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.pass4itsure.com/70-767.html

100% Passing Guarantee
100% Money Back Assurance

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

QUESTION 1

You are creating a Data Quality Services (DQS) solution. You must provide statistics on the accuracy of the data.

You need to use DQS profiling to obtain the required statistics.

Which DQS activity should you use?

A. Cleansing

B. Knowledge Discovery

C. SQL Profiler

D. Matching Rule Definition

Correct Answer: A

QUESTION 2

You install a SQL Server 2016 database engine instance on a production server. A month later, you install SQL Server 2016 Integration Services (SSIS).

You must develop an SSIS project and deploy it to the server by using the Project Deployment model.

Operations Log records that are outside the configured retention period must be cleaned automatically.

You need to create the SSIS catalog on the production server.

What should you do? (Each correct answer presents part of the solution. Choose all that apply.)

A. Enable XP Command Shell.

B. Enable CLR Integration.

C. Enable OLE Automation.

D. Start the SQL Server Browser service.

E. Enable Cross Database Ownership Chaining

F. Start the SQL Server Agent service

G. Enable Ad Hoc Remote Queries.

Correct Answer: BF

Explanation: Ref: http://msdn.microsoft.com/en-us/library/gg471509(v=sql.110).aspx

QUESTION 3

You are designing a SQL Server Integration Services (SSIS) package to execute 12 Transact-SQL (T-SQL) statements on a SQL Azure database. The T-SQL statements may be executed in any order. The T-SQL statements have

unpredictable execution times.

You have the following requirements:

The package must maximize parallel processing of the T-SQL statements.

After all the T-SQL statements have completed, a Send Mail task must notify administrators.

You need to design the SSIS package. Which three actions should you perform in sequence? (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:

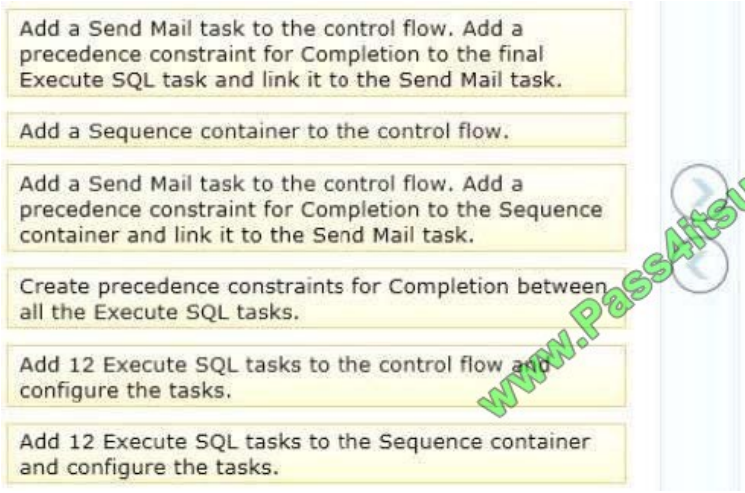

Correct Answer:

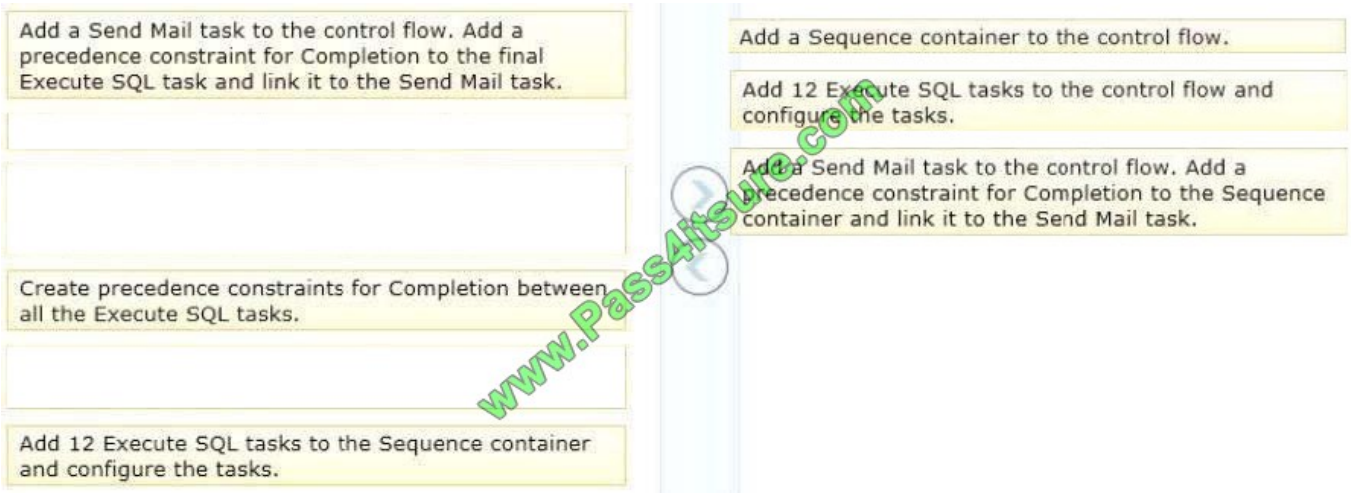

QUESTION 4

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).

Each day, data from the table OnlineOrder in DB2 must be exported by partition. The tables must not be locked during the process.

You need to write a Microsoft SQL Server Integration Services (SSIS) package that performs the data export. 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: E

The Union All transformation combines multiple inputs into one output. For example, the outputs from five different Flat File sources can be inputs to the Union All transformation and combined into one output. References: https://docs.microsoft.com/en-us/sql/integration-services/data- flow/transformations/union-all-transformation

QUESTION 5

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 6

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 7

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

Correct Answer: D

Explanation: 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 Ref: http://msdn.microsoft.com/en-us/library/hh758656.aspx

QUESTION 8

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?

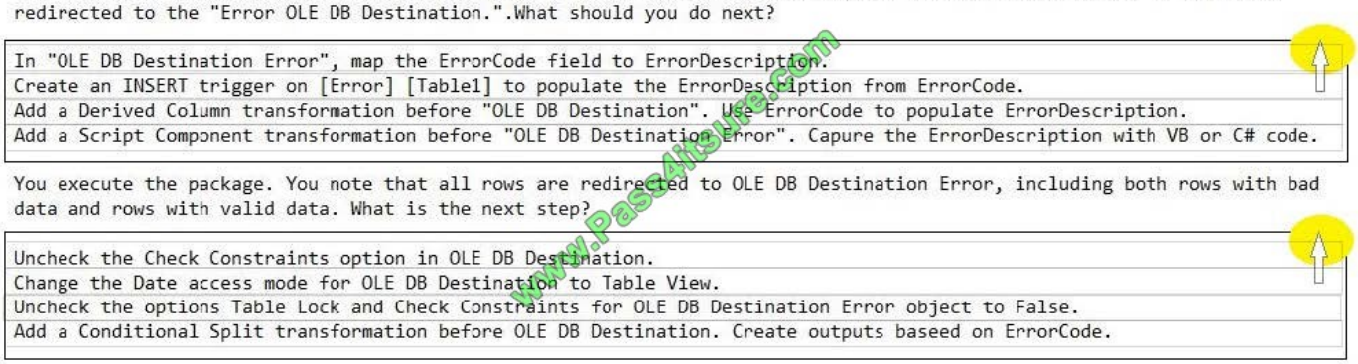

You execute the package. You note that all rows are redirected to OLE DB Destination Error, including both rows with bad data and rows with valid data. What is the next step?

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?

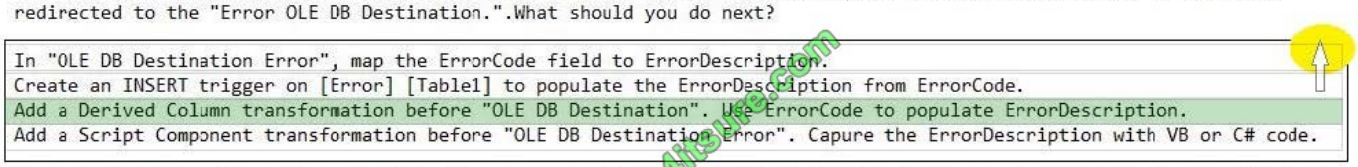

You execute the package. You note that all rows are redirected to OLE DB Destination Error, including both rows with bad data and rows with valid data. What is the next step?

