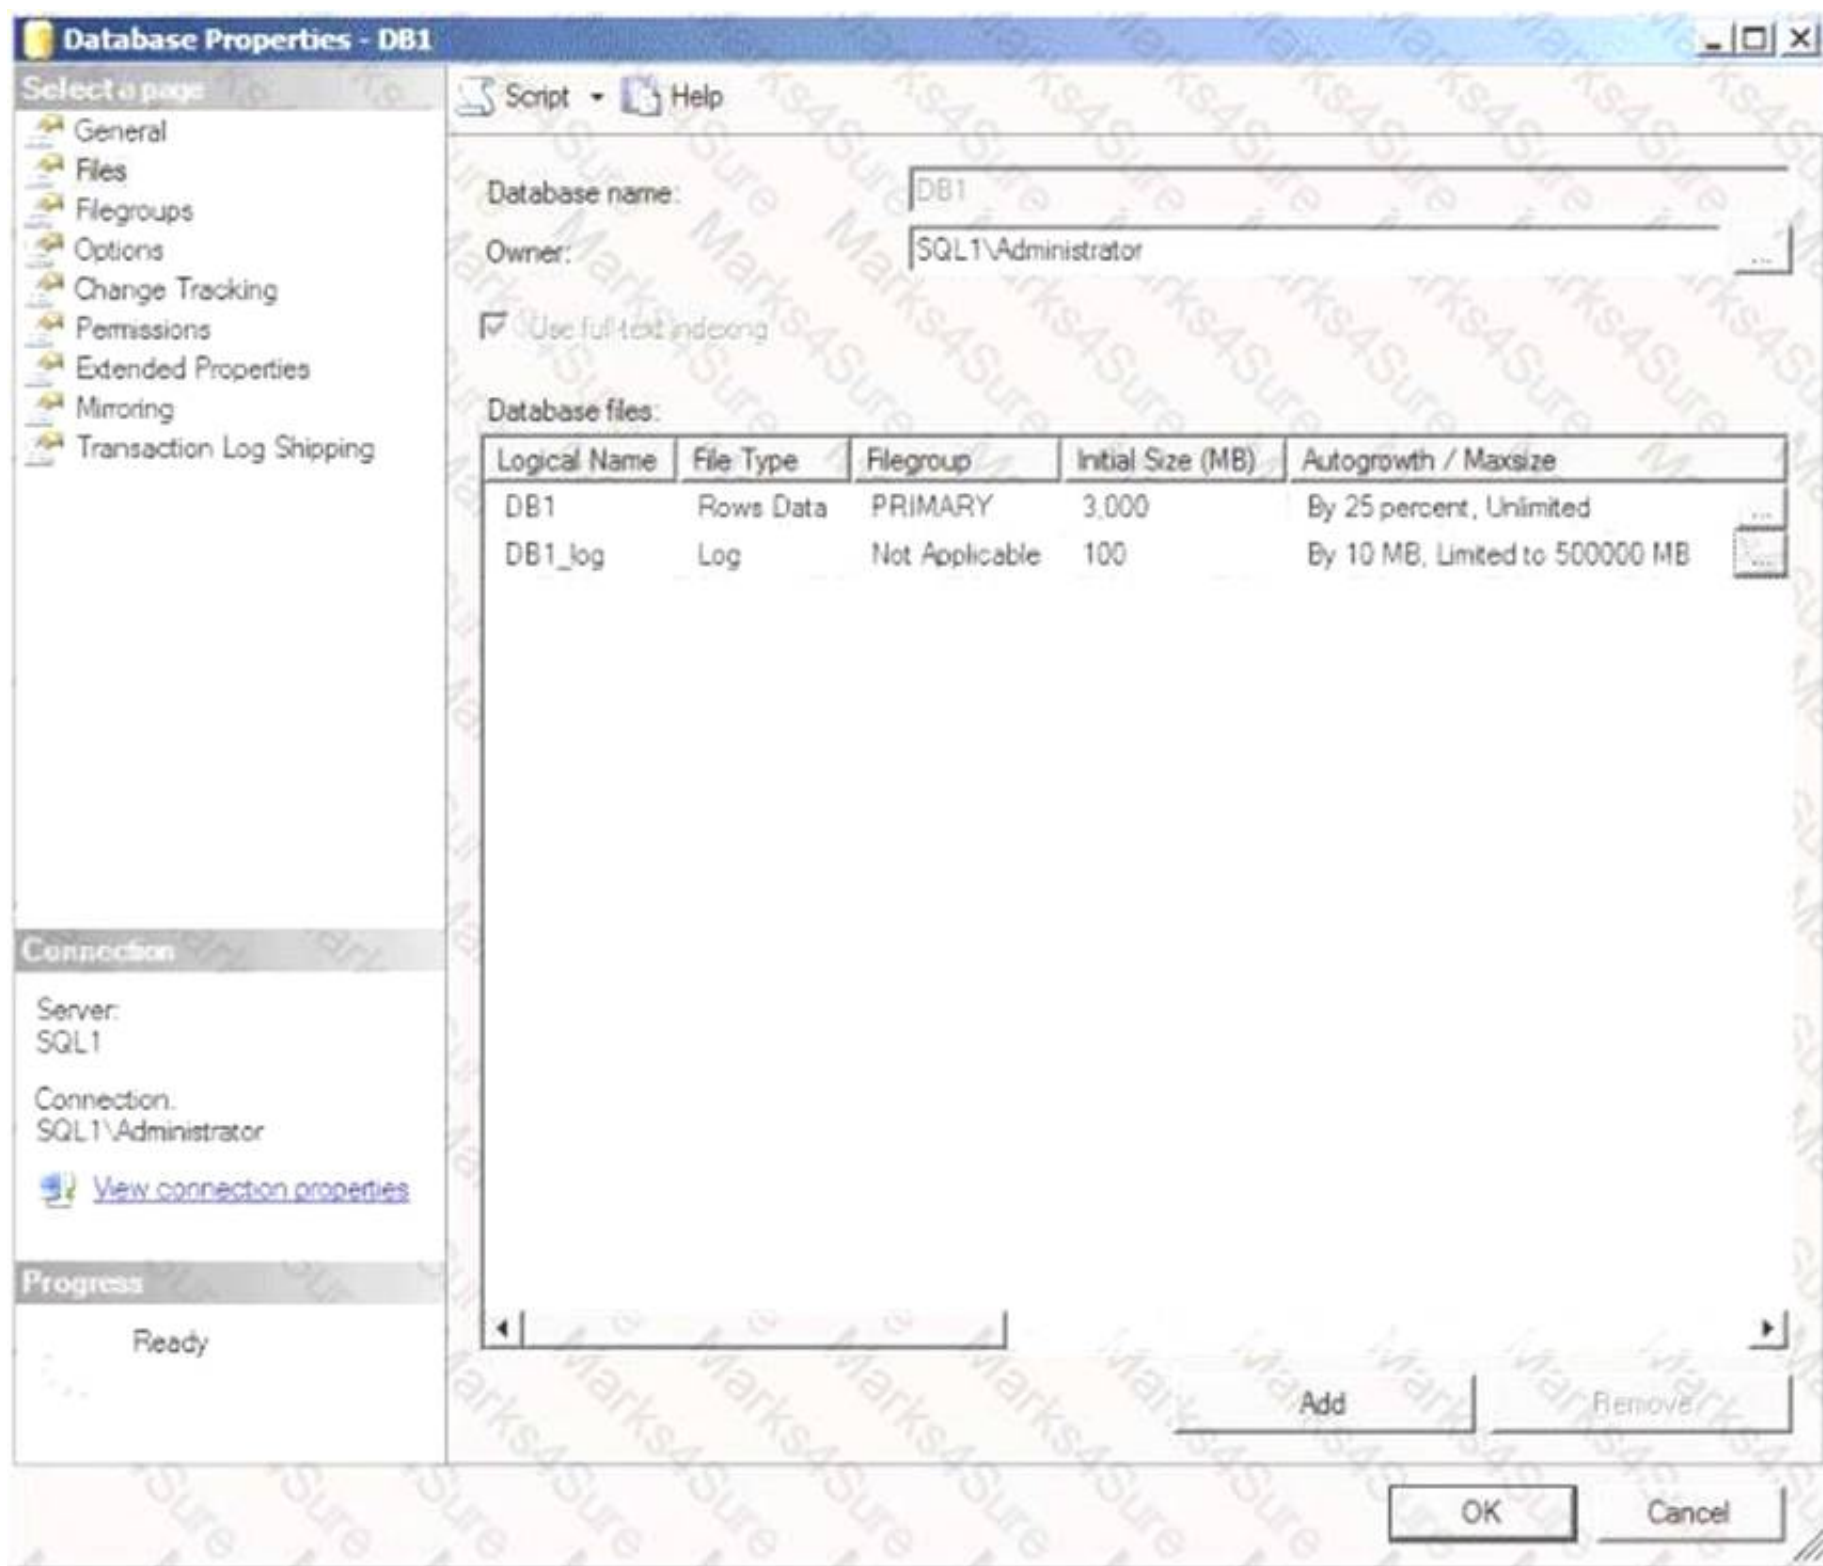
NEW QUESTION 101

- (Exam Topic 8)

You have a database named DB1.

You plan to configure change data capture on the existing tables in DB1.

The database file settings for the DB1 are shown in the exhibit. (Click the Exhibit button.)



You need to minimize disk contention caused by change data capture. What should you do?
More than one answer choice may achieve the goal. Select the BEST answer.

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

Answer: D

NEW QUESTION 106

- (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	Answer Area
DepartmentId	CREATE INDEX IX_Index3 ON dbo.Employee
EmployeeId	(Code element , Code element)
FirstName	INCLUDE (Code element , Code element)
LastName	
OfficeId	

- 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 108

- (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 113

- (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 114

- (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 118

- (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 120

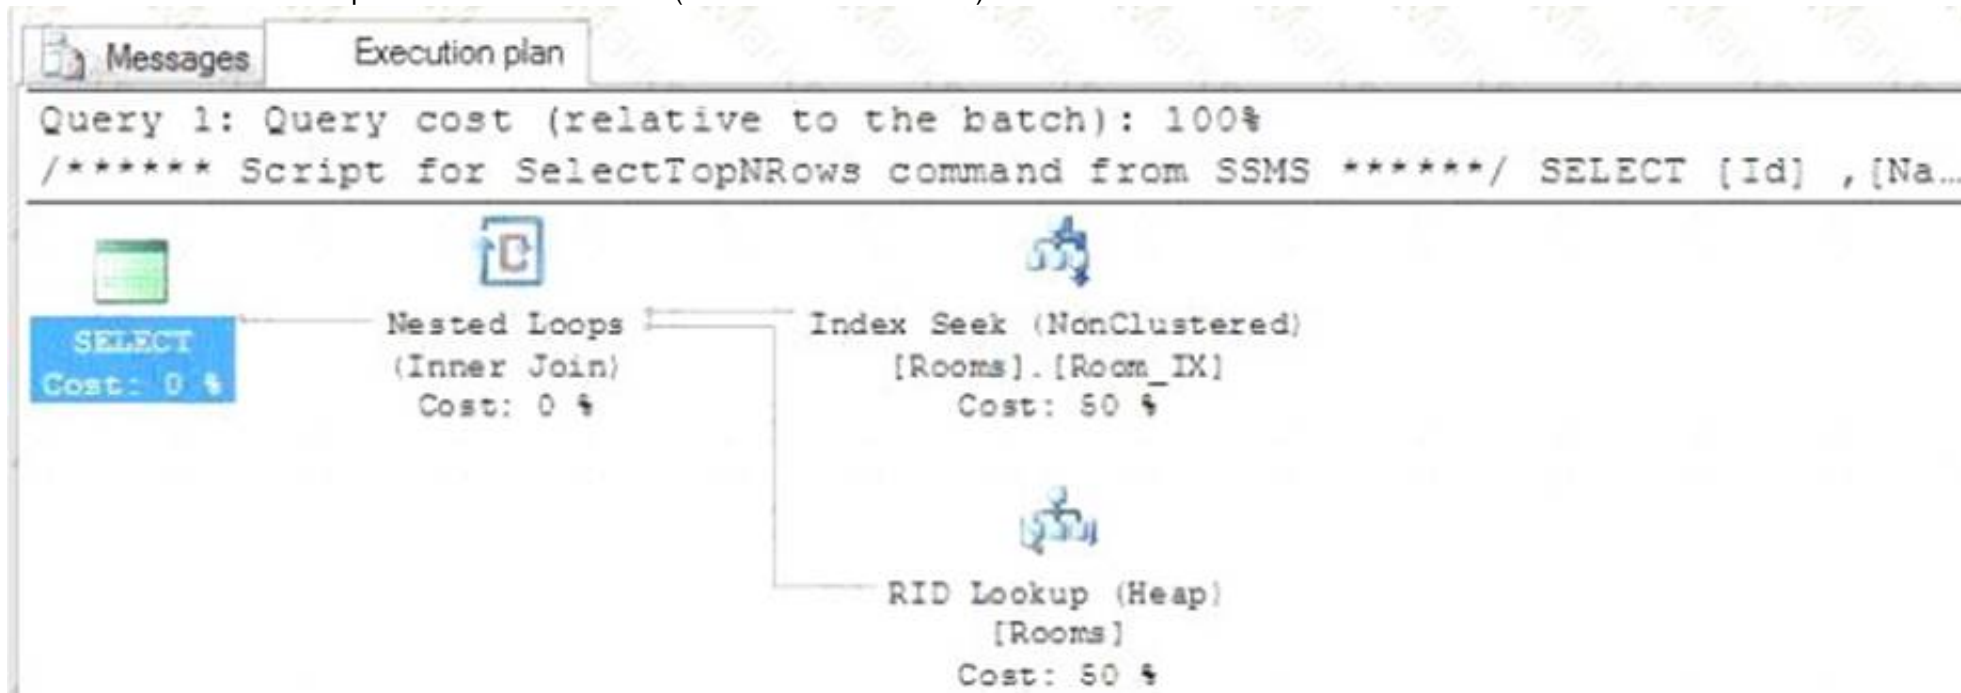
- (Exam Topic 8)

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 122

- (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 125

- (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
Extensible Markup Language (XML)	Procedure2
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 126

- (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 129

- (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 131

- (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 134

- (Exam Topic 8)

You plan to execute the following code:

```
01 CREATE TABLE dbo.Table1
02 (
03     datavalue varchar(20)
04 );
05 GO
06 BEGIN TRANSACTION;
07 INSERT INTO Table1 VALUES('entry1');
08     BEGIN TRANSACTION;
09         INSERT INTO Table1 VALUES('entry2');
10     COMMIT TRANSACTION;
11 INSERT INTO Table1 VALUES('entry3');
12 ROLLBACK TRANSACTION;
13 Go
```

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

- A. 1
B. 2
C. 3

Answer: A

NEW QUESTION 137

- (Exam Topic 8)

You have a SQL Server 2012 database named Database1. Database1 has a data file named Database1_data.mdf and a transaction log named Database1log.ldf. Database1_data.mdf is 1.5 GB. Database1log.ldf is 1.5 terabytes.

A full backup of Database1 is performed every day.

You need to reduce the size of the log file. The solution must ensure that you can perform transaction log backups in the future.

Which code segment should you execute?

To answer, move the appropriate code segments from the list of code segments to the answer area and arrange them in the correct order.

DBCC SHRINKFILE (database1_log,1)

ALTER DATABASE database1 SET RECOVERY FULL

ALTER DATABASE database1 SET RECOVERY SIMPLE

BACKUP LOG database1 WITH TRUNCATE_ONLY

DBCC SHRINKFILE (database1_data,1)

- A. Mastered
B. Not Mastered

Answer: A

Explanation:



NEW QUESTION 139

- (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 140

- (Exam Topic 8)

You are planning two stored procedures named SProc1 and SProc2. You identify the following requirements:

- ▶ SProc1 must return a table.
- ▶ SProc2 must return a status code.

You need to identify which options must be implemented to meet each stored procedure requirement. Which options should you identify?

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

Options	Answer Area
a raise error	SProc1 Option
a return value	SProc2 Option
a SELECT statement	
a table-valued parameter (TVP)	

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Options	Answer Area
a raise error	SProc1 a SELECT statement
a return value	SProc2 a return value
a SELECT statement	
a table-valued parameter (TVP)	

NEW QUESTION 145

- (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 146

- (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 151

- (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: 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.

Identify which indexes should be created.

Function

Function

- A. Mastered
- B. Not Mastered

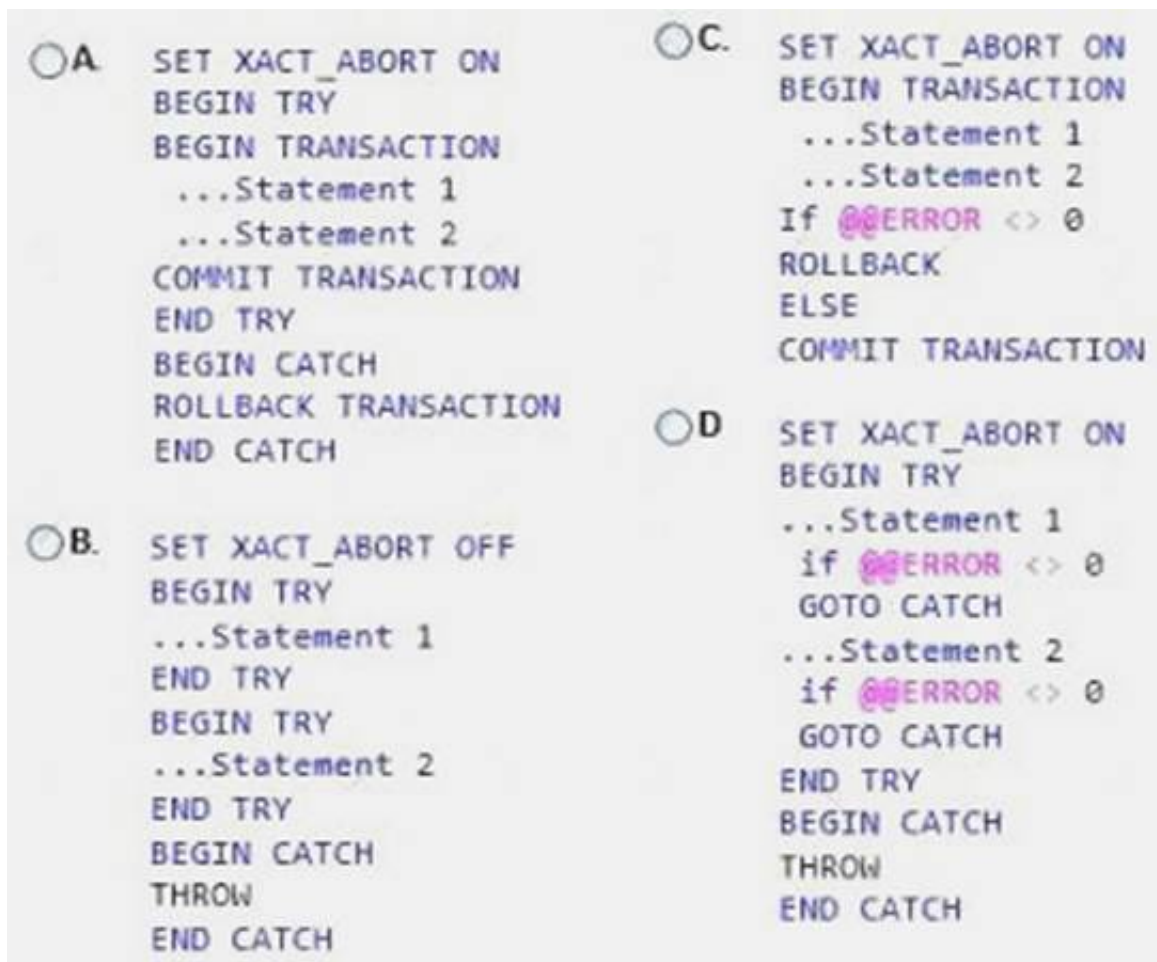
Answer: A

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.dm_db_missing_index_details
sys.dm_db_missing_index_details returns detailed information about a missing index; for example, it returns the name and identifier of the table where the index is missing, and the columns and column types that should make up the missing index.
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-dynamic-management-views/sys-dm-db-index-usage-stats)
<https://docs.microsoft.com/en-us/sql/relational-databases/system-catalog-views/sys-indexes-transact-sql> [https://technet.microsoft.com/en-us/library/ms345524\(v=sql.105\).aspx](https://technet.microsoft.com/en-us/library/ms345524(v=sql.105).aspx)

NEW QUESTION 153

- (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?



- 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 155

- (Exam Topic 8)

You plan to create a new table that has the following requirements:

- ▶ Uses a GUID data type as the primary key.
- ▶ Uses a clustered index as the primary key.
- ▶ Minimizes fragmentation.

You need to recommend which option to include in the CREATE statement. Which option should you include?

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

- A. NEWID
- B. @@IDENTITY
- C. NEWSEQUENTIALID
- D. IDENTITY

Answer: C

NEW QUESTION 159

- (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 162

- (Exam Topic 8)

You have a table named Customers that has a clustered index defined on the ID column. You write a script to create a stored procedure.

