



Vendor: Microsoft

Exam Code: 70-463

Exam Name: Implementing a Data Warehouse with
Microsoft SQL Server 2012 Exam

Version: DEMO

QUESTION 1

You are developing a SQL Server Integration Services (SSIS) package to implement an incremental data load strategy. The package reads data from a source system that uses the SQL Server change data capture (CDC) feature.

You have added a CDC Source component to the data flow to read changed data from the source system.

You need to add a data flow transformation to redirect rows for separate processing of insert, update, and delete operations.

Which data flow transformation should you use?

- A. Audit
- B. Merge Join
- C. Merge
- D. CDC Splitter

Answer: C

Explanation:

MERGE with Incremental Loads The change data capture functionality that is new in SQL Server 2008 makes it easier to perform incremental loads reliably to a data warehouse. As an alternative to using parameterized OLE DB Command transformations to perform the inserts and the updates, you can use the MERGE statement to combine both operations.

QUESTION 2

You are implementing the indexing strategy for a fact table in a data warehouse. The fact table is named Quotes. The table has no indexes and consists of seven columns:

```
[ID]
[QuoteDate]
[Open]
[Close]
[High]
[Low]
[Volume]
```

Each of the following queries must be able to use a columnstore index:

```
SELECT AVG ([Close]) AS [AverageClose] FROM Quotes WHERE [QuoteDate]
BETWEEN '20100101' AND '20101231'.
SELECT AVG([High] - [Low]) AS [AverageRange] FROM Quotes WHERE
[QuoteDate] BETWEEN '20100101' AND
'20101231'.
SELECT SUM([Volume]) AS [SumVolume] FROM Quotes WHERE [QuoteDate]
BETWEEN '20100101' AND '20101231'.
```

You need to ensure that the indexing strategy meets the requirements.

The strategy must also minimize the number and size of the indexes.

What should you do?

- A. Create one columnstore index that contains [ID], [Close], [High], [Low], [Volume], and [QuoteDate].
- B. Create three columnstore indexes:
 - One containing [QuoteDate] and [Close]
 - One containing [QuoteDate], [High], and [Low]
 - One containing [QuoteDate] and [Volume]
- C. Create one columnstore index that contains [QuoteDate], [Close], [High], [Low], and [Volume].

D. Create two columnstore indexes:

- One containing [ID], [QuoteDate], [Volume], and [Close]
- One containing [ID], [QuoteDate], [High], and [Low]

Answer: C

QUESTION 3

Drag and Drop Questions

You are loading a dataset into SQL Server. The dataset contains numerous duplicates for the Artist and Song columns.

The values in the Artist column in the dataset must exactly match the values in the Artist domain in the knowledge base. The values in the Song column in the dataset can be a close match with the values in the Song domain.

You need to use SQL Server Data Quality Services (DQS) to define a matching policy rule to identify duplicates.

How should you configure the Rule Editor? (To answer, drag the appropriate answers to the answer area.)

The screenshot shows the Rule Editor interface with the following configuration:

Domain	Similarity	Weight	Prerequisite
Artist	Similar		
Song	Similar		

On the left side, there are input fields for '0', '50', and '100', and a 'Prerequisite Choices' section with 'No' and 'Yes' options.

Answer:

The screenshot shows the Rule Editor interface with the following configuration:

Domain	Similarity	Weight	Prerequisite
Artist	Similar	0	Yes
Song	Similar	100	No

On the left side, there are input fields for '0', '50', and '100', and a 'Prerequisite Choices' section with 'No' and 'Yes' options.

QUESTION 4

You are developing a SQL Server Integration Services (SSIS) package.

The package uses a data flow task to source data from a SQL Server database for loading into a dimension table in a data warehouse.

You need to create a separate data flow path for data that has been modified since it was last processed.

Which data flow components should you use to identify modified data? (Each correct answer presents a complete solution. Choose all that apply.)

- A. Multicast
- B. Data Conversion
- C. Lookup
- D. Slowly Changing Dimension
- E. Aggregate

Answer: AC

Explanation:

A: The transformation that distributes data sets to multiple outputs. The transformation that distributes data sets to multiple outputs.

C: Lookup Transformation

The transformation that looks up values in a reference table using an exact match.

Note:

* SQL Server Integration Services provides three different types of data flow components: sources, transformations, and destinations. Sources extract data from data stores such as tables and views in relational databases, files, and Analysis Services databases. Transformations modify, summarize, and clean data. Destinations load data into data stores or create in-memory datasets.

Incorrect:

Not B: Data Conversion Transformation

The transformation that converts the data type of a column to a different data type.

Not D: Slowly Changing Dimension Transformation

The transformation that configures the updating of a slowly changing dimension.

Not E: The Aggregate transformation applies aggregate functions, such as Average, to column values and copies the results to the transformation output. Besides aggregate functions, the transformation provides the GROUP BY clause, which you can use to specify groups to aggregate across.

QUESTION 5

Which of the following statements are true? (Choose all that apply.)

- A. An SSIS package can contain one or more SSDT solutions, each performing a specific data management operation.
- B. An SSIS project can contain one or more SSIS packages.
- C. An SSIS project can contain exactly one SSIS package.
- D. SSIS packages contain programmatic logic used in data movements and data transformation operations.

Answer: BD

QUESTION 6

Hotspot Question

You are developing a SQL Server Integration Services (SSIS) package to load data into a SQL Server 2012 database.

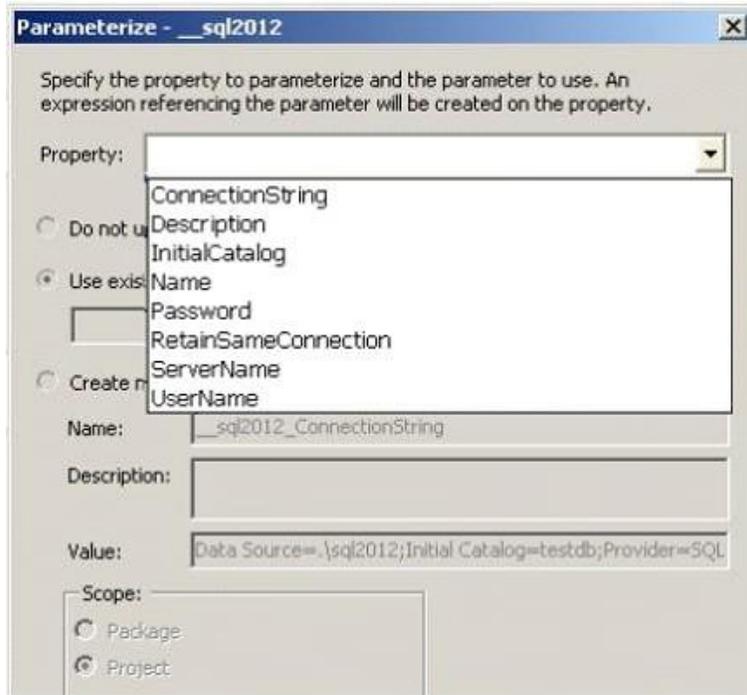
The package is allowed to connect to only one database.

An Environment variable contains the name of the database.

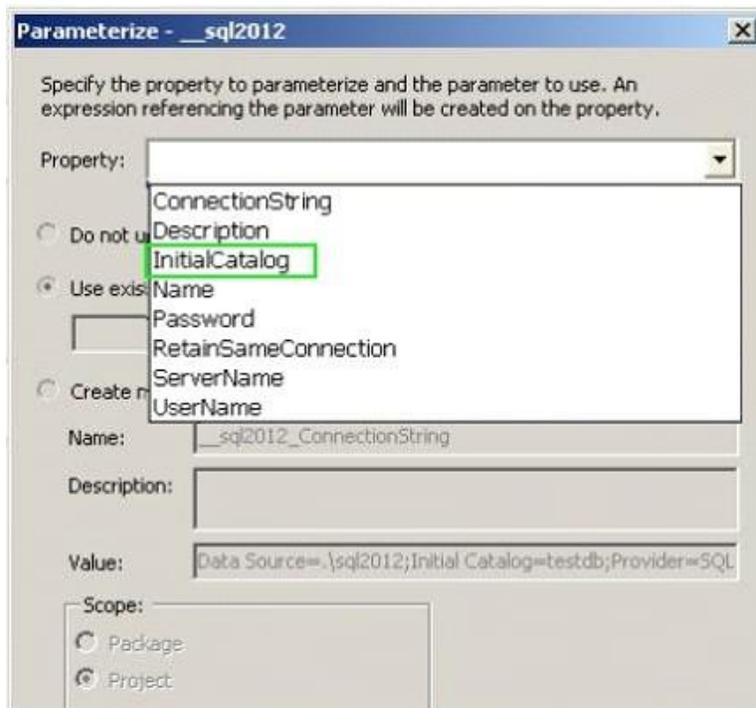
The OLE DB project connection manager has been parameterized.

You need to configure the connection manager property to accept the value of the Environment variable.

Which property should you use? (To answer, configure the appropriate option or options in the dialog box in the answer area.)