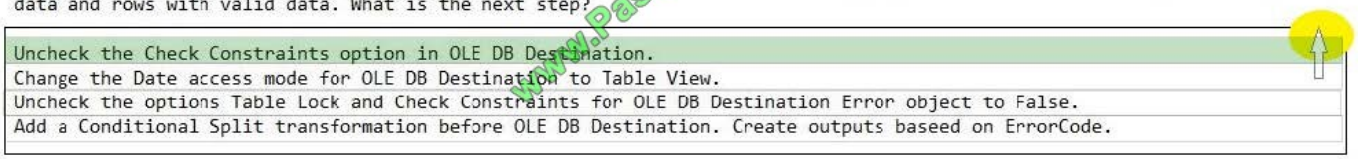

QUESTION 9

A SQL Server Integration Services (SSIS) package imports daily transactions from several files into a SQL Server table named Transaction. Each file corresponds to a different store and is imported in parallel with the other files. The data flow tasks use OLE DB destinations in fast load data access mode.

The number of daily transactions per store can be very large and is growing. The Transaction table does not have any indexes.

You need to minimize the package execution time.

What should you do?

A. Partition the table by day and store.

B. Create a clustered index on the Transaction table.

C. Run the package in Performance mode.

D. Increase the value of the Row per Batch property.

Correct Answer: D

Explanation: * Data Access Mode ?This setting provides the \\'fast load\\' option which internally uses a BULK INSERT

statement for uploading data into the destination table instead of a simple INSERT statement (for each single row) as in the case for other options.

* BULK INSERT parameters include:

ROWS_PER_BATCH =rows_per_batch

Indicates the approximate number of rows of data in the data file. By default, all the data in the data file is sent to the server as a single transaction, and the number of rows in the batch is unknown to the query optimizer. If you specify

ROWS_PER_BATCH (with a value > 0) the server uses this value to optimize the bulk- import operation. The value specified for ROWS_PER_BATCH should approximately the same as the actual number of rows.

QUESTION 10

Note: This question is part of a series of questions that use the same scenario. For your convenience, the scenario is repeated in each question. Each question presents a different goal and answer choices, but the text of the scenario is

exactly the same in each question in this series.

You have a Microsoft SQL Server data warehouse instance that supports several client applications.

The data warehouse includes the following tables: Dimension.SalesTerritory, Dimension.Customer, Dimension.Date, Fact.Ticket, and Fact.Order. The Dimension.SalesTerritory and Dimension.Customer tables are frequently updated. The

Fact.Order table is optimized for weekly reporting, but the company wants to change it daily. The Fact.Order table is loaded by using an ETL process. Indexes have been added to the table over time, but the presence of these indexes slows

data loading.

All data in the data warehouse is stored on a shared SAN. All tables are in a database named DB1. You have a second database named DB2 that contains copies of production data for a development environment. The data warehouse has

grown and the cost of storage has increased. Data older than one year is accessed infrequently and is considered historical.

You have the following requirements:

You are not permitted to make changes to the client applications.

You need to optimize the storage for the data warehouse.

What change should you make?

A. Partition the Fact.Order table, and move historical data to new filegroups on lower-cost storage.

B. Create new tables on lower-cost storage, move the historical data to the new tables, and then shrink the database.

C. Remove the historical data from the database to leave available space for new data.

D. Move historical data to new tables on lower-cost storage.

E. Implement row compression for the Fact.Order table.

F. Move the index for the Fact.Order table to lower-cost storage.

Correct Answer: A

Create the load staging table in the same filegroup as the partition you are loading. Create the unload staging table in the same filegroup as the partition you are deleteing.

From scenario: Data older than one year is accessed infrequently and is considered historical.

References: https://blogs.msdn.microsoft.com/sqlcat/2013/09/16/top-10-best-practices-for- building-a-large-scale-relational-data-warehouse/

QUESTION 11

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 view on the sales table.

B. Partition the fact tables by day.

C. Create a surrogate key for the time dimension.

D. Change the level of granularity in both fact tables to be the same.

Correct Answer: D

QUESTION 12

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 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 Change Type values should you select? (To answer, drag the appropriate value from the list of values to the correct location or locations in the answer area.)

Select and Place:

Correct Answer:

QUESTION 14

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while

others might not have a correct solution.

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

You have the following line-of-business solutions:

If a change is made to the ReferenceNr column in any of the sources, set the value of IsDisabled to True and create a new row in the Products table. If a row is deleted in any of the sources, set the value of IsDisabled to True in the data

warehouse.

One or more Microsoft SQL Server instances support each solution. Each solution has its own product catalog. You have an additional server that hosts SQL Server Integration Services (SSIS) and a data warehouse. You populate the data

warehouse with data from each of the line-of-business solutions. The data warehouse does not store primary key values from the individual source tables.

The database for each solution has a table named Products that stored product information. The Products table in each database uses a separate and unique key for product records. Each table shares a column named ReferenceNr between

the databases. This column is used to create queries that involve more than once solution.

You need to load data from the individual solutions into the data warehouse nightly. The following requirements must be met:

Enable the Change Tracking for the Product table in the source databases. Query the cdc.fn_cdc_get_all_changes_capture_dbo_products function from the sources for updated rows.

Set the IsDisabled column to True for rows with the old ReferenceNr value.

Create a new row in the data warehouse Products table with the new ReferenceNr value.

Solution: Perform the following actions:

Does the solution meet the goal?

A. Yes

B. No

Correct Answer: B

We must also handle the deleted rows, not just the updated rows.

References: https://solutioncenter.apexsql.com/enable-use-sql-server-change-data- capture/

QUESTION 15

You are the administrator for a Data Quality Server. You are adding a user who must have permission to:

Edit and execute a project

View the activity monitoring data

This user must not be able to:

Perform any kind of knowledge management

Create or change a knowledge base

Terminate an activity or perform administrative duties

You need to develop a Transact-SQL (T-SQL) script to meet these requirements.

What should you do? (To answer, drag the appropriate code segment or segments to the correct location or locations in the answer area.)

Select and Place:

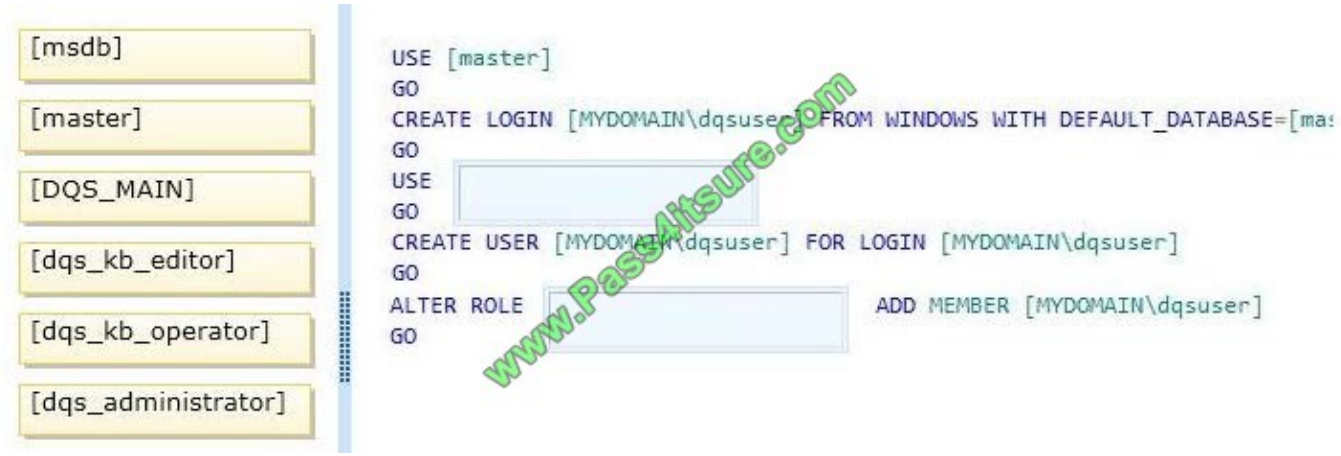

Correct Answer:

[Latest 70-767 Dumps](#) [70-767 Study Guide](#) [70-767 Braindumps](#)