


[100% Pass PassLeader Share New Real Microsoft 70-467 191q Exam Questions With 100 Percent Passing Ensure (111-140)]

100% Valid 70-467 Exam Pass Tips: PassLeader's 70-467 191q exam dumps were updated in recent days with the change of new 70-467 exam questions, PassLeader ensure the 70-467 191q braindumps are the newest and the most valid, our 70-467 191q practice test will help you 100 percent pass 70-467 exam. Visit passleader.com and get the 70-467 191q exam questions with PDF and VCE. New version vce player is also free now. keywords: 70-467 exam,70-467 191q exam dumps,70-467 191q exam questions,70-467 pdf dumps,70-467 vce dumps,70-467 braindump,Designing Business Intelligence Solutions with Microsoft SQL Server 2012 Exam



Case Study: 6 - Tailspin Toys Case B (QUESTION 111 - QUESTION 120) Overview Tailspin Toys is a manufacturing company that has offices across the United States, Europe, and Asia. Tailspin Toys plans to implement a business intelligence (BI) solution for its US-based headquarters to manage the sales data, including information on customer transactions, products, sales quotas, and bonuses. Existing Environment Data Sources Tailspin Toys currently stores data in line-of-business applications, relational databases, flat files, and the following:

- A Microsoft Excel spreadsheet named MarketResearch.xlsx. The spreadsheet is stored on a network drive in a directory owned by an analyst.
- A tabular model named Research.xlsx used in PowerPivot for Excel. Research.xlsx uses MarketResearch.xlsx as one of its data sources.

Network The network contains an Active Directory forest named tailspintoys.com. The forest contains a Microsoft SharePoint Server 2013 server farm. Implementation Plans Databases Tailspin Toys plans to build a star schema data warehouse named DB1. DB1 will be loaded from several different sources and will be updated nightly to contain new sales data. DB1 will contain the following table types:

- A fact table to store transactional data, including transaction date, productID, customerID, quantity, and sales amounts.
- Dimension tables to store information about each customer, each product, each date, and each sales department user.

BI Semantic Models Tailspin Toys plans to deploy the following BI semantic models:

- A multidimensional cube named CUBE1 that will store sales data. CUBE1 will be based on DB1 and will be hosted in SQL Server Analysis Services (SSAS). CUBE1 will contain two distinct count measures named UniqueCustomers and UniqueProducts. The measures are expected to aggregate hundreds of millions of rows from DB1.
- A tabular model named SalesCommission that will contain information about sales department user quotas and commissions.
- A tabular model named Research that will contain the migrated model from Research.xlsx.
- An instance of SSAS in tabular mode named Tabular.

Planned Reports and Queries Tailspin Toys plans to implement the following reports and queries:

- Power View reports that use data from the Research model.
- Reports for each year the company recorded sales data that used the SalesCommission model.
- The reports will use the Dates_Between() and the DatesInPeriod() DAX functions in queries.

Reports that use CUBE1 that contain the following query statements:

```
01 SELECT [Measures].[UniqueCustomers] ON 0,
02 [Date].[Date].[Date] ON 1
03 FROM [CUBE1]
04 WHERE
05 [Date].[Calendar Month].[Calendar Month].&[2012]&[1]
06 SELECT [Measures].[UniqueProducts] ON 0,
07 [Date].[Date].[Date] ON 1
08 FROM [CUBE1]
09 WHERE
10 [Date].[Calendar Month].[Calendar Month].&[2012]&[1]
```

A report named SalesByCategory that uses CUBE1 and the following query statement: (Line numbers are included for reference only.)

```
01 SELECT
02 {[Measures].[SalesAmount]} on 0
03 ,{(
04 [Date].[CalendarYear].[&2012]
05 ,
06 [Product].[Categories].[Category].[Category1]
07 ,{(
08 [Product].[Categories].[Category].[Category2]
09 ,
10 [Date].[CalendarYear].[&2012]
11 )}} ON 1
12 from CUBE1
```

Self-Service Reporting Tailspin Toys plans to deploy the following self-service reports: - Reports created by sales department specialists that use CUBE1 and contain drillthroughs, maps, sparklines, and Key Performance Indicators (KPIs). The reports will be stored in a SharePoint Server document library named Library1. - Reports created by sales department managers that use the SalesCommission model. The reports will contain visualizations that show sales department users their current sales as compared to their quota. - Power Pivot models stored in a SharePoint Server document library that is configured as a PowerPivot Gallery named Gallery1.

