VCE & PDF
GeekCert.com

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.geekcert.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 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 2

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

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_batchIndicates 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 3

A SQL Server Integration Services (SSIS) package currently downloads sales data from a Microsoft Azure SQL Database database.

To improve sales data accuracy, exchange rates must be downloaded daily from a public HTTP web service instead of from a weekly flat file.

You need to implement the change to the existing package while minimizing the development effort.

What should you use to call the web service to retrieve the daily exchange rates?

A. a Script task

B. a Web Service task

C. a Script component

D. a Web Service source

Correct Answer: B

References: https://docs.microsoft.com/en-us/sql/integration-services/control-flow/web-service-task?view=sql-server-2017

https://www.mssqltips.com/sqlservertip/3272/example-using-web-services-with-sql-server-integration-services/

QUESTION 4

You work as a database developer at ABC.com. The ABC.com network consists of a single domain named ABC.com. ABC.com makes use of Microsoft SQL Server 2012 in their environment.

You are in the process of creating a SQL Server Integration Services (SSIS) package. The package will be configured to

make use of the GenerateAndPersistNewIndex Fuzzy Lookup transformation option.

Which of the following is the purpose of this option?

A. It allows for a current index to be reused.

B. It allows for a new index to be created, without saving it.

C. It allows for a new index to be created and saved, without maintaining it.

D. It allows for a new index to be created, saved, and maintained.

Correct Answer: C

QUESTION 5

You are designing a data warehouse for a software distribution business that stores sales by software title. It stores sales targets by software category. Software titles are classified into subcategories and categories. Each software title is

included in only a single software 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 software title

factTarget,

 used to record the monthly sales targets by software category

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

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

What should you do?

A.

Create three software tables, dimSoftware, dimSoftwareCategory, and dimSoftwareSubcategory and a fourth bridge table that joins software titles to their appropriate category and subcategory table records with foreign key constraints. Direct the cube developer to use key granularity attributes.

B.

Create three software tables, dimSoftware, dimSoftwareCategory, and dimSoftwareSubcategory. Connect factSales to all three tables and connect factTarget to dimSoftwareCategory with foreign key constraints. Direct the cube developer to use key granularity attributes.

C.

Create one table, dimSoftware, which contains Software Detail, Category, and Subcategory columns. Connect factSales to dimSoftware with a foreign key constraint. Direct the cube developer to use a non-key granularity attribute for

factTarget.

D.

Create two tables, dimSoftware and dimSoftwareCategory. Connect factSales to dimSoftware and factTarget to dimSoftwareCategory with foreign key constraints. Direct the cube developer to use key granularity attributes.

Correct Answer: C

[70-463 PDF Dumps](https://www.geekcert.com/70-463.html)   [70-463 VCE Dumps](https://www.geekcert.com/70-463.html)   [70-463 Practice Test](https://www.geekcert.com/70-463.html)

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.geekcert.com/allproducts

Need Help

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