VCE & PDF
GeekCert.com

70-767^{Q&As}

Implementing a Data Warehouse using SQL

Pass Microsoft 70-767 Exam with 100% Guarantee

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

https://www.geekcert.com/70-767.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 are the Master Data Services (MDS) administrator at your company.

An existing user needs to integrate data from other systems into MDS.

You need to give the user permissions to integrate data.

Which user management menu item should you select? (To answer, configure the appropriate option or options in the dialog box in the answer area.)

Hot Area:

Correct Answer:

QUESTION 2

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

You need to design a package to change a variable value during package execution by using the least amount of development effort.

What should you use?

A. Expression task

B. Data Cleansing transformation

C. Fuzzy Lookup transformation

D. Term Lookup transformation

E. Data Profiling task

Correct Answer: A

QUESTION 3

You maintain a SQL Server Integration Services (SSIS) package. The package was developed by using SQL Server 2008 Business Intelligence Development Studio (BIDS).

The package includes custom scripts that must be upgraded. You need to upgrade the package to SQL Server 2016.

Which tool should you use?

A. SSIS Upgrade Wizard in SQL Server 2008 BIDS

B. SSIS Upgrade Wizard in SQL Server Data Tools

C. SQL Server DTExecUI utility (dtexecui.exe)

D. SQL Server dtexec utility (dtexec.exe)

Correct Answer: B

QUESTION 4

Note: This question is part of a series of questions that use the same or similar answer choices. An answer choice may be correct for more than one question in the series. Each question is independent of the other questions in this series. Information and details provided in a question apply only to that question.

You are a database administrator for an e-commerce company that runs an online store. The company has the databases described in the following table.

Database	Description
DB1	This database supports the online store.
DB2	This is the data warehouse for the company. **DB2** contains a table named OnlineOrder that is partitioned in hourly increments. The LOCK_ESCALATION option is set to **AUTO**. The data flow contains 24 OLE DB destinations, one for each partition.
DB3	This database runs Master Data Services (MDS).

Product prices are updated and are stored in a table named Products on DB1. The Products table is deleted and refreshed each night from MDS by using a Microsoft SQL Server Integration Services (SSIS) package. None of the data sources are sorted.

You need to update the SSIS package to add current prices to the Products table.

What should you use?

A. Lookup transformation

B. Merge transformation

C. Merge Join transformation

D. MERGE statement

E. Union All transformation

F. Balanced Data Distributor transformation

G. Sequential container

H. Foreach Loop container

Correct Answer: D

In the current release of SQL Server Integration Services, the SQL statement in an Execute SQL task can contain a MERGE statement. This MERGE statement enables you to accomplish multiple INSERT, UPDATE, and DELETE operations in a single statement.

References: https://docs.microsoft.com/en-us/sql/integration-services/control-flow/merge- in-integration-services-packages

---

QUESTION 5

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 6

You manage Master Data Services (MDS). You plan to create entities and attributes and load them with the data. You also plan to match data before loading it into Data Quality Services (DQS).

You need to recommend a solution to perform the actions.

What should you recommend?

A. MDS Add-in for Microsoft Excel

B. MDS Configuration Manager

C. Data Quality Matching

D. MDS repository

Correct Answer: A

In the Master Data Services Add-in for Excel, matching functionality is provided by Data Quality Services (DQS). This

functionality must be enabled to be used.

1.

 To enable Data Quality Services integration

2.

 Open Master Data Services Configuration Manager.

3.

 In the left pane, click Web Configuration.

4.

 On the Web Configuration page, select the website and web application.

5.

 In the Enable DQS Integration section, click Enable integration with Data Quality Services.

6.

 On the confirmation dialog box, click OK.

References: https://docs.microsoft.com/en-us/sql/master-data-services/install-windows/enable-data-quality-services-integration-with-master-data-services

---

QUESTION 7

You plan to deploy several Microsoft SQL Server Integration Services (SSIS) packages to a highly available SQL Server instance. The instance is configured to use an AlwaysOn availability group that has two replicas.

