

100% Money Back
Guarantee

Vendor:Microsoft

Exam Code:70-464

Exam Name:Developing Microsoft SQL Server
Databases

Version:Demo

QUESTION 1

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

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

You need to prevent duplicate values in the SKU field.

Which five code segments should you use?

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

Select and Place:

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

Correct Answer:

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

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

Box 3: CK_SKU Box 4: UNIQUE Box 5: (SKU)

Note: The SQL command is: ALTER TABLE ADD CONSTRAINT UNIQUE ()

QUESTION 2

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.

Select and Place:

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

Correct Answer:

Tasks	Maintenance Steps
an identity for UserActivityID starting at the next value	Convert UserActivity to use <input type="text" 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 <input type="text" value="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 <input type="text" value="UserActivity and UserActivity_Archive"/>
UserActivity, UserActivity_Staging, and UserActivity_Archive	
Alter the partition function and UserActivity_Staging constraints	After switching a new partition from UserActivity_Staging into UserActivity_Archive, <input type="text" value="Alter the partition function and UserActivity_Archive constraints"/>

QUESTION 3

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

```
01 CREATE PROC up_employees
02     @ID int,
03     @Name nvarchar(50)
04 AS
05
06 SELECT Name AS OriginalName
07 FROM HR.Employees
08
09 WHERE ID = @ID;
10
11 UPDATE HR.Employees
12 SET Name = @Name
13
14 WHERE ID = @ID;
```

You discover that the stored procedure periodically fails to update HR.Employees.

You need to ensure that HR.Employees is always updated when up_employees executes.

The solution must minimize the amount of time required for the stored procedure to execute and the number of locks held.

What should you do?

- A. Add the following line of code to line 05: SET TRANSACTION ISOLATION LEVEL SNAPSHOT
- B. Add the following line of code to line 13: WITH (UPDLOCK)
- C. Add the following line of code to line 05: SET TRANSACTION ISOLATION LEVEL SERIALIZABLE
- D. Add the following line of code to line 08: WITH (UPDLOCK)

Correct Answer: D

QUESTION 4

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.

Select and Place:

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

```
CREATE SYMMETRIC KEY Key1 WITH ALGORITHM = S
HA1
ENCRYPTION BY CERTIFICATE Cert1;
```

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

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

```
CLOSE SYMMETRIC KEY;
```

```
CREATE CREDENTIAL Cred1 WITH IDENTITY = 'Use
r1', SECRET = 'P@sswOrd';
```

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

```
CREATE SYMMETRIC KEY Key1 WITH ALGORITHM = A
ES_256
ENCRYPTION BY CERTIFICATE Cert1;
```

Correct Answer:

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

```
CREATE CREDENTIAL Cred1 WITH IDENTITY = 'Use
r1', SECRET = 'P@sswOrd';
```

```
CREATE SYMMETRIC KEY Key1 WITH ALGORITHM = A
ES_256
ENCRYPTION BY CERTIFICATE Cert1;
```

```
CREATE SYMMETRIC KEY Key1 WITH ALGORITHM = S
HA1
ENCRYPTION BY CERTIFICATE Cert1;
```

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

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

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

```
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/>

QUESTION 5

You have a table named Table1 that stores customer data.

Each customer has a credit limit that can only be discovered by querying multiple tables.

You need to ensure that the value of the credit limit is returned by executing a query on Table1.

What should you create?

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

Correct Answer: C

QUESTION 6

You use SQL Server 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. SET XACT_ABORT ON
BEGIN TRY
BEGIN TRANSACTION
...Statement 1
...Statement 2
COMMIT TRANSACTION
END TRY
BEGIN CATCH
ROLLBACK TRANSACTION
END CATCH
- B. SET XACT_ABORT OFF
BEGIN TRY
...Statement 1
END TRY
BEGIN TRY
...Statement 2
END TRY
BEGIN CATCH
THROW
END CATCH
- C. SET XACT_ABORT OFF
BEGIN TRANSACTION
BEGIN TRY
...Statement 1
...Statement 2
END TRY
BEGIN CATCH
IF @@ERROR > 0
BEGIN
ROLLBACK TRANSACTION
THROW
END
ELSE
COMMIT TRANSACTION
END CATCH
- D. SET XACT_ABORT OFF
BEGIN TRY
BEGIN TRANSACTION
...Statement 1
...Statement 2
END TRY
BEGIN CATCH
ROLLBACK TRANSACTION
THROW
END CATCH

A. Option A

B. Option B

C. Option C

D. Option D

Correct Answer: A

* 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.

QUESTION 7

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.

Select and Place:

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

Correct Answer:

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

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)

QUESTION 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.)

Select and Place:

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

Correct Answer:

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

QUESTION 9

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.

Correct Answer: B

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

QUESTION 10

You need to implement a solution that addresses the bulk insert requirements. What should you add to line 91 in usp_ImportOrderDetails?

- A. TABLOCK,
- B. MAXERRORS=0,
- C. BATCHSIZE=1000,
- D. LASTROW=0,

Correct Answer:

BATCHSIZE specifies the number of rows in the BULK INSERT operation to be copied to the server as one transaction at a time. If this transaction fails, SQL Server commits or rolls back the transaction for every batch. Incorrect Answers:

A: TABLOCK specifies that a table-level lock is acquired for the duration of the bulk-import operation.

B: MAXERRORS specifies the maximum number of syntax errors allowed in the data before the entire bulk-import operation is canceled. Each row that cannot be imported by the bulk-import operation is ignored and counted as one error.

D: LASTROW specifies the number of the last row to load. The default is 0, which indicates the last row in the specified data file.

QUESTION 11

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.)

Select and Place:

Functions	Answer Area
sys.dm_cb_index_usage_stats	Identify unused indexes. Function
sys.dm_cb_index_operational_stats	Identify indexes that need to be defragmented. Function
sys.dm_cb_index_physical_stats	
sys.dm_cb_missing_index_columns	
sys.dm_cb_missing_index_details	
sys.dm_cb_missing_index_groups	

Correct Answer:

Functions	Answer Area
	Identify unused indexes. sys.dm_cb_index_usage_stats
sys.dm_cb_index_operational_stats	Identify indexes that need to be defragmented. sys.dm_cb_index_physical_stats
sys.dm_cb_missing_index_columns	
sys.dm_cb_missing_index_details	
sys.dm_cb_missing_index_groups	

*

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.

QUESTION 12

You have a table named ORDERS that contains 10,514,003 Orders. The ORDERS table has an IDENTITY(1,1) column named ORDER_ID. ORDER_ID is the UNIQUE CLUSTERED INDEX and PRIMARY KEY for the table. The first ORDER_ID is 1. There are no missing ORDER_IDs in the set.

Based on table usage patterns, you decide to use partitioning on this table based off of the ORDER_ID column.

You need to create the following partitions:

Partition	Values
1	Orders \leq 7,500,000
2	Orders $>$ 7,500,000 and \leq 10,000,000
3	Orders $>$ 10,000,000

Which Transact-SQL statement should you run?

- A. CREATE PARTITION FUNCTION pfOrderIDRange (int) AS RANGE RIGHT FOR VALUES (7500000,10000000)
- B. CREATE PARTITION FUNCTION pfOrderIDRange (int) AS RANGE LEFT FOR VALUES (7500000,10000000)
- C. CREATE PARTITION FUNCTION pfOrderIDRange (int) AS RANGE RIGHT FOR VALUES (0,7500000,10000000)
- D. CREATE PARTITION FUNCTION pfOrderIDRange (int) AS RANGE LEFT FOR VALUES (0,7500000,10000000)

Correct Answer: B

We need to specify two boundary values that separate the partitions. And we must use RANGE LEFT as the boundary value must belong to the partition to the left. Incorrect Answers:

A: With RANGE RIGHT the boundary value belongs to the partition to the right. It is the equivalent of \geq . C, D: Providing three boundary values will result in four partitions.

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

To Read the [Whole Q&As](#), please purchase the [Complete Version](#) from [Our website](#).

Try our product !

100% Guaranteed Success

100% Money Back Guarantee

365 Days Free Update

Instant Download After Purchase

24x7 Customer Support

Average 99.9% Success Rate

More than 800,000 Satisfied Customers Worldwide

Multi-Platform capabilities - [Windows](#), [Mac](#), [Android](#), [iPhone](#), [iPod](#), [iPad](#), [Kindle](#)

Need Help

Please provide as much detail as possible so we can best assist you.

To update a previously submitted ticket:



 <p>One Year Free Update Free update is available within One Year after your purchase. After One Year, you will get 50% discounts for updating. And we are proud to boast a 24/7 efficient Customer Support system via Email.</p>	 <p>Money Back Guarantee To ensure that you are spending on quality products, we provide 100% money back guarantee for 30 days from the date of purchase.</p>	 <p>Security & Privacy We respect customer privacy. We use McAfee's security service to provide you with utmost security for your personal information & peace of mind.</p>
---	---	--

Any charges made through this site will appear as Global Simulators Limited.

All trademarks are the property of their respective owners.