70-767Q&As

Implementing a Data Warehouse using SQL

Pass Microsoft 70-767 Exam with 100% Guarantee

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

https://www.lead4pass.com/70-767.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Microsoft Official Exam Center

QUESTION 1

You are developing a SQL Server Integration Services (SSIS) project to read and write data from a Windows Azure SQL Database database to a server that runs SQL Server 2016. The connection will be used by data flow tasks in multiple SSIS packages. The address of the target Windows Azure SQL Database database will be provided by a project parameter. You need to create a solution to meet the requirements by using the least amount of administrative effort.

What should you do?

A. Add a SQLMOBILE connection manager to each package.

B. Add an ADO.NET project connection manager.

C. Add a SQLMOBILE project connection manager.

D. Add an ADO.NET connection manager to each data flow task.

E. Add a SQLMOBILE connection manager to each data flow task.

F. Add an ADO.NET connection manager to each package.

Correct Answer: B

Ref: http://www.databasejournal.com/features/mssql/windows-azure-sql-database- uploading-data-by-using-sql-server-integration-services.html

QUESTION 2

You have a data warehouse named DW1 that contains 20 years of data. DW1 contains a very large fact table. New data is loaded to the fact table monthly.

Many reports query DW1 for the past year of data.

Users frequently report that the reports are slow.

You need to modify the fact table to minimize the amount of time it takes to run the reports. The solution must ensure that other reports can continue to be generated from DW1.

What should you do?

A. Move the historical data to SAS disks and move the data from the past year to SSD disks. Run the ALTER TABLE statement.

B. Move all the data to SSD disks. Load and archive the data by using partition switching.

C. Move all the data to SAS disks. Load and archive the data by using partition switching.

D. Move the historical data to SAS disks and move the data for the past year to SSD disks. Create a distributed partitioned view.

Correct Answer: A

We use ALTER TABLE to partition the table. Incorrect Answers:

D: A Distributed Partitioned View contains participating tables from multiple SQL Server instances, which can be used to distribute the data processing load across multiple servers. Another advantage for the SQL Server Partitioned Views is that the underlying tables can participate in more than one Partitioned View, which could be helpful in some implementations.

QUESTION 3

Note: 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 have a database named DB1 that has change data capture enabled.

A Microsoft SQL Server Integration Services (SSIS) job runs once weekly. The job loads changes from DB1 to a data warehouse by querying the change data capture tables.

You discover that the job loads changes from the previous three days only.

You need re ensure that the job loads changes from the previous week.

Which stored procedure should you execute?

A. catalog.deploy_project

B. catalog.restore_project

C. catalog.stop.operation

D. sys.sp_cdc.addJob

E. sys.sp.cdc.changejob

F. sys.sp_cdc_disable_db

G. sys.sp_cdc_enable_db

H. sys.sp_cdc.stopJob

Correct Answer: A

catalog.deploy_project deploys a project to a folder in the Integration Services catalog or updates an existing project that has been deployed previously. References: https://docs.microsoft.com/en-us/sql/integration-services/system-stored-procedures/catalog-deploy-project- ssisdb-database

QUESTION 4

You have a database Ihat includes a table named dbo.sales. The table contains two billion rows. You created the table by running the following Transact-SQL statement:

```
CREATE TABLE dbo.Sales (
  SaleId BIGINT PRIMARY KEY,
  StoreId INT,
  EmployeeId INT,
  SaleAmount MONEY,
  TaxAmount MONEY,
  SubTotalAmount MONEY,
  LineItems XML,
  Refund BIT,
  SaleData DATE,
  SaleTime TIME
)
```

You run the following queries against the dbo.sales table. All of the queries perform poorly.

Query name	Query text
Query1	SELECT StoreId, SUM(SaleAmount) SaleTotal, SUM(TaxAmount) TaxTotal FROM dbo. Sales WHERE SaleDate BETWEEN '1/1/2015' AND '1/1/2016' GROUP BY StoreId
Query2	SELECT StoreId, datepart(hh, SaleTime) SaleHour,count (*) FROM dbo. Sales WHERE SaleDate = convert (varchar(10), getdate ( ) -1, 111) GROUP BY StoreId, datepart (hh,SaleTime)
Query3	SELECT SaleId, StoreId, EmployeeId, SaleAmount FROM dbo. Sales WHERE Refund = 1 AND SaleDate = convert (varchar(10), getdate ( ) -1, 111)

The ETL process that populates the table uses bulk insert to load 10 million rows each day. The process currently takes six hours to load the records.

The value of the Refund column is equal to 1 for only 0.01 percent of the rows in the table. For all other rows, the value of the Refund column is equal to 0.