You need to identify which deployment method must be used to ensure that the packages are always accessible from all the nodes in the availability group.

Which deployment method should you use for the packages?

A. Deploy to the msdb database on the secondary replica.

B. Deploy to the msdb database on the primary replica.

C. Deploy to a file on the hard drive of the primary replica.

D. Deploy to a shared folder on a file server.

Correct Answer: A

Before you can configure SSIS to enable support of AlwaysOn on the new added secondary Replicas, you must connect to all new added secondary replicas.
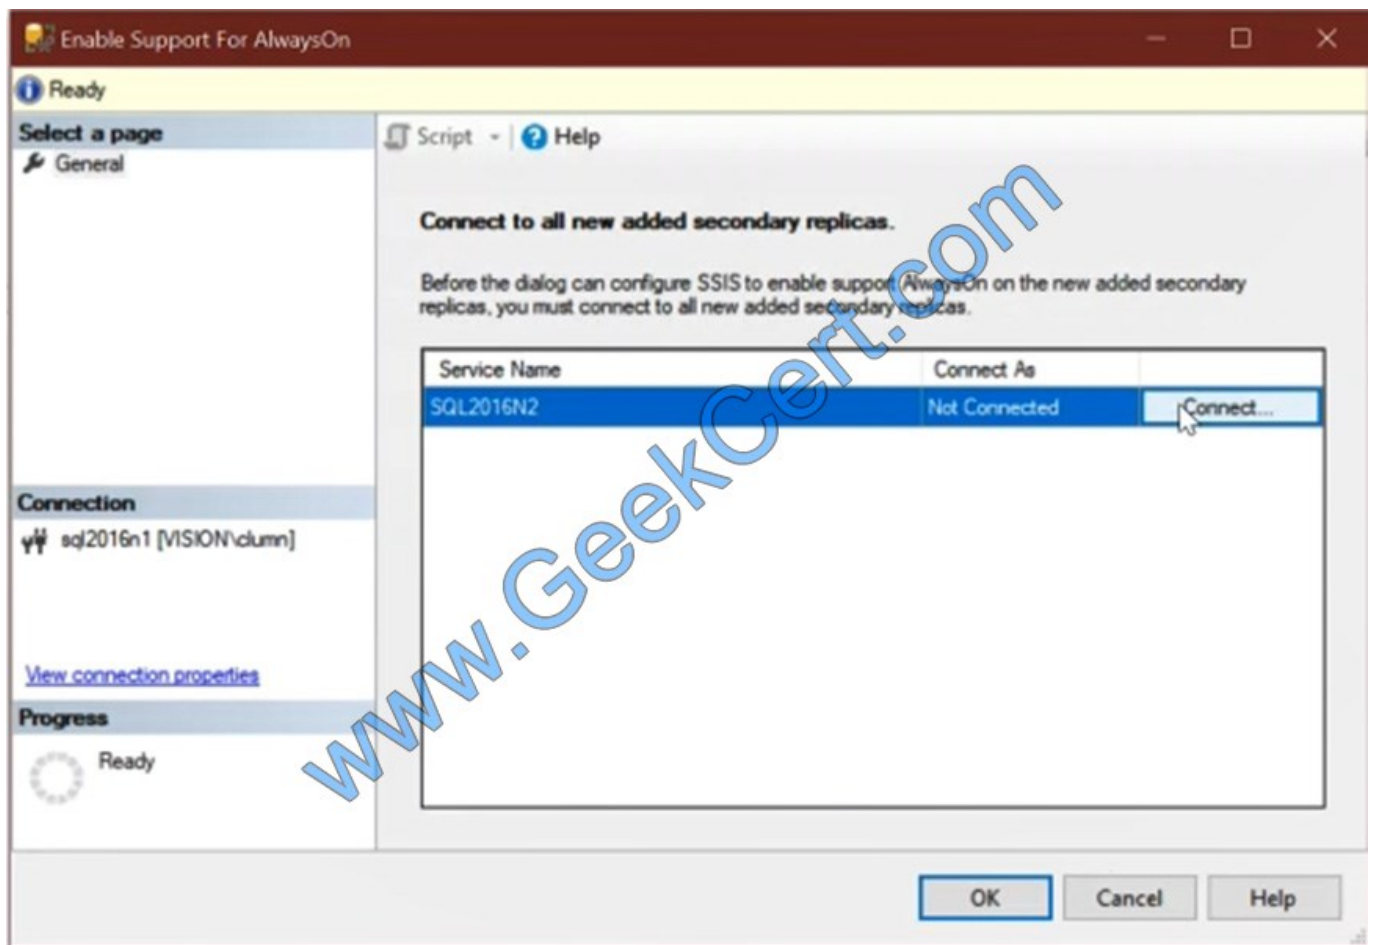
Note: To use SSIS with AlwaysOn, you'll need to add the SSIS Catalog (SSISDB) into an Availability Group. You'll need to do the following steps:

