



Vendor: Microsoft

Exam Code: 70-767

Exam Name: Implementing a Data Warehouse using SQL

Version: DEMO

QUESTION 1

A SQL Server Integration Services (SSIS) package on a computer is failing periodically in production. The package was created less than one year ago and was deployed to the SSIS catalog. Sometimes the package is started on a SQL Agent schedule; sometimes the package is started manually by an SSIS developer by using the Object Explorer in SQL Server Management Studio.

You need to identify the authenticated user responsible for starting the package each time it failed in the past.

Where can you find this information?

- A. the SQL Server Log
- B. the SSISDB.[catalog].[executions] view
- C. the SSISDB.[catalog].[event_messages] view
- D. the SQL Agent Job History
- E. the SQL Agent Error Log

Answer: B

QUESTION 2

You are reviewing the design of a customer dimension table in an existing data warehouse hosted on SQL Azure. The current dimension design does not allow the retention of historical changes to customer attributes such as Postcode.

You need to redesign the dimension to enable the full historical reporting of changes to multiple customer attributes including Postcode.

What should you do?

- A. Add StartDate and EndDate columns to the customer dimension.
- B. Add an IsCurrent column to the customer dimension.
- C. Enable Snapshot Isolation on the data warehouse.
- D. Add CurrentValue and PreviousValue columns to the customer dimension.

Answer: A

QUESTION 3

You are designing an enterprise star schema that will consolidate data from three independent data marts. One of the data marts is hosted on SQL Azure. Most of the dimensions have the same structure and content. However, the geography dimension is slightly different in each data mart.

You need to design a consolidated dimensional structure that will be easy to maintain while ensuring that all dimensional data from the three original solutions is represented.

What should you do?

- A. Create a junk dimension for the geography dimension.
- B. Implement change data capture.
- C. Create a conformed dimension for the geography dimension.
- D. Create three geography dimensions.

Answer: C

QUESTION 4

You are designing an extract, transform, load (ETL) process for loading data from a SQL Server

database into a large fact table in a data warehouse each day with the prior day's sales data. The ETL process for the fact table must meet the following requirements:

- . Load new data in the shortest possible time.
- . Remove data that is more than 36 months old.
- . Ensure that data loads correctly.
- . Minimize record locking.
- . Minimize impact on the transaction log.

You need to design an ETL process that meets the requirements.

What should you do? (More than one answer choice may achieve the goal. Select the BEST answer.)

- A. Partition the destination fact table by date.
Insert new data directly into the fact table and delete old data directly from the fact table.
- B. Partition the destination fact table by date.
Use partition switching and staging tables both to remove old data and to load new data.
- C. Partition the destination fact table by customer.
Use partition switching both to remove old data and to load new data into each partition.
- D. Partition the destination fact table by date.
Use partition switching and a staging table to remove old data.
Insert new data directly into the fact table.

Answer: B

QUESTION 21

You are designing a fact table in a SQL Server database.

The fact table must meet the following requirements:

- . Include a column store index.
- . Allow users to choose up to 10 dimension tables and up to five facts at one time.
- . Maximize performance of queries that aggregate measures by using any of the 10 dimensions.
- . Support billions of rows.
- . Use the most efficient design strategy.

You need to design the fact table to meet the requirements.

What should you do? (More than one answer choice may achieve the goal. Select the BEST answer.)

- A. Design a fact table with 5 dimensional key columns and 10 measure columns.
Place the column store index on the dimensional key columns.
- B. Design a fact table with 5 dimensional key columns and 10 measure columns.
Place the column store index on the measure columns.
- C. Design a fact table with 10 dimensional key columns and 5 measure columns.
Place the column store index on the dimensional key columns and the measure columns.
- D. Design a fact table with 10 dimensional key columns and 5 measure columns.
Place the column store index on only the measure columns.

Answer: C

QUESTION 6

You are designing a SQL Server Integration Services (SSIS) solution that will load multiple Online Transactional Processing (OLTP) data sources into a SQL Server data mart.

You have the following requirements:

- . Ensure that the process supports the creation of an exception report that details possible duplicate key values, null ratios within columns, and column-length distributions of values.
- . Ensure that users can generate the exception report in an XML format.
- . Use the minimum development effort.

You need to design the SSIS solution to meet the requirements.

What should you do? (More than one answer choice may achieve the goal. Select the BEST answer.)

- A. Use a Data Profiling task. Use a Data Flow task to extract the XML output of the Data Profiling task into a SQL Server table. Query the table to view the exceptions.
- B. Use Data Flow tasks to process the clean data.
- C. Use a Data Profiling task. Read the exceptions in Data Profile Viewer.
- D. Design a stored procedure that examines data for common dirty data patterns. Use an Execute SQL task.

Answer: C

QUESTION 7

You are designing a partitioning strategy for a large fact table in a data warehouse. Tens of millions of new records are loaded into the data warehouse weekly, outside of business hours. Most queries are generated by reports and by cube processing. Data is frequently queried at the day level and occasionally at the month level.

You need to partition the table to maximize the performance of queries.

What should you do? (More than one answer choice may achieve the goal. Select the BEST answer.)

- A. Partition the fact table by month, and compress each partition.
- B. Partition the fact table by week.
- C. Partition the fact table by year.
- D. Partition the fact table by day, and compress each partition.

Answer: D

Thank You for Trying Our Product

Braindump2go Certification Exam Features:

- ★ More than **99,900** Satisfied Customers Worldwide.
- ★ Average **99.9%** Success Rate.
- ★ **Free Update** to match latest and real exam scenarios.
- ★ **Instant Download** Access! No Setup required.
- ★ Questions & Answers are downloadable in **PDF** format and **VCE** test engine format.
- ★ Multi-Platform capabilities - **Windows, Laptop, Mac, Android, iPhone, iPod, iPad**.
- ★ **100%** Guaranteed Success or **100%** Money Back Guarantee.
- ★ **Fast**, helpful support **24x7**.



View list of all certification exams: <http://www.braindump2go.com/all-products.html>



Microsoft



ORACLE



CITRIX



JUNIPER
NETWORKS



EMC²
where information lives[®]

10% Discount Coupon Code: BDN2014