

[28/Mar/2019 Updated Pass 70-767 Exam By Training PassLeader Free 70-767 Study Materials

New Updated 70-767 Exam Questions from PassLeader 70-767 PDF dumps! Welcome to download the newest PassLeader 70-767 VCE dumps: <https://www.passleader.com/70-767.html> (350 Q&As) Keywords: 70-767 exam dumps, 70-767 exam questions, 70-767 VCE dumps, 70-767 PDF dumps, 70-767 practice tests, 70-767 study guide, 70-767 braindumps, Implementing a Data Warehouse using SQL Exam P.S. New 70-767 dumps PDF:

https://drive.google.com/open?id=0B-ob6L_QjGLpcXZXWU14dHhIUvk >> New 70-761 dumps PDF:

https://drive.google.com/open?id=0B-ob6L_QjGLpaEZzRVFnOE9OenM >> New 70-762 dumps PDF:

https://drive.google.com/open?id=0B-ob6L_QjGLpN3RVQ25sVUM5dkU >> New 70-764 dumps PDF:

https://drive.google.com/open?id=0B-ob6L_QjGLpN3N6eHJ6Z2EzZWc >> New 70-765 dumps PDF:

https://drive.google.com/open?id=0B-ob6L_QjGLpZHIHSG5KM09xUms >> New 70-768 dumps PDF:

https://drive.google.com/open?id=0B-ob6L_QjGLpeXAXaUJkWEZnVIU NEW QUESTION 333 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. Answer: A Explanation: We use ALTER TABLE to partition the table.

Incorrect: Not 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. NEW QUESTION 334 You have a database named DB1. You need to track auditing data for four tables in DB1 by using change data capture. Which stored procedure should you execute first? A. catalog.deploy_project B. catalog.restore_project C. catalog.stop_operation D. sys.sp_cdc_add_job E. sys.sp_cdc_change_job F. sys.sp_cdc_disable_db Answer: D

Explanation: Because the cleanup and capture jobs are created by default, the sys.sp_cdc_add_job stored procedure is necessary only when a job has been explicitly dropped and must be recreated. Note: sys.sp_cdc_add_job creates a change data capture cleanup or capture job in the current database. A cleanup job is created using the default values when the first table in the database is enabled for change data capture. A capture job is created using the default values when the first table in the database is enabled for change data capture and no transactional publications exist for the database. When a transactional publication exists, the transactional log reader is used to drive the capture mechanism, and a separate capture job is neither required nor allowed. Note: sys.sp_cdc_change_job. <https://docs.microsoft.com/en-us/sql/relational-databases/track-changes/track-data-changes-sql-server>

NEW QUESTION 335 You have a data warehouse named DW1. All data files are located on drive E. You expect queries that pivot hundreds of millions of rows for each report. You need to modify the data files to minimize latency. What should you do? A. Add more data files to DW1 on drive E. B. Add more data files to tempdb on drive E. C. Remove data files from tempdb. D. Remove data files from DW1. Answer: B Explanation: The number of files depends on the number of (logical) processors on the machine. As a general rule, if the number of logical processors is less than or equal to eight, use the same number of data files as logical processors. If the number of logical processors is greater than eight, use eight data files and then if contention continues, increase the number of data files by multiples of 4 until the contention is reduced to acceptable levels or make changes to the workload/code.

<https://docs.microsoft.com/en-us/sql/relational-databases/databases/tempdb-database> NEW QUESTION 336 You are the administrator of a Microsoft SQL Server Master Data Services (MDS) model. The model was developed to provide consistent and validated snapshots of master data to the ETL processes by using subscription views. A new model version has been created. You need to ensure that the ETL processes retrieve the latest snapshot of master data. What should you do? A. Add a version flag to the new version, and create new subscription views that use this version flag. B. Create new subscription views for the new version. C. Update the subscription views to use the new version. D. Update the subscription views to use the last committed version. Answer: A Explanation: When a

version is ready for users or for a subscribing system, you can set a flag to identify the version. You can move this flag from version to version as needed. Flags help users and subscribing systems identify which version of a model to use.

<https://docs.microsoft.com/en-us/sql/master-data-services/versions-master-data-services>

NEW QUESTION 337 You are developing a Microsoft SQL Server Master Data Services (MDS) solution. The model contains an entity named Product. The Product entity has three user-defined attributes named Category, Subcategory, and Price, respectively. You need to ