



Vendor: Microsoft

Exam Code: 70-463

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

Version: Demo

QUESTION 1

You have a server named SQL1 that has SQL Server Integration Services (SSIS) installed. SQL1 has eight CPUs, 8 GB of RAM, RAID-1 storage, and a 10-gigabit Ethernet connection.

A package named Package1 runs on SQL1. Package1 contains 10 data tasks. Package1 queries 3 GB of data from the source system.

You review the performance statistics on SQL1 and discover that Package1 executes slower than expected.

You need to identify the problem that is causing package1 to execute slowly.

Which performance monitor counter data should you review?

- A. SQLServer:SSIS Pipeline 10.0:Buffers in use
- B. Processor\% Idle Time
- C. SQLServer:SSIS Pipeline 10.0:Buffers spooled
- D. SQLServer:SSIS Pipeline 10.0:Private buffer memory

Correct Answer: B

QUESTION 2

You are designing a SQL Server Integration Services (SSIS) data flow to load sales transactions from a source system into a data warehouse hosted on Windows Azure SQL Database. One of the columns in the data source is named ProductCode.

Some of the data to be loaded will reference products that need special processing logic in the data flow.

You need to enable separate processing streams for a subset of rows based on the source product code.

Which Data Flow transformation should you use?

- A. Script Component
- B. Audit
- C. Destination Assistant
- D. Data Conversion

Correct Answer: A

Explanation:

The Script component can be used as a source, a transformation, or a destination. This

component supports one input and multiple outputs. Depending on how the component is used, it supports either an input or outputs or both. The script is invoked by every row in the input or output.

If used as a source, the Script component supports multiple outputs. If used as a transformation, the Script component supports one input and multiple outputs.

If used as a destination, the Script component supports one input.

QUESTION 3

DRAG DROP

You have a SQL Server data warehouse that contains all of the customer data for your company. You plan to deploy an archiving solution.

You need to move a table named OLDCUSTOMERS to a new storage device. OLDCUSTOMERS contains a clustered index.

Which four actions should you perform in sequence?

The screenshot shows a SQL Server Enterprise Manager interface. On the left, there is a list of actions: "Rebuild the clustered index", "Create a file group", "Create a data file", "Enable Page compression", "Execute the MERGE statement", and "Remove the clustered index". On the right, there are four empty rectangular boxes arranged vertically, intended for the user to drag and drop the selected actions into them in the correct sequence.

Correct Answer:

The screenshot shows the same SQL Server Enterprise Manager interface as above, but with the correct sequence of actions highlighted. The actions "Rebuild the clustered index", "Create a file group", "Create a data file", "Enable Page compression", "Execute the MERGE statement", and "Remove the clustered index" are highlighted with a green border. The actions "Enable Page compression", "Execute the MERGE statement", "Remove the clustered index", and "Create a data file" are highlighted with a red border, indicating they are the correct sequence of actions to perform in order.

QUESTION 4

You have Master Data Services (MDS) in a SQL Server deployment. You need to deploy a model that contains data to the MDS repository. What should you use?

- A. the Model Deployment Wizard
- B. the MDSModelDeploy tool
- C. the Master Data Services Configuration Manager tool
- D. the Integration Services Deployment Wizard

Correct Answer: B

Explanation:

<https://msdn.microsoft.com/en-us/library/ff486956.aspx>

QUESTION 5

You are using SQL Server Data Tools to develop a SQL Server Integration Services (SSIS) project.

The first package that you create in this project contains a package connection that accesses a Microsoft Excel file. Additional packages in the project must also access this file.

You need to define and reuse the Excel file connection in all project packages.

What should you do?

- A. Copy the package Connection Manager and paste it in the second package.
- B. Set the RetainSameConnection property of the package Connection Manager to True.
- C. Convert the package Connection Manager in the first package to a project Connection Manager.
- D. Convert the project to the Package Deployment model.

Correct Answer: C

QUESTION 6

You are developing a SQL Server Integration Services (SSIS) package that loads data from a source transactional tables in a data warehouse.

You discover that one of the data flows loads incorrect data to one of the destination tables.

You need to identify the cause of the incorrect data flow. Which tool should you use?

- A. SSIS logging
- B. a breakpoint

- C. a data viewer
- D. the locals window

Correct Answer: C

QUESTION 7

You plan to deploy a package to a server that has SQL Server installed. The server contains a SQL Server Integration Services Catalog.

You need to ensure that the package credentials are encrypted.

Which protection level should you set for the package?

- A. EncryptSensitivewithUserKey
- B. EncryptAllWithUserKey
- C. ServerStorage
- D. EncryptSensitivewithPassword

Correct Answer: D

QUESTION 8

You are designing a SQL Server Integration Services (SSIS) data flow to load sales transactions from a source system into a data warehouse hosted on Windows Azure SQL Database. One of the columns in the data source is named ProductCode.

Some of the data to be loaded will reference products that need special processing logic in the data flow.

You need to enable separate processing streams for a subset of rows based on the source product code.

Which Data Flow transformation should you use?

- A. Multicast
- B. Conditional Split
- C. Destination Assistant
- D. Script Task

Correct Answer: B

Explanation:

We use Conditional Split to split the source data into separate processing streams.

A Script Component (Script Component is the answer to another version of this question) could be used but this is not the same as a Script Task.

QUESTION 9

You are completing the installation of the Data Quality Server component of SQL Server Data Quality Services (DQS).

You need to complete the post-installation configuration.

What should you do?

- A. Run the Data Quality Server Installer.
- B. Install the data providers that are used for data refresh.
- C. Run the dbimpexp.exe command.
- D. Install the Analysis Services OLE DB Provider.

Correct Answer: A

QUESTION 10

You have a database named Sales and a data warehouse named DataW.

From Sales, you plan to bulk insert to a table in DataW that is more than 2TB. The process will insert a minimum of 11 million rows simultaneously.

You need to identify which data storage strategy must be used to minimize load times.

Which data storage strategy should you identify?

- A. a Clustered Columnstore
- B. a file table
- C. a durable In-Memory OLTP
- D. a page-compressed heap

Correct Answer: A

QUESTION 11

Your company has several line-of-business applications. The applications use a server that has SQL Server installed and contains several databases.

You need that the business applications can access curated and validated data from the databases.

Which features should you deploy to the server?

- A. Data Quality Services (DQS)
- B. SQL Server Analysis Services (SSAS) dimension hierarchies.
- C. SQL Server Integration Services (SSIS)
- D. Master Data Services (MDS) subscriptions.

Correct Answer: C

QUESTION 12

You are designing a data warehouse with two fact tables. The first table contains sales per month and the second table contains orders per day.

Referential integrity must be enforced declaratively.

You need to design a solution that can join a single time dimension to both fact tables.

What should you do?

- A. Create a time mapping table.
- B. Partition the fact tables by day.
- C. Create a time dimension that can join to both fact tables at their respective granularity.
- D. Join the two fact tables.

Correct Answer: C

QUESTION 13

You have a SQL server integration Services (SSIS) package named package1.

You discover that the Data Flow task for Package1 runs slower than expected.

You need to reduce the amount of time required to run the Data Flow task.

Which two actions should you perform? (Select Two)

- A. Configure the package to run by using the CallerInfo option.
- B. Modify the DefaultBufferSize package setting.
- C. Modify the ForceExecutionResult Package setting.
- D. Modify the DefaultBufferMaxRows package setting.
- E. Configure the package to run by using MaxConcurrent option.

Correct Answer: DE

QUESTION 14

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

The package sources data from an HTML web page that lists product stock levels.

You need to implement a data flow task that reads the product stock levels from the HTML web page.

Which data flow sources should you use? (Select Two)

- A. Raw File source
- B. XML source
- C. Custom source component
- D. Flat File source
- E. Script component

Correct Answer: CE

QUESTION 15

You have a secured database that contains all of the customer data for your company. You plan to use a project deployment model.

You need to create a SQL Server Integration Services (SSIS) package that connects to the database by using a dedicated username and password. The solution must ensure that the password used for the connection string is encrypted.

Which two actions should you perform? (Select Two)

- A. Select the Sensitive check box for the catalog environment.
- B. Set the sensitive property of the package parameter to True.
- C. Modify the set_execution_parameter_value stored procedure.
- D. Set the package protection level to EncrypSensitiveWithPassword.

Correct Answer: CD

QUESTION 16

You are designing a data warehouse with two fact tables. The first table contains sales per month and the second table contains orders per day.

Referential integrity must be enforced declaratively.

You need to design a solution that can join a single time dimension to both fact tables.

What should you do?

- A. Join the two fact tables.
- B. Merge the fact tables.
- C. Change the level of granularity in both fact tables to be the same.
- D. Partition the fact tables by day.

Correct Answer: D