Make sure you meet the prerequisites for using AlwaysOn

Connect to every node and create the SSISDB catalog. We need to create the catalog even on secondary nodes to create the other server-level objects (cleanup jobs, keys, accounts etc) that are used by SSIS.

Delete the SSISDB databases on secondary nodes.

Create an availability group, specifying SSISDB as the user database

Specify secondary replicas.

References: https://chrislumnah.com/2017/05/09/enabling-alwayson-for-ssisdb/

QUESTION 8

You are developing a Microsoft SQL Server Integration Services (SSIS) package to incrementally load new and changed records from a data source.

The SSIS package must load new records into Table1 and updated records into Table1_Updates. After loading records, the package must call a Transact-SQL statement to process updated rows according to existing business logic.
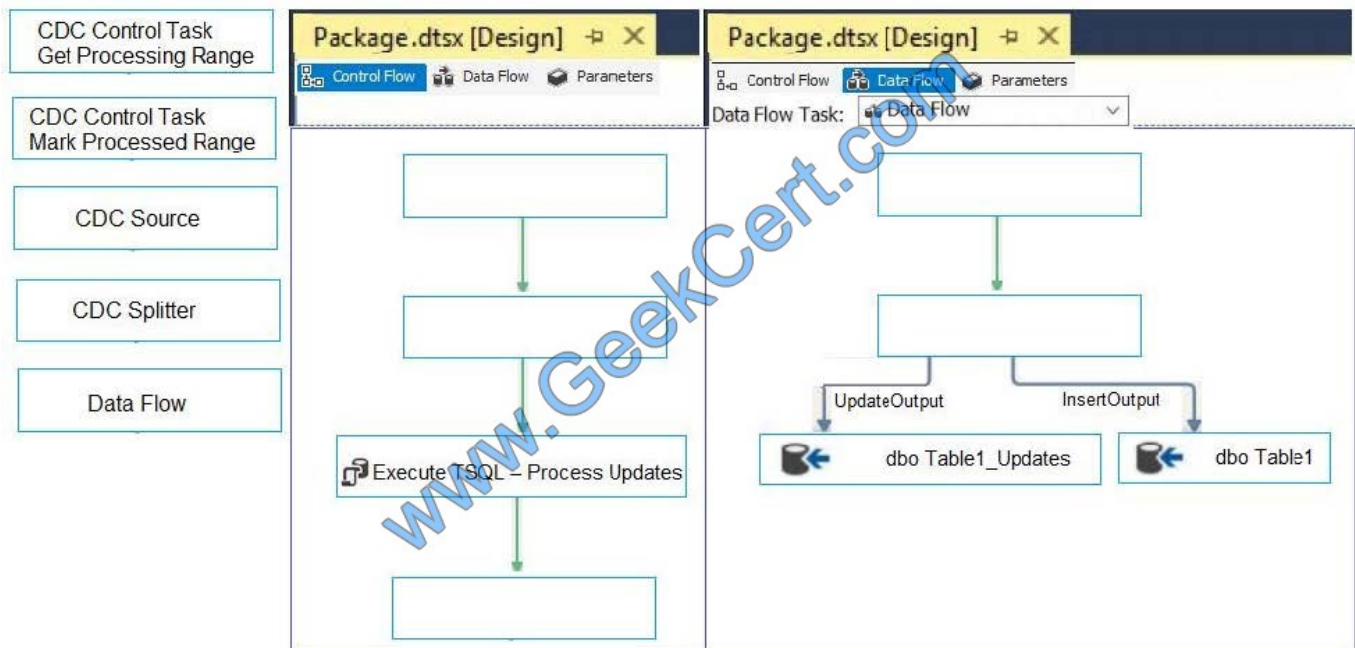
You need to complete the design of the SSIS package.

