

[2017 New Easily Pass 70-461 Exam With Lead2pass Updated Microsoft 70-461 Dumps (169-177)]

2017 May **Microsoft** Official New Released **70-461** Q&As in Lead2pass.com!] 100% Free Download! 100% Pass Guaranteed!

Good news, Lead2pass has updated the 70-461 exam dumps. With all the questions and answers in your hands, you will pass the Microsoft 70-461 exam easily. Following questions and answers are all new published by Microsoft Official Exam Center:

<http://www.lead2pass.com/70-461.html> QUESTION 169 You have a database that contains a table named Customer. The customer table contains a column named LastName that has a column definition of varchar(50). An application named App1 reads from the table frequently. You need to change the column definition to nvarchar(100). The solution must minimize the amount of time it takes for App1 to read the data. Which statement should you execute? A. B. C. D. Answer: C Explanation: To change the data type of a column in a SQL Server (or Microsoft access) table, use the following syntax: ALTER TABLE table_name ALTER COLUMN column_name datatype References: https://www.w3schools.com/SQL/sql_alter.asp QUESTION 170 You are maintaining a SQL Server database that uses the default settings. The database contains a table that is defined by the following Transact-SQL statement: You must write a query that returns the AddressLine1, AddressLine2, and Region fields separated by carriage returns. You must return an empty string for any null values. A. B. C. D. Answer: A Explanation: Char(13) is a carriage return. Use the IIF construct to return an empty string for NULL values of the AddressLine2 column. IIF returns one of two values, depending on whether the Boolean expression evaluates to true or false in SQL Server. References:

<https://msdn.microsoft.com/en-us/library/hh213574.aspx> QUESTION 171 A table named Profits stores the total profit made each year within a territory. The Profits table has columns named Territory, Year, and Profit. You need to create a report that displays the profits made by each territory for each year and its previous year. Which Transact-SQL query should you use? A. SELECT Territory, Year, Profit, LEAD(Profit, 1, 0) OVER (PARTITION BY Territory ORDER BY Year) AS PreviousYearProfit FROM Profits B. SELECT Territory, Year, Profit, LAG(Profit, 1, 0) OVER (PARTITION BY Year ORDER BY Territory) AS PreviousYearProfit FROM Profits C. SELECT Territory, Year, Profit, LAG(Profit, 1, 0) OVER (PARTITION BY Territory ORDER BY Year) AS PreviousYearProfit FROM Profits D. SELECT Territory, Year, Profit, LEAD(Profit, 1, 0) OVER (PARTITION BY Year ORDER BY Territory) AS PreviousYearProfit FROM Profits Answer: C Explanation: LAG accesses data from a previous row in the same result set without the use of a self-join in SQL Server 2016. LAG provides access to a row at a given physical offset that comes before the current row. Use this analytic function in a SELECT statement to compare values in the current row with values in a previous row. Use ORDER BY Year, not ORDER BY Territory. Example: The following example uses the LAG function to return the difference in sales quotas for a specific employee over previous years. Notice that because there is no lag value available for the first row, the default of zero (0) is returned. USE AdventureWorks2012; GO SELECT BusinessEntityID, YEAR(QuotaDate) AS SalesYear, SalesQuota AS CurrentQuota, LAG(SalesQuota, 1, 0) OVER (ORDER BY YEAR(QuotaDate)) AS PreviousQuota FROM Sales.SalesPersonQuotaHistory WHERE BusinessEntityID = 275 and YEAR(QuotaDate) IN ('2005','2006'); Incorrect Answers: A, D: LEAD accesses data from a subsequent row in the same result set without the use of a self-join in SQL Server 2016. LEAD provides access to a row at a given physical offset that follows the current row. Use this analytic function in a SELECT statement to compare values in the current row with values in a following row. B: Use ORDER BY Year, not ORDER BY Territory. References: <https://msdn.microsoft.com/en-us/library/hh231256.aspx> QUESTION 172 Your database contains a table named Products that has columns named ProductID and Name. You want to write a query that retrieves data from the Products table sorted by Name listing 15 rows at a time. You need to view rows 31 through 45. Which Transact-SQL query should you create? A. B. C. D. Answer: C Explanation: The OFFSET-FETCH clause provides you with an option to fetch only a window or page of results from the result set. OFFSET-FETCH can be used only with the ORDER BY clause. Example: Skip first 10 rows from the sorted resultset and return next 5 rows. SELECT First Name + ' ' + Last Name FROM Employees ORDER BY First Name OFFSET 10 ROWS FETCH NEXT 5 ROWS ONLY; References:

[https://technet.microsoft.com/en-us/library/gg699618\(v=sql.110\).aspx](https://technet.microsoft.com/en-us/library/gg699618(v=sql.110).aspx) QUESTION 173 A database named AdventureWorks contains two tables named Production.Product and Sales.SalesOrderDetail. The tables contain data on the available products and a detailed order history. The Production.Product table contains the following two columns: - ProductID - Name The Sales.SalesOrderDetail table contains the following three columns: - SalesOrderID - ProductID - OrderQty You need to create a query listing all of the products that were never ordered. Which statements should you execute? A. B. C. D. Answer: A Explanation: EXCEPT and INTERSECT returns distinct rows by comparing the results of two queries. EXCEPT returns distinct rows from the left input query that aren't output by the right input query. Column names or aliases in ORDER BY clauses must reference column names returned by the left-side query. Example: The following query returns any distinct values from the query to

the left of the EXCEPT operator that are not also found on the right query. -- Uses AdventureWorks SELECT CustomerKey FROM FactInternetSales EXCEPT SELECT CustomerKey FROM DimCustomer WHERE DimCustomer.Gender = 'F' ORDER BY CustomerKey; --Result: 9351 Rows (Sales to customers that are not female.) Incorrect Answers: B: Column names or aliases in ORDER BY clauses must reference column names returned by the left-side query. C, D: INTERSECT returns distinct rows that are output by both the left and right input queries operator. References: <https://msdn.microsoft.com/en-us/library/ms188055.aspx>

QUESTION 174 You plan to write a query for a new business report that will contain several nested queries. You need to ensure that a nested query can call a table-valued function for each row in the main query. Which query operator should you use in the nested query? A. CROSS APPLY B. INNER JOIN C. OUTER JOIN D. PIVOT Answer: A Explanation: The APPLY operator allows you to invoke a table-valued function for each row returned by an outer table expression of a query. The table-valued function acts as the right input and the outer table expression acts as the left input. The right input is evaluated for each row from the left input and the rows produced are combined for the final output. The list of columns produced by the APPLY operator is the set of columns in the left input followed by the list of columns returned by the right input. There are two forms of APPLY: CROSS APPLY and OUTER APPLY. CROSSAPPLY returns only rows from the outer table that produce a result set from the table-valued function. OUTER APPLY returns both rows that produce a result set, and rows that do not, with NULL values in the columns produced by the table-valued function. References: [https://technet.microsoft.com/en-us/library/ms175156\(v=sql.105\).aspx](https://technet.microsoft.com/en-us/library/ms175156(v=sql.105).aspx)

QUESTION 175 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. Answer: B Explanation: 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 176 You are a database developer of a Microsoft SQL Server database. You are designing a table that will store Customer data from different sources. The table will include a column that contains the CustomerID from the source system and a column that contains the SourceID. A sample of this data is as shown in the following table. You need to ensure that the table has no duplicate CustomerID within a SourceID. You also need to ensure that the data in the table is in the order of SourceID and then CustomerID. Which Transact-SQL statement should you use? A. CREATE TABLE Customer (SourceID int NOT NULL, CustomerID int NOT NULL, CustomerName varchar(255) NOT NULL CONSTRAINT UQ_Customer UNIQUE (SourceID, CustomerID)); B. CREATE TABLE Customer (SourceID int NOT NULL UNIQUE, CustomerID int NOT NULL UNIQUE, CustomerName varchar(255) NOT NULL); C. CREATE TABLE Customer (SourceID int NOT NULL PRIMARY KEY CLUSTERED, CustomerID int NOT NULL UNIQUE, CustomerName varchar(255) NOT NULL); D. CREATE TABLE Customer (SourceID int NOT NULL, CustomerID int NOT NULL, CustomerName varchar(255) NOT NULL, CONSTRAINT PK_Customer PRIMARY KEY CLUSTERED (SourceID, CustomerID)); Answer: D Explanation: A PRIMARY KEY is a constraint that enforces entity integrity for a specified column or columns by using a unique index. Only one PRIMARY KEY constraint can be created for each table. We need to use both SourceID and CustomerID, in that order, in the PRIMARY KEY constraint. References: <https://msdn.microsoft.com/en-us/library/ms188066.aspx>

QUESTION 177 Drag and Drop Question You maintain a SQL Server database that is used by a retailer to track sales and salespeople. The database includes two tables and two triggers that is defined by the following Transact-SQL statements: During days with a large sales volume, some new sales transaction fail and report the following error: Arithmetic overflow error converting expression to data type int. You need to ensure that the two triggers are applied once per sale, and that they do not interfere with each other. How should you complete the relevant Transact-SQL statement? To answer, drag the appropriate Transact-SQL segments to the correct location or locations. Each Transact-SQL 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. Answer: Explanation: Box 1: OFF Box 2: 0 Only direct recursion of AFTER triggers is prevented when the RECURSIVE_TRIGGERS database option is set to OFF. To disable indirect recursion of AFTER triggers, also

set the nested triggers server option to 0. Note: Both DML and DDL triggers are nested when a trigger performs an action that initiates another trigger. These actions can initiate other triggers, and so on. DML and DDL triggers can be nested up to 32 levels. You can control whether AFTER triggers can be nested through the nested triggers server configuration option. If nested triggers are allowed and a trigger in the chain starts an infinite loop, the nesting level is exceeded and the trigger terminates. References: <https://msdn.microsoft.com/en-us/library/ms190739.aspx> Once there are some changes on 70-461 exam questions, we will update the study materials timely to make sure that our customer can download the latest edition. 70-461 new questions on Google Drive: <https://drive.google.com/open?id=0B3Syig5i8gpDQ25HZW9OcFRldDQ> 2017 Microsoft 70-461 exam dumps (All 200 Q&As) from Lead2pass: <http://www.lead2pass.com/70-461.html> [100% Exam Pass Guaranteed!!!]