

Exam Questions 70-464

Developing Microsoft SQL Server 2012 Databases

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



NEW QUESTION 1

- (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 2

- (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 3

- (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 4

- (Exam Topic 2)

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

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

Answer: A

NEW QUESTION 5

- (Exam Topic 2)

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

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

Answer: B

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/

NEW QUESTION 6

- (Exam Topic 2)

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

- A. CREATE XML SCHEMA COLLECTION SessionDataTable
- B. CREATE TYPE SessionDataTable AS Table
- C. CREATE SCHEMA SessionDataTable
- D. CREATE TABLE SessionDataTable

Answer: B

NEW QUESTION 7

- (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 8

- (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 9

- (Exam Topic 2)

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

- A. CREATE SCHEMA SessionDataTable
- B. CREATE TYPE SessionDataTable AS Table
- C. CREATE TABLE SessionDataTable
- D. CREATE XML SCHEMA COLLECTION SessionDataTable

Answer: A

NEW QUESTION 10

- (Exam Topic 3)

You need to add a new column named Confirmed to the Employees table. The solution must meet the following requirements:

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

Which code segment should you use?

- A. ALTER TABLE Employees
ADD Confirmed char(1) DEFAULT '1';
- B. ALTER TABLE Employees
ADD Confirmed char(1) DEFAULT '0';
- C. ALTER TABLE Employees
ADD Confirmed bit DEFAULT 0;
- D. ALTER TABLE Employees
ADD Confirmed bit DEFAULT 1;

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

Answer: D

NEW QUESTION 10

- (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 12

- (Exam Topic 3)

You need to provide referential integrity between the Offices table and Employees table.

Which code segment or segments should you add at line 27 of Tables.sql? (Each correct answer presents part of the solution. Choose all that apply.)

- A. ALTER TABLE dbo.Offices ADD CONSTRAINT
PK_Offices_EmployeeID PRIMARY KEY (EmployeeID);
- B. ALTER TABLE dbo.Employees ADD CONSTRAINT
FK_Employees_Offices FOREIGN KEY (OfficeID)
REFERENCES dbo.Offices (OfficeID);
- C. ALTER TABLE dbo.Employees ADD CONSTRAINT
PK_Employees_EmployeeID PRIMARY KEY (EmployeeID);
- D. ALTER TABLE dbo.Offices ADD CONSTRAINT
FK_Offices_Employees FOREIGN KEY (EmployeeID)
REFERENCES dbo.Employees (EmployeeID);

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

Answer: CD

Explanation:

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

NEW QUESTION 15

- (Exam Topic 3)

You are evaluating the index design.

You need to recommend a change to Indexes.sql that minimizes the amount of time it takes for usp_SelectEmployeesByName to execute.

Which line of code should you use to replace line 08 of Indexes.sql?

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

Answer: A

Explanation:

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/indexes/create-indexes-with-included-columns?view=sql-server-2017>

NEW QUESTION 16

- (Exam Topic 4)

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

- A. CREATE FUNCTION Production.fnProductsByProductType (@ProductType varchar(11))
RETURNS @tblInvoices TABLE (ProductID bigint, ProductType varchar(11), CreationDate date)
AS
INSERT INTO @tblInvoices
- B. CREATE FUNCTION Production.fnProductsByProductType (@ProductType varchar(11))
RETURNS TABLE
AS
RETURN
- C. CREATE FUNCTION Production.fnProductsByProductType (@ProductType varchar(11))
RETURNS @TblInvoices TABLE (ProductID bigint, ProductType varchar(11), CreationDate date)
AS
- D. CREATE FUNCTION Production.fnProductsByProductType (@ProductType varchar(11))
RETURNS xml
AS
RETURN

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

Answer: B

Explanation:

<http://msdn.microsoft.com/en-us/library/ms191320.aspx> <http://msdn.microsoft.com/en-us/library/ms186755.aspx>

NEW QUESTION 18

- (Exam Topic 4)

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

- A. OPEN PRODUCTSCERT;
ALTER PROCEDURE Production.ProductDetails_Insert
WITH ENCRYPTION;
CLOSE PRODUCTSCERT;
- B. OPEN DBCERT;
ALTER PROCEDURE Production.ProductDetails_Insert
WITH ENCRYPTION;
CLOSE DBCERT;
- C. ADD SIGNATURE TO Production.ProductDetails_Insert
BY CERTIFICATE DBCERT;
- D. ADD SIGNATURE TO Production.ProductDetails_Insert
BY CERTIFICATE PRODUCTSCERT;

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

Answer: C

Explanation:

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

NEW QUESTION 20

- (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 25

- (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)

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 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 35

- (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 40

- (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 42

- (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 46

- (Exam Topic 5)

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

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

Answer: C

NEW QUESTION 47

- (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 48

- (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 49

- (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 50

- (Exam Topic 7)

You need to optimize the index and table structures for POSTransaction.

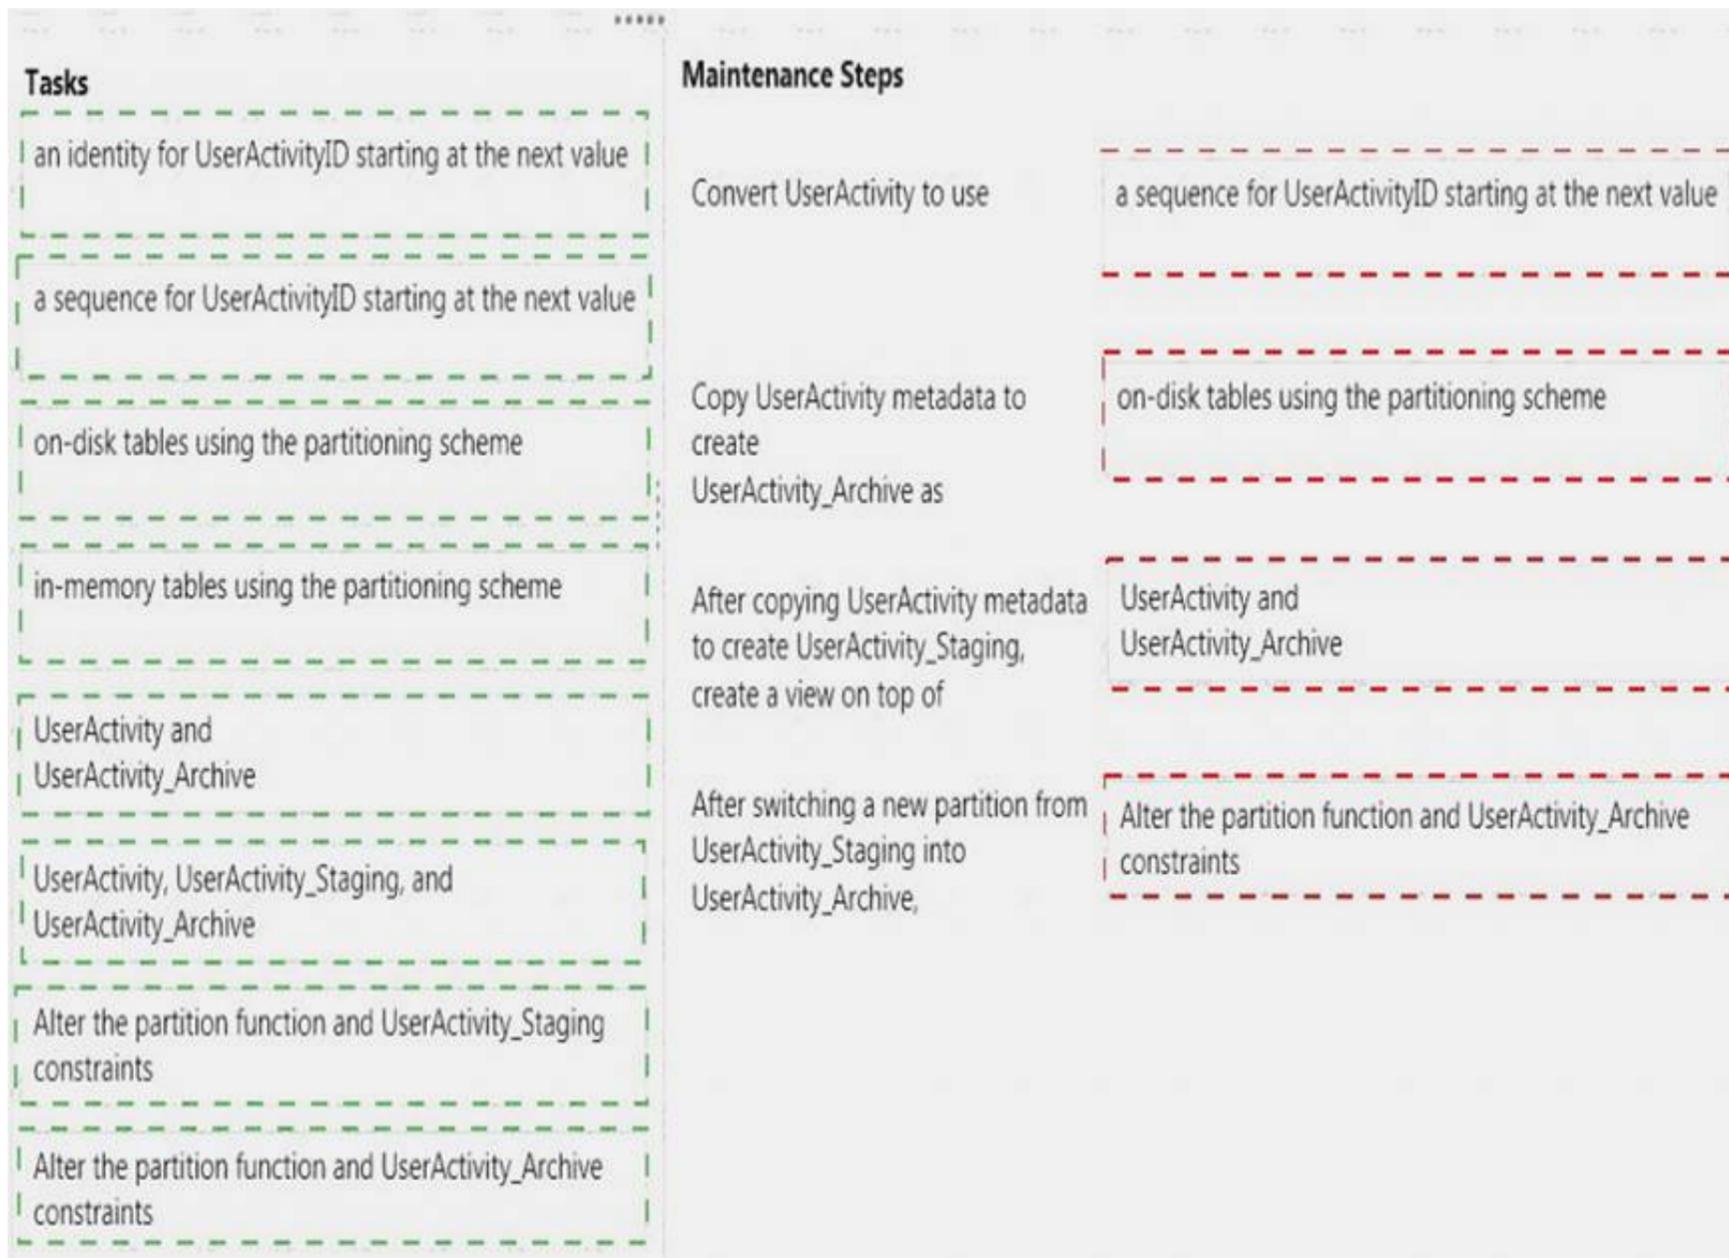
Which task should you use with each maintenance step? To answer, drag the appropriate tasks to the correct maintenance steps. Each task may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:



NEW QUESTION 53

- (Exam Topic 7)

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

- A. Query system DMVs to monitor avg_chain_length and max_chain_lengt
- B. Create alerts to notify you when these values converge.
- C. Create a SQL Agent alert when the File Table: Avg time per file I/O request value is increasing.
- D. Query system DMVs to monitor total_bucket_coun
- E. Create alerts to notify you when this value increases.
- F. Query system DMVs to monitor total_bucket_coun
- G. Create alerts to notify you when this value decreases.

Answer: A

Explanation:

From scenario:

- * You need to anticipate when POSTransaction table will need index maintenance.
- * The index maintenance strategy for the UserActivity table must provide the optimal structure for both maintainability and query performance.

NEW QUESTION 54

- (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 56

- (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=@POSTransactionId
```

Box 6:

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

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 it 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, whereas the journey of THROW statement has just begun with Sql Server 2012. Obviously, Microsoft is 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 61

- (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 65

- (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 66

- (Exam Topic 8)

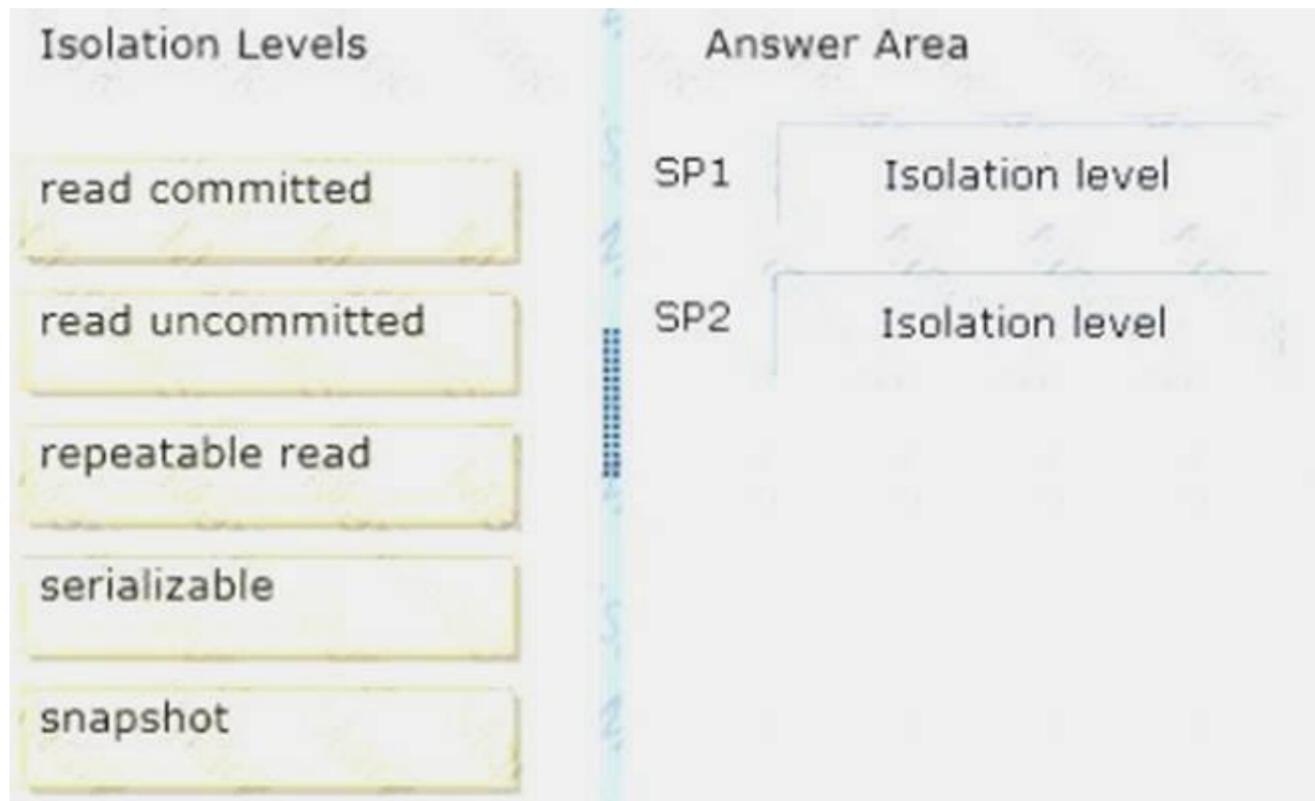
You plan to deploy two stored procedures named SP1 and SP2 that read data from the database. Your company identifies the following requirements for each stored procedure:

- SP1 must allow dirty reads.
- SP2 must place range locks on the data to ensure read consistency.

You need to identify which isolation level you must set for each stored procedure. The solution must minimize the number of locks.

Which isolation level should you identify?

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



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Note:

* READ UNCOMMITTED

Specifies that statements can read rows that have been modified by other transactions but not yet committed. Transactions running at the READ UNCOMMITTED level do not issue shared locks to prevent other transactions from modifying data read by the current transaction. READ UNCOMMITTED transactions are also not blocked by exclusive locks that would prevent the current transaction from reading rows that have been modified but not committed by other transactions. When this option is set, it is possible to read uncommitted modifications, which are called dirty reads. Values in the data can be changed and rows can appear or disappear in the data set before the end of the transaction. This option has the same effect as setting NOLOCK on all tables in all SELECT statements in a transaction. This is the least restrictive of the isolation levels.

* SERIALIZABLE

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.

Range locks are placed in the range of key values that match the search conditions of each statement executed in a transaction. This blocks other transactions from updating or inserting any rows that would qualify for any of the statements executed by the current transaction. This means that if any of the statements in a transaction are executed a second time, they will read the same set of rows. The range locks are held until the transaction completes. This is the most restrictive of the isolation levels because it locks entire ranges of keys and holds the locks until the transaction completes. Because concurrency is lower, use this option only when necessary.

Reference: SET TRANSACTION ISOLATION LEVEL (Transact-SQL)

NEW QUESTION 71

- (Exam Topic 8)

You plan to design an application that temporarily stores data in a SQL Azure database.

You need to identify which types of database objects can be used to store data for the application. The solution must ensure that the application can make changes to the schema of a temporary object during a session.

Which type of objects should you identify?

- A. Common table expressions (CTEs)
- B. Temporary stored procedures
- C. Temporary tables
- D. Table variables

Answer: C

Explanation:

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

NEW QUESTION 72

- (Exam Topic 8)

You use the following statement to create a table.

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

You have the following queries.

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

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

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

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

Code Elements

- DepartmentId
- EmployeeId
- FirstName
- LastName
- OfficeId

Answer Area

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: LastName

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

Box 2: FirstName

Box 3: DepartmentID

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

Box 4: OfficeID

NEW QUESTION 75

- (Exam Topic 8)

You have a Microsoft SQL Azure database named DBAzure1. DBAzure1 contains a table named Orders that stores sales data.

Each order has a sales total that can only be discovered by querying multiple tables.

You need to ensure that the value of the sales total is returned by executing a query on Orders. What should you create?

- A. A calculated column that uses a scalar function

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

Answer: C

Explanation:

A table-valued parameter is scoped to the stored procedure, function, or dynamic Transact-SQL text, exactly like other parameters. Similarly, a variable of table type has scope like any other local variable that is created by using a DECLARE statement. You can declare table-valued variables within dynamic Transact-SQL statements and pass these variables as table-valued parameters to stored procedures and functions.

Table-valued parameters offer more flexibility and in some cases better performance than temporary tables or other ways to pass a list of parameters.

Incorrect:

Not A: A scalar function would only be able to use other columns from the same table.

NEW QUESTION 77

- (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 81

- (Exam Topic 8)

You have a SQL Server 2012 database that contains a table named Users. The Users table contains usernames and passwords.

You need to ensure that all new records have a password. Which code segment should you use?

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

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

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

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

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

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

**Answer:** D

**NEW QUESTION 83**

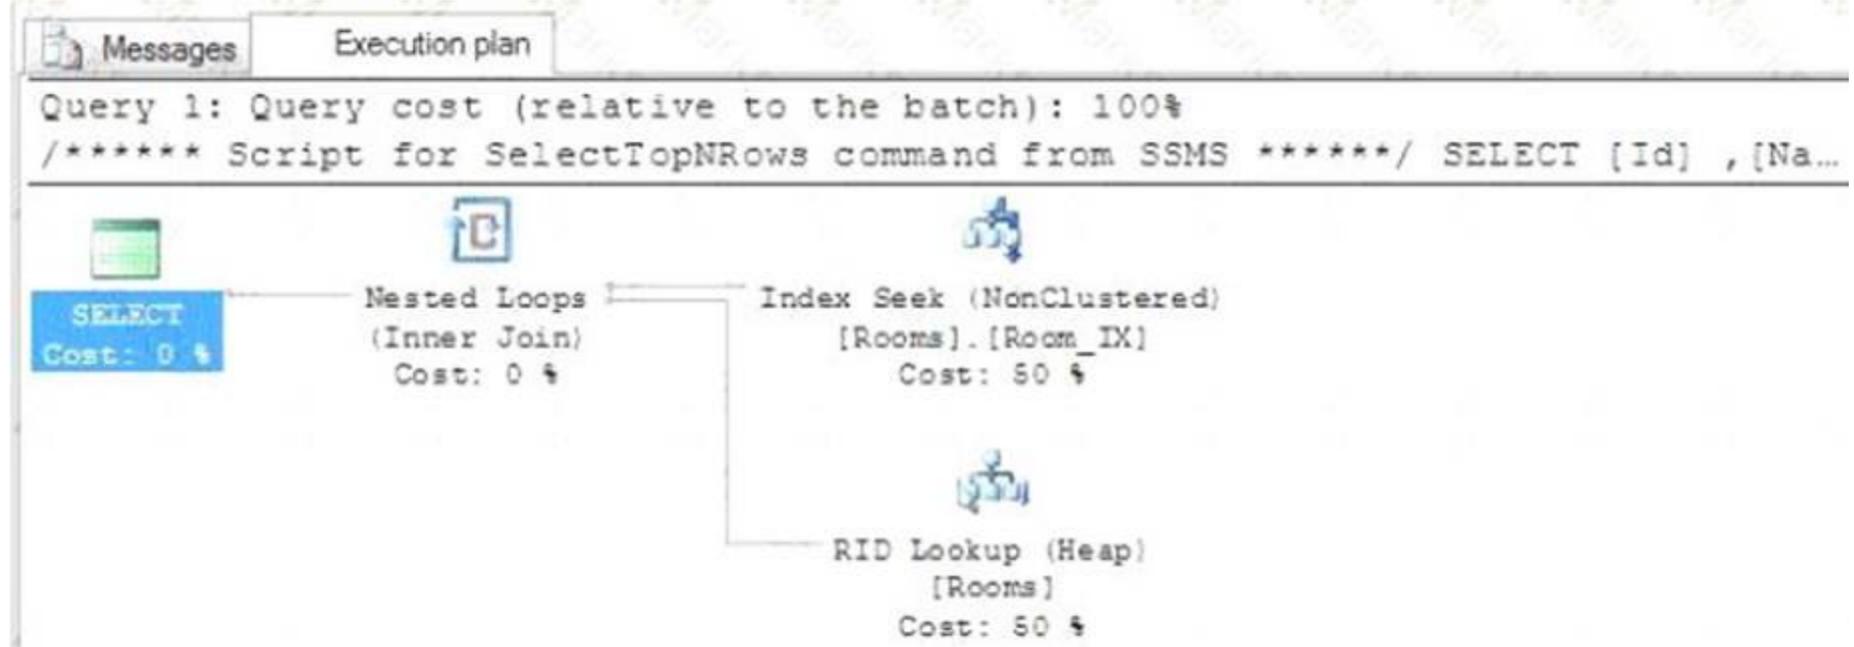
- (Exam Topic 8)

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 86**

- (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 88**

- (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 90**

- (Exam Topic 8)

Your network contains a server named SQL1 that has SQL Server 2012 installed. SQL1 contains a database name DB1 and a table named Customers.

You add an additional server named SQL2 that runs SQL Server 2012.

You need to create a distributed partitioned view. The solution must minimize the amount of network traffic. What should you do? (Each correct answer presents

part of the solution. Choose all that apply.)

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

**Answer:** BCDF

### NEW QUESTION 93

- (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 98

- (Exam Topic 8)

You have a table named Table1 that contains 1 million rows. Table1 contains a column named Column1 that stores sensitive information. Column1 uses the nvarchar(16) data type.

You have a certificate named Cert1.

You need to replace Column1 with a new encrypted column that uses two-way encryption. Which code segment should you execute before you remove Column1? To answer, move the appropriate code segments from the list of code segments to the answer area and arrange them in the correct order.

Code segments	Answer Area
<pre>ALTER TABLE Table1 ADD Column2 varbinary(256);</pre>	
<pre>CLOSE SYMMETRIC KEY;</pre>	
<pre>CREATE SYMMETRIC KEY Key1 WITH ALGORITHM = SHA1 ENCRYPTION BY CERTIFICATE Cert1;</pre>	
<pre>CREATE CREDENTIAL Cred1 WITH IDENTITY = 'User1', SECRET = 'P@ssw0rd';</pre>	
<pre>ALTER TABLE Table1 ADD Column2 nvarchar(256);</pre>	
<pre>OPEN SYMMETRIC KEY Key1 DECRYPTION BY CERTIFICATE Cert1;</pre>	
<pre>UPDATE table1 SET Column2 = EncryptByKey(Key_GUID ('Key1'),Column1);</pre>	
<pre>CREATE SYMMETRIC KEY Key1 WITH ALGORITHM = AES_256 ENCRYPTION BY CERTIFICATE Cert1;</pre>	

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Note:  