You need to complete the script for the stored procedure. The solution must minimize the number of locks and deadlocks.

What should you do?

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

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

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Note:

* Optimized bulk load operations on heaps block queries that are running under the following isolation levels: SNAPSHOT

READ UNCOMMITTED

READ COMMITTED using row versioning

* READ COMMITTED

Specifies that statements cannot read data that has been modified but not committed by other transactions. This prevents dirty reads. Data can be changed by other transactions between individual statements within the current transaction, resulting in nonrepeatable reads or phantom data. This option is the SQL Server default.

* SERIALIZABLE (more locks) Specifies the following:

Statements cannot read data that has been modified but not yet committed by other transactions.

No other transactions can modify data that has been read by the current transaction until the current transaction completes.

Other transactions cannot insert new rows with key values that would fall in the range of keys read by any statements in the current transaction until the current transaction completes.

* 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.

* XLOCK

Specifies that exclusive locks are to be taken and held until the transaction completes. If specified with ROWLOCK, PAGLOCK, or TABLOCK, the exclusive locks apply to the appropriate level of granularity.

Reference: Table Hints (Transact-SQL)

NEW QUESTION 163

- (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 166

- (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 169

- (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.)

Values	Answer Area
<input type="button" value="ON"/> <input type="button" value="OFF"/>	ALTER INDEX ALL ON Customers
	REBUILD WITH
	(ONLINE = <input style="width: 100px;" type="text" value="Value"/> ,
	STATISTICS_NORECOMPUTE <input style="width: 100px;" type="text" value="Value"/>);

- 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 173

- (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.

The screenshot shows a drag-and-drop interface with two columns. The left column is divided into two sections: 'Encryption Types' and 'Data Types'. The 'Encryption Types' section contains four items: 'Transparent data encryption (TDE)', 'Encrypting File System (EFS)', 'Cell-level encryption', and 'BitLocker Drive Encryption (BitLocker)'. The 'Data Types' section contains four items: 'decimal', 'varchar', 'varbinary', and 'money'. The right column is titled 'Answer Area' and contains two empty boxes labeled 'Encryption Type' and 'Data Type'.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

Cell-level encryption

Data Type

varbinary

NEW QUESTION 176

- (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.

