



70-463^{Q&As}

Implementing a Data Warehouse with Microsoft SQL Server 2012/2014

Pass Microsoft 70-463 Exam with 100% Guarantee

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

<https://www.passapply.com/70-463.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 developing a SQL Server Integration Services (SSIS) package to implement an incremental data load strategy. The package reads rows from a source system and compares them to rows in a destination system. New rows will be inserted and changed rows will be updated.

You have used a Lookup transformation and a Conditional Split transformation. The Lookup transformation joins the source and destination table on the business key, and includes all columns from the destination table in the data flow output.

The Conditional Split transformation inspects the destination columns and directs data flow to either insert new records or update existing records.

You need to configure the Lookup transformation to ensure that all records flow to the Conditional Split transformation, regardless of whether the rows match an existing row in the destination table. Which setting should you select? (To answer, select the appropriate option in the answer area.)

Hot Area:

Cache connection manager

OLE DB connection manager

Specify how to handle rows with no matching entries

- Ignore Failure
- Redirect rows to error output
- Fail component
- Redirect rows to no match output

Correct Answer:

Cache connection manager

OLE DB connection manager

Specify how to handle rows with no matching entries

- Ignore Failure
- Redirect rows to error output
- Fail component
- Redirect rows to no match output



QUESTION 2

You are designing a data warehouse hosted on SQL Azure. The data warehouse currently includes the dimUser and dimDistrict dimension tables and the factSales fact table. The dimUser table contains records for each user permitted to run

reports against the warehouse; and the dimDistrict table contains information about sales districts.

The system is accessed by users from certain districts, as well as by area supervisors and users from the corporate headquarters.

You need to design a table structure to ensure that certain users can see sales data for only certain districts. Some users must be permitted to see sales data from multiple districts.

What should you do?

- A. Add a district column to the dimUser table.
- B. Partition the factSales table on the district column.
- C. Create a userDistrict table that contains primary key columns from the dimUser and dimDistrict tables.
- D. For each district, create a view of the factSales table that includes a WHERE clause for the district.

Correct Answer: C

QUESTION 3

You are designing a data warehouse for a fresh food distribution business that stores sales by individual product. It stores sales targets by product category. Products are classified into subcategories and categories. Each product is included in only a single product subcategory, and each subcategory is included in only a single category.

The data warehouse will be a data source for an Analysis Services cube. The data warehouse contains two fact tables:

factSales, used to record daily sales by product

factProductTarget, used to record the monthly sales targets by product category

Reports must be developed against the warehouse that reports product sales by product, category and subcategory, and product sales targets.

You need to design the product dimension. The solution should use as few tables as possible while supporting all the requirements.

What should you do?

- A. Create two product tables, dimProduct and dimProductCategory. Connect factSales to dimProduct and factProductTarget to dimProductCategory with foreign key constraints. Direct the cube developer to use key granularity attributes.
- B. Create one product table, dimProduct, which contains product detail, category, and subcategory columns. Connect



factSales to dimProduct with a foreign key constraint. Direct the cube developer to use a non-key granularity attribute for factProductTarget.

C. Create three product tables, dimProduct, dimProductCategory, and dimProductSubcategory, and a fourth bridge table that joins products to their appropriate category and subcategory table records with foreign key constraints. Direct the cube developer to use key granularity attributes.

D. Create three product tables, dimProduct, dimProductCategory, and dimProductSubcategory. Connect factSales to all three product tables and connect factProductTarget to dimProductCategory with foreign key constraints. Direct the cube developer to use key granularity attributes.

Correct Answer: B

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.

Correct Answer: B

QUESTION 5



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. Join the two fact tables.
- B. Partition the fact tables by day.
- C. Create a surrogate key for the time dimension.
- D. Create a time dimension that can join to both fact tables at their respective granularity.

Correct Answer: D

Microsoft SQL Server Analysis Services, a time dimension is a dimension type whose attributes represent time periods, such as years, semesters, quarters, months, and days. The periods in a time dimension provide time-based levels of granularity for analysis and reporting. The attributes are organized in hierarchies, and the granularity of the time dimension is determined largely by the business and reporting requirements for historical data

References: <https://docs.microsoft.com/en-us/sql/analysis-services/multidimensional-models/database-dimensions-create-a-date-type-dimension>

[Latest 70-463 Dumps](#)

[70-463 Study Guide](#)

[70-463 Braindumps](#)



To Read the [Whole Q&As](#), please purchase the [Complete Version](#) from [Our website](#).

Try our product !


100% Guaranteed Success
100% Money Back Guarantee
365 Days Free Update
Instant Download After Purchase
24x7 Customer Support
Average 99.9% Success Rate
More than 800,000 Satisfied Customers Worldwide
Multi-Platform capabilities - [Windows](#), [Mac](#), [Android](#), [iPhone](#), [iPod](#), [iPad](#), [Kindle](#)

We provide exam PDF and VCE of Cisco, Microsoft, IBM, CompTIA, Oracle and other IT Certifications.
You can view Vendor list of All Certification Exams offered:

<https://www.passapply.com/allproducts>

Need Help

Please provide as much detail as possible so we can best assist you.
To update a previously submitted ticket:



| | | |
|---|---|--|
|  <p>One Year Free Update Free update is available within One Year after your purchase. After One Year, you will get 50% discounts for updating. And we are proud to boast a 24/7 efficient Customer Support system via Email.</p> |  <p>Money Back Guarantee To ensure that you are spending on quality products, we provide 100% money back guarantee for 30 days from the date of purchase.</p> |  <p>Security & Privacy We respect customer privacy. We use McAfee's security service to provide you with utmost security for your personal information & peace of mind.</p> |
|---|---|--|

Any charges made through this site will appear as Global Simulators Limited.
All trademarks are the property of their respective owners.
Copyright © passapply, All Rights Reserved.