

[2018-March-NewFree Microsoft 287Q 70-767 Exam Dumps PDF Braindump2go Offers[248-258

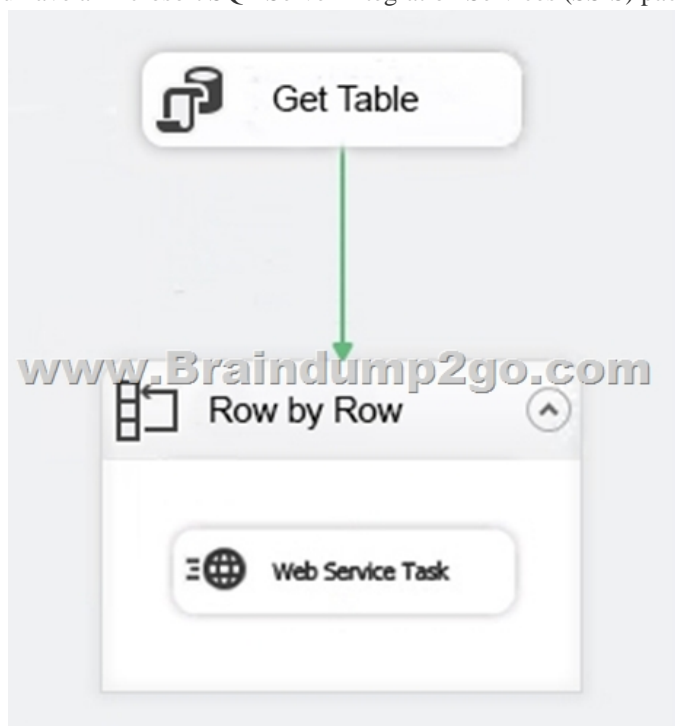
2018 March New Microsoft 70-697 Exam Dumps with PDF and VCE Free Updated Today! Following are some new 70-697 Real Exam Questions:1.[2018 Latest 70-697 Exam Dumps (PDF & VCE) 287Q&As Download:

<https://www.braindump2go.com/70-767.html>2.[2018 Latest 70-697 Exam Questions & Answers Download:

<https://drive.google.com/drive/folders/0B75b5xYLjSSNN1RSdIN6Z0VwRjg?usp=sharing>
QUESTION 248Note: 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 a Microsoft Azure SQL Data Warehouse instance that must be available six months a day for reporting.You need to pause the compute resources when the instance is not being used.Solution: You use the Azure portal.Does the solution meet the goal?A. YesB. NoAnswer: A**QUESTION 249**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.Answer: AExplanation: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 deleting.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 250You have a Microsoft SQL Server Integration Services (SSIS) package that includes the control flow shown in the following diagram.



You need to choose the enumerator for the Foreach Loop container. Which enumerator should you use? A. Foreach SMO Enumerator B. Foreach Azure Blob Enumerator C. Foreach NodeList Enumerator D. Foreach ADO Enumerator
Answer: D
Explanation: Use the Foreach ADO enumerator to enumerate rows in tables. For example, you can get the rows in an ADO recordset.
QUESTION 251
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: ERP system Online WebStore Partner extranet 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: 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.
Solution: Perform the following actions: Enable the Change Tracking feature for the Products table in the three source databases. Query the CHANGETABLE function from the sources for the deleted rows. Set the IsDisabled column to True on the data warehouse Products table for the listed rows.
Does the solution meet the goal? A. Yes B. No
Answer: B
Explanation: We must check for updated rows, not just deleted rows.
References:

<https://www.timmitchell.net/post/2016/01/18/getting-started-with-change-tracking-in-sql-server/>
QUESTION 252
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 a data warehouse that stores information about products, sales, and orders for a manufacturing company. The instance contains a database that has two tables named SalesOrderHeader and SalesOrderDetail. SalesOrderHeader has 500,000 rows and SalesOrderDetail has 3,000,000 rows. Users report performance degradation when they run the following stored procedure:

```
CREATE PROCEDURE Sales.GetRecentSales (@date datetime)
AS BEGIN
    IF @date is NULL
        SET @date = GETDATE()
    SELECT h.SalesOrderID, d.OrderDate, d.ProductID
    FROM SalesOrderHeader h
    JOIN SalesOrderDetail d ON h.SalesOrderID = d.SalesOrderID
    WHERE h.SalesOrderID > 1000000
    AND h.OrderDate > @date
END
```

You need to optimize performance. Solution: You run the following Transact-SQL statement:

```
CREATE STATISTICS Stat1
ON SalesOrderHeader (SalesOrderID)
WITH SAMPLE 5 PERCENT
```

Does the solution meet the goal? A. Yes B. No
Answer: A
Explanation: You can specify the sample size as a percent. A 5% statistics sample size would be helpful.
References:

<https://docs.microsoft.com/en-us/azure/sql-data-warehouse/sql-data-warehouse-tables-statistics>
QUESTION 253
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: Implement table partitioning to improve

the manageability of the data warehouse and to avoid the need to repopulate all transactional data each night. Use a partitioning strategy that is as granular as possible. Partition the Fact.Order table and retain a total of seven years of data. Partition the Fact.Ticket table and retain seven years of data. At the end of each month, the partition structure must apply a sliding window strategy to ensure that a new partition is available for the upcoming month, and that the oldest month of data is archived and removed. Optimize data loading for the Dimension.SalesTerritory, Dimension.Customer, and Dimension.Date tables. Incrementally load all tables in the database and ensure that all incremental changes are processed. Maximize the performance during the data loading process for the Fact.Order partition. Ensure that historical data remains online and available for querying. Reduce ongoing storage costs while maintaining query performance for current data. You are not permitted to make changes to the client applications. You need to implement the data partitioning strategy. How should you partition the Fact.Order table?

A. Create 17,520 partitions.

B. Create 17,520 partitions.

Use a granularity of two days.C. Create 2,557 partitions.D. Create 730 partitions.
 Answer: C
 Explanation: We create one partition for each day. 7 years times 365 days is 2,555. Make that 2,557 to provide for leap years.
 From scenario: Partition the Fact.Order table and retain a total of seven years of data. Maximize the performance during the data loading process for the Fact.Order

partition. **QUESTION 254** 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 a Microsoft Azure SQL Data Warehouse instance that must be available six months a day for reporting. You need to pause the compute resources when the instance is not being used. Solution: You

use SQL Server Configuration Manager.Does the solution meet the goal?A. YesB. NoAnswer: BExplanation:To pause a SQL Data Warehouse database, use any of these individual methods.Pause compute with Azure portalPause compute with PowerShellPause compute with REST APIsReferences:

<https://docs.microsoft.com/en-us/azure/sql-data-warehouse/sql-data-warehouse-manage-compute-overview>QUESTION 255Note:

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 three databases as described in the following table.

Database	Description
----------	-------------

Database	Description
DB1	This database supports the online store.
DB2	This is the data warehouse for the company. DB2 contains a table named <code>SALES</code> with 24 columns. The <code>SALES</code> table has a <code>PARTITION</code> option set to AUTO . The data flow contains 24 <code>DB</code> destinations, one for each partition.
DB3	This database runs Master Data Services (MDS).