Which tasks should you use? To answer, drag the appropriate SSIS objects to the correct targets. Each SSIS object 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.

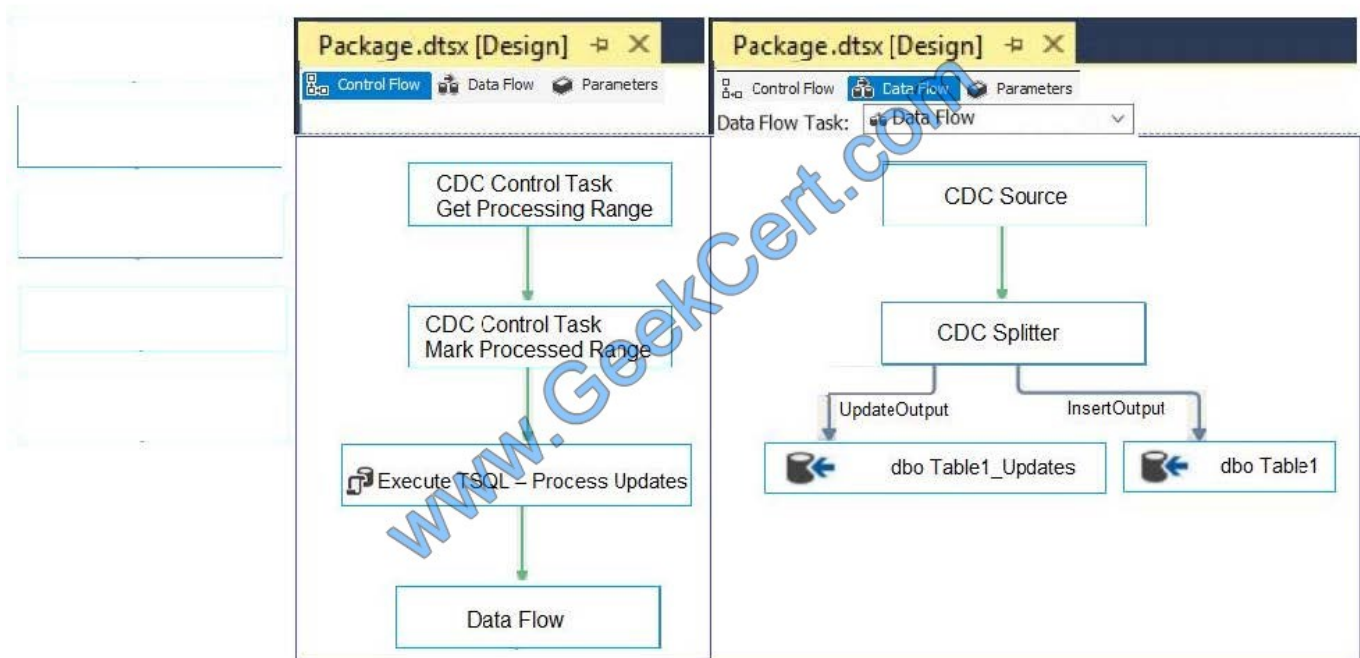NOTE: Each correct selection is worth one point.

