70-461 Q&As

Querying Microsoft SQL Server 2012

Pass Microsoft 70-461 Exam with 100% Guarantee

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

https://www.pass4lead.com/70-461.html

100% Passing Guarantee
100% Money Back Assurance

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

QUESTION 1

You administer a Microsoft SQL Server instance that will support several databases.

You need to ensure that every new database created has a data type named postalcode that contains the same attributes.

What should you do?

A. Create a user-defined type on the model database.

B. Create a user-defined type on the master database.

C. Create a user-defined data type on the master database.

D. Create a user-defined data type on the model database.

Correct Answer: D

One option is to create SQL Server user defined data types.

One trick with new databases is to create the objects in the model database, so as new databases are created the user defined data types will automatically be available.

References: https://www.mssqltips.com/sqlservertip/1628/sql-server-user-defined-data-types-rules-and-defaults/

QUESTION 2

How many clustered indexes can you create on a table?

A. 999

B. 16

C. 1

D. 900

Correct Answer: C

QUESTION 3

You create the following stored procedure. (Line numbers are included for reference only.)

```
01 CREATE PROCEDURE dbo.InsertCountryRegion
02  @CountryRegionCode nvarchar(3),
03  @Name nvarchar(50)
04 AS
05 BEGIN
06  SET NOCOUNT ON;
07  ...
08 END;
```


You need to ensure that the stored procedure performs the following tasks:

If a record exists, update the record.

If no record exists, insert a new record.

Which four Transact-SQL statements should you insert at line 07? (To answer, move the appropriate statements from the list of statements to the answer area and arrange them in the correct order.)

Select and Place:

Correct Answer:

QUESTION 4

You develop a database application for a university. You need to create a view that will be indexed that meets the following requirements:

Displays the details of only students from Canada.

Allows insertion of details of only students from Canada.

Which four Transact-SQL statements should you use? (To answer, move the appropriate SQL statements from the list of statements to the answer area and arrange them in the correct order.)

Select and Place:

Correct Answer:

QUESTION 5

You use Microsoft SQL Server to develop a database application.

Your application sends data to a VARCHAR(50) variable named @var.

You need to write a Transact-SQL statement that will return information on a successful or unsuccessful cast to an integer in a table.

Which Transact-SQL statement should you run?

A.
```
SELECT
    CASE
    WHEN CONVERT(int, @var) IS NULL
    THEN 'True'
    ELSE 'False'
    END
 As BadCast
```

B.
```
SELECT
    IIF(TRY_PARSE(@var AS int) IS NULL,
   'True',
   'False'
)
As BadCast
```

C.
```
TRY(
  SELECT CONVERT(int, @var)
  SELECT 'True' As BadCast
)
CATCH
  SELECT 'False' As BadCast
)
```

D.
```
BEGIN TRY
SELECT
  CONVERT(int, @var) as value,
  'True' As BadCast
END TRY
BEGIN CATCH
SELECT
  CONVERT(int, @var) as Value,
'False' As BadCast
END CATCH
```


A. Option A

B. Option B

C. Option C

D. Option D

Correct Answer: B

TRY_PARSE returns the result of an expression, translated to the requested data type, or null if the cast fails in SQL Server. Use TRY_PARSE only for converting from string to date/time and number types. References: https://docs.microsoft.com/en-us/sql/t-sql/functions/try-parse-transact-sql?view=sql-server-2017

QUESTION 6

Which of the following is invalid? (Choose all that apply.)

A. Referring to an attribute that you group by in the WHERE clause

B. Referring to an expression in the GROUP BY clause; for example, GROUP BY YEAR(orderdate)

C. In a grouped query, referring in the SELECT list to an attribute that is not part of

D. the GROUP BY list and not within an aggregate function Referring to an alias defined in the SELECT clause in the HAVING clause

Correct Answer: CD

QUESTION 7

You work as a database developer at ABC.com. ABC has an in-house application named ABCApp3 that runs a Transact-SQL query against a SQL Server 2012 database.

You want to run an execution plan against the query that will provide detailed information on missing indexes.

How would you accomplish this task?

A. You should make use of the READPAST hint in the queries.

B. You should make use of the READCOMMITTED hint in the queries.

C. You should make use of the SET SHOWPLAN_XML ON statement in the query.

D. You should make use of the SET STATISTICS XML ON statement in the query.

E. You should make use of the SET XACT_ABORT OFF statement in the query.

F. You should make use of the SET CONTEXT_INFO statement in the query.

Correct Answer: C

QUESTION 8

You have a database that contains the tables as shown in the exhibit. (Click the Exhibit button.)

You need to create a query that returns a list of products from Sales.ProductCatalog. The solution must meet the following requirements:

UnitPrice must be returned in descending order.

The query must use two-part names to reference the table.

The query must use the RANK function to calculate the results.

The query must return the ranking of rows in a column named PriceRank.

The list must display the columns in the order that they are defined in the table.

PriceRank must appear last.

Which code segment should you use?

To answer, type the correct code in the answer area.

Correct Answer: Please review the explanation part for this answer

SELECT ProductCatalog.CatID, ProductCatalog.CatName, ProductCatalog.ProductID, ProductCatalog.ProdName, ProductCatalog.UnitPrice, RANK() OVER (ORDER BY ProductCatalog.UnitPrice DESC) AS PriceRank FROM

Sales.ProductCatalog ORDER BY ProductCatalog.UnitPrice DESC

QUESTION 9

You work as a database administrator at ABC.com. ABC.com has a SQL Server 2012 database named SalesDB. The SalesDB is shown in the following database diagram: You need to write a Transact-SQL query that display a single row in the following XML format:

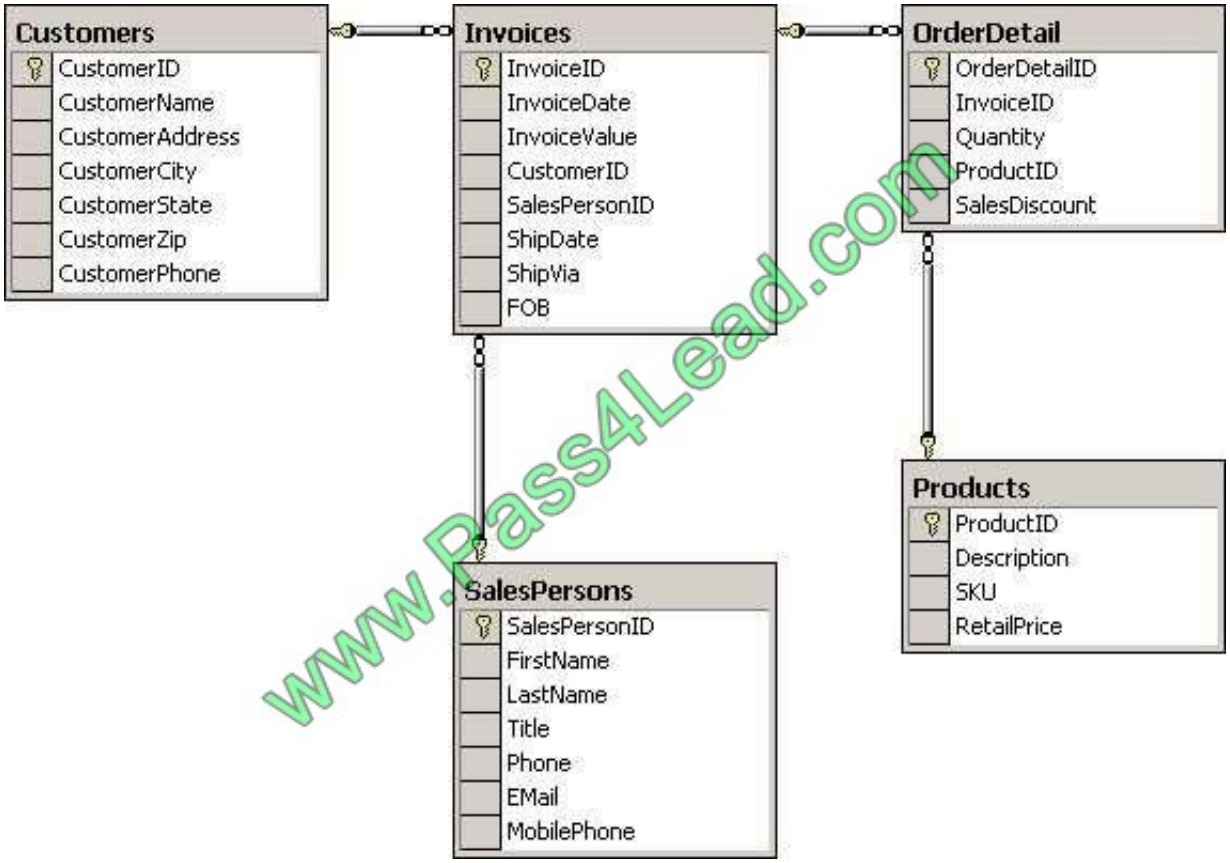

Which of the following SELECT statement would you write?

A. SELECT in.InvoiceID, in.InvoiceDate AS [Date], in.InvoiceValue AS [Value], cu.CustomerName AS [Name], cu.CustomerCity AS [ShippedTo] FROM Invoices AS in INNER JOIN Customers AS cu ON in.CustomerID = cu.CustomerID WHERE cu.CustomerID = 1001 FOR XML RAW

B. SELECT InvoiceID, InvoiceDate AS [Date], InvoiceValue AS [Value], CustomerName AS [Name], CustomerCity AS [ShippedTo] FROM Invoices INNER JOIN Customers ON Invoices.CustomerID = Customers.CustomerID WHERE Customers.CustomerID = FOR XML

C. SELECT Invoices.InvoiceID, Invoices.InvoiceDate AS [Date], Invoices.InvoiceValue AS [Value], Customers.CustomerName AS [Name], Customers.CustomerCity AS [ShippedTo] FROM Invoices INNER JOIN Customers ON Invoices.CustomerID = Customers.CustomerID WHERE Customers.CustomerID = 1001 FOR XML AUTO

D. SELECT InvoiceID, InvoiceDate AS [Date], InvoiceValue AS [Value], CustomerName AS [Name], CustomerCity AS [ShippedTo] FROM Invoices INNER JOIN Customers ON Invoices.CustomerID = Customers.CustomerID WHERE

Customers.CustomerID = FOR XML AUTO, RAW

Correct Answer: C

QUESTION 10

Which of the following are true about table-valued UDFs?

A. Table-valued UDFs can return scalar values or tables.

B. Table-valued UDFs always involve multiple T-SQL statements.

C. Table-valued UDFs can be invoked in a SELECT list or a WHERE clause.

D. Table-valued UDFs can be invoked in the FROM clause of a SELECT statement.

Correct Answer: D

The FROM clause requires a table and table-valued UDFs return tables.

QUESTION 11

You want to add a new GUID column named BookGUID to a table named dbo.Book that already contains data.

BookGUID will have a constraint to ensure that it always has a value when new rows are inserted into dbo.Book.

You need to ensure that the new column is assigned a GUID for existing rows.

Which four Transact-SQL statements should you use? (To answer, move the appropriate SQL statements from the list of statements to the answer area and arrange them in the correct order.)

Select and Place:

Correct Answer:

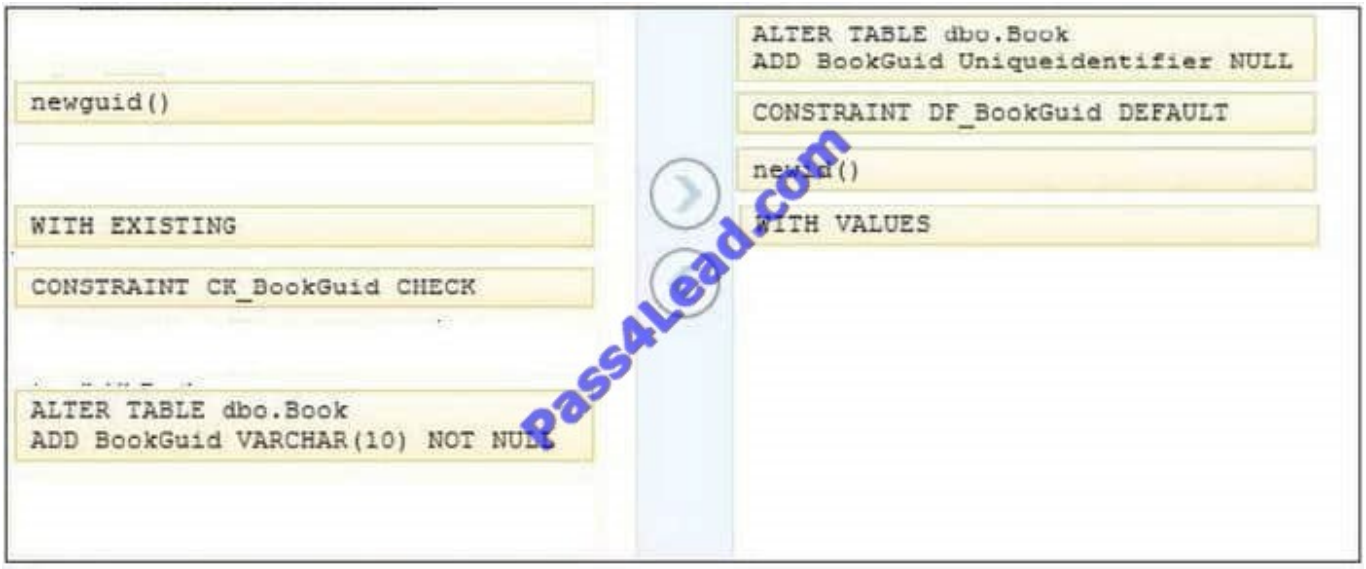

Actually, in the real world, you don\\'t have to use WITH VALUES at the end of the statement and it works just as well. But because the question specifically states which FOUR TSQL statements to use, we have to include it.

QUESTION 12

You create a stored procedure that will update multiple tables within a transaction. You need to ensure that if the stored procedure raises a run-time error, the entire transaction is terminated and rolled back. Which Transact-SQL statement should you include at the beginning of the stored procedure?

A. SET XACT_ABORT ON

B. SET ARITHABORT ON

C. TRY

D. BEGIN

E. SET ARITHABORT OFF

F. SET XACT_ABORT OFF

Correct Answer: A

http://msdn.microsoft.com/en-us/library/ms190306.aspx http://msdn.microsoft.com/en-us/library/ms188792.aspx

QUESTION 13

You develop a database for a travel application. You need to design tables and other database objects. You create a view that displays the dates and times of the airline schedules on a report. You need to display dates and times in several

international formats.

What should you do?