Requirements Data Security Requirements Sales department users browsing CUBE1 must be able to view the sales data that relates to their respective customers only. Access to reports must be controlled by using SharePoint permissions.

ETL Requirements Tailspin Toys identifies the following extract, transformation, and load (ETL) requirements: - Nightly updates of DB1 must support the incremental load of dimension and fact tables on separate schedules. Fact data may be loaded before dimension data. - ETL processes must be able to update dimension attributes without losing context for historical facts. - Referential integrity between dimension and fact tables must be maintained at all times.

Cube Performance Requirements The design of CUBE1 must minimize the processing time of the UniqueCustomers and Unique Products measures. The time required to process CUBE1 each night must be minimized.

Data Refresh Requirements The Research model must be refreshed nightly without interrupting the workflow of the analyst.

QUESTION 111 You need to recommend a cube architecture for CUBE1. The solution must meet the performance requirements for CUBE1. Which two partitions should you recommend creating? Each Answer presents part of the solution. A. Partitions based on the values of the customerID column in the dimension table B. Partitions based on the values of the customerID column in the fact table C. Partitions based on the values of the productID column in the fact table D. Partitions based on the values of the productID column in the dimension table Answer: AD

QUESTION 112 You execute the SalesbyCategory report and receive the following error message: "Members, tuples, or sets must use the same hierarchies in the function." You need to ensure that the query executes successfully. Which two actions should you perform? Each Answer presents part of the solution. A. Move the Product clause from line 08 to line 10. B. Move the Date and Product clauses on line 11 to axis 0. C. Move the Date clause from line 10 to line 08. D. Move the Measures clause on line 02 to axis 1. Answer: AC

QUESTION 113 You need to implement the SalesCommission model to support the planned reports and queries. What should you do? A. Create a date table that contains only one row for each date on which a sale is recorded. B. Use the existing transaction date column in the sales table for date calculations. C. Create a date table that contains a row for every date since data started being recorded. D. Create a new calculated date column in the sales table for date calculations. Answer: C

QUESTION 114 You need to deploy a solution for the planned self-service reports that will be used by the sales department managers. What is the best solution you should deploy? More than one answer choice may achieve the goal. Select the BEST answer. A. A filter B. A KPI C. A calculated column D. A measure Answer: B

QUESTION 115 You need to recommend a partitioning strategy that meets the performance requirements for CUBE1. What should you include in the recommendation? A. Create separate measure groups for each distinct count measure. B. Create one measure group for all distinct count measures. C. Create a separate dimension for each distinct count attribute. D. Create one dimension for all distinct count attributes. Answer: A

QUESTION 116 You need to prepare the infrastructure for the planned implementation of Gallery1. Which three actions should you perform? Each Answer presents part of the solution, A. Install a Database Engine instance. B. Run the PowerPivot Configuration Tool. C. Install the SQL Server Reporting Services add-in for SharePoint. D. Install SQL Server PowerPivot for SharePoint. E. Install the SQL Server Reporting Services - SharePoint feature. F. Run the Install-SPUserSolution cmdlet. Answer: BCD

QUESTION 117 You are deploying the Research model. You need to ensure that the data contained in the model can be refreshed. What should you do? A. Import MarketResearch.xlsx to a new tabular database on the Tabular instance. B. Assign the Tabular instance service account permissions to the MarketResearch.xlsx network location. C. Create a SQL Server Integration Services (SSIS) package that imports data from MarketResearch.xlsx nightly. Load the data to CUBE1. D. Upload MarketResearch.xlsx to Library1. Answer: B

QUESTION 118 You need to configure Library1 to support the planned self-service reports. What is the best configuration you should add to Library1? More than one answer choice may achieve the goal. Select the BEST answer. A. The Report Builder report content type B. The PowerPivot Gallery Document content type C. The Report Builder Model content type D. The Report content type Answer: A

QUESTION 119 You need to recommend a SQL

Server Integration Services (SSIS) package design that meets the ETL requirements. What should you include in the recommendation? A. Add new rows for changes to existing dimension members and enable inferred members. B. Update non-key attributes in the dimension tables to use new values. C. Update key attributes in the dimension tables to use new values. D. Add new rows for changes to existing dimension members and disable inferred members. Answer: A QUESTION 120 You need to recommend a solution to implement the data security requirements for CUBE1. Which three actions should you recommend performing in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions	Answer Area
Create a factless fact table.	
Create a perspective.	
Write an MDX expression.	
Enable Visual Totals.	
Create a SQL Server login.	