Select and Place:

Answer Area

Correct Answer:

Answer Area

Step 1: CDC Control Task Get Processing Range

Step 2: Mark Processed Range

Step 3: Data Flow

The Data Flow task encapsulates the data flow engine that moves data between sources and destinations, and lets the user transform, clean, and modify data as it is moved. Addition of a Data Flow task to a package control flow makes it

possible for the package to extract, transform, and load data.

Step 4: CDC Source

The CDC source reads a range of change data from SQL Server 2017 change tables and delivers the changes downstream to other SSIS component.

Step 5: CDC Splitter

The CDC splitter splits a single flow of change rows from a CDC source data flow into different data flows for Insert, Update and Delete operations.

https://docs.microsoft.com/en-us/sql/integration-services/control-flow/cdc-control-task

https://docs.microsoft.com/en-us/sql/integration-services/control-flow/data-flow-task

https://docs.microsoft.com/en-us/sql/integration-services/data-flow/cdc-splitter?view=sql-server-2017

QUESTION 9

Note: This question is part of a series of questions that use the same or similar answer choices. An answer choice may be correct for more than one question in the series. Each question is independent of the other questions in this series.

Information and details provided in a question apply only to that question.

You have a database named DB1 that has change data capture enabled.

A Microsoft SQL Server Integration Services (SSIS) job runs once weekly. The job loads changes from DB1 to a data warehouse by querying the change data capture tables.

Users report that an application that uses DB1 is suddenly unresponsive.

You discover that the Integration Services job causes severe blocking issues in the application. You need to ensure that the users can run the application as quickly as possible. Your SQL Server login is a member of only the ssis.admin

database role.

Which stored procedure should you execute?

A. catalog.deploy_project

B. catalog.restore_project

C. catalog.stop.operation

D. sys.sp.cdc.addjob

E. sys.sp.cdc.changejob

F. sys.sp_cdc_disable_db

G. sys.sp_cdc_enable_db

H. sys.sp_cdc.stopJob

Correct Answer: E

sys.sp_cdc_change_job modifies the configuration of a change data capture cleanup or capture job in the current database.

References: https://docs.microsoft.com/en-us/sql/relational-databases/system-stored-procedures/sys-sp-cdc-change-job-transact-sql? view=sql-server-2017

QUESTION 10

You are developing a SQL Server Integration Services (SSIS) project by using the Project Deployment model.

The project contains many packages. It is deployed on a server named SQLTest1. The project will be deployed to several servers that run SQL Server 2016.

The project accepts one required parameter. The data type of the parameter is a string.

A SQL Agent job is created that will call the Loading.dtsx package in the project. A job step is created for the SSIS package.

The job must pass the value of an SSIS Environment Variable to the project parameter.

The value of the Environment Variable must be configured differently on each server that runs SQL Server. The value of

the Environment Variable must provide the server name to the project parameter.

You need to configure SSIS on the SQLTest1 server to pass the Environment Variable to the package.

Which four actions should you perform in sequence by using SQL Server Management Studio?

(To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.)

Select and Place:

In the SSIS catalog, create an Environment named **EnvTest**.
In the **EnvTest** Environment, create an Environment Variable. Name the Environment Variable **EnvServer**. Set the Environment Variable value to **SQLTest1**.
In the SQL Agent job, choose the **EnvTest** Environment in the Configuration tab of the SQL Server Integration Services job step.
In the SSIS catalog, create a folder named **EnvTest**.
In the project configuration, create an Environment reference to **EnvTest**. Assign the value of the **EnvServer** Environment Variable to the project parameter.
In the project configuration, change the value of the project parameter to **EnvTest**.

Correct Answer:

	In the SSIS catalog, create an Environment named **EnvTest**.
	In the **EnvTest** Environment, create an Environment Variable. Name the Environment Variable **EnvServer**. Set the Environment Variable value to **SQLTest1**.
	In the project configuration, create an Environment reference to **EnvTest**. Assign the value of the **EnvServer** Environment Variable to the project parameter.
In the SSIS catalog, create a folder named **EnvTest**.	In the SQL Agent job, choose the **EnvTest** Environment in the Configuration tab of the SQL Server Integration Services job step.
In the project configuration, change the value of the project parameter to **EnvTest**.	

