

(New Version) Real Microsoft 70-467 Practice Tests For Free Share From PassLeader (21-30)

Attention: These 70-467 Exam Questions Were Updated After 2014/7/1 With The Change Of New Microsoft Exam. 100 Percent Valid And 100 Percent Pass Ensure. Visit Our PassLeader Website And Get All Valid Exam Questions With PDF And VCE.

QUESTION 21 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.

QUESTION 22 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 23 You have a business intelligence (BI) infrastructure that contains three servers. The servers are configured as shown in the following table.

| Server name | |
|-------------|--------------------|
| Server1 | Front-end |
| Server2 | Application Server |
| Server3 | SQL Server |
| | SQL Server |
| | SQL Server |

You need to recommend a health monitoring solution for the BI infrastructure. The solution must meet the following requirements:
 - Monitor the status of the Usage Data Collection feature.
 - Monitor the number of end-users accessing the solution.
 - Monitor the amount of cache used when the users query data.

Which health monitoring solution should you recommend using on each server? To answer, drag the appropriate monitoring solutions to the correct servers. Each monitoring solution 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.

The screenshot shows a drag-and-drop interface. On the left, under 'Monitoring Solutions', there are three items: 'PowerPivot server health rules', 'Web Service performance object', and 'Dynamic management views'. On the right, under 'Answer Area', there is a table with three rows for 'Server1', 'Server2', and 'Server3', each with an empty 'Monitoring Solution' column.

Answer:

The screenshot shows the same interface as above, but with the following configurations in the Answer Area table:
 - Server1: PowerPivot server health rules
 - Server2: Dynamic management views
 - Server3: PowerPivot server health rules

QUESTION 24 Drag and Drop Questions You are validating whether a SQL Server Integration Services (SSIS) package named Master.dtsx in the SSIS catalog is executing correctly. You need to display the number of rows in each buffer passed between each data flow component of the package. 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.)

Execute a SQL statement with a package name of Master.dtsx against the **catalog.executions** view and return its execution ID.

Run the Master.dtsx package with the logging level set to **Performance**.

Execute a SQL statement with the execution ID equal to the previously retrieved execution ID against the **catalog.execution_data_statistics** view and return the **rows_sent** column values for all the rows.

Run the Master.dtsx package with the logging level set to **Verbose**.

Execute a SQL statement with a package name of Master.dtsx against the **msdb..sysssislog** table and return its execution ID.

Answer:

Run the Master.dtsx package with the logging level set to **Performance**.

Execute a SQL statement with a package name of Master.dtsx against the **catalog.executions** view and return its execution ID.

Execute a SQL statement with the execution ID equal to the previously retrieved execution ID against the **catalog.execution_data_statistics** view and return the **rows_sent** column values for all the rows.

Run the Master.dtsx package with the logging level set to **Verbose**.

Execute a SQL statement with a package name of Master.dtsx against the **msdb..sysssislog** table and return its execution ID.

QUESTION 25 You are creating a Multidimensional Expressions (MDX) calculation for Projected Revenue in a cube. For Customer A, Projected Revenue is defined as 150 percent of the Total Sales for the customer. For all other customers, Projected Revenue is defined as 110 percent of the Total Sales for the customer. You need to calculate the Projected Revenue as efficiently as possible. Which calculation should you use? (More than one answer choice may achieve the goal. Select the BEST answer.)

- A. CREATE MEMBER AS [Measures] SCOPE ([Customer].[Customer]) IF [Customer] = 'A' THEN [Measures].[TotalSales] * 1.5 ELSE [Measures].[TotalSales] * 1.1 END SCOPE;
- B. CREATE MEMBER AS CASE WHEN [Customer] = 'A' THEN [Measures].[TotalSales] * 1.5 ELSE [Measures].[TotalSales] * 1.1 END SCOPE;
- C. CREATE MEMBER AS [Measures] SCOPE ([Customer].[Customer]) THIS = [Measures].[TotalSales] * 1.5 END SCOPE;
- D. CREATE MEMBER AS [Measures] SCOPE ([Customer].[Customer]) [Customer] = 'A' THEN [Measures].[TotalSales] * 1.5 ELSE [Measures].[TotalSales] * 1.1 END SCOPE;

A. Option A B. Option B C. Option C D. Option D Answer: C



<http://www.passleader.com/70-467.html>] QUESTION 26 You are developing the database schema for a SQL Server Analysis Services (SSAS) BI Semantic Model (BISM). The BISM will be based on the schema displayed in the following graphic.

| Student | |
|---------|-------------|
| PK | StudentID |
| | StudentName |
| | StudentAge |

You have the following requirements: - Ensure that queries of the data model correctly display average student age by class and average class level by student. - Minimize development effort. You need to design the data model. What should you do? (More than one answer choice may achieve the goal. Select the BEST answer.)