 \* Use AES\_256 for two-way encryption.  
 \* Use varbinary to store key.  
 \* CLOSE SYMMETRIC KEY (Transact-SQL)  
 Closes a symmetric key, or closes all symmetric keys open in the current session.  
 \* Example:  
 CREATE SYMMETRIC KEY CreditCards\_Key11 WITH ALGORITHM = AES\_256  
 ENCRYPTION BY CERTIFICATE Sales09; GO  
 -- Create a column in which to store the encrypted data. ALTER TABLE Sales.CreditCard  
 ADD CardNumber\_Encrypted varbinary(128); GO  
 -- Open the symmetric key with which to encrypt the data. OPEN SYMMETRIC KEY CreditCards\_Key11 DECRYPTION BY CERTIFICATE Sales09;  
 -- Encrypt the value in column CardNumber using the  
 -- symmetric key CreditCards\_Key11.  
 -- Save the result in column CardNumber\_Encrypted. UPDATE Sales.CreditCard  
 SET CardNumber\_Encrypted = EncryptByKey(Key\_GUID('CreditCards\_Key11')  
 , CardNumber, 1, HashBytes('SHA1', CONVERT( varbinary  
 , CreditCardID))); GO

**NEW QUESTION 102**

- (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-usage-stats>

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

**NEW QUESTION 103**

- (Exam Topic 8)

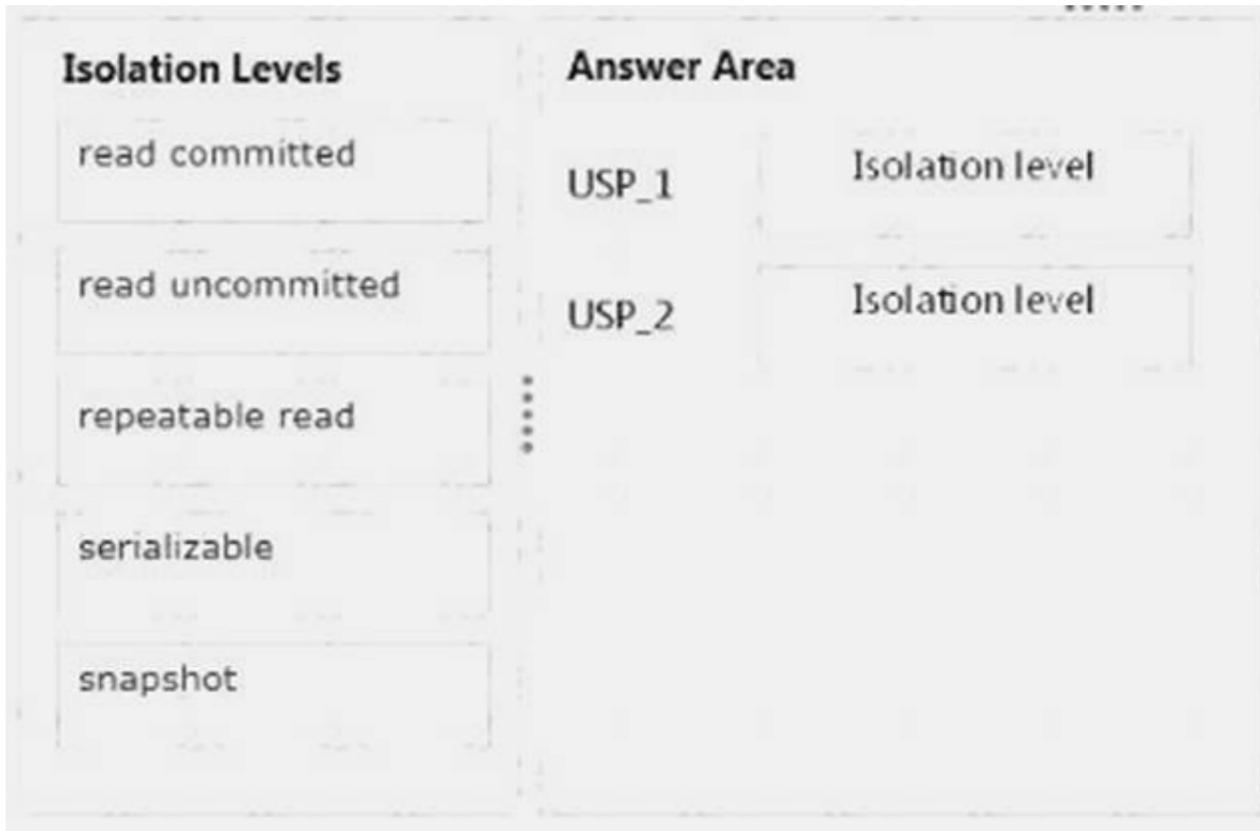
You plan to deploy two stored procedures name USP\_1 and USP\_2 that read data from a database. Your company identifies the following requirements for each stored procedure:

- 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 108**

- (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 Schema1.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 109**

- (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 111**

- (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 113**

- (Exam Topic 8)

You have a database for a mission-critical web application. The database is stored on a SQL Server 2012 instance and is the only database on the instance.

The application generates all T-SQL statements dynamically and does not use stored procedures. You need to maximize the amount of memory available for data caching.

Which advanced server option should you modify?

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

**Answer:** A

**NEW QUESTION 117**

- (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 120**

- (Exam Topic 8)

You have the following stored procedure.

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

The stored procedure takes much longer to execute than expected.

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

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

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

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

**Answer:** A

**Explanation:**

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

References:

[https://sqlserverperformance.wordpress.com/2009/02/02/beware-of-convert\\_implicit-in-a-sql-execution-plan/](https://sqlserverperformance.wordpress.com/2009/02/02/beware-of-convert_implicit-in-a-sql-execution-plan/)

**NEW QUESTION 124**

- (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	Answer Area
sys.dm_db_index_usage_stats	Identify unused indexes.
sys.dm_db_index_operational_stats	Identify which indexes should be created.
sys.dm_db_index_physical_stats	
sys.dm_db_missing_index_columns	
sys.dm_db_missing_index_details	
sys.dm_db_missing_index_groups	

- A. Mastered
- B. Not Mastered

Answer: A

**Explanation:**

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-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 127**

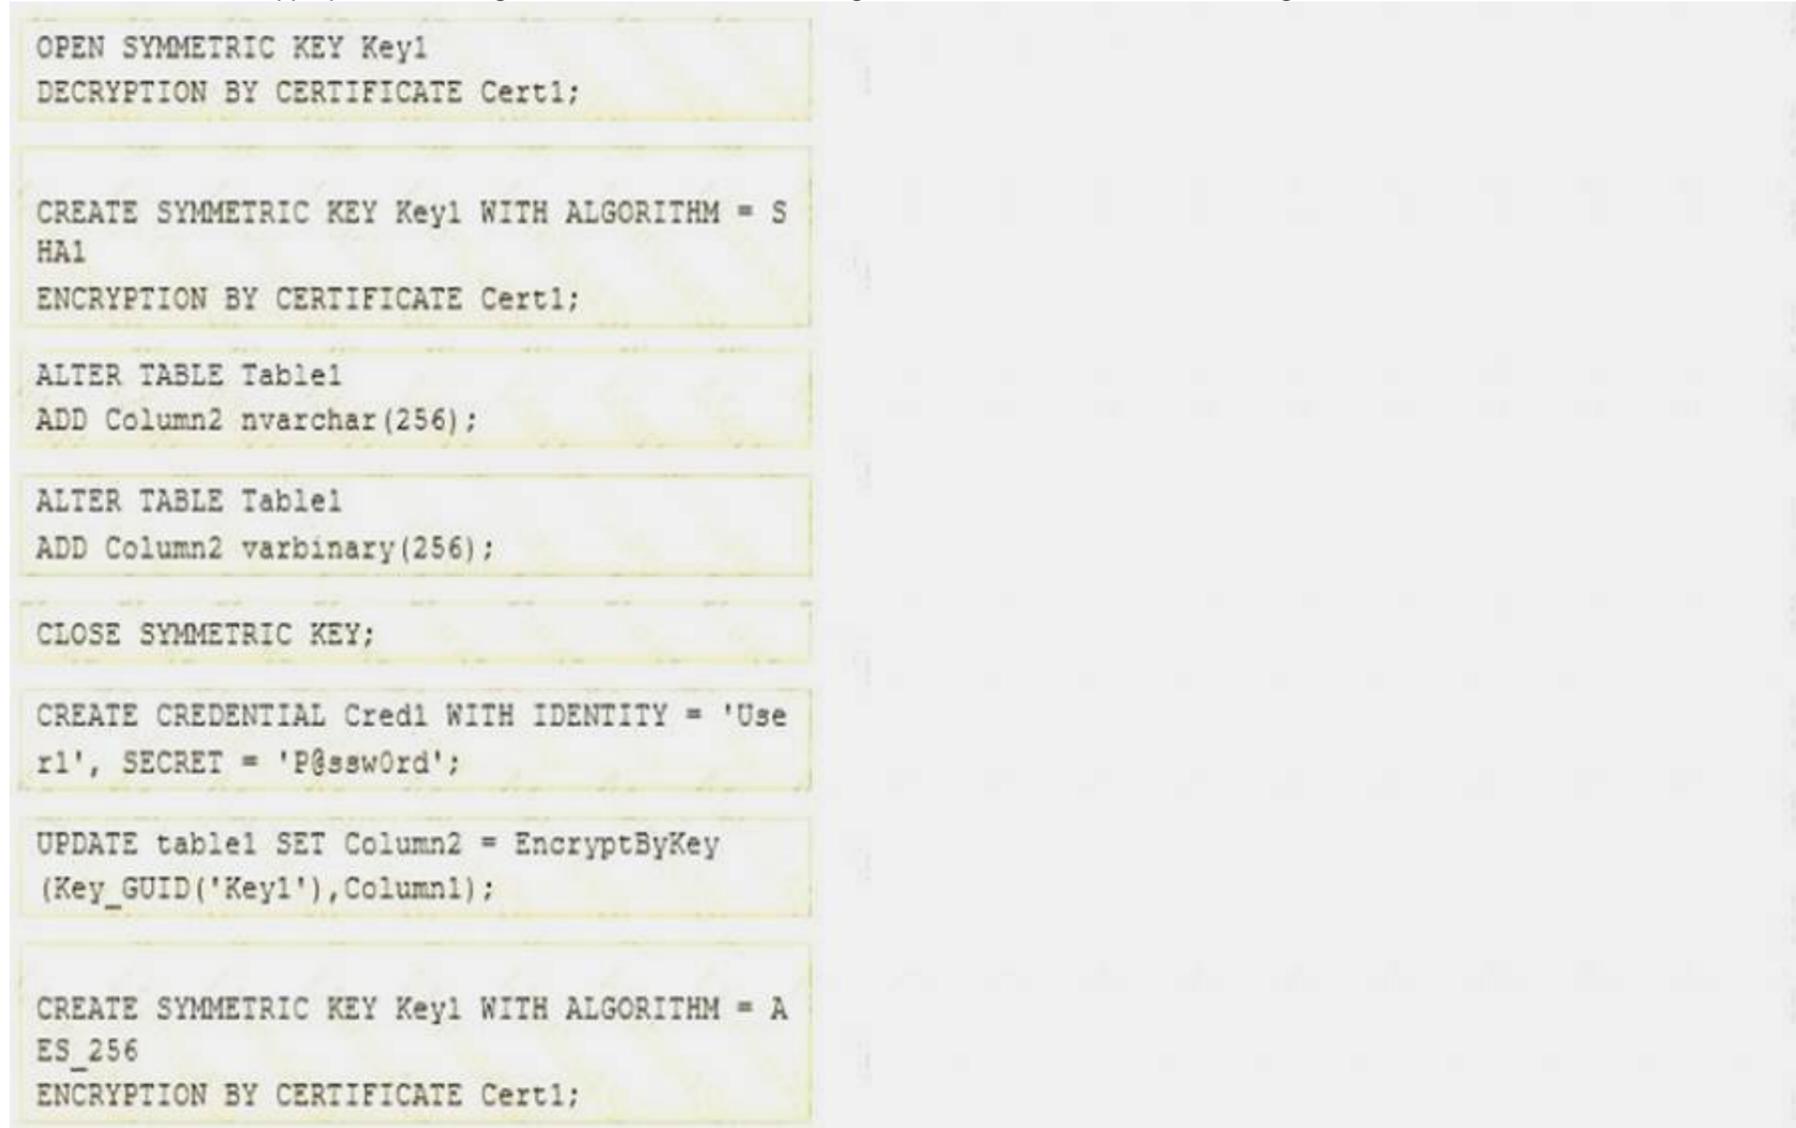
- (Exam Topic 8)

You have a table named Table1 that contains 1 million rows. Table1 contains a column named Column1 that stores sensitive information. Column1 uses the nvarchar (16) data type.

You have a certificate named Cert1.

You need to replace Column1 with a new encrypted column named Column2 that uses one-way hashing. Which code segment should you execute before you remove Column1?

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



- A. Mastered
- B. Not Mastered

Answer: A

**Explanation:**

Box 1:

```
CREATE SYMMETRIC KEY Key1 WITH ALGORITHM = SHA1
ENCRYPTION BY CERTIFICATE Cert1;
```

First create a hash key using the certificate. Not AES: AES is not based on hashing. Box 2:

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

Add a column with varbinary data type. Box 3:

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

Box 4:

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

Box 5:

```
CLOSE SYMMETRIC KEY;
```

Note:

\* There are a few different hashing algorithms available in SQL Server 2005: MD2, MD4, MD5, SHA, SHA1, with each having pros and cons.

\* In cryptography, SHA-1 is a cryptographic hash function designed by the United States National Security Agency and published by the United States NIST as a US Federal Information Processing Standard. SHA stands for "secure hash algorithm". The four SHA algorithms are structured differently and are distinguished as SHA-0, SHA-1, SHA-2, and SHA-3. SHA-1 is very similar to SHA-0, but corrects an error in the original SHA hash specification that led to significant weaknesses. The SHA-0 algorithm was not adopted by many applications. SHA-2 on the other hand significantly differs from the SHA-1 hash function. SHA-1 is the most widely used of the existing SHA hash functions, and is employed in several widely used applications and protocols.

\* To encrypt a column of data using a simple symmetric encryption in Object Explorer, connect to an instance of Database Engine.

On the Standard bar, click New Query.

Copy and paste the following example into the query window and click Execute. USE AdventureWorks2012;

-- If there is no master key, create one now. IF NOT EXISTS

```
(SELECT * FROM sys.symmetric_keys WHERE symmetric_key_id = 101) CREATE MASTER KEY ENCRYPTION BY
PASSWORD = '23987hxJKL95QYV4369#ghf0%lekjg5k3fd117r$$#1946kcj$N44ncjhdj' GO
```

```
CREATE CERTIFICATE Sales09
```

```
WITH SUBJECT = 'Customer Credit Card Numbers'; GO
```

```
CREATE SYMMETRIC KEY CreditCards_Key11 WITH ALGORITHM = AES_256
```

```
ENCRYPTION BY CERTIFICATE Sales09; GO
```

```
-- Create a column in which to store the encrypted data. ALTER TABLE Sales.CreditCard
```

```
ADD CardNumber_Encrypted varbinary(128); GO
```

```
-- Open the symmetric key with which to encrypt the data. OPEN SYMMETRIC KEY CreditCards_Key11 DECRYPTION BY CERTIFICATE Sales09;
```

```
-- Encrypt the value in column CardNumber using the
```

```
-- symmetric key CreditCards_Key11.
```

```
-- Save the result in column CardNumber_Encrypted. UPDATE Sales.CreditCard
```

```
SET CardNumber_Encrypted = EncryptByKey(Key_GUID('CreditCards_Key11')
```

```
, CardNumber, 1, HashBytes('SHA1', CONVERT(varbinary
```

```
, CreditCardID))); GO
```

Reference: SQL Server 2012, Encrypt a Column of Data

Ref:

<http://www.mssqltips.com/sqlservertip/2431/sql-server-column-level-encryption-example-using-symmetric-keys>

### NEW QUESTION 128

- (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 131

- (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
SERIALIZABLE
WITH(UPDLOCK)
WITH(XLOCK)

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

```

- 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 136**

- (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 141**

- (Exam Topic 8)

You have a database named Database1. Database1 has two stored procedures named Proc1 and Proc2 and a table named Table1. Table1 has millions of rows. Proc1 updates data in Table1. Proc2 reads data from Table1.

You discover that when Proc1 is executed to update more than 4,000 rows, Proc2 is blocked. The block affects all rows, including those that are not being updated by Proc1.

You need to ensure that when Proc1 is executing, Proc2 can access the data in Table1 that Proc1 is not updating.

What should you change Proc1 to do?

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

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

**Answer: A**

**NEW QUESTION 143**

- (Exam Topic 8)

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

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

You create a non-clustered index named IX\_Name on the name column.