QUESTION 11

Note: This question is part of a series of questions that use the same or similar answer choices. An answer choice may be correct for more than one question in the series. Each question is independent of the other questions in this series.

Information and details provided in a question apply only to that question.

You are implementing a Microsoft SQL Server data warehouse with a multi-dimensional data model. You have a fact table that includes sales data for all products. The model includes a dimension named Geography that stores all

geographies. You create a dimension that has a foreign key and provides the ability to analyze sales by the following sales channels: Internet or retail store.

You need to update the data model to allow business users to analyze Internet sales by geography without changing the overall structure of the data model. What should you do?

A. star schema

B. snowflake schema

C. conformed dimension

D. slowly changing dimension (SCD)

E. fact table

F. semi-additive measure

G. non-additive measure

H. dimension table reference relationship
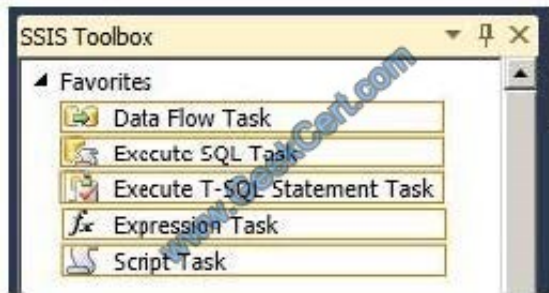
Correct Answer: D

QUESTION 12

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

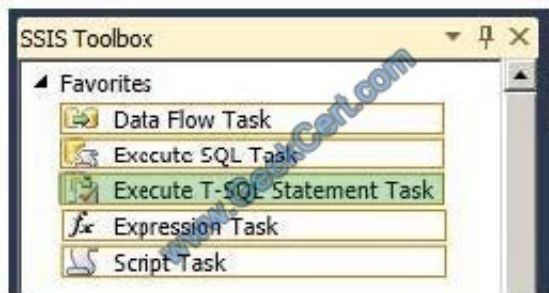The package must run a parameterized query against a Windows Azure SQL Database database.

You need to use the least amount of development effort to meet the package requirement.

Which task should you use? (To answer, select the appropriate task in the answer area.)

Hot Area:

Correct Answer:

QUESTION 13

You are developing a SQL Server Integration Services (SSIS) package that imports data into a data warehouse. You are developing the part of the SSIS package that populates the ProjectDates dimension table. The business key of the ProjectDates table is the ProjectName column. The business user has given you the dimensional attribute behavior for each of the four columns in the ProjectDates table:

ExpectedStartDate - New values should be tracked over time.

ActualStartDate - New values should not be accepted.

ExpectedEndDate - New values should replace existing values.

ActualEndDate - New values should be tracked over time.

You use the SSIS Slowly Changing Dimension Transformation.

You must configure the Change Type value for each source column.

Which settings should you select? (To answer, select the appropriate setting or settings in the answer area. Each Change Type may be used once, more than once, or not at all.)

Select and Place:

Dimension Columns	Change Type
ExpectedStartDate	⌄
ActualStartDate	⌄
ExpectedEndDate	⌄
ActualEndDate	⌄

Fixed Attribute ⌄

Changing Attribute ⌄

Historical Attribute ⌄

Correct Answer:

Dimension Columns	Change Type
ExpectedStartDate	Historical Attribute ⌄
ActualStartDate	Fixed Attribute ⌄
ExpectedEndDate	Changing Attribute ⌄
ActualEndDate	Historical Attribute ⌄

Fixed Attribute ⌄

Changing Attribute ⌄

Historical Attribute ⌄

QUESTION 14

You are developing a Microsoft SQL Server Integration Services (SSIS) package. You create a data flow that has the following characteristics:

The package moves data from the table [source].Tabid to DW.Tablel.

All rows from [source].Table1 must be captured in DW.Tablel for error.Tablel.