A. Use the CAST function.

B. Use the DATE data type.

C. Use the FORMAT function.

D. Use an appropriate collation.

E. Use a user-defined table type.

F. Use the VARBINARY data type.

G. Use the DATETIME data type.

H. Use the DATETIME2 data type.

I. Use the DATETIMEOFFSET data type.

J. Use the TODATETIMEOFFSET function.

Correct Answer: C

http://msdn.microsoft.com/en-us/library/hh213505.aspx

QUESTION 14

You are developing a database that will contain price information. You need to store the prices that include a fixed precision and a scale of six digits. Which data type should you use?

A. Float

B. Money

C. Smallmoney

D. Numeric

Correct Answer: D

Numeric is the only one in the list that can give a fixed precision and scale.

http://msdn.microsoft.com/en-us/library/ms179882.aspx

QUESTION 15

Your application contains a stored procedure for each country. Each stored procedure accepts an employee identification number through the @EmpID parameter. You plan to build a single process for each employee that will execute the

stored procedure based on the country of residence. Which approach should you use?

A. a recursive stored procedure

B. Trigger

C. An UPDATE statement that includes CASE

D. Cursor

E. The foreach SQLCLR statement

Correct Answer: D

QUESTION 16

You are designing an order entry system that uses an SQL Server database. The database includes the following tables:

You need to ensure that Orders are added to the Orders table only for customers that have an account balance of zero.

How should you complete the relevant Transact-SQL statement? To answer, select the correct Transact-SQL statement from each list in the answer area.

Hot Area:

Answer Area

```
CREATE [ RULE | TRIGGER | FUNCTION | NOTIFICATION EVENT ] Purchasing.ZeroBalance ON Purchasing.Orders
AFTER [ INSERT | ON INSERT | AFTER INSERT | BEFORE INSERT ]
AS
IF EXISTS (SELECT *
           FROM Purchasing.Orders AS o
           JOIN inserted AS i
           ON o.OrderId = i.OrderId
           JOIN Purchasing.Customers AS c
           ON c.CustomerID = o.CustomerID
           WHERE [ c.AccountBalance > 0 | c.AccountBalance = 0 | c.AccountBalance <0 ]
          )
BEGIN
    [ DUMP TRANSACTION; | ROLLBACK TRANSACTION; ]
   RETURN
END;
GO
```


Correct Answer:

Answer Area

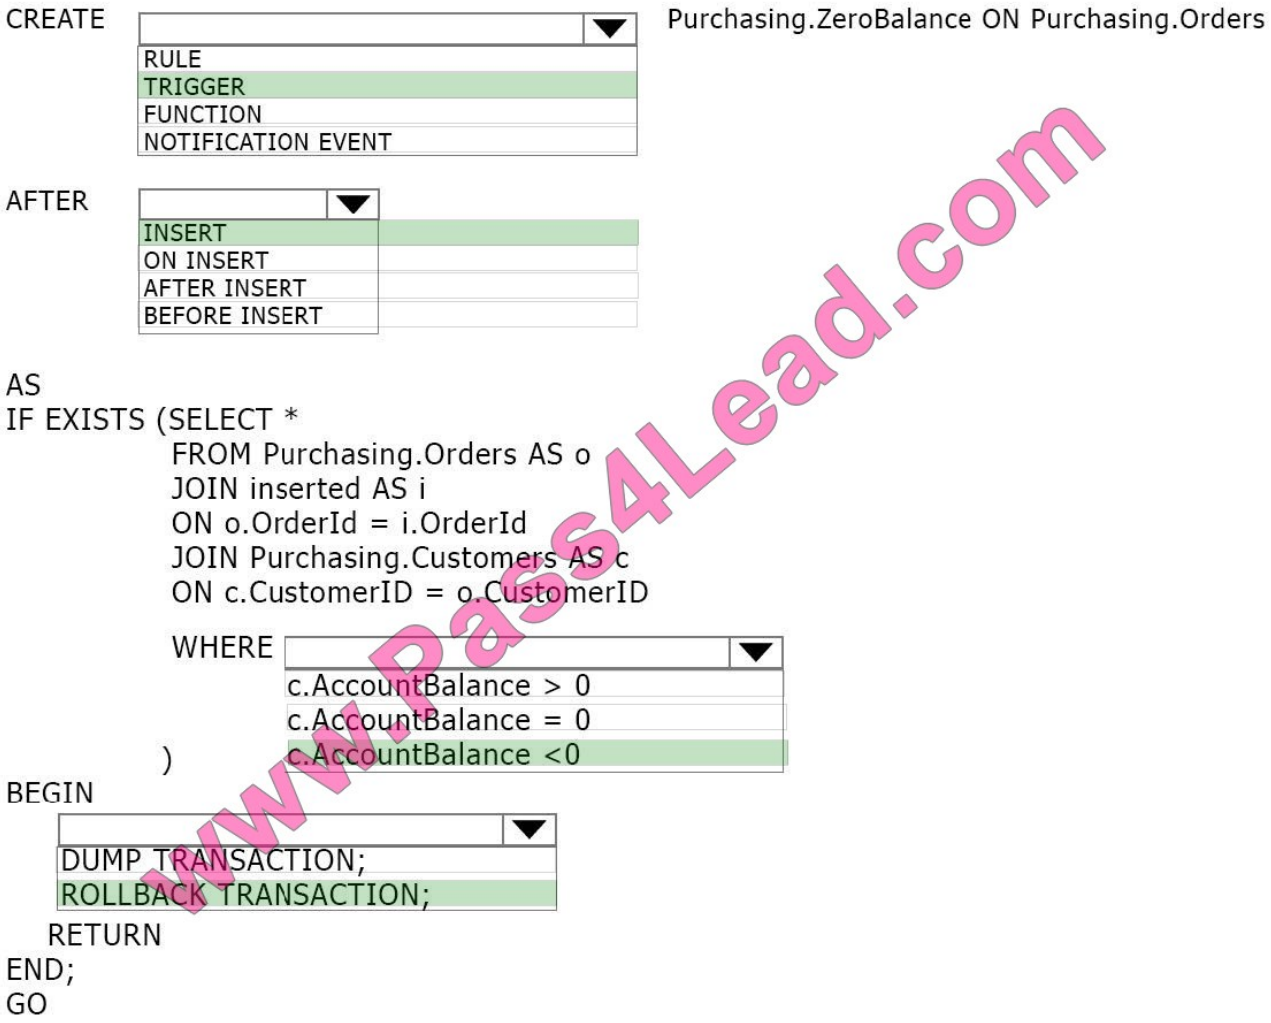

The Transact SQL CREATE TRIGGER command creates a DML, DDL, or logon trigger. A trigger is a special kind of stored procedure that automatically executes when an event occurs in the database server. DML triggers execute when a

user tries to modify data through a data manipulation language (DML) event. DML events are INSERT, UPDATE, or DELETE statements on a table or view. These triggers fire when any valid event is fired, regardless of whether or not any

table rows are affected.

Partial syntax is:

CREATE TRIGGER [schema_name .]trigger_name

ON { table | view }

[WITH [,...n]]

{ FOR | AFTER | INSTEAD OF }

{ [INSERT] [,] [UPDATE] [,] [DELETE] }

QUESTION 17

Which of the following strategies can help reduce blocking and deadlocking by reducing shared locks? (Choose all that apply.)

A. Add the READUNCOMMITTED table hint to queries.

B. Use the READ COMMTTED SNAPSHOT option.

C. Use the REPEATABLE READ isolation level.

D. Use the SNAPSHOT isolation level.

Correct Answer: ABD

QUESTION 18

You have a database named Sales that contains the tables as shown in the exhibit. (Click the Exhibit button.)