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

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

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

- A. Replace LEFT(name,3) = 'smi' by using name like 'smi%'
- B. Replace LEFT(name,3) = 'smi' by using substring(name,1,3) = 'smi'
- C. Recreate IX\_Name as a unique index
- D. Recreate IX Name as a clustered index

**Answer: A**

**NEW QUESTION 145**

- (Exam Topic 8)

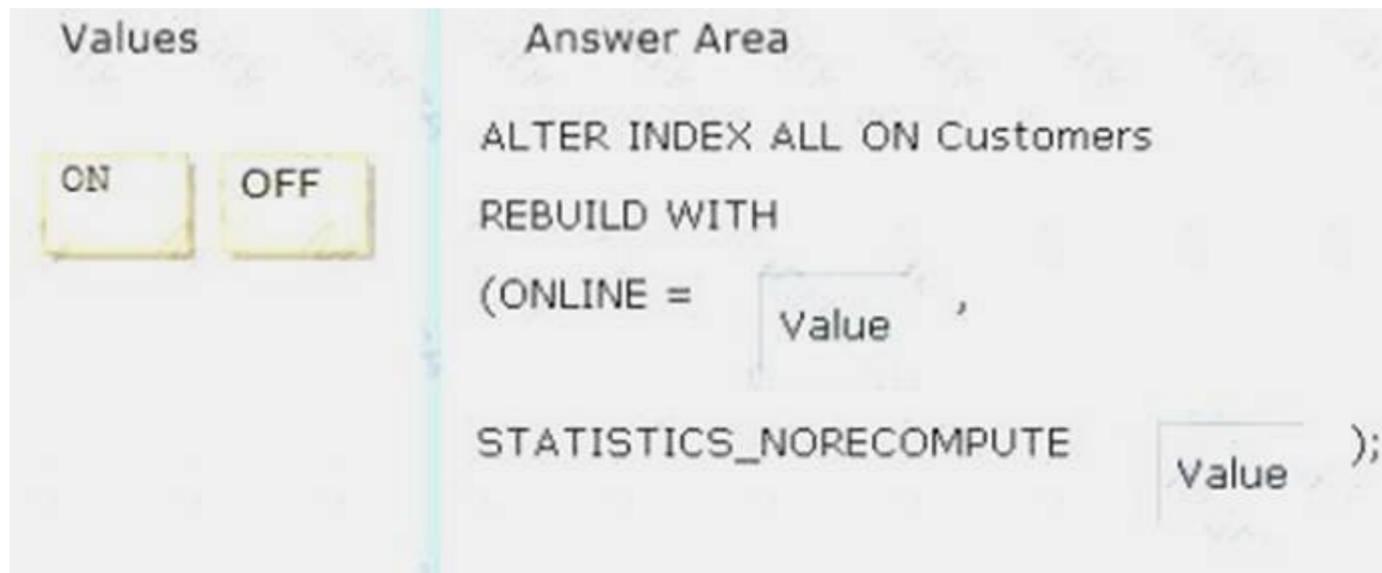
You run the following code segment:

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

After you add 10,000 rows to Customers, you discover that the index is fragmented. You need to defragment the index in the least amount of time.

Which code segment should you execute?

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



- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Note:

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

\* Online = OFF

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

\* STATISTICS\_NORECOMPUTE = ON

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

**NEW QUESTION 150**

- (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 153**

- (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 157**

- (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 160**

- (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 164

- (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 169

- (Exam Topic 8)

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

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

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

**Answer:** A

#### NEW QUESTION 173

- (Exam Topic 8)

You plan to create a custom aggregation function named Median.

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

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

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

What three tasks should you perform next?

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

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

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: Execute the CREATE ASSEMBLY statement.

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

Note:

\* CREATE AGGREGATE

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

**NEW QUESTION 174**

- (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 179**

- (Exam Topic 8)

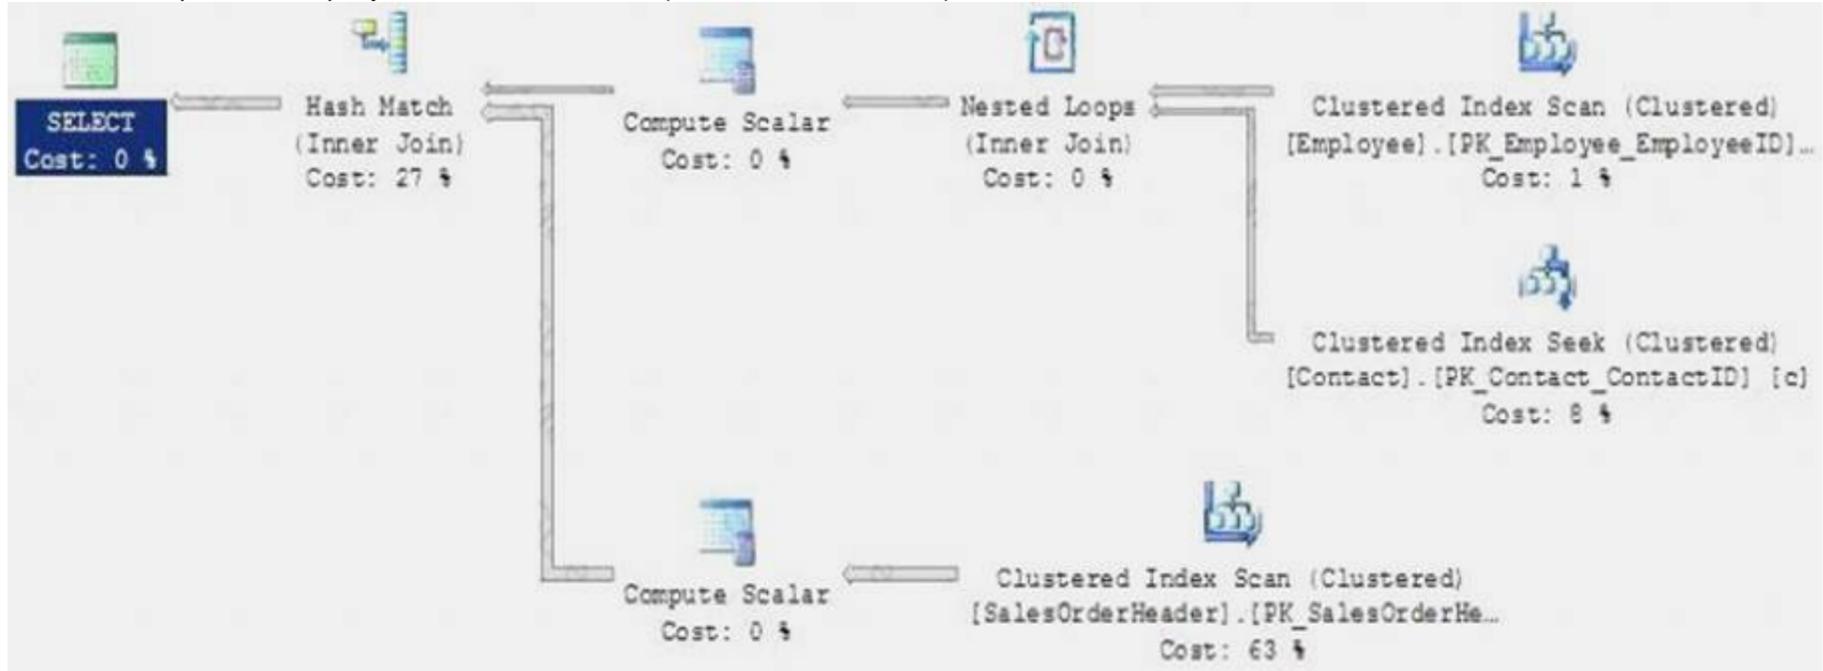
You have a database that contains three tables. The tables are configured as shown in the following table.

Table	Primary key index
SalesOrderHeader	PK_SalesOrderHeader_SalesOrderID
Employee	PK_Employee_EmployeeID
Contact	PK_Contact_ContactID

You have the following query:

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

The execution plan for the query is shown in the exhibit. (Click the Exhibit button.)



You need to create one index to minimize the amount of time it takes to execute the query.

What should you do?

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

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

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Note:

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

**NEW QUESTION 184**

- (Exam Topic 8)

You use SQL Server 2012 to store data used by an e-commerce application.

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

You need to ensure that sp1 can read data even while another transaction is modifying the price of a product. Sp1 must only read committed data.

Which transaction isolation level should you use in sp1?

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

**Answer: B**

#### NEW QUESTION 189

- (Exam Topic 8)

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

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

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

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

**Answer: B**

#### Explanation:

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

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

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

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

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

#### NEW QUESTION 191

- (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 192

- (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.

Options

- Sparse
- Computed
- Persisted computed

Answer Area

Column1  Option

Column2  Option

- 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 196

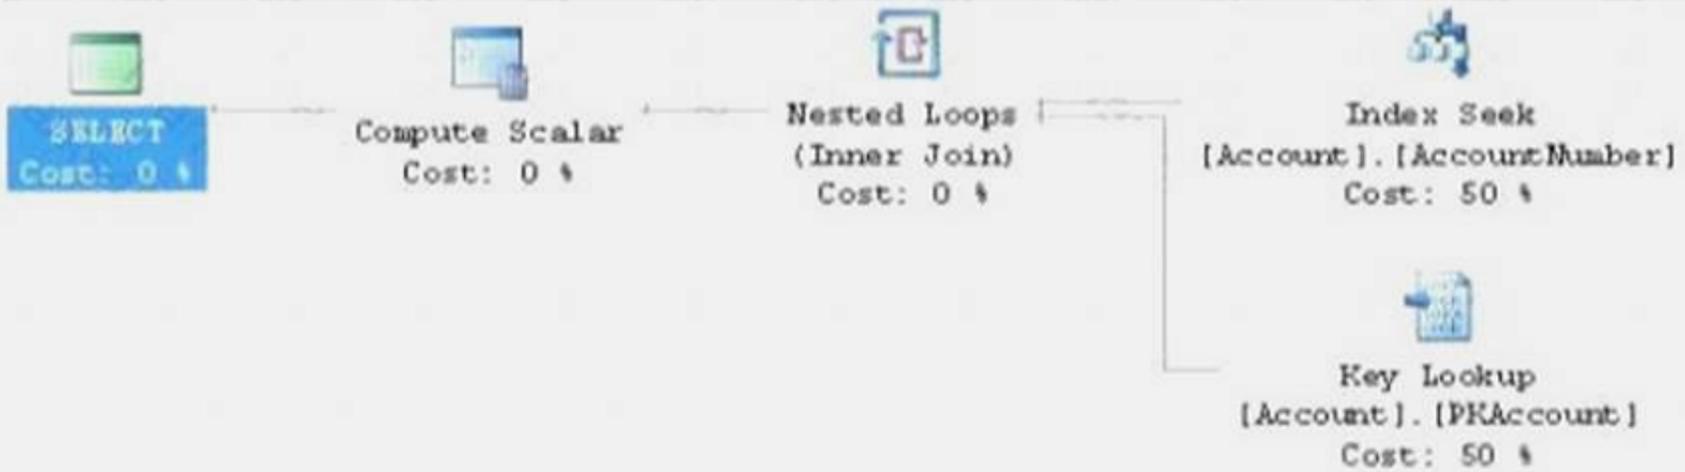
- (Exam Topic 8)

You administer an instance of SQL Server 2014.

You are tasked with tuning a common set of queries. You have the results of several test executions, along with query plans. The schema and the data for all database object(s) used remain unchanged between executions. The QueryTime column is defined as a computed column that uses the GETDATE() system function. The query plans and results are shown below:

```
SELECT *
FROM dbo.Account
WHERE AccountNumber = 'A10000001'
```

Query 1: Query cost (relative to the batch): 100%  
 SELECT \* FROM [dbo].[Account] WHERE [AccountNumber]=@1



AccountID	AccountNumber	Name	QueryTime
-----	-----	-----	-----
-----	-	-	----
0F638176-7257-4480-9D0E-126C45 CEFFF1	A10000001	Don Hall	2014-01-29 18:01:50.923

```
SELECT *
FROM dbo.Account
WHERE AccountNumber IN('A10000001', 'Q88700323', 'R00000012')
GO
```

Query 1: Query cost (relative to the batch): 100%

```
SELECT * FROM dbo.Account WHERE AccountNumber IN('A10000001', 'Q88700323', 'R00000012')
```

AccountID	AccountNumber	Name	QueryTime
----- -----	-	-	----
0F63B176-7257-4480-9D0E-126C45 CEFFF1	A10000001	Don Hall	2014-01-29 20:14:05.660
337227AA-3A4B-4B28-8E02-0ADEAD 06EA10	Q88700323	Darren Parker	2014-01-29 20:14:05.660
C4980E64-874E-4640-8826- BAF35D8FB845	R00000012	Carol Philips	2014-01-29 20:14:05.660

You need to make an initial diagnosis of the situation, based solely on this input

Which two statements can you make about the performance characteristics of this query? Each correct answer presents a complete solution. Choose two.

- A. The queries would perform better if the index named AccountNumber included the Name and QueryTime column.
- B. The queries would perform worse if the index named AccountNumber included the NameColumn.
- C. The queries would perform better if the index named AccountNumber included the Name column.
- D. The object Account is a table, with an index having a leading column of AccountNumber and a Clustered Index named PKAccount.
- E. The object Account is an indexed view, with an index having a leading column of AccountNumber and a Clustered Index named PKAccount.
- F. The object Account is a view, joining the Account-AccountNumber and Account.PKAccount objects together.

Answer: BD

**NEW QUESTION 199**

- (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 200**

- (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 many page splits occur when records are inserted or updated in the table. You need to minimize the number of page splits.

What should you set from the index options?

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

**Answer:** D

**Explanation:**

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

**NEW QUESTION 202**

- (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 207**

- (Exam Topic 8)

You have a database named database1.

Database developers report that there are many deadlocks.

You need to implement a solution to monitor the deadlocks. The solution must meet the following requirements:

- Support real-time monitoring.
- Be enabled and disabled easily.
- Support querying of the monitored data. What should you implement?

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

- A. Log errors by using trace flag 1222
- B. Log errors by using trace flag 1204
- C. A SQL Server Profiler template
- D. An Extended Events session

**Answer:** D

**Explanation:**

<http://www.sqlservercentral.com/blogs/james-sql-footprint/2012/08/12/monitor-deadlock-in-sql-2012/>

[http://blogs.technet.com/b/mspfe/archive/2012/06/28/how\\_2d00\\_to\\_2d00\\_monitor\\_2d00\\_deadlocks\\_2d00\\_in\\_2](http://blogs.technet.com/b/mspfe/archive/2012/06/28/how_2d00_to_2d00_monitor_2d00_deadlocks_2d00_in_2)

**NEW QUESTION 208**

- (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 211**

.....

## 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