The table error.Tablel must accept rows that fail upon insertion into DW.Tablel due to violation of nullability or data type errors such as an invalid date, or invalid characters in a number.

The behavior for the Error Output on the";OLE DB Destinatio"; object is Redirect. ?The data types for all columns in [sourceJ.Tablel are VARCHAR. Null values are allowed.

The Data access mode for both OLE DB destinations is set to Table or view - fast load.

The table definitions are as follows:

```
CREATE TABLE [source].Table1
(
    ID INT NULL,
    CreateDate VARCHAR(100) NULL,
    Date1 DATETIME2(7) NULL,
    Number1 VARCHAR(100) NULL
)

CREATE TABLE error.Table
(
    ID INT NULL,
    CreateDate VARCHAR(100) NULL,
    Date1 DATETIME2(7) NULL,
    Number1 VARCHAR(100) NULL,
    ErrorDescription VARCHAR(255) NULL
)
```

Use the drop-down menus to select the answer choice that answers each question.

Hot Area:

The ErrorDescription column is not yet populated in error.Table1.You must capture the error description for any rows redirected to the "Error OLE DB Destination.".What should you do next?

In "OLE DB Destination Error", map the ErrorCode field to ErrorDescription.
Create an INSERT trigger on [Error] [Table1] to populate the ErrorDescription from ErrorCode.
Add a Derived Column transformation before "OLE DB Destination". Use ErrorCode to populate ErrorDescription.
Add a Script Component transformation before "OLE DB Destination Error". Capure the ErrorDescription with VB or C# code.

You execute the package. You note that all rows are redirected to OLE DB Destination Error, including both rows with bad data and rows with valid data. What is the next step?

Uncheck the Check Constraints option in OLE DB Destination.
Change the Date access mode for OLE DB Destination to Table View.
Uncheck the options Table Lock and Check Constraints for OLE DB Destination Error object to False.
Add a Conditional Split transformation before OLE DB Destination. Create outputs baseed on ErrorCode.

Correct Answer:

The ErrorDescription column is not yet populated in error.Table1.You must capture the error description for any rows redirected to the "Error OLE DB Destination.".What should you do next?

In "OLE DB Destination Error", map the ErrorCode field to ErrorDescription.
Create an INSERT trigger on [Error] [Table1] to populate the ErrorDescription from ErrorCode.
Add a Derived Column transformation before "OLE DB Destination". Use ErrorCode to populate ErrorDescription.
Add a Script Component transformation before "OLE DB Destination Error". Capure the ErrorDescription with VB or C# code.

You execute the package. You note that all rows are redirected to OLE DB Destination Error, including both rows with bad data and rows with valid data. What is the next step?

Uncheck the Check Constraints option in OLE DB Destination.
Change the Date access mode for OLE DB Destination to Table View.
Uncheck the options Table Lock and Check Constraints for OLE DB Destination Error object to False.
Add a Conditional Split transformation before OLE DB Destination. Create outputs baseed on ErrorCode.

QUESTION 15

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You create a set of Microsoft SQL Server Integration Services (SSIS) packages to support an ETL process. You need to deploy the SSIS packages and meet the following requirements:

1.

Maximize ease of use for backup processes.

2.

Manage security of packages from a single place.

3.

Maintain versioning pf packages.

4.

Implement parameters to assign values to packages.

5.

Use stored procedures to run SSIS packages.

6.

Use SQL Server Agent to run packages.

Solution: You use the Project Deployment model with an SSIS catalog to store and manage packages.

Does this meet the goal?

A. Yes

B. No

Correct Answer: A

Built-in stored procedures only work for with an SSIS catalog to store.

References: https://docs.microsoft.com/en-us/sql/integration-services/packages/run-integration-services-ssis-packages

[70-767 PDF Dumps](https://www.geekcert.com/70-767.html)   [70-767 Practice Test](https://www.geekcert.com/70-767.html)   [70-767 Exam Questions](https://www.geekcert.com/70-767.html)