Answer:



QUESTION 7

To facilitate the troubleshooting of SQL Server Integration Services (SSIS) packages, a logging methodology is put in place. The methodology has the following requirements:

- * The deployment process must be simplified.
- * All the logs must be centralized in SQL Server.
- * Log data must be available via reports or T-SQL.
- * Log archival must be automated.

You need to configure a logging methodology that meets the requirements while minimizing the amount of deployment and development effort.

What should you do?

- A. Open a command prompt and run the gacutil command.
- B. Open a command prompt and execute the package by using the SQL Log provider and running the dtexecui.exe utility.
- C. Add an OnError event handler to the SSIS project.
- D. Use an msi file to deploy the package on the server.
- E. Configure the output of a component in the package data flow to use a data tap.
- F. Run the dtutil command to deploy the package to the SSIS catalog and store the configuration in SQL Server.
- G. Open a command prompt and run the dtexec /rep /conn command.
- H. Open a command prompt and run the dtutil /copy command.
- I. Open a command prompt and run the dtexec /dumperror /conn command.
- J. Configure the SSIS solution to use the Project Deployment Model.
- K. Create a reusable custom logging component and use it in the SSIS project.

Answer: J

QUESTION 8

You are designing a data warehouse that uses SQL Server 2012.

The data warehouse contains a table named factSales that stores product sales.

The table has a clustered index on the primary key, four foreign keys to dimension tables, and an aggregate column for sales totals.

All key columns use the INT data type, and the aggregate column uses the MONEY data type.

You need to increase the speed of data retrieval from the factSales table.

Which index type should you add to the table?

- A. Full text
- B. Spatial
- C. Non-clustered
- D. Clustered

Answer: C

QUESTION 9

You are administering SQL Server Integration Services (SSIS) permissions on a production server that runs SQL Server 2012.

Quality Assurance (QA) testers in the company must have permission to perform the following tasks on specific projects:

- View and validate projects and packages

- View Environments and Environment variables
- Execute packages

You need to grant the minimum possible privileges to the QA testers.
What should you do? (Each correct answer presents part of the solution. Choose all that apply.)

- A. In the SSISDB database, add QA Tester logons to the ssis_admin role.
- B. In the msdb database, add QA Tester logons to the db_ssisoperator role.
- C. Grant Modify permission in the projects to the QA Tester logons.
- D. Grant Read permission in the SSIS catalog folder, the projects, and the Environments to the QA Tester logons.
- E. Grant Execute permission in the projects to the QA Tester logons.
- F. In the msdb database, add QA Tester logons to the db_ssisltduser role.

Answer: BD

QUESTION 10

You are implementing a SQL Server Integration Services (SSIS) package that imports Microsoft Excel workbook data into a Windows Azure SQL Database database. The package has been deployed to a production server that runs Windows Server 2008 R2 and SQL Server 2012. The package fails when executed on the production server. You need to ensure that the package can load the Excel workbook data without errors. You need to use the least amount of administrative effort to achieve this goal. What should you do?

- A. Create a custom SSIS source component that encapsulates the 32-bit driver and compile it in 64-bit mode.
- B. Install a 64-bit ACE driver and execute the package by using the 64-bit run-time option.
- C. Execute the package by using the 32-bit run-time option.
- D. Replace the SSIS Excel source with a SSIS Flat File source.

Answer: C

Explanation:

* See step 3 below.

To publish an Excel worksheet to Azure SQL Database, your package will contain a Data Flow Task, Excel Source task, and ADO NET Destination.

- 1) Create an SSIS project.
- 2) Drop a Data Flow Task onto the Control Flow design surface, and double click the Data Flow Task.
- 3) Drop an Excel Source onto the Data Flow design surface. Note When using the Excel Source task on a 64-bit machine, set Run64BitRuntime to False.

*Incorrect:

Not D: The Flat File source reads data from a text file. The text file can be in delimited, fixed width, or mixed format.

Thank You for Trying Our Product

PassLeader Certification Exam Features:

- ★ More than **99,900** Satisfied Customers Worldwide.
- ★ Average **99.9%** Success Rate.
- ★ **Free Update** to match latest and real exam scenarios.
- ★ **Instant Download** Access! No Setup required.
- ★ Questions & Answers are downloadable in **PDF** format and **VCE** test engine format.
- ★ Multi-Platform capabilities - **Windows, Laptop, Mac, Android, iPhone, iPod, iPad.**
- ★ **100%** Guaranteed Success or **100%** Money Back Guarantee.
- ★ **Fast**, helpful support **24x7**.



View list of all certification exams: <http://www.passleader.com/all-products.html>



Microsoft



ORACLE



CITRIX



JUNIPER
NETWORKS



EMC²
where information lives®

10% Discount Coupon Code: STNAR2014