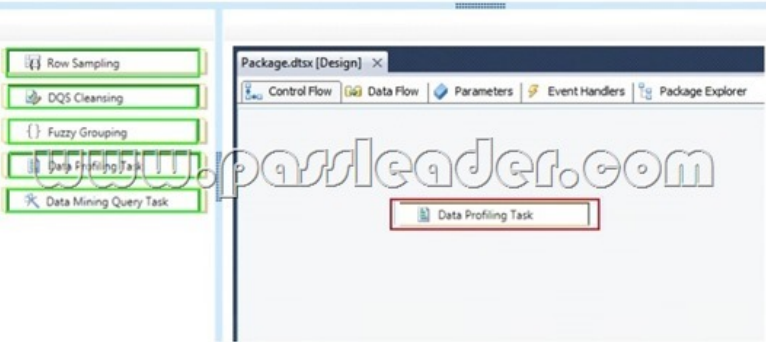
Answer:

Actions	Answer Area
Create a factless fact table.	Create a SQL Server login.
Create a perspective.	Write an MDX expression.
	Enable Visual Totals.

QUESTION 121 A company has a data warehouse that is rapidly increasing in size. You plan to improve query performance by partitioning a very large fact table. You need to identify the columns that are good partitioning candidates and the range that each partition should include. Which component should you add to the package? To answer, drag the appropriate component from the list of items to the answer area.


Available Components	Answer Area
Row Sampling	
DQS Cleansing	
Fuzzy Grouping	
Data Mining Task	
Data Mining Query Task	

Answer:



QUESTION 122 A company runs SQL Server Database Engine and SQL Server Reporting Services (SSRS) in native mode. Reports are based on data that is cached in multiple shared datasets. Source data is archived each day at midnight for regulatory compliance purposes. The shared datasets may continue to cache data that should not be used in reports. Shared report schedules are often paused during nightly server maintenance windows. Reports must not return archived data. You need to create a fully automated solution to ensure that reports do not deliver archived data. What should you do? (More than one answer choice may achieve the goal. Select the BEST answer.) A. Write a script that calls the flushcache method to clear individual items from the SSRS cache. Create a SQL Server Agent job that runs rs.exe with the script as an input file, and schedule the job to run every day after the archive process completes. B. Create a SQL Server Agent job that uses a Transact-SQL (T-SQL) step to truncate the dbo. ExecutionCache table in the ReportServerTempDB database. Schedule the job to run every day after the archive process completes. C. Create a SQL Server Agent job that restarts the SQL Server Reporting Services service. Schedule the job to run every day after the archive process completes. D. Create a shared schedule. Configure the datasets to expire on the shared schedule. Answer: A

QUESTION 123 You are the administrator of a SQL Server Integration Services (SSIS) catalog. You have access to the original password that was used to create the SSIS catalog. A full database backup of the SSISDB database on the production server is made each day. The server used for disaster recovery has an operational SSIS catalog. Information that is encrypted in the SSISDB database must not be lost. You need to document the disaster recovery solution that restores the production SSIS catalog to the disaster recovery server. Which three steps should you document in sequence? (To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.



Answer:



QUESTION 124 You are setting role permissions for a SQL Server Analysis Service (SSAS) cube. You plan to grant Read access for specific attribute members only to specific Active Directory security groups. You need to enter a Multidimensional Expressions (MDX) expression to configure the attribute member access. On which tab should you enter the MDX expression? To answer, select the appropriate tab in the answer area.



