

Exam Questions 70-464

Developing Microsoft SQL Server 2012 Databases

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



NEW QUESTION 1

- (Exam Topic 1)

Which data type should you use for CustomerID?

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

Answer: D

Explanation:

Invoices.xml

All customer IDs are 11 digits. The first three digits of a customer ID represent the customer's country. The remaining eight digits are the customer's account number.

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

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

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

NEW QUESTION 2

- (Exam Topic 1)

You execute IndexManagement.sql and you receive the following error message: "Msg 512, Level 16, State 1, Line 12 Subquery returned more than 1 value. This is not permitted when the subquery follows =, !=, <, <=, >, >= or when the subquery is used as an expression."

You need to ensure that IndexManagement.sql executes properly. Which WHILE statement should you use at line 18?

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

Answer: B

NEW QUESTION 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 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 6

- (Exam Topic 2)

You execute usp_TestSpeakers.

You discover that usp_SelectSpeakersByName uses inefficient execution plans.

You need to update usp_SelectSpeakersByName to ensure that the most efficient execution plan is used. What should you add at line 30 of Procedures.sql?

- A. OPTION (FORCESCAN)
- B. OPTION (FORCESEEK)
- C. OPTION (OPTIMIZE FOR UNKNOWN)
- D. OPTION (OPTIMIZE FOR (@LastName= 'Anderson'))

Answer: C

Explanation:

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

NEW QUESTION 7

- (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.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 8

- (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 9

- (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 10

- (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 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 14

- (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 19

- (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 21

- (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 25

- (Exam Topic 4)

You execute IndexManagement.sql and you receive the following error message: "Msg 512, Level 16, State 1, Line 12 Subquery returned more than 1 value. This is not permitted when the subquery follows =, !=, <, <=, >, >= or when the subquery is used as an expression."

You need to ensure that IndexManagement.sql executes properly. Which WHILE statement should you use at line 18?

- A. WHILE SUM(@RowNumber) < (SELECT @counter FROM @indextable)

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

Answer: D

NEW QUESTION 30

- (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 35

- (Exam Topic 5)

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

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

Answer: D

NEW QUESTION 37

- (Exam Topic 6)

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

- A. ALTER INDEX IX_Orders_ShipDate ON OrdersREBUILD WITH (PAD_INDEX=OFF, DROP_EXISTING = ON);
- B. ALTER INDEX IX_Orders_ShipDate ON OrdersREBUILD WITH (FILLFACTOR=50, DROP_EXISTING = ON);
- C. ALTER INDEX IX_Orders_ShipDate ON OrdersREBUILD WITH (FILLFACTOR = 0, DROP_EXISTING = ON);
- D. ALTER INDEX IX_Orders_ShipDate ON OrdersREBUILD WITH (PAD_INDEX=ON, DROP_EXISTING = ON);

Answer: B

NEW QUESTION 41

- (Exam Topic 6)

You need to implement a solution that solves the performance issues of usp_GetOrdersAndItems.

Which statements should you execute?

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

Answer: B

NEW QUESTION 42

- (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 43

- (Exam Topic 6)

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

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

Answer: D

NEW QUESTION 45

- (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 46

- (Exam Topic 7)

You need to redesign the system to meet the scalability requirements of the application. Develop the solution by selecting and arranging the required code blocks in the correct order. You may not need all of the code blocks.

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1:

```

ALTER DATABASE CoffeeTransactions
ADD FILEGROUP [CoffeeTransactions_inmem
] CONTAINS MEMORY_OPTIMIZED_DATA

```

Box 2:

```

CREATE TABLE dbo.POSTransaction (

```

Box 3:

```

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

```

Box 4:

```

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

```

Box 5:

```

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

```

Box 6:

```

    WITH (MEMORY_OPTIMIZED=ON,
    DURABILITY=SCHEMA_ONLY)

```

Box 7:

```

    ON [CoffeeTransactions_inmem]

```

Note:

* MEMORY_OPTIMIZED_DATA

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

* You must specify a value for the BUCKET_COUNT parameter when you create the memory-optimized table. In most cases the bucket count should be between 1 and 2 times the number of distinct values in the index key.

* Example:

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

NEW QUESTION 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

an identity for UserActivityID starting at the next value

a sequence for UserActivityID starting at the next value

on-disk tables using the partitioning scheme

in-memory tables using the partitioning scheme

UserActivity and UserActivity_Archive

UserActivity, UserActivity_Staging, and UserActivity_Archive

Alter the partition function and UserActivity_Staging constraints

Alter the partition function and UserActivity_Archive constraints

Maintenance Steps

Convert UserActivity to use

Copy UserActivity metadata to create UserActivity_Archive as

After copying UserActivity metadata to create UserActivity_Staging, create a view on top of

After switching a new partition from UserActivity_Staging into UserActivity_Archive,

Task

Task

Task

Task

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

NEW QUESTION 52

- (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 54

- (Exam Topic 7)

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

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

Answer: A

Explanation:

The fraud detection service will need to meet the following requirement (among others):

* Detect micropayments that are flagged with a StatusId value that is greater than 3 and that occurred within the last minute.

NEW QUESTION 55

- (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 60

- (Exam Topic 8)

Your company has a main office in London and a branch office in New York.

Your network contains a server named Server5 that has SQL Server 2012 installed. Server5 contains a database name ContentDB and a table named ContentTable.

You add an additional server named Server9 that runs SQL Server 2012.

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

- A. Create the view on Server5.
- B. Add Server9 as a linked server.
- C. Create the view on Server9.
- D. Add the Customers table to Server9.
- E. Add Server9 as a Distributor.
- F. Remove the Customers table from Server5.

Answer: ABCD

NEW QUESTION 62

- (Exam Topic 8)

Your network contains a server named Server1 that runs SQL Server 2012. Server1 contains an instance named Instance1. Instance1 contains a database named ContentDatabase.

ContentDatabase uses transaction log backups.

The recovery model of ContentDatabase is set to FULL.

You need to shrink the ContentDatabase_Log log file to 10 MB. The solution must ensure that you can continue to back up the transaction log.

Which three code segments should you execute?

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

Answer Area

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

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

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

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

```
BACKUP LOG ContentDatabase
WITH TRUNCATE_ONLY
```

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1:

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

Box 2:

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

Box 3:

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

Note:

* Shrinking a log file to a specified target size

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

```
USE AdventureWorks2012; GO
```

```
-- Truncate the log by changing the database recovery model to SIMPLE. ALTER DATABASE AdventureWorks2012
SET RECOVERY SIMPLE; GO
```

```
-- Shrink the truncated log file to 1 MB.
```

```
DBCC SHRINKFILE (AdventureWorks2012_Log, 1); GO
```

```
-- Reset the database recovery model. ALTER DATABASE AdventureWorks2012 SET RECOVERY FULL;
GO
```

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

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

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

Arguments include: target_size

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

NEW QUESTION 63

- (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 67

- (Exam Topic 8)

You administer a SQL Server 2014 instance.

The server is capable of 10000 IO/second (IOPS). During the time period when the second process executes, the disk IO can reach 7000 IOPS, and CPU use can average 30% over the eight processors.

The first process summarizes the day's activity executed by a login of [SummaryReportLogin]. The second process submits transactions executed by a login of [ETLLogin].

A Resource Governor classifier function has been created to return WG_Low for connections from the [ETLLogin] and [SummaryReportLogin].

You need to set up the Resource Group and Workgroup Pools on the instance. You have the following requirements:

- ☒ Both processes must never use more than 50 percent of the CPU at any one time.
- ☒ The number of active queries that these processes can execute simultaneously should be limited to a maximum of 10.
- ☒ The SummaryReportLogin process must always achieve the minimum IOPS required to be minimally affected during executing the ETLLogin processes.

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

```
MAX_IOPS_PER_VOLUME=3000
)
```

```
CREATE WORKLOAD GROUP WG_Low
WITH
(
    MAX_DOP = 4
)
USING RP_Low
```

```
CREATE WORKLOAD GROUP WG_Low
WITH
(
    GROUP_MAX_REQUESTS=10
)
USING RP_Low
```

```
CREATE WORKLOAD GROUP WG_Low
WITH
(
    REQUEST_MAX_CPU_TIME_SEC = 100,
    MAX_DOP = 4
)
USING RP_Low
```

```
CREATE RESOURCE POOL RP_Low
WITH
(
    CAP_CPU_PERCENT=50,
    MAX_CPU_PERCENT=30,
```

```
CREATE RESOURCE POOL RP_Low
WITH
(
    AFFINITY_SCHEDULER = (0 to 50),
    MAX_CPU_PERCENT=30,
```

```
CREATE RESOURCE POOL RP_Low
WITH
(
    MAX_CPU_PERCENT=50,

    MAX_IOPS_PER_VOLUME=30
)
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1:

```
CREATE RESOURCE POOL RP_Low
WITH
(
    CAP_CPU_PERCENT=50,
    MAX_CPU_PERCENT=30,
```

Box 2:

```
MAX_IOPS_PER_VOLUME=3000
)
```

Box 3:

```
CREATE WORKLOAD GROUP WG_Low
WITH
(
GROUP_MAX_REQUESTS=10
)
USING RP_Low
```

Note:

CREATE WORKLOAD RESOURCE POOL

* Resource pools. A resource pool, represents the physical resources of the server. You can think of a pool as a virtual SQL Server instance inside of a SQL Server instance.

* Workload groups. A workload group serves as a container for session requests that have similar classification criteria. A workload allows for aggregate monitoring of the sessions, and defines policies for the sessions. Each workload group is in a resource pool.

* CAP_CPU_PERCENT =value

Specifies a hard cap on the CPU bandwidth that all requests in the resource pool will receive. Limits the maximum CPU bandwidth level to be the same as the specified value. value is an integer with a default setting of 100. The allowed range for value is from 1 through 100.

* MIN_IOPS_PER_VOLUME =value

Specifies the minimum I/O operations per second (IOPS) per disk volume to reserve for the resource pool.

* GROUP_MAX_REQUESTS =value

Specifies the maximum number of simultaneous requests that are allowed to execute in the workload group. value must be a 0 or a positive integer.

NEW QUESTION 72

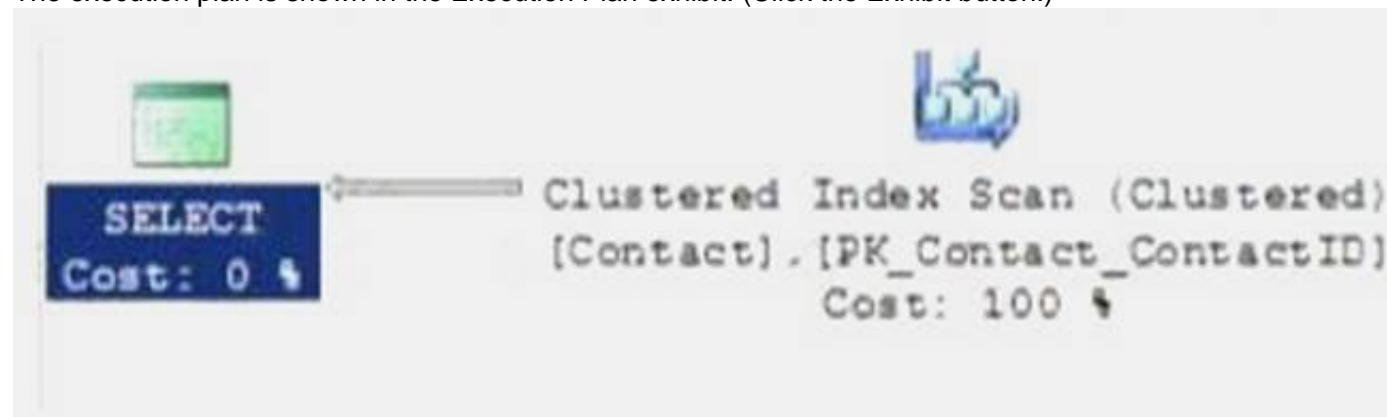
- (Exam Topic 8)

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 76

- (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 80

- (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 85

- (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 90

- (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 94

- (Exam Topic 8)

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

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

Answer: BE

Explanation:

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

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

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

NEW QUESTION 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 98

- (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 101

- (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 106

- (Exam Topic 8)

You have a database hosted on SQL Azure.

You are developing a script to create a view that will be used to update the data in a table.

The following is the relevant portion of the script. (Line numbers are included for reference only.)

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

You need to ensure that the view can update the data in the table, except for the data in Column1. Which code segment should you add at line 06?

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

Answer: A

Explanation:

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

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

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

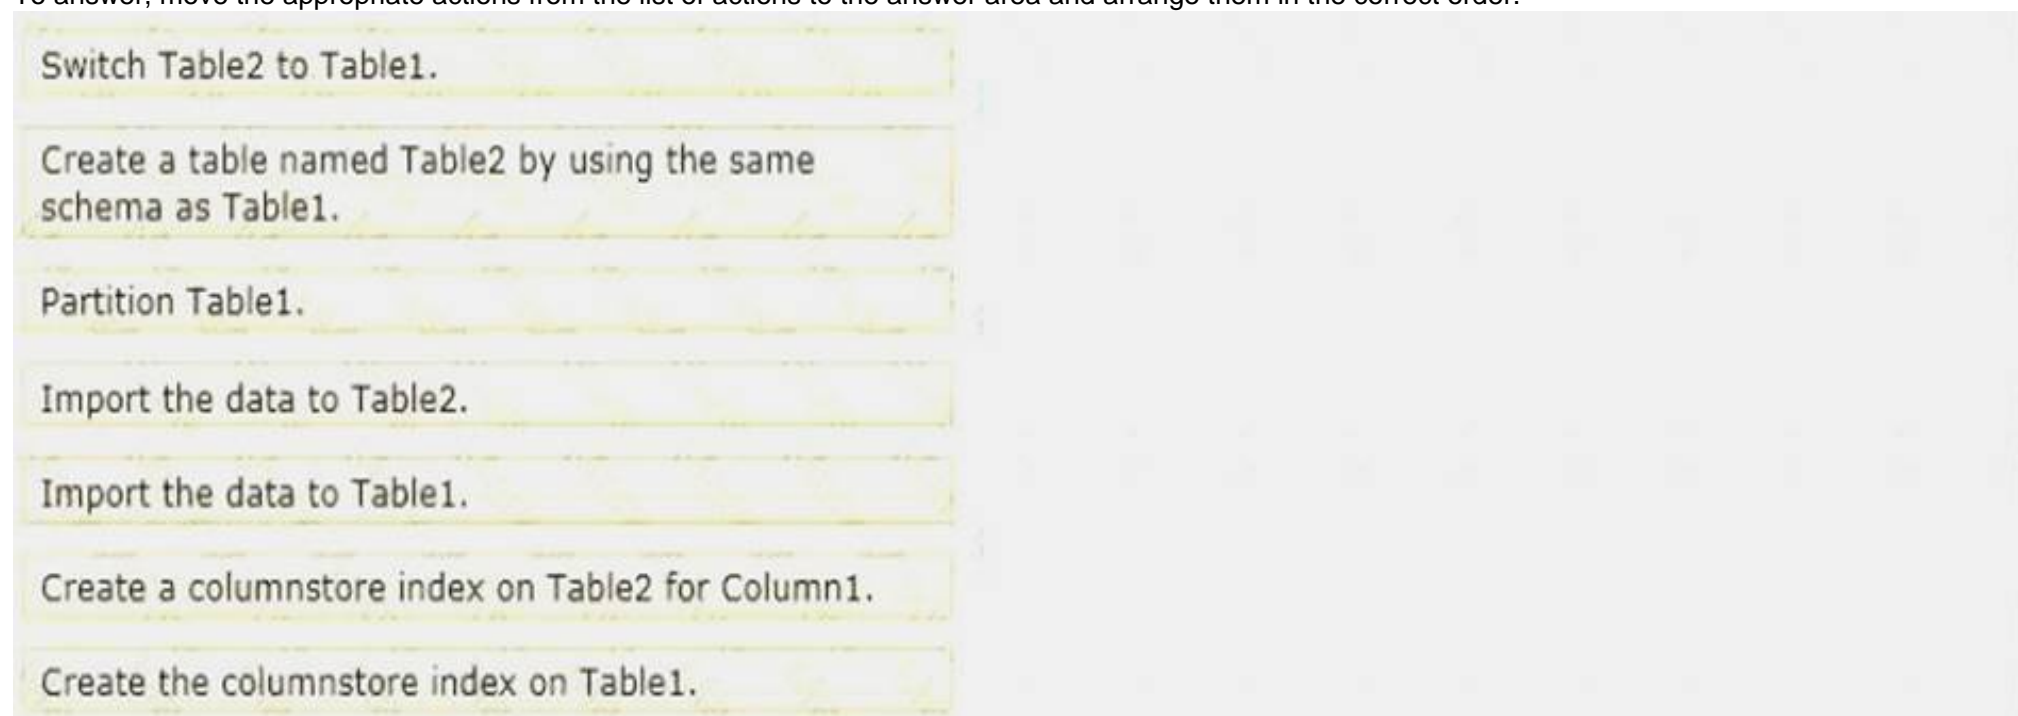
NEW QUESTION 111

- (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.



- 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 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 121

- (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 122

- (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 123

- (Exam Topic 8)
You have two existing tables, one named COUNTRY and the other named STATES. The tables are defined as follows:

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

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

Code Blocks

REFERENCES STATES (Country_Abbr)

REFERENCES COUNTRY (Country_Abbr)

ON STATES

FOREIGN KEY (Country_Abbr)

ON COUNTRY

ADD CONSTRAINT FK_StateCountry

ON COUNTRY_ABBR

ALTER TABLE COUNTRY

ADD FOREIGN KEY FK_StateCountry

ALTER TABLE STATES

Answer Area

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Note:
 To allow naming of a FOREIGN KEY constraint, and for defining a FOREIGN KEY constraint on multiple columns, use the following SQL syntax:
 MySQL / SQL Server / Oracle / MS Access:
 ALTER TABLE Orders ADD CONSTRAINT fk_PerOrders FOREIGN KEY (P_Id) REFERENCES
 Persons(P_Id)

NEW QUESTION 127

- (Exam Topic 8)

You have a database named Database1. You execute the following code:

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

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

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

Users report that the query takes a long time to complete. You create a full-text index for the Notes column.

You need to recommend changes to the query to reduce the amount of time it takes for the query to complete. Which code segment should you use to replace line 03?

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

Answer: A

NEW QUESTION 129

- (Exam Topic 8)

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

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

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

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

You insert 3 million rows into Sales.

You need to reduce the amount of time it takes to execute Proc1. What should you do?

☐ A. Change the query inside Proc1 to:

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

☐ B. Change the definition of Proc1 to:

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

☐ C. Change the query inside Proc1 to:

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

☐ D. Change the definition of Proc1 to:

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

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

Answer: A

NEW QUESTION 134

- (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

ALTER SCHEMA

ALTER TABLE

dbo

dbo.Customers

EXEC sp_rename

TRANSFER

Schema2

Answer Area

Code

Code

Code

Code

- A. Mastered
- B. Not Mastered

Answer: A

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

NEW QUESTION 135

- (Exam Topic 8)
You need to identify which nonclustered indexes are unused by queries.
How should you complete the statement? To answer, drag the appropriate values to the correct locations. Each value may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

Values

sys.dm_db_index_operational_stats

sys.dm_db_index_physical_stats

sys.dm_db_index_usage_stats

sys.indexes

sys.tables

••••

Answer Area

SELECT database_id, a.object_id, a.index_id, b.name, a.user_seeks, a.user_scans, a.user_updates

FROM Value a
join Value b on a.object_id = b.object_id

and a.index_id = b.index_id
where a.index_id >= 2
and a.database_id = db_id()
and (a.user_seeks = 0
or a.user_scans = 0)

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: sys.dm_db_index_usage_stats

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

Box 2: sys.indexes

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

References:

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

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

NEW QUESTION 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

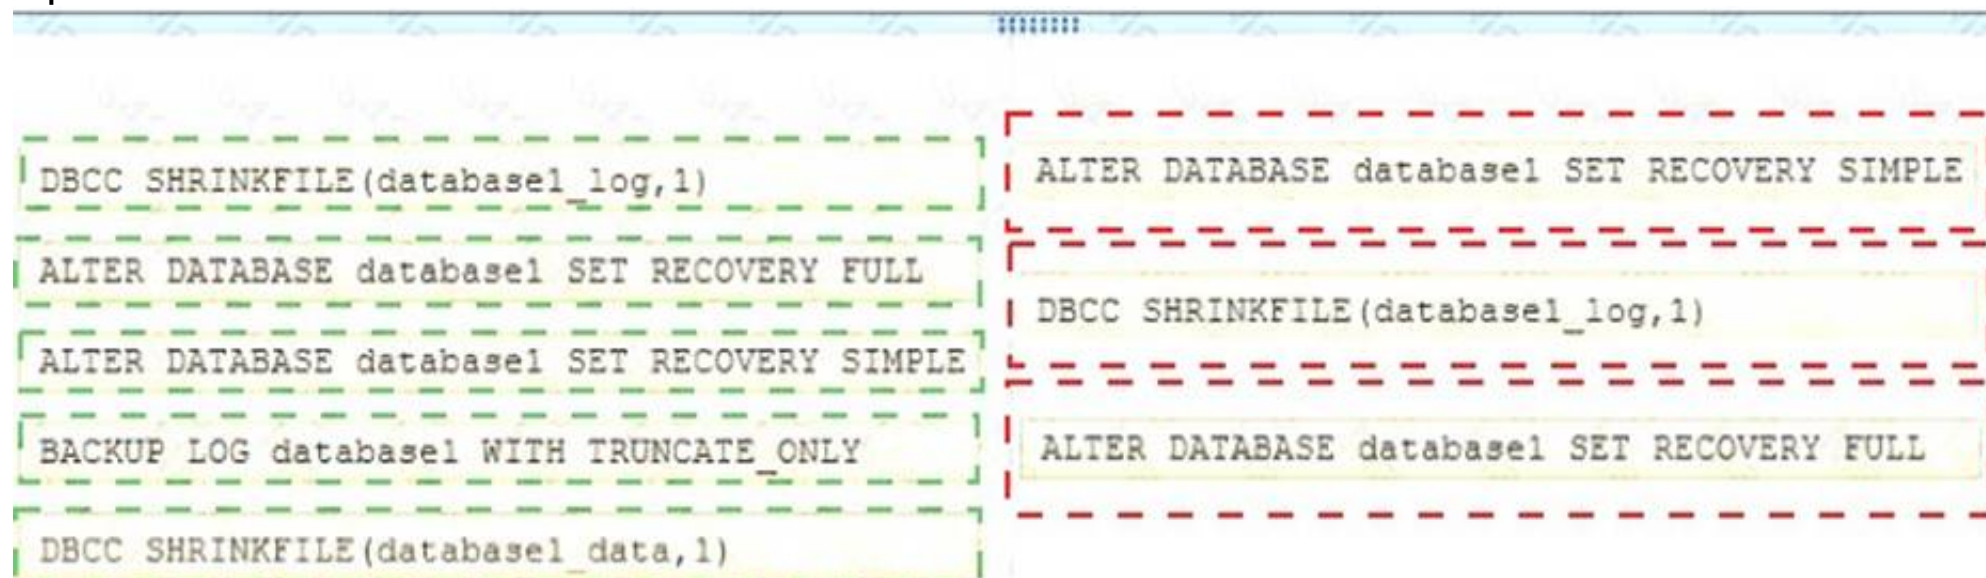
BACKUP LOG database1 WITH TRUNCATE_ONLY

DBCC SHRINKFILE (database1_data,1)

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:



NEW QUESTION 138

- (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 142

- (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.)

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:

* 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 143

- (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 147

- (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 150

- (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 152

- (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 156

- (Exam Topic 8)

You have the following stored procedure.

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

The stored procedure takes much longer to execute than expected.

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

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

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

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

Answer: A

Explanation:

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

References:

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

NEW QUESTION 160

- (Exam Topic 8)

You have a SQL Server 2012 instance that hosts a single-user database.

The database does not contain user-created stored procedures or user-created functions. You need to minimize the amount of memory used for query plan caching.

Which advanced server option should you modify?

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

Answer: C

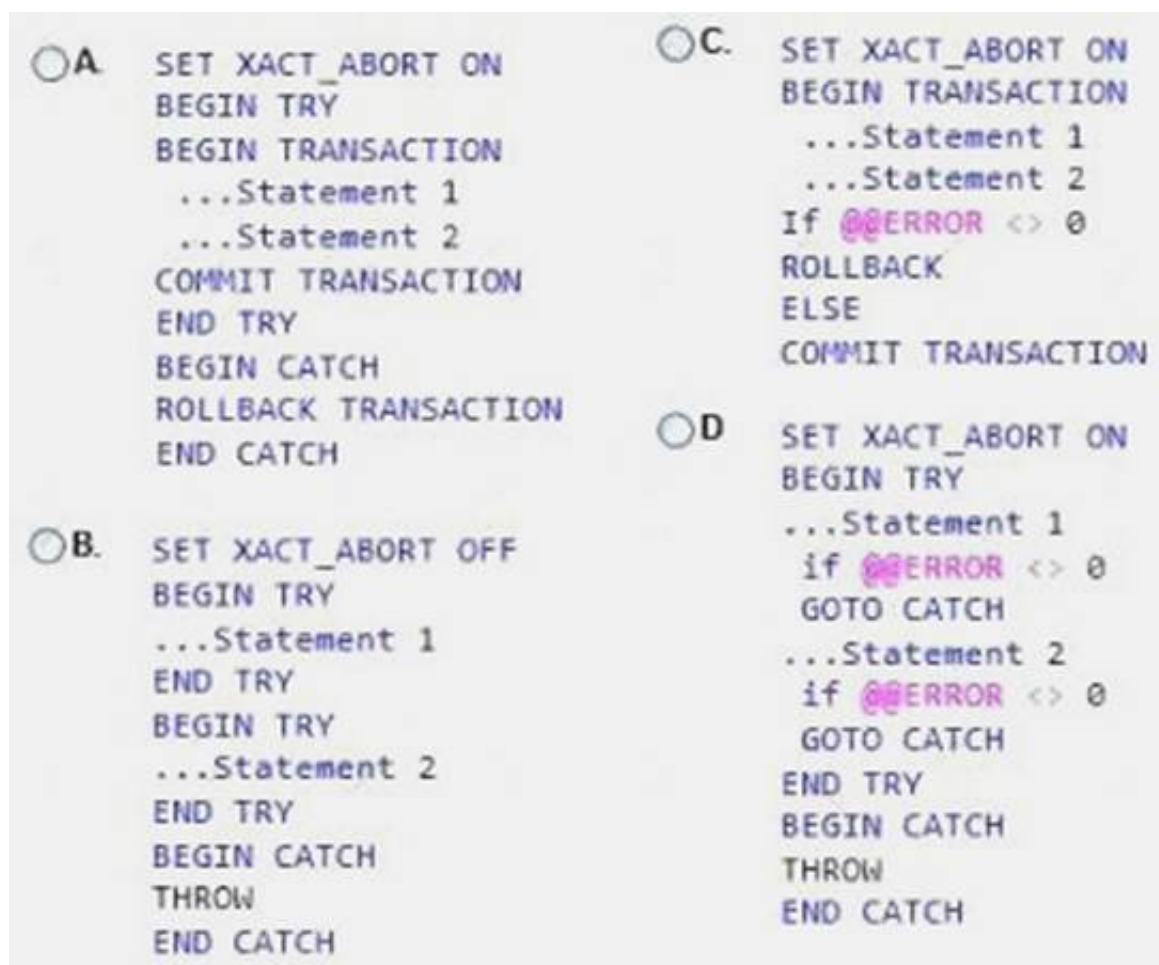
NEW QUESTION 165

- (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 169

- (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 170

- (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$44ncjhdj' 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 171

- (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 174

- (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 179

- (Exam Topic 8)

You execute the following code:

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

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

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

- ☐ A.

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

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

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

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

A. Option A

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

Answer: B

NEW QUESTION 183

- (Exam Topic 8)

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

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

You discover that the stored procedure periodically fails to update Person.Contact.

You need to ensure that Person.Contact is always updated when UpdateContact executes. The solution must minimize the amount of time required for the stored procedure to execute and the number of locks held.

What should you do?

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

Answer: C

Explanation:

* Overall, you should use UPDLOCK when you read a value that you plan to update later in the same transaction to prevent the value from changing.

* UPDLOCK

Specifies that update locks are to be taken and held until the transaction completes. UPDLOCK takes update locks for read operations only at the row-level or page-level. If UPDLOCK is combined with TABLOCK, or a table-level lock is taken for some other reason, an exclusive (X) lock will be taken instead.

When UPDLOCK is specified, the READCOMMITTED and READCOMMITTEDLOCK isolation level hints are ignored. For example, if the isolation level of the session is set to SERIALIZABLE and a query specifies (UPDLOCK, READCOMMITTED), the READCOMMITTED hint is ignored and the transaction is run using the SERIALIZABLE isolation level.

NEW QUESTION 188

- (Exam Topic 8)

You have a SQL Server 2012 database named database1.

Users report that queries that usually take less than one second to execute, take more than 30 seconds to execute.

You need to view the server resource consumption when the queries are executed. What should you do?

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



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Start a SQL Server Profiler trace.

Box 2: Start a data collection by using Performance monitor. Box 3: Save the SQL Profiler trace.

Box 4: Save the Performance Monitor data.

Box 5: Import the performance data into SQL Server Profiler. Note:

* (step1, step 2) Both the Profiler trace and the Performance Monitor logs should be started and stopped at about the same time.

* (step 3, step 4) Once you have completed capturing the data for both tools, you are ready to perform the correlation analysis.

* (step 5) How to Correlate SQL Server Profiler Data with Performance Monitor Data

Correlating Performance Monitor and Profiler data is a straightforward process that simply involves importing both sets of data into Profiler. Start Profiler and load the trace file you want to correlate.

From the main menu of Profiler, select File | Import Performance Data,

* With SQL Server Profiler, we have the tools to identify the causes of such spikes. We can import Performance Monitor log data and compare it directly with Profiler activity. If we see a spike in CPU utilization, we can identify which statement or statements were running at the same time, and diagnose potential problems.

NEW QUESTION 192

- (Exam Topic 8)

You use SQL Server 2014. The physical server is a dedicated database server that has 120GB of RAM available. There is approximately 50GB of storage space available on a slow local disk.

You create a new stored procedure. You decide you need to temporarily hold approximately 300,000 rows from two tables, from which you will compute two complex business scores.

The stored procedure will use temporary storage defined as follows:

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

The code will make several passes through the data, applying complex calculations before writing the data to a permanent disk-based table in the same database from which it reads the data.

For this stored procedure, you need to deal with temporary data in the most efficient way to minimize physical disk pressure.

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

- ☐ A.

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

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

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

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

A. Option A

B. Option B

- C. Option C
D. Option D

Answer: A

Explanation:

* You must specify a value for the BUCKET_COUNT parameter when you create the memory-optimized table. In most cases the bucket count should be between 1 and 2 times the number of distinct values in the index key. If the index key contains a lot of duplicate values, on average there are more than 10 rows for each index key value, use a nonclustered index instead

You may not always be able to predict how many values a particular index key may have or will have. Performance should be acceptable if the BUCKET_COUNT value is within 5 times of the actual number of key values.

NEW QUESTION 195

- (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 196

- (Exam Topic 8)

You have a database named database1. Each table in database1 has one index per column. Users often report that creating items takes a long time.

You need to perform the following maintenance tasks:

- ▶ Identify unused indexes.
- ▶ Identify indexes that need to be defragmented. What should you use?

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

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Note:

* sys.dm_db_index_usage_stats

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

* sys.dm_db_index_physical_stats

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

NEW QUESTION 200

- (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 205

- (Exam Topic 8)

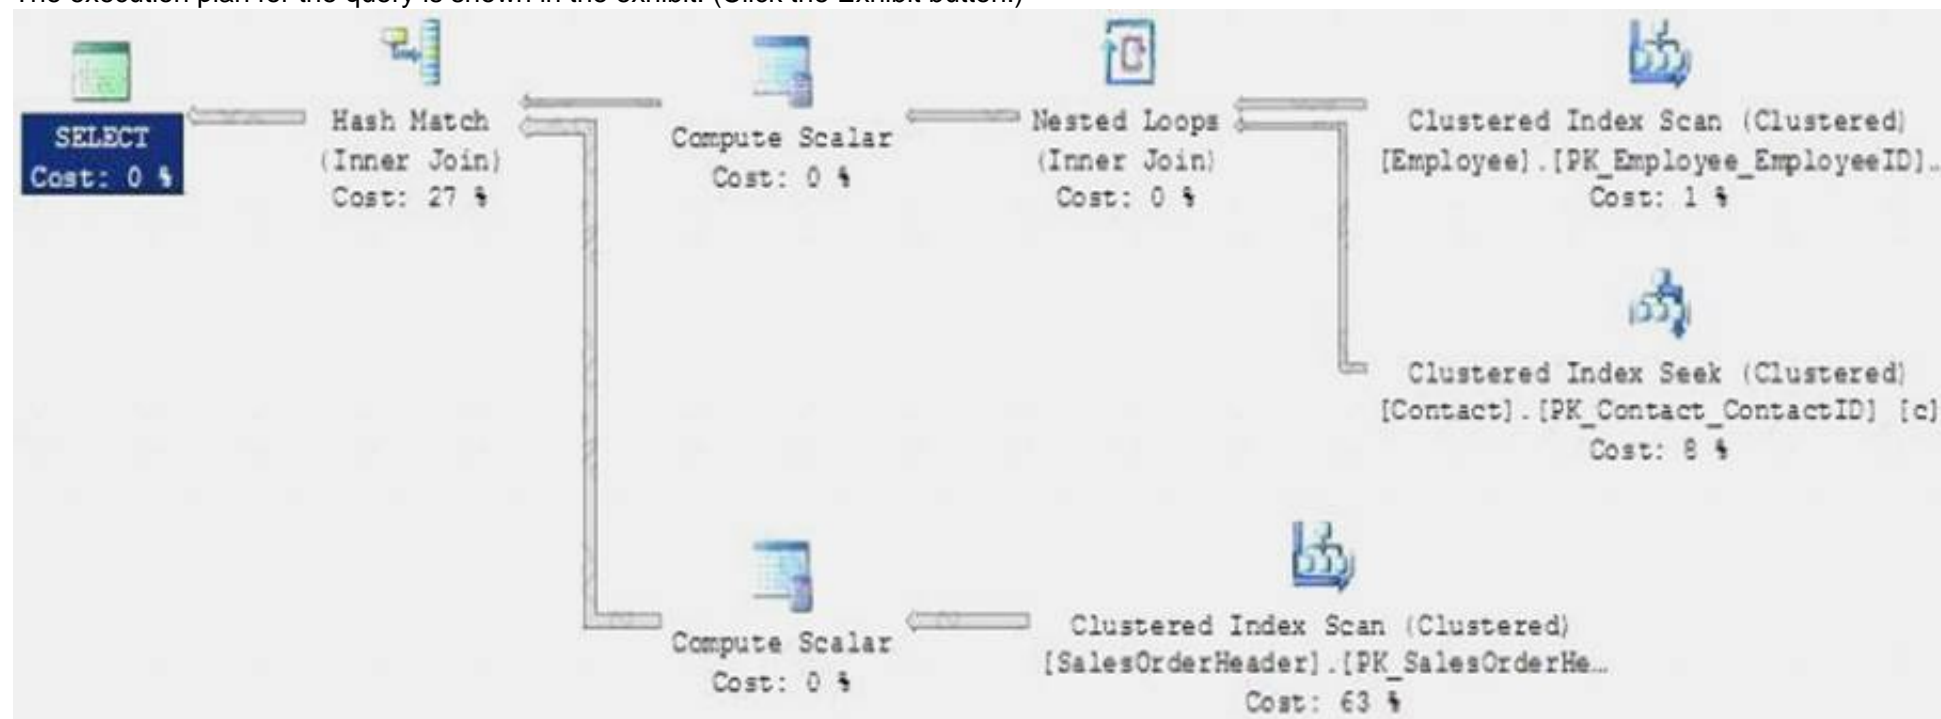
You have a database that contains three tables. The tables are configured as shown in the following table.

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

Contact.EmailAddress

Contact.FirstName

Contact.LastName

Employee.Title

SalesOrderHeader.OrderDate

SalesOrderHeader.SalesPersonID

SalesOrderHeader.SubTotal

Answer Area

Indexed Columns

Column

Included Columns

Column

Column

- 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 208

- (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 209

- (Exam Topic 8)

Your network contains a SQL Server 2012 instance named SQL1. SQL1 contains a database named DB1. DB1 contains three tables.

The tables are configured as shown in the following table.

Table name	Configuration
Table1	<ul style="list-style-type: none"> Table data will not be updated. The table will contain historical calculations. The table will contain 10 million records.
Table2	<ul style="list-style-type: none"> 20% of the table data will be updated weekly. The table will contain 25 million records.
Table3	<ul style="list-style-type: none"> 40% of the table data will be updated weekly. The table will contain 1 million records.

You plan to create indexes for the tables.

You need to identify which type of index must be created for each table. The solution must minimize the amount of time required to return information from the tables.

Which type of index should you create for each table? To answer, drag the appropriate index type to the correct table in the answer area.

Index Types	Answer Area	
Columnstore Index	Table1	
Nonclustered Index	Table2	
	Table3	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Index Types	Answer Area	
Columnstore Index	Table1	Columnstore Index
Nonclustered Index	Table2	Nonclustered Index
	Table3	Nonclustered Index

NEW QUESTION 214

- (Exam Topic 8)

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

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

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

- A. Repeatable read

- B. Read committed
- C. Serializable
- D. Snapshot

Answer: C

NEW QUESTION 219

- (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 222

- (Exam Topic 8)

You are creating a stored procedure named usp1. Usp1 will create a table that will be used during the execution of usp1. Only usp1 will be allowed to access the table.

You need to write the code required to create the table for usp1. The solution must minimize the need to recompile the stored procedure.

Which code segment should you use to create the table?

- A. CREATE TABLE oneTable
- B. CREATE TABLE ##oneTable
- C. CREATE TABLE #oneTable
- D. DECLARE oneTable TABLE

Answer: B

NEW QUESTION 225

- (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 227

- (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 230

- (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

NEW QUESTION 234

.....

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