

[2018-March-New100% Real 70-767 PDF and VCE Dumps 287Q-Braindump2go[270-280

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 270Hotspot QuestionYou have a series of analytic data models and reports that provide insights into the participation rates for sports at different schools. Users enter information about sports and participants into a client application. The application stores this transactional data in a Microsoft SQL Server database. A SQL Server Integration Services (SSIS) package loads the data into the models.When users enter data, they do not consistently apply the correct names for the sports. The following table shows examples of the data entry issues. You need to create a new knowledge base to improve the quality of the sport name data.How should you configure the knowledge base? To answer, select the appropriate options in the dialog box in the answer area. Answer: Explanation:Spot 1: Create Knowledge base from: NoneSelect None if you do not want to base the new knowledge base on an existing knowledge base or data file.QUESTION 271Drag and Drop QuestionYou have a series of analytic data models and reports that provide insights into the participation rates for sports at different schools. Users enter information about sports and participants into a client application. The application stores this transactional data in a Microsoft SQL Server database. A SQL Server Integration Services (SSIS) package loads the data into the models.When users enter data, they do not consistently apply the correct names for the sports. The following table shows examples of the data entry issues. You need to improve the quality of the data.Which four 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. Answer: Explanation:<https://docs.microsoft.com/en-us/sql/data-quality-services/perform-knowledge-discovery>QUESTION 272Drag and Drop QuestionNote: 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 configure the Fact.Order table.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. Answer: Explanation: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.Step 1: Create a partition function.Using CREATE PARTITION FUNCTION is the first step in creating a partitioned table or index.Step 2: Create a partition scheme based on the partition function. To migrate SQL Server partition definitions to SQL Data Warehouse simply:Step 3: Execute an ALTER TABLE command to specify the partition function.References: <https://docs.microsoft.com/en-us/azure/sql-data-warehouse/sql-data-warehouse-tables-partition>QUESTION 273Hotspot QuestionYou manage a data warehouse in a Microsoft SQL Server instance. Company employee information is imported from the human resources system to a table named Employee in the data warehouse instance. The Employee table was created by running the query shown in the Employee Schema exhibit. (Click the Exhibit button.) The personal identification number is stored in a column named EmployeeSSN. All values in the EmployeeSSN column must be unique.When

importing employee data, you receive the error message shown in the SQL Error exhibit. (Click the Exhibit button.) You determine that the Transact-SQL statement shown in the Data Load exhibit in the cause of the error. (Click the Exhibit button.) You remove the constraint on the EmployeeSSN column. You need to ensure that values in the EmployeeSSN column are unique. For each of the following statements, select Yes if the statement is true. Otherwise, select No. NOTE: Each correct selection is worth one point.

Answer: Explanation: With the ANSI standards SQL:92, SQL:1999 and SQL:2003, a UNIQUE constraint must disallow duplicate non-NULL values but accept multiple NULL values. In the Microsoft world of SQL Server however, a single NULL is allowed but multiple NULLs are not. From SQL Server 2008, you can define a unique filtered index based on a predicate that excludes NULLs.

References: <https://stackoverflow.com/questions/767657/how-do-i-create-a-unique-constraint-that-also-allows-nulls> QUESTION 274

Hotspot Question Your company has a Microsoft SQL Server data warehouse instance. The human resources department assigns all employees a unique identifier. You plan to store this identifier in a new table named Employee. You create a new dimension to store information about employees by running the following Transact-SQL statement: You have not added data to the dimension yet. You need to modify the dimension to implement a new column named [EmployeeKey]. The new column must use unique values. How should you complete the Transact-SQL statements? To answer, select the appropriate Transact-SQL segments in the answer area.

Answer: QUESTION 275 Drag and Drop Question You have a Microsoft SQL Server Integration Services (SSIS) package that loads data into a data warehouse each night from a transactional system. The package also loads data from a set of Comma-Separated Values (CSV) files that are provided by your company's finance department. The SSIS package processes each CSV file in a folder. The package reads the file name for the current file into a variable and uses that value to write a log entry to a database table. You need to debug the package and determine the value of the variable before each file is processed. Which four 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. Answer: Explanation: You debug control flows. The Foreach Loop container is used for looping through a group of files. Put the breakpoint on it. The Locals window displays information about the local expressions in the current scope of the Transact-SQL debugger. References:

<https://docs.microsoft.com/en-us/sql/integration-services/troubleshooting/debugging-control-flow>

<http://blog.pragmaticworks.com/looping-through-a-result-set-with-the-foreach-loop> QUESTION 276 Hotspot Question You create a Microsoft SQL Server Integration Services (SSIS) package as shown in the SSIS Package exhibit. (Click the Exhibit button.) The package uses data from the Products table and the Prices table. Properties of the Prices source are shown in the OLE DB Source Editor exhibit (Click the Exhibit Button.) and the Advanced Editor for Prices exhibit (Click the Exhibit button.) You join the Products and Prices tables by using the ReferenceNr column. You need to resolve the error with the package. For each of the following statements, select Yes if the statement is true. Otherwise, select No. NOTE: Each correct selection is worth one point.

Answer: Explanation: There are two important sort properties that must be set for the source or upstream transformation that supplies data to the Merge and Merge Join transformations: The Merge Join Transformation requires sorted data for its inputs. If you do not use a Sort transformation to sort the data, you must set these sort properties manually on the source or the upstream transformation. References:

<https://docs.microsoft.com/en-us/sql/integration-services/data-flow/transformations/sort-data-for-the-merge-and-merge-join-transformations> QUESTION 277

Drag and Drop Question You deploy a Microsoft Server database that contains a staging table named EmailAddress_Import. Each night, a bulk process will import customer information from an external database, cleanse the data, and then insert it into the EmailAddress table. Both tables contain a column named EmailAddressValue that stores the email address. You need to implement the logic to meet the following requirements: Email addresses that are present in the EmailAddress_Import table but not in the EmailAddress table must be inserted into the EmailAddress table. Email addresses that are not in the EmailAddress_Import but are present in the EmailAddress table must be deleted from the EmailAddress table. How should you complete the Transact-SQL statement? To answer, drag the appropriate Transact-SQL segments to the correct locations. Each Transact-SQL segment 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. Answer: Explanation: Box 1: EmailAddress The EmailAddress table is the target. Box 2: EmailAddress_import The EmailAddress_import table is the source. Box 3: NOT MATCHED BY TARGET Box 4: NOT MATCHED BY SOURCE

References: <https://docs.microsoft.com/en-us/sql/t-sql/statements/merge-transact-sql> QUESTION 278 Drag and Drop Question 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 partitioning for the Fact.Ticket table. Which three actions should you perform in sequence? To answer, drag the appropriate actions to the correct locations. Each action 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: More than one combination of answer choices is correct. You will receive credit for any of the correct combinations you select. Answer: Explanation: From scenario: - 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. The detailed steps for the recurring partition maintenance tasks are: References:

<https://docs.microsoft.com/en-us/sql/relational-databases/tables/manage-retention-of-historical-data-in-system-versioned-temporal-tables> QUESTION 279 Drag and Drop Question 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 optimize data loading for the Dimension.Customer table. Which three Transact-SQL segments should you use to develop the solution? To answer, move the appropriate Transact-SQL segments from the list of Transact-SQL segments to the answer area and arrange them in the correct order. NOTE: You will not need all of the Transact-SQL segments. Answer: Explanation: Step 1: USE DB1 From Scenario: 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. Step 2: EXEC sys.sp_cdc_enable_db Before you can enable a table for change data capture, the database must be enabled. To enable the database, use the sys.sp_cdc_enable_db stored procedure. sys.sp_cdc_enable_db has no parameters. Step 3: EXEC sys.sp_cdc_enable_table @source_schema = N 'schema' etc. Sys.sp_cdc_enable_table enables change data capture for the specified source table in the current database. Partial syntax: sys.sp_cdc_enable_table [@source_schema =] 'source_schema', [@source_name =] 'source_name', [[@capture_instance =] 'capture_instance'] [[@supports_net_changes =] supports_net_changes] Etc. References:

<https://docs.microsoft.com/en-us/sql/relational-databases/system-stored-procedures/sys-sp-cdc-enable-table-transact-sql>

<https://docs.microsoft.com/en-us/sql/relational-databases/system-stored-procedures/sys-sp-cdc-enable-db-transact-sql>QUESTION 280Hotspot QuestionYou have a Microsoft SQL Server Integration Services (SSIS) package that contains a Data Flow task as shown in the Data Flow exhibit. (Click the Exhibit button.) You install Data Quality Services (DQS) on the same server that hosts SSIS and deploy a knowledge base to manage customer email addresses. You add a DQS Cleansing transform to the Data Flow as shown in the Cleansing exhibit. (Click the Exhibit button.) You create a Conditional Split transform as shown in the Splitter exhibit. (Click the Exhibit button.) You need to split the output of the DQS Cleansing task to obtain only Correct values from the EmailAddress column.For each of the following statements, select Yes if the statement is true. Otherwise, select No. Answer: Explanation:The DQS Cleansing component takes input records, sends them to a DQS server, and gets them back corrected. The component can output not only the corrected data, but also additional columns that may be useful for you. For example - the status columns. There is one status column for each mapped field, and another one that aggregated the status for the whole record. This record status column can be very useful in some scenarios, especially when records are further processed in different ways depending on their status. In such cases, it is recommended to use a Conditional Split component below the DQS Cleansing component, and configure it to split the records to groups based on the record status (or based on other columns such as specific field status).References:
<https://blogs.msdn.microsoft.com/dqs/2011/07/18/using-the-ssis-dqs-cleansing-component/>!!!RECOMMEND!!!1.|2018 Latest 70-697 Exam Dumps (PDF & VCE) 287Q&As Download:<https://www.braindump2go.com/70-767.html>2.|2018 Latest 70-697 Study Guide Video: YouTube Video: [YouTube.com/watch?v=di0FBePt_-w](https://www.youtube.com/watch?v=di0FBePt_-w)