You plan to load at least one million rows of data each night from DB1 into the OnlineOrder table. You must load data into the correct partitions using a parallel process. You create 24 Data Flow tasks. You must place the tasks into a component to allow parallel load. After all of the load processes compete, the process must proceed to the next task. You need to load the data for the OnlineOrder 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 containerAnswer: HExplanation:The Parallel Loop Task is an SSIS Control Flow task, which can execute multiple iterations of the standard ForEach Loop Container concurrently.References: <http://www.cozyroc.com/ssis/parallel-loop-task>

QUESTION 256Note: 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	
Product	ProductID, ProductName, ProductDescription, ProductPrice, ProductCategory, ProductStatus
Customer	CustomerID, CustomerName, CustomerAddress, CustomerEmail, CustomerPhone, CustomerStatus
Order	OrderID, OrderDate, OrderStatus, OrderTotal, OrderItems, OrderCustomerID, OrderProductID
Inventory	InventoryID, InventoryName, InventoryQuantity, InventoryLocation, InventoryStatus

Database	Description
DB1	This database supports the online store.
DB2	This is the data warehouse for the company. All data in the table is compressed. The compression option is set to AUTO . The data flow contains one row per partition.
DB3	This database runs Master Data Services (MDS).

Each week, you import a product catalog from a partner company to a staging table in DB2. You need to create a stored procedure that will update the staging table by inserting new products and deleting discontinued products. 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

Answer: G
QUESTION 257
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: ERP system Online WebStore Partner extranet 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: 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.

Solution: Perform the following actions: Enable the Change Tracking for the Product table in the source databases. Query the CHANGETABLE function from the sources for the updated rows. Set the IsDisabled column to True for the listed rows that have the old ReferenceNr value. Create a new row in the data warehouse Products table with the new ReferenceNr value.

Does the solution meet the goal?

A. Yes B. No

Answer: B

Explanation: We must check for deleted rows, not just updates rows.

References:

<https://www.timmitchell.net/post/2016/01/18/getting-started-with-change-tracking-in-sql-server/>QUESTION 258Note: 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.

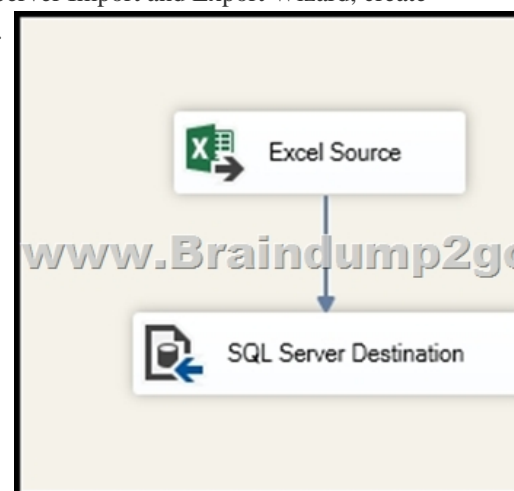
Database	Description
----------	-------------

Database	Description
DB1	This database supports the online store.
DB2	This is the data warehouse for the company. DB2 contains a table named <code>SALES</code> that is partitioned by month. The <code>TABLESPACE</code> option is set to AUTO . The data flow contains 24 <code>OLE DB</code> destinations, one for each partition.
DB3	This database runs Master Data Services (MDS).

Each day, you publish a Microsoft Excel workbook that contains a list of product names and current prices to an external website. Suppliers update pricing information in the workbook. Each supplier saves the workbook with a unique name. Each night, the Products table is deleted and refreshed from MDS by using a Microsoft SQL Server Integration Services (SSIS) package. All files must be loaded in sequence. You need to add a data flow in an SSIS package to perform the Excel files import in the data warehouse. 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

Answer: A
Explanation: If you're familiar with SSIS and don't want to run the SQL Server Import and Export Wizard, create an SSIS package that uses the Excel Source and the SQL Server Destination in the data flow.



References: <https://docs.microsoft.com/en-us/sql/integration-services/import-export-data/import-data-from-excel-to-sql>
 !!!RECOMMEND!!!1.2018 Latest 70-697 Exam Dumps (PDF & VCE) 2870&As Download:

<https://www.braindump2go.com/70-767.html2>.|2018 Latest 70-697 Study Guide Video: YouTube Video:
[YouTube.com/watch?v=di0FBePt_-w](https://www.youtube.com/watch?v=di0FBePt_-w)