Answer:



QUESTION 125 You have a SQL Server Reporting Services (SSRS) 2008 instance integrated with an Internet-facing Microsoft SharePoint 2010 farm. PerformancePoint Services and Excel Services are installed and configured. You need to upgrade to SSRS 2012 with the minimum downtime and hardware requirements. What should you do? (More than one answer choice may achieve the goal. Select the BEST answer.) A. Migrate SSRS content to a separate instance of SSRS 2012 on the same server. B. Perform an in-place upgrade of the SSRS environment. C. Migrate SSRS content to new servers that run SSRS 2012. D. Install SSRS 2012 on a new SharePoint farm and migrate content to the new farm. Answer: C



<http://www.passleader.com/70-467.html> **QUESTION 126** You are redesigning a SQL Server Analysis Services (SSAS) database that contains a cube named Sales. Before the initial deployment of the cube, partition design was optimized for processing time. The cube currently includes five partitions named FactSales1 through FactSales5. Each partition contains from 1 million to 2 million rows. A. Use the Usage-Based Optimization Wizard to create aggregations for the current partitions. B. Use the Aggregation Design Wizard to create aggregations for the current partitions. C. Combine all the partitions into a single partition. Use the Usage-Based Optimization Wizard to create aggregations. D. Combine all the partitions into a single partition. Use the Aggregation Design Wizard to create aggregations. Answer: A **QUESTION 127** You have a SQL Server 2012 database named DB1. You have a backup device named Device1. You discover that the log file for the database is full. You need to ensure that DB1 can complete transactions. The solution must not affect the chain of log sequence numbers (LSNs). Which code segment should you execute? A. BACKUP LCG DB1 TO Device1 WITH COPY_ONLY B. BACKUP LOG DB1 TO Device1 C. BACKUP LOG DB1 TO Device1 WITH NCRECCVERY D. BACKUP LOG D31 TO Device1 WITH TRUNCATE ONLY Answer: B **QUESTION 128** A group of report writers develop reports. The report writers currently use Report Builder 1.0 and report models. The report writers are not trained on SQL Server query writing. Their report assignments come from various departments within the organization. The company is upgrading to SQL Server 2012 with SQL Server Reporting Services (SSRS) in native mode. After the upgrade, reports will retrieve data from a large new data warehouse that will reside on an instance of the SQL Server 2012 Database Engine. Reports for each department will retrieve data from only a subset of the data warehouse tables. You are designing the data access strategy. You have the following requirements: - Ensure that report writers can create only reports that display data which they have database permissions to view. - Minimize the effort required to update connection strings for all reports developed by the report writers. - Minimize the number of fields visible in a specific report for report writers who work for multiple departments. You need to design the data access strategy. What should you do? (More than one answer choice may achieve the goal. Select the BEST answer.) A. Create report models by using SQL Server Data Tools (SSDT). Include data sources and data source views for each department's data requirements. Configure each data source to use integrated security. B. Create one shared data source that uses integrated security. Create an embedded dataset for each report. C. Create one shared data source that uses integrated security. Create one shared dataset that includes all tables required by the report writers. D. Create one shared data source that uses integrated security. Create a shared dataset for each department's data requirements. Answer: D **QUESTION 129** You have a database for a mission-critical web application. The database is stored on a SQL Server 2012 instance and is the only database on the instance. The application generates all T-SQL statements dynamically and does not use stored procedures. You need to maximize the amount of memory available for data caching. Which advanced server option should you modify? A. scan for Startup Procs B. Allow Triggers to Fire Others C. Enable Contained Databases

D. Optimize for Ad hoc Workloads Answer: C QUESTION 130 You have a server that has SQL Server 2012 installed. You need to identify which parallel execution plans are running in serial. Which tool should you use?
A. Data Profile Viewer B. Database Engine Tuning Advisor
C. Performance Monitor D. Extended Events Answer: D QUESTION 131 You have a database named database1. Database developers report that there are many deadlocks. You need to implement a solution to monitor the deadlocks. The solution must meet the following requirements: - Support real-time monitoring. - Be enabled and disabled easily. - Support querying of the monitored data. What should you implement? More than one answer choice may achieve the goal. Select the BEST answer. A. an Extended Events session B. a SQL Server Profiler template C. log errors by using trace flag 1204 D. log errors by using trace flag 1222 Answer: A QUESTION 132 You are designing a SQL Server 2012 Integration Services (SSIS) deployment strategy. You currently have many SQL Server 2008 SSIS packages that require upgrading. The production environment includes SSIS 2012 and SSIS 2008. The environment includes existing command shell scripts that invoke the dtutil command-line utility. You need to design a deployment strategy that supports existing deployment strategies and requires the minimum amount of effort. What should you do? (More than one answer choice may achieve the goal. Select the BEST answer.) A. Use a project deployment model. Change the command shell scripts to reference the SQL Server 2012 path to dtutil. B. Use a package deployment model. Use the Integration Services Deployment Wizard. C. Use a package deployment model. Change the command shell scripts to reference the SQL Server 2012 path to dtutil. D. Use a project deployment model. Use the Integration Services Deployment Wizard. Answer: C QUESTION 133 You are creating a table named Orders. You need to ensure that every time a new row is added to the Orders table, a user-defined function is called to validate the row before the row is added to the table. What should you use? More than one answer choice may achieve the goal. Select the BEST answer. A. a FOREIGN KEY constraint B. a data manipulation language (DML) trigger C. a DEFAULT constraint D. a CHECK constraint E. a Data Definition Language (DDL) trigger Answer: D QUESTION 134 You plan to design an application that temporarily stores data in a SQL Azure database. You need to identify which types of database objects can be used to store data for the application. The solution must ensure that the application can make changes to the schema of a temporary object during a session. Which type of objects should you identify? A. common table expressions (CTEs) B. table variables C. temporary tables D. temporary stored procedures Answer: C QUESTION 135 You have a text file that contains an XML Schema Definition (XSD). You have a table named Schemal.Table1. You have a stored procedure named Schemal.Proc1 that accepts an XML parameter named Param1. You need to store validated XML data in Schemal.Table1. The solution must ensure that only valid XML data is accepted by Param1. What should you do? (Each correct answer presents part of the solution. Choose all that apply.) A. Define an XML column in Table1 by using an XML schema collection. B. Create an XML schema collection in the database from the text file. C. Declare Param1 var1 as type XML and associate the variable to the XML schema collection. D. use the modify method to insert the XML schema into each row of the XML column in Table1. Answer: ABD QUESTION 136 You are developing a SQL Server Analysis Services (SSAS) tabular project. The model includes a table named DimEmployee. The table contains employee details, including the sales territory for each employee. The table also defines a column named EmployeeAlias which contains the Active Directory Domain Services (AD DS) domain and logon name for each employee. You create a role named Employees. You need to configure the Employees roles so that users can query only sales orders for their respective sales territory. What should you do? A. Add a row filter that implements the LOOKUPVALUE and CUSTOMDATA functions. B. Add a row filter that implements only the USERNAME function. C. Add a row filter that implements the LOOKUPVALUE and USERNAME functions. D. Add a row filter that implements only the CUSTOMDATA function. Answer: C QUESTION 137 You are developing a SQL Server Analysis Services (SSAS) tabular project. A column named City must be added to the table named Customer. The column will be used in the definition of a hierarchy. The City column exists in the Geography table that is related to the Customer table. You need to add the City column to the Customer table. How should you write the calculation? A. =RELATED(Geography[City]) B. City:=LOOKUPVALUE(Geography[City],Geography[GeographyKey],[GeographyKey]) C. City:=VALUES(Geography[City]) D. =VALUES(Geography[City]) E. City:=RELATED(Geography.City) F. City:=

LOCKUP(Geography[City],Geography[GeographyKey],[GeographyKey]) Answer: A QUESTION 138 You are developing a SQL Server Analysis Services (SSAS) tabular project. You need to grant the minimum permissions necessary to enable users to query data in a data model. Which role permission should you use? A. Browser B. Process C. Read D. Administrator E. Explorer F. Select Answer: C QUESTION 139 You have a server that has SQL Server 2012 installed. You need to identify which parallel execution plans are running in serial. Which tool should you use? A. Data Profile Viewer B. Database Engine Tuning Advisor C. Performance Monitor D. Extended Events Answer: D QUESTION 140 You have a database named database1. Database developers report that there are many deadlocks. You need to implement a solution to monitor the deadlocks. The solution must meet the following requirements: - Support real-time monitoring. - Be enabled and disabled easily. - Support querying of the monitored data. What should you implement? More than one answer choice may achieve the goal. Select the BEST answer. A. an Extended Events session B. a SQL Server Profiler template C. log errors by using trace flag 1204 D. log errors by using trace flag 1222 Answer:

A

Compare And Choose The Best **PassLeader** 70-467 Brain Dumps

Pass4sure	PL PassLeader®	TEST KING
	 Leader of IT Certifications	
		
Banned By Microsoft Not Available	191 Q&As Price: \$99.99 Coupon Code -- CELEB	50 Q&As Price: \$124.99

<http://www.passleader.com/70-467.html>