A. Create a multidimensional project and define measures and a reference relationship. B. Create a tabular project and define calculated columns. C. Create a multidimensional project and define measures and a many-to-many dimensional

relationship. D. Create a tabular project and define measures. Answer: C QUESTION 27 Drag and Drop Questions You are designing a self-service business intelligence and reporting environment. Business analysts will create and publish PowerPivot for Microsoft Excel workbooks and create reports by using SQL Server Reporting Services (SSRS) and Power View. When the data models become more complex and the data volume increases, the data models will be replaced by IT-hosted server-based models. You have the following requirements: - Maintain the self-service nature of the reporting environment. - Reuse existing reports. - Add calculated columns to the data models. You need to create a strategy for implementing this process. What should you do? To answer, drag the appropriate term or terms to the correct location or locations in the answer area. (Answer choices may be used once, more than once, or not all.)

Answer:

QUESTION 28 You are modifying a star schema data mart that feeds order data from a SQL Azure database into a SQL Server Analysis Services (SSAS) cube. The data mart contains two large tables that include flags and indicators for some orders. There are 100 different flag columns, each with 10 different indicator values. Some flags reuse indicators. The tables both have a granularity that matches the fact table. You have the following requirements: - Allow users to slice data by all flags and indicators. - Modify the date dimension table to include a surrogate key of a numeric data type and add the surrogate key to the fact table. - Use the most efficient design strategy for cube processing and queries. You need to modify the schema. What should you do? (More than one answer choice may achieve the goal. Select the BEST answer.) A. Define the surrogate key as an INT data type. Combine the distinct flag/indicator combinations into a single dimension. B. Define the surrogate key as an INT data type. Create a single fact dimension in each table for its flags and indicators. C. Define the surrogate key as a BIGINT data type. Combine the distinct flag/indicator combinations into a single dimension.

D. Define the surrogate key as a BIGINT data type. Create a single fact dimension in each table for its flags and indicators. Answer: A QUESTION 29 You are defining a named set by using Multidimensional Expressions (MDX) in a sales cube. The cube includes a Product dimension that contains a Category hierarchy and a Color attribute hierarchy. You need to return only the blue products in the Category hierarchy. Which set should you use? (More than one answer choice may achieve the goal. Select the BEST answer.)

- A. CrossJoin


```
{
    [Product].[Product Category].[Product Name].Members,
    [Product].[Color].&[Blue]
}
```
- B. Filter


```
{
    [Product].[Product Category].[Product Name].Members,
    ([Product].[Color].&[Blue], [Measures].[Sales Amount]) > 0
}
```
- C. Exists


```
{
    [Product].[Product Category].[Product Name].Members,
    [Product].[Color].&[Blue]
}
```
- D. Generate


```
{
    [Product].[Color].&[Blue],
    [Product].[Model Name].[Model Name].Members, ALL
}
```

A. Option A B. Option B C. Option C
D. Option D Answer: C QUESTION 30 An existing cube dimension that has 30 attribute hierarchies is performing very poorly. You have the following requirements: - Implement drill-down browsing. - Reduce the number of attribute hierarchies but ensure that the information contained within them is available to users on demand. - Optimize performance. You need to redesign the cube dimension to meet the requirements. What should you do? (More than one answer choice may achieve the goal. Select the BEST answer.) A. set the AggregateFunction property to Sum on all measures. Use the SCOPE statement in a Multidimensional Expressions (MDX) calculation to tune the aggregation types. B. Set the AttributeHierarchyOptimizedState property to FullyOptimized on the attribute hierarchies. C. Create user-defined hierarchies. For the attributes sourced by the levels of the user-defined hierarchies, set the RelationshipType property to Rigid. Run incremental processing. D. Remove as many attribute hierarchies as possible from the dimension. Reintroduce the information in the attribute hierarchies as properties. Implement natural hierarchies and set the AttributeHierarchyVisible property to False for attributes used as levels in the natural hierarchies. Answer: D

(New Version) Real Microsoft 70-467 Practice Tests For Free Share From PassLeader [Click Here To Get The New Update And 100 Percent Valid & Pass 70-467 Exam Questions -- http://www.passleader.com/70-467.html](#)]