Code Snippets

PRIMARY KEY CLUSTERED

UNIQUE NONCLUSTERED

(Lastname, FirstName) INCLUDE (AccountNumber)

(Lastname, FirstName) INCLUDE (AccountId)

(FirstName, Lastname)

/* No Change To Structure */

CREATE TABLE Statement

CREATE TABLE dbo.Account
(
AccountNumber nchar(10) NOT NULL

AccountId int NOT NULL

GovernmentId nvarchar(11) NOT NULL

FirstName nvarchar(20) NOT NULL,
MiddleInitial nvarchar(1) NULL,
LastName nvarchar(20) NOT NULL
)
Go

CREATE NONCLUSTERED INDEX X1 ON dbo.Account

Code Snippet

Code Snippet

Code Snippet

Code Snippet

- A. Mastered
- B. Not Mastered

Answer: A

The Leader of IT Certification

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

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 178

- (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 182

- (Exam Topic 8)

You are creating a table to support an application that will cache data outside of SQL Server. The application will detect whether cached values were changed before it updates the values. You need to create the table, and then verify that you can insert a row into the table.

Which code segment should you use?

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

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

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

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

- A. Option A
- B. Option B

- C. Option C
- D. Option D

Answer: C

Explanation:





<http://msdn.microsoft.com/en-us/library/ms182776.aspx>

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

NEW QUESTION 184

- (Exam Topic 8)

You have a database that uses the following management views:

-  Sys.dm_os_volume_stats
-  Sys.dm_db_partition_stats
-  Sys.dm_db_file_space_usage
-  Sys.fulltext_indexes

You plan to migrate the database to Microsoft SQL Azure. You need to identify which view can be used in SQL Azure. Which view should you identify?

- A. sys.fulltext_indexes
- B. sys.dm_db_file_space_usage
- C. sys.dm_os_volume_stats
- D. sys.dm_db_partition_stats

Answer: D

NEW QUESTION 185

- (Exam Topic 8)

You have an application that queries a database.

Users report that the application runs more slowly than expected.

You need to identify which queries take the most time to execute. The solution must minimize the use of processor resources on the server.

What should you use?

- A. a SQL Server Profiler trace
- B. a Central Management Server
- C. Query Store
- D. Distributed Replay

Answer: C

Explanation:

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/performance/monitoring-performance-by-using-theque>

NEW QUESTION 189

- (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 192

- (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 193

- (Exam Topic 8)

You run the following code:

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

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

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

Answer: E

Explanation:

<http://msdn.microsoft.com/en-us/library/ms187856.aspx>

NEW QUESTION 198

- (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/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 201

- (Exam Topic 8)

You have a SQL Server 2012 database named Database1. Database1 has a data file named database1_data.mdf and a transaction log file named database1_Log.ldf. Database1_Data.mdf is

1.5 GB.

Database1_Log.ldf is 1.5 terabytes. A full backup of Database1 is performed every day.

You need to reduce the size of the log file. The solution must ensure that you can perform transaction log

backups in the future. Which code segment should you execute? To answer, move the appropriate code segments from the list of code segments to the answer area and arrange them in the correct order.

Ordered List Title	Answer Choices Title
<div><div>▲▼</div><div></div></div>	<div>ALTER DATABASE Database1 SET RECOVERY FULL ; ALTER DATABASE Database1 SET RECOVERY SIMPLE ; DBCC SHRINKFILE (Database1_Log); DBCC SHRINKFILE (Database1_Log, TRUNCATEONLY); BACKUP LOG Database1 TO Database1_Log_Backup; BACKUP LOG Database1 TO Database1_Log_Backup WITH NO_TRUNCATE; USE Database1; USE master;</div>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

<http://msdn.microsoft.com/en-gb/library/ms189493.aspx>

NEW QUESTION 202

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your 70-464 Exam with Our Prep Materials Via below:

<https://www.certleader.com/70-464-dumps.html>