You need to create a query that meets the following requirements:

References columns by using one-part names only.

Groups aggregates only by SalesTerritoryID, and then by ProductID.

Orders the results in descending order by SalesTerritoryID and then by ProductID in descending order for both.

Part of the correct T-SQL statement has been provided in the answer area. Provide the complete code.

```
SELECT SalesTerritoryID,
  ProductID,
  AVG(UnitPrice),
  MAX(OrderQty),
  MAX(DiscountAmount)
FROM Sales.Details
```

Correct Answer: Please review the explanation part for this answer

SELECT SalesTerritoryID, ProductID, AVG(UnitPrice), MAX(OrderQty) MAX(DiscountAmount) FROM Sales.Details GROUP BY SalesTerritoryID, ProductID ORDER BY SalesTerritoryID DESC, ProductID DESC

QUESTION 19

You develop a Microsoft SQL Server 2012 database. You create a view from the Orders and OrderDetails tables by using the following definition. You need to improve the performance of the view by persisting data to disk. What should you do?

```
CREATE VIEW vOrders
WITH SCHEMABINDING
AS
SELECT o.ProductID,
  o.OrderDate,
  SUM(od.UnitPrice * od.OrderQty) AS Amount
FROM OrderDetails AS od INNER JOIN
  Orders AS o ON od.OrderID = o.OrderID
WHERE od.SalesOrderID = o.SalesOrderID
GROUP BY o.OrderDate, o.ProductID
GO
```

A. Create an INSTEAD OF trigger on the view.

B. Create an AFTER trigger on the view.

C. Modify the view to use the WITH VIEW_METADATA clause.

D. Create a clustered index on the view.

Correct Answer: D

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

QUESTION 20

You have a vendor application that uses a scalar function.

You discover that the queries for the application run slower than expected.

You need to gather the runtime information of the scalar function.

What are two possible ways to achieve this goal? Each correct answer presents a complete solution.

A. Enable the Include Actual Execution Plan option.

B. Enable the Display Estimated Execution Planoption.

C. Create and then enable a profiler trace.

D. Create and then enable an extended events trace.

E. Run the Database Engine Tuning Advisor.

Correct Answer: AD

A: An execution plan is the result of the query optimizer\\'s attempt to calculate the mostefficient way to implement the request represented by the T-SQL query you submitted. To generate the first execution plan, you can enable the Include Actual Execution Plan option.

D: SQL Server Extended Events can be used to capture User Defined Function(UDF) counts
References:https://www.mssqltips.com/sqlservertip/4100/how-to-find-udfs-causing-sql-server-performance-issues/

QUESTION 21

You use a Microsoft SQL Server 2012 database that contains two tables named SalesOrderHeader and SalesOrderDetail. The indexes on the tables are as shown in the exhibit.

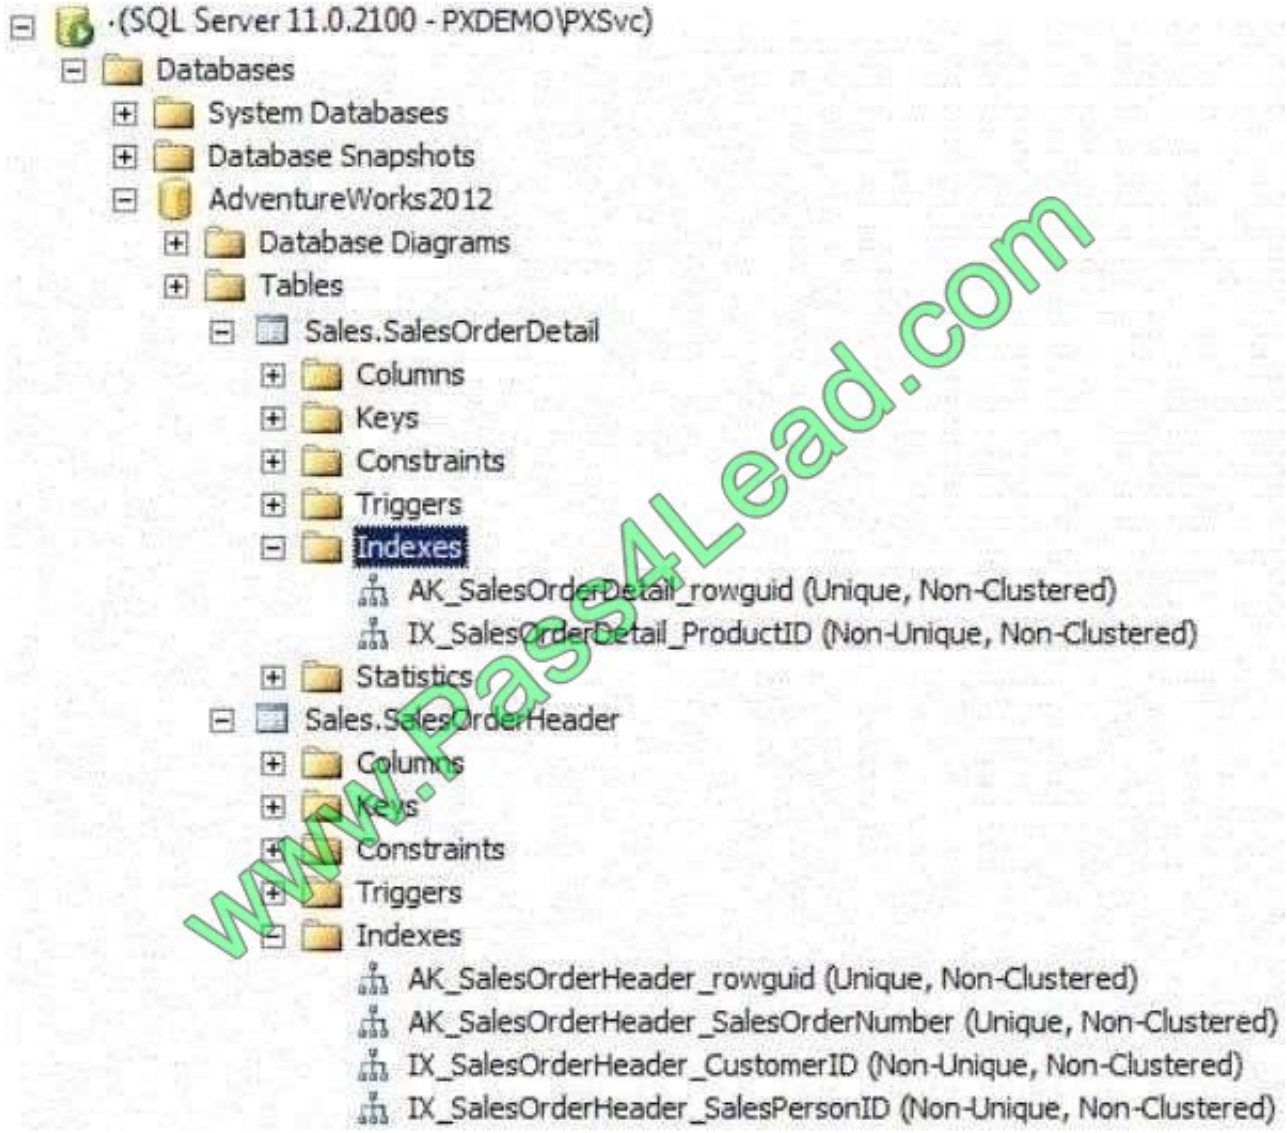

You write the following Transact-SQL query: You discover that the performance of the query is slow. Analysis of the query plan shows table scans where the estimated rows do not match the actual rows for SalesOrderHeader by using an unexpected index on SalesOrderDetail. You need to improve the performance of the query. What should you do?

```
SELECT h.SalesOrderID, h.TotalDue, d.OrderQty
FROM Sales.SalesOrderHeader AS h
  INNER JOIN Sales.SalesOrderDetail AS d
  ON h.SalesOrderID = d.SalesOrderID
WHERE h.TotalDue > 100
AND (d.OrderQty > 5 OR d.LineTotal < 1000.00);
```