You need to maximize the performance of queries and the ETL process.

Which index type should you use for each query? To answer, select the appropriate index types in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

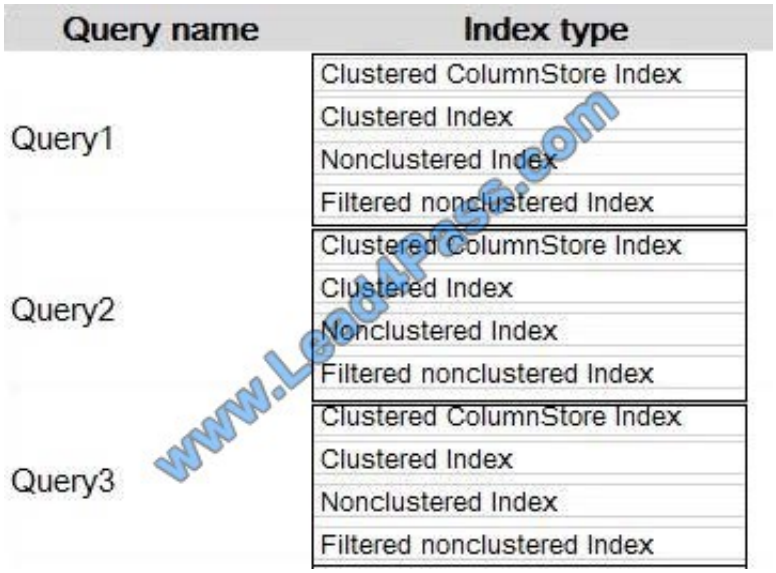

Correct Answer:

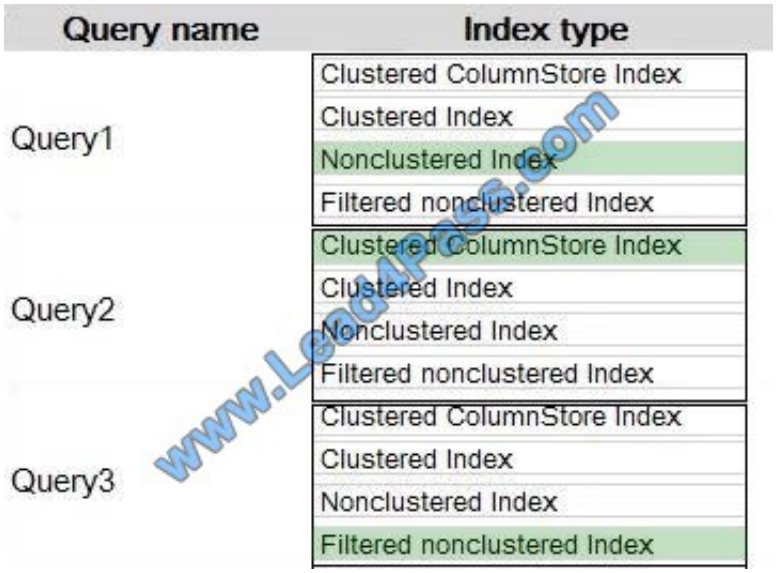

Query1: Nonclustered Index

The query include a date range.

If you have included columns in your index, then the leaf level page of your non-clustered index contains the columns as defined in the nonclustered index the clustering key column(s) all those additional columns as defined in your INCLUDE

statement.

Query2: Clustered columnstore index

Columnstore index is a new type of index introduced in SQL Server 2012. It is a column-based non-clustered index geared toward increasing query performance for workloads that involve large amounts of data, typically found in data

warehouse fact tables.

Query3: Filtered nonclustered index

* When a column only has a small number of relevant values for queries, you can create a filtered index on the subset of values. For example, when the values in a column are mostly NULL and the query selects only from the non-NULL values, you can create a filtered index for the non-NULL data rows. The resulting index will be smaller and cost less to maintain than a full-table nonclustered index defined on the same key columns.

When a table has heterogeneous data rows, you can create a filtered index for one or more categories of data. This can improve the performance of queries on these data rows by narrowing the focus of a query to a specific area of the table. Again, the resulting index will be smaller and cost less to maintain than a full-table nonclustered index.

References: https://docs.microsoft.com/en-us/sql/relational-databases/indexes/create-filtered-indexes

https://logicalread.com/sql-server-columnstore-index-w02/#.XRo06egzaUk

QUESTION 5

You are editing a SQL Server Integration Services (SSIS) package.

You need to edit the Script task that has an expression defined.

Which task should you select? To answer, select the appropriate setting or settings in the answer area.

