



70-464^{Q&As}

Developing Microsoft SQL Server Databases

Pass Microsoft 70-464 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

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

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Microsoft
Official Exam Center

-  **Instant Download** After Purchase
-  **100% Money Back** Guarantee
-  **365 Days** Free Update
-  **800,000+** Satisfied Customers





QUESTION 1

You need to ensure that `usp_AddXMLOrder` can be used to validate the XML input from the retailers.

Which parameters should you add to `usp_AddXMLOrder` on line 04 and line 05? (Each correct answer presents part of the solution. Choose all that apply.)

- A. `@schema varbinary(100)`.
- B. `@items varchar(max)`.
- C. `@schema sysname`.
- D. `@items varbinary(max)`.
- E. `@items xml`.
- F. `@schema xml`.

Correct Answer: CE

QUESTION 2

A company has a main office in London and a branch office in New York. The company's network contains a server named `Server5` that has SQL Server installed. `Server5` contains a database name `ContentDB` and a table named

`ContentTable`.

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

You need to create a distributed partitioned view. The solution must minimize the amount of network traffic.

Which three actions should you perform? (Each correct answer presents part of the solution. Choose all that apply.)

- A. Create the view on `Server5` and `Server9`.
- B. Remove `ContentTable` from `Server5`.
- C. Add `Server9` as a Distributor.
- D. Add `ContentTable` to `Server9`.
- E. Add `Server9` as a linked server.

Correct Answer: ADE

D: Creating Distributed Partitioned Views Before you implement a partitioned view, you must first partition a table horizontally. In designing a partitioning scheme, it must be clear what data belongs to each member table. The original table is replaced with several smaller member tables. Each member table has the same number of columns as the original table, and each column has the same attributes as the corresponding column in the original table, such as data type, size, and collation. If you are creating a distributed partitioned view, each member table is on a separate member



server.

AE: Defining Distributed Partition Views

After you create the member tables, you define a distributed partitioned view on each member server, with each view having the same name.

You build the distributed partitioned views by performing the following tasks:

Adding linked server definitions on each member server that contains the connection information required to run distributed queries on the other member servers. This gives a distributed partitioned view access to data on the other servers.

Setting the lazy schema validation option.

Creating a distributed partitioned view on each member server.

References: [https://technet.microsoft.com/en-us/library/ms188299\(v=sql.105\).aspx](https://technet.microsoft.com/en-us/library/ms188299(v=sql.105).aspx)

QUESTION 3

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
<pre>REFERENCES STATES (Country_Abbr)</pre>	<p>www.PassApply.com</p>
<pre>REFERENCES COUNTRY (Country_Abbr)</pre>	
<pre>ON STATES</pre>	
<pre>FOREIGN KEY (Country_Abbr)</pre>	
<pre>ON COUNTRY</pre>	
<pre>ADD CONSTRAINT FK_StateCountry</pre>	
<pre>ON COUNTRY_ABBR</pre>	
<pre>ALTER TABLE COUNTRY</pre>	
<pre>ADD FOREIGN KEY FK_StateCountry</pre>	
<pre>ALTER TABLE STATES</pre>	

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 4

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

Correct Answer: A

QUESTION 5

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

Correct Answer: DE

Note:



* Move micropayments to dbo.POSException table by using a stored procedure named ups_DetectSuspiciousActivity.

QUESTION 6

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

Correct Answer: A

QUESTION 7



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

Correct Answer: C

QUESTION 8

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\\'`,

Correct Answer: C

QUESTION 9

You use SQL Server. 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

Correct Answer: A

* 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.

QUESTION 10

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

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

You need to create a distributed partitioned view. The solution must minimize the amount of network traffic.

What should you do? (Each correct answer presents part of the solution. Choose all that apply.)

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

Correct Answer: BCDF

QUESTION 11

You have a SQL Server database.

The recovery model is set to full. The transaction log is backed up every night.

You discover that the transaction log never decreases in size.

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

You need to ensure that you can reduce the size of the transaction log.

What should you do first?

- A. Truncate the transaction log



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

Correct Answer: A

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

C: The transaction log can be truncated by setting the database recovery model to SIMPLE, or by backing up the log and then running the DBCC SHRINKFILE operation again.

References: <https://docs.microsoft.com/en-us/sql/t-sql/database-console-commands/dbcc-shrinkfile-transact-sql?view=sql-server-2017>

QUESTION 12

Which data type should you use for CustomerID?

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

Correct Answer: D

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>

[70-464 VCE Dumps](#)

[70-464 Study Guide](#)

[70-464 Exam Questions](#)



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](#)

We provide exam PDF and VCE of Cisco, Microsoft, IBM, CompTIA, Oracle and other IT Certifications. You can view Vendor list of All Certification Exams offered:

<https://www.passapply.com/allproducts>

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.
Copyright © passapply, All Rights Reserved.