A. Use a FORCESCAN hint in the query.

B. Add a clustered index on SalesOrderId in SalesOrderHeader.

C. Use a FORCESEEK hint in the query.

D. Update statistics on SalesOrderId on both tables.

Correct Answer: D

http://msdn.microsoft.com/en-us/library/ms187348.aspx

QUESTION 22

Which FOR XML options can you use to manually format the XML returned? (Choose all that apply.)

A. FOR XML AUTO

B. FOR XML EXPLICIT

C. FOR XML RAW

D. FOR XML PATH

Correct Answer: BD

A. Incorrect: FOR XML AUTO automatically formats the XML retuned. B. Correct: FOR XML EXPLICIT allows you to manually format the XML returned. C. Incorrect: FOR XML RAW automatically formats the XML retuned. D. Correct: FOR XML PATH allows you to manually format the XML returned.

QUESTION 23

You develop a database for a travel application. You need to design tables and other database objects. You create the Airline_Schedules table. You need to store the departure and arrival dates and times of flights along with time zone information. What should you do?

A. Use the CAST function.

B. Use the DATE data type.

C. Use the FORMAT function.

D. Use an appropriate collation.

E. Use a user-defined table type.

F. Use the VARBINARY data type.

G. Use the DATETIME data type.

H. Use the DATETIME2 data type.

I. Use the DATETIMEOFFSET data type.

J. Use the TODATETIMEOFFSET function.

Correct Answer: I

http://msdn.microsoft.com/en-us/library/ff848733.aspx http://msdn.microsoft.com/en-us/library/bb630289.aspx

QUESTION 24

You are designing a table for a SQL Server database. The table uses horizontal partitioning.

You have the following requirements:

Each record in the table requires a unique key.

You must minimize table fragmentation as the table grows.

You need to choose the appropriate data type for the key value.

What should you do?

A. Use the NEWID function to create a unique identifier.

B. Use the NEWSEQUENTIALID function to create a unique identifier.

C. Generate a random value that uses the bigint datatype.

D. Generate a random value that uses the char(16) data type.

Correct Answer: B

Horizontal partitioning divides a table into multiple tables. Each table then contains the same number of columns, but fewer rows. For example, a table that contains 1 billion rows could be partitioned horizontally into 12 tables, with each smaller table representing one month of data for a specific year. Any queries requiring data for a specific month only reference the appropriate table.

NEWSEQUENTIALID creates a GUID that is greater than any GUID previously generated by this function on a specified computer since Windows was started. After restarting Windows, the GUID can start again from a lower range, but is still globally unique. When a GUID column is used as a row identifier, using NEWSEQUENTIALID can be faster than using the NEWID function. This is because the NEWID function causes random activity and uses fewer cached data pages. Using NEWSEQUENTIALID also helps to completely fill the data and index pages.

References:https://msdn.microsoft.com/en-us/library/ms189786.aspx

QUESTION 25

Which of the following datatypes has a fixed precision and a scale of six digits?

A. Double

B. Money

C. Int

D. Numeric

E. SmallInt

F. VarInt

G. Float

Correct Answer: D

QUESTION 26

You have a Microsoft SQL Server database that includes two tables named EmployeeBonus and BonusParameters. The tables are defined by using the following Transact-SQL statements.

```
CREATE TABLE [dbo].[EmployeeBonus](
[EmpNumber] [int] NOT NULL,
[Quarterly] [tinyint] NULL,
[HalfYearly] [tinyint] NULL,
[Yearly] [tinyint] NULL
) ON [PRIMARY]

CREATE TABLE [dbo].[BonusParameters](
[AvailableBonus] [money] NOT NULL,
[CompanyPerformance] [tinyint] NOT NULL
) ON [PRIMARY]
```


The tables are used to compute a bonus for each employee. The EmployeeBonus table has a non- null value in either the Quarterly, HalfYearly or Yearly column. This value indicates which type of bonus an employee receives. The BonusParameters table contains one row for each calendar year that stores the amount of bonus money available and a company performance indicator for that year. You need to calculate a bonus for each employee at the end of a calendar year. Which Transact-SQL statement should you use?

A. SELECT CAST(CHOOSE((Quarterly * AvailableBonus * CompanyPerformance)/40, (HalfYearly * AvailableBonus * CompanyPerformance)/20, (Yearly * AvailableBonus * CompanyPerformance)/10) AS money) AS `Bonus\\' FROM EmployeeBonus, BonusParameters

B. SELECT "Bonus" = CASE EmployeeBonus WHEN Quarterly=1 THEN (Quarterly * AvailableBonus *CompanyPerformance)/40 WHEN HalfYearly=1 THEN (HalfYearly * AvailableBonus * CompanyPerformance)/20 WHEN Yearly=1 THEN (Yearly * AvailableBonus * CompanyPerformance)/10 END FROM EmployeeBonus,BonusParameters

C. SELECT CAST(COALESCE((Quarterly * AvailableBonus * CompanyPerformance)/40, (HalfYearly * AvailableBonus * CompanyPerformance)/20, (Yearly * AvailableBonus * CompanyPerformance)/10) AS money) AS `Bonus\\' FROM EmployeeBonus, BonusParameters

D. SELECT NULLIF(NULLIF((Quarterly * AvailableBonus * CompanyPerformance)/40,(HalfYearly * AvailableBonus * CompanyPerformance)/20), (Yearly * AvailableBonus * CompanyPerformance)/10) AS `Bonus\\' FROM EmployeeBonus, BonusParameters

Correct Answer: B

QUESTION 27

You\\'re designing a new SQL Server 2012 query for the HR department. The query will find records from the persons table for people whose name starts with \\'Ja\\'. Which WHERE statement would be the correct choice?

A. where (name like \\'Ja%\\')

B. where (name = \\'Ja%\\')

C. where (name > \\'Ja\\')

D. where (name like \\'%Ja%\\')

Correct Answer: A

QUESTION 28

You need to create a query that calculates the total sales of each OrderID from a table named Sales.Details. The table contains two columns named OrderID and ExtendedAmount. The solution must meet the following requirements:

Use one-part names to reference columns.

Order the results by OrderID with the smallest value first.

NOT depend on the default schema of a user.

Use an alias of TotalSales for the calculated ExtendedAmount.

Display only the OrderID column and the calculated TotalSales column.

Provide the correct code in the answer area.