Hot Area:

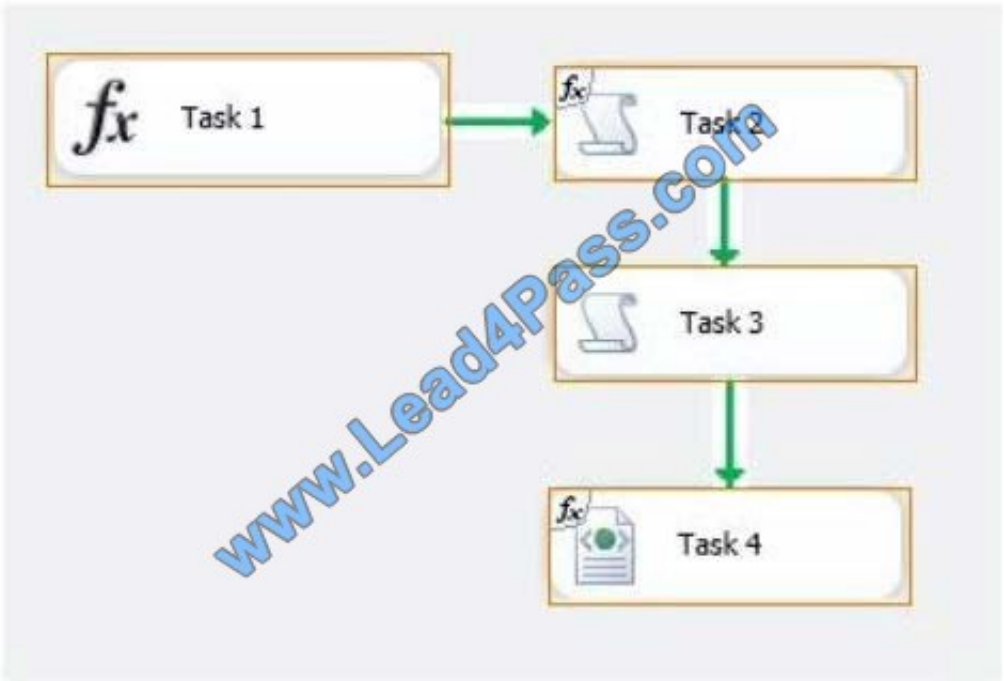

Correct Answer:

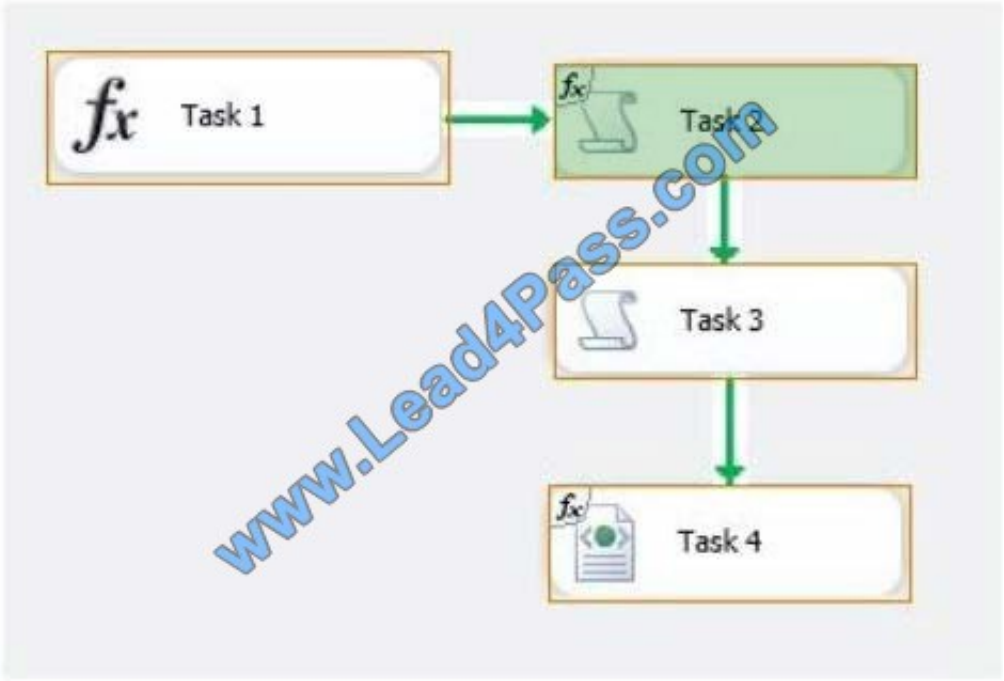

Latest 70-767 Dumps 70-767 Study Guide 70-767 Braindumps