Key Words:

ADD	CREATE	EXTERNAL	LIKE	PUBLIC	TABLE
ALL	CROSS	FETCH	LINENO	RAISERROR	TABLESAMPLE
ALTER	CURRENT	FILE	LOAD	READ	TEXTSIZE
AND	CURRENT_DATE	FILLFACTOR	MERGE	READTEXT	THEN
ANY	CURRENT_TIME	FOR	NATIONAL	RECONFIGURE	TO
AS	CURRENT_TIMESTAMP	FOREIGN	NOCHECK	REFERENCES	TOP
ASC	CURRENT_USER	FREETEXT	NONCLUSTERED	REPLICATION	TRAN
AUTHORIZATION	CURSOR	FREETEXTTABLE	NOT	RESTORE	TRANSACTION
BACKUP	DATABASE	FROM	NULL	RESTRICT	TRIGGER
BEGIN	DBCC	FULL	NULLIF	RETURN	TRUNCATE
BETWEEN	DEALLOCATE	FUNCTION	OF	REVERT	TRY_CONVERT
BREAK	DECLARE	GOTO	OFFSETS	REVOKE	TSEQUAL
BROWSE	DEFAULT	GRANT	ON	RIGHT	UNION
BULK	DELETE	GROUP	OPEN	ROLLBACK	UNIQUE
BY	DENY	HAVING	OPENDATASOURCE	ROWCOUNT	UNPIVOT
CASCADE	DESC	HOLDLOCK	OPENQUERY	ROWGUIDCOL	UPDATE
CASE	DISK	IDENTITY	OPENROWSET	RULE	UPDATETEXT
CHECK	DISTINCT	IDENTITY_INSERT	OPENXML	SAVE	USE
CHECKPOINT	DISTRIBUTED	IDENTITYCOL	OPTION	SCHEMA	USER
CLOSE	DOUBLE	IF	OR	SECURITYAUDIT	VALUES
CLUSTERED	DROP	IN	ORDER	SELECT	VARYING
COALESCE	DUMP	INDEX	OUTER	SEMANTICKEYPHRASETABLE	VIEW
COLLATE	ELSE	INNER	OVER	SEMANTICSIMILARITYDETAILSTABLE	WAITFOR
COLUMN	END	INSERT	PERCENT	SEMANTICSIMILARITYTABLE	WHEN
COMMIT	ERRL VL	INTERSECT	PIVOT	SESSION_USER	WHERE
COMPUTE	ESCAPE	INTO	PLAN	SET	WHILE
CONSTRAINT	EXCEPT	IS	PRECISION	SETUSER	WITH
CONTAINS	EXEC	JOIN	PRIMARY	SHUTDOWN	WITHIN GROUP
CONTAINSTABLE	EXECUTE	KEY	PRINT	SOME	WRITETEXT
CONTINUE	EXISTS	KILL	PROC	STATISTICS	
CONVERT	EXIT	LEFT	PROCEDURE	SYSTEM_USER	

Correct Answer:

SELECT OrderId, SUM(ExtendedAmount) AS TotalSales FROM Sales.Details ORDER BY OrderID ASC

QUESTION 29

You use Microsoft SQL Server 2012 database to develop a shopping cart application. You need to rotate the unique values of the ProductName field of a table-valued expression into multiple columns in the output. Which Transact-SQL operator should you use?

A. CROSS JOIN

B. CROSS APPLY

C. PIVOT

D. UNPIVOT

Correct Answer: C

http://technet.microsoft.com/en-us/library/ms177634.aspx

QUESTION 30

You use a Microsoft Azure SQL DataBase instance. The instance contains a table named Customers that has columns

named Id, Name, and IsPriority.

You need to create a view named VwPriorityCustomers that:

returns rows from Customer that have a value of True in the IsPriority column, and does not allow columns to be altered or dropped in the underlying table.

Which Transact-SQL statement shoul you run?

A. CREATE VIEW VwPriorityCustomers AS SELECT Id, Name FROM dbo.Customers WHERE IsPriority=1 WITH CHECK OPTION

B. CREATE VIEW VwPriorityCustomers WITH VIEW_METADATA AS SELECT Id, Name FROM dbo.Customers WHERE IsPriority=1

C. CREATE VIEW VwPriorityCustomers WITH ENCRYPTION AS SELECT Id, Name FROM dbo.Customers WHERE IsPriority=1

D. CREATE VIEW VwPriorityCustomers WITH SCHEMABINDING AS SELECT Id, Name FROM dbo.Customers WHERE IsPriority=1

Correct Answer: D

SCHEMABINDING binds the view to the schema of the underlying table or tables. When SCHEMABINDING is specified, the base table or tables cannot be modified in a way that would affect the view definition. References: https://docs.microsoft.com/en-us/sql/t-sql/statements/create-view-transact-sql?view=sql-server-2017

QUESTION 31

Which of the following is not a FLWOR clause?

A.for

B. let

C. where

D. over

E. return

Correct Answer: D

FLWOR (pronounced "flower") is an acronym for

"For,

Let,

Where,

Order by,

Return".

QUESTION 32

You need to create a stored procedure that enters values into multiple tables. The solution must ensure that if a single insert fails, none of the values are inserted into the tables.

How should you complete the stored procedure? 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.

Select and Place:

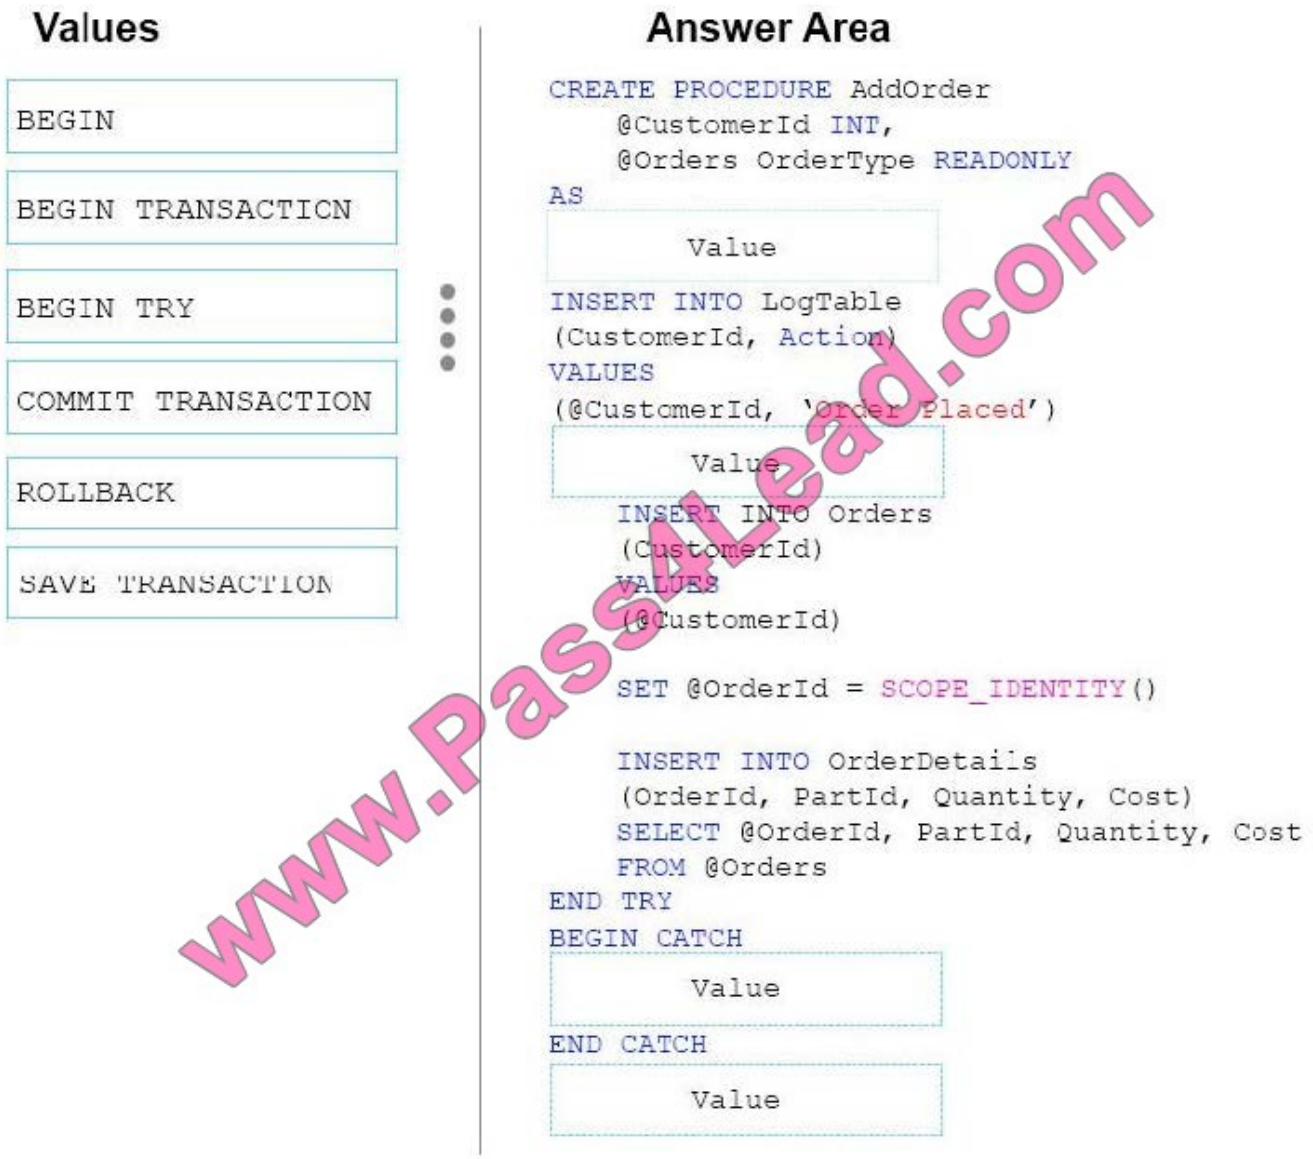

Correct Answer:

Box 1: BEGIN TRY Box 2:SAVE TRANSACTION Box 3: ROLLBACK Box 4: COMMIT TRANSACTION
References:https://msdn.microsoft.com/en-us/library/ms188378.aspx

QUESTION 33

You use a Microsoft SQL Server database that contains a table. The table has records of customer orders.

Your company has three divisions that have the following names:

East

Central West

You need to create a query that displays the following information: The number of sales for each product (ProductName) grouped by the division (Division) that sold the product A column for each division

Which Transact-SQL query should you use?

A.

```
DECLARE @Results TABLE (
  ProductName VARCHAR(255),
  East INT,
  Central INT,
  West INT)

INSERT INTO @Results (ProductName, East)
SELECT ProductName, COUNT(OrderId)
FROM Orders
WHERE Division = "East"

UPDATE @Results
SET Central = COUNT(OrderId)
FROM Orders o INNER JOIN @Results r ON o.ProductName = r.ProductName
WHERE Division = "Central"

UPDATE @Results
SET West = COUNT(OrderId)
FROM Orders o INNER JOIN @Results r ON o.ProductName = r.ProductName
WHERE Division = "West"

SELECT
  ProductName,
  East,
  Central,
  West
FROM
  @Results
```

B.

```
SELECT
  ProductName,
  SUM(CASE WHEN Division = "East" THEN 1 ELSE 0 END) AS East,
  SUM(CASE WHEN Division = "Central" THEN 1 ELSE 0 END) AS Central,
  SUM(CASE WHEN Division = "West" THEN 1 ELSE 0 END) AS West
FROM
  Orders
GROUP BY
  Division
ORDER BY
  ProductName
```

A. B.

C.

```
SELECT
  ProductName,
  East,
  Central,
  West
FROM
  (SELECT OrderId, Division, ProductName FROM Orders) o
   PIVOT (
    COUNT (o.OrderId)
    FOR o.Division IN (East, Central, West)
  ) AS pvt
ORDER BY
  pvt.ProductName
```

D.

```
SELECT
  ProductName,
  Server,
  Requests
FROM
  (SELECT OrderId, Division, ProductName FROM Orders) o
   UNPIVOT (
    Requests FOR Division IN ([Division])
  ) AS pvt
ORDER BY
  pvt.ProductName
```


C. D.

Correct Answer: D

You can use the PIVOT and UNPIVOT relational operators to change a table-valued expression into another table. PIVOT rotates a table-valued expression by turning the unique values from one column in the expression into multiple columns in the output, and performs aggregations where they are required on any remaining column values that are wanted in the final output. UNPIVOT performs the opposite operation to PIVOT by rotating columns of a table-valued expression into column values.

References: https://technet.microsoft.com/en-us/library/ms177410(v=sql.105).aspx

QUESTION 34

You administer a Microsoft SQL Server database that has Trustworthy set to On. You create a stored procedure that returns database-level information from Dynamic Management Views. You grant User1 access to execute the stored

procedure. You need to ensure that the stored procedure returns the required information when User1 executes the stored procedure. You need to achieve this goal by granting the minimum permissions required. What should you do? (Each correct answer presents a complete solution. Choose all that apply.)

A. Create a SQL Server login that has VIEW SERVER STATE permissions. Create an application role and a secured password for the role.

B. Modify the stored procedure to include the EXECUTE AS OWNER statement. Grant VIEW SERVER STATE permissions to the owner of the stored procedure.

C. Create a SQL Server login that has VIEW SERVER STATE permissions. Modifythe stored procedure to include the EXECUTE AS {newlogin} statement.

D. Grant the db_owner role on the database to User1.

E. Grant the sysadmin role on the database to User1.

Correct Answer: DE

QUESTION 35

Which FOR XML options are valid? (Choose all that apply.)

A. FOR XML AUTO

B. FOR XML MANUAL

C. FOR XML DOCUMENT

D. FOR XML PATH

Correct Answer: AD

QUESTION 36

You develop an SQL Server database. The database contains a table that is defined by the following T-SQL statements:

```
CREATE TABLE Employees
(employeeNumber INT,
  surName VARCHAR(100),
  givenName VARCHAR(25),
  dateOfBirth DATE,
  workPhone VARCHAR(12));
```

The table contains duplicate records based on the combination of values in the surName, givenName, and dateOfBirth fields.

You need to remove the duplicate records.

How should you complete the relevant Transact-SQL statements? To answer, drag the appropriate code segment or segments to the correct location or locations in the answer area. Each code segment 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:

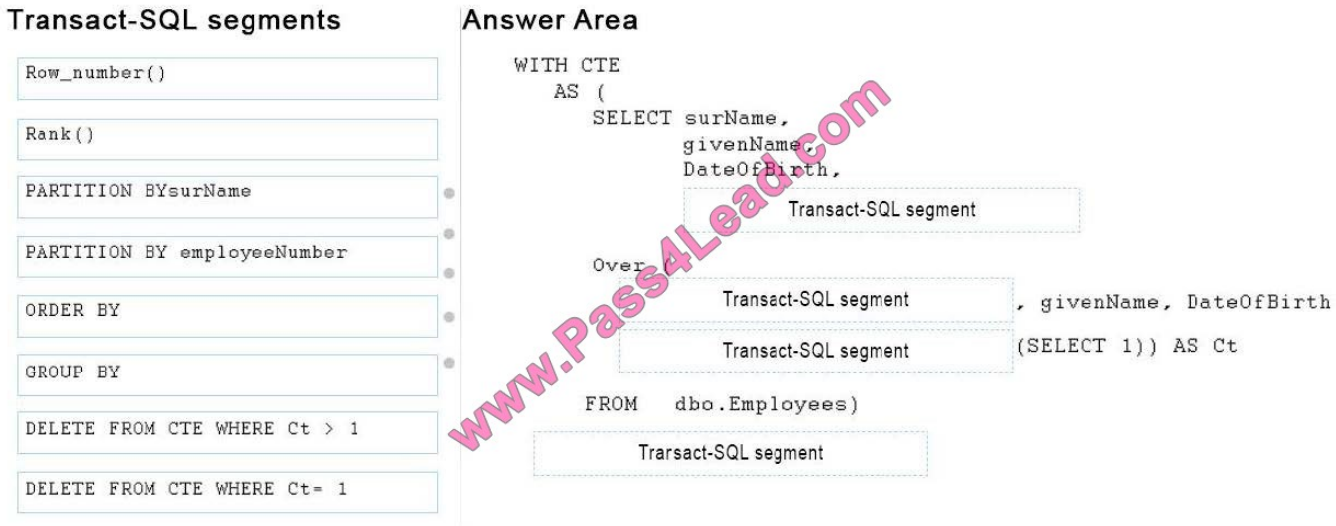

Correct Answer:

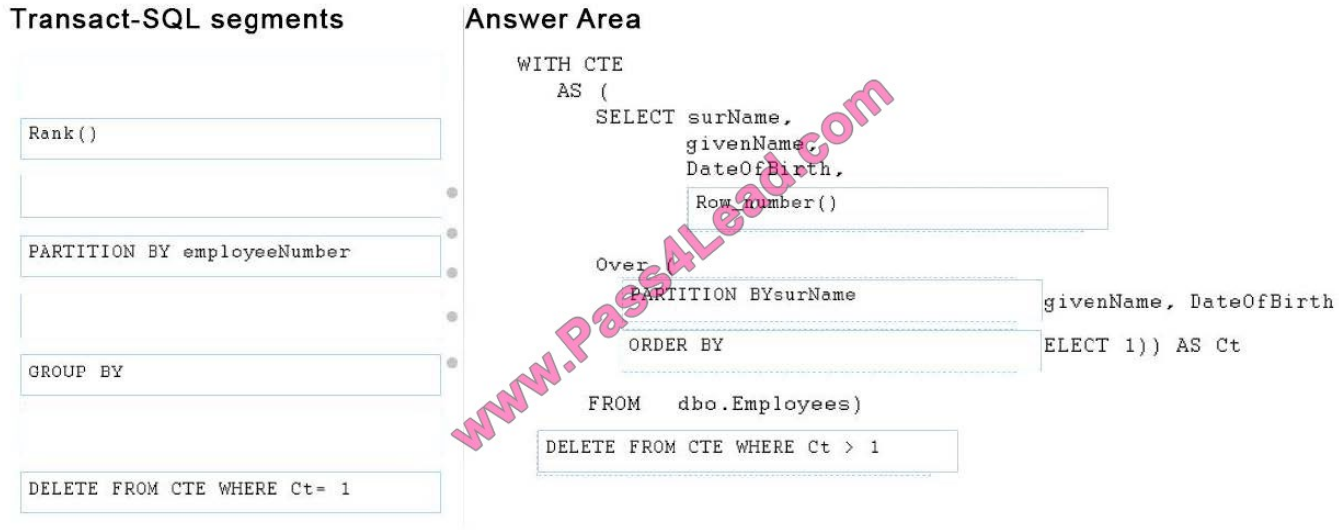

let us write a query which will delete all duplicate data in one shot. We will use a CTE (Common Table Expression) for this purpose. We will read in future posts what a CTE is and why it is used. On a lighter note, CTE\\'s can be imagined as equivalent to temporary result sets that can be used only in an underlying SELECT, INSERT, UPDATE, DELETE or CREATE VIEW statement. ;WITH CTE AS (SELECT Name , City , [State] , ROW_NUMBER() OVER(PARTITION BY Name, City, [State] ORDER BY [Name]) AS Rnum FROM Persons) DELETE FROM CTE WHERE Rnum 1 In the code by saying WHERE Rnum 1, we are asking SQL Server to keep all the records with Rank 1, which are not duplicates, and delete any other record. After executing this query in SQL Server Management Studio, you will end up with no duplicates in your table. To confirm that just run a simple query against your table.

QUESTION 37

You administer a Microsoft SQL Server database that contains a table named OrderDetail. You discover that the NCI_OrderDetail_CustomerID non-clustered index is fragmented. You need to reduce fragmentation.

You need to achieve this goal without taking the index offline. Which Transact-SQL batch should you use?

A. CREATE INDEX NCI_OrderDetail_CustomerID ON OrderDetail.CustomerID WITH DROP EXISTING

B. ALTER INDEX NCI_OrderDetail_CustomerID ON OrderDetail.CustomerID REORGANIZE

C. ALTER INDEX ALL ON OrderDetail REBUILD

D. ALTER INDEX NCI_OrderDetail_CustomerID ON OrderDetail.CustomerID REBUILD

Correct Answer: B

Reference: http://msdn.microsoft.com/en-us/library/ms188388.aspx

QUESTION 38

You need to create a table named OrderDetails on a new server. OrderDetails must meet the following requirements:

Contain a new column named LineItemTotal that stores the product of ListPrice and Quantity for each row.

The calculation for a line item total must not be run every time the table is queried.

The code must NOT use any object delimiters.

The solution must ensure that LineItemTotal is stored as the last column in the table.

Part of the correct T-SQL statement has been provided in the answer area. Provide the complete code.

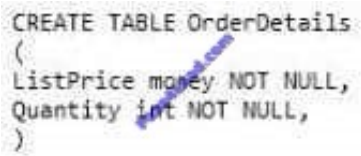

Correct Answer: Please review the explanation part for this answer

CREATE TABLE OrderDetails (ListPrice money NOT NULL, Quantity int NOT NULL, LineItemTotal AS (ListPrice * Quantity) PERSISTED)

QUESTION 39

Which of the following are true about the SET QUOTED_IDENTIFIER statement? (Choose all that apply.)

A. When set to ON, QUOTED_IDENTIFIER allows you to use double quotation marks to delimit T-SQL identifiers such as table and column names.

B. When set to OFF, QUOTED_IDENTIFIER allows you to use double quotation marks to delimit T-SQL identifiers such as table and column names.

C. When set to ON, QUOTED_IDENTIFIER allows you to use double quotation marks to delimit strings.

D. When set to OFF, QUOTED_IDENTIFIER allows you to use double quotation marks to delimit strings.

Correct Answer: AD

QUESTION 40

You have an XML schema collection named Sales.InvoiceSchema.

You need to declare a variable of the XML type named XML1. The solution must ensure that XML1 is validated by using Sales.InvoiceSchema.

Which code segment should you use?

To answer, type the correct code in the answer area.

Correct Answer: DECLARE @XML1 XML(Sales.InvoiceSchema)

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.pass4lead.com/allproducts

Need Help

Please provide as much detail as possible so we can best assist you.
To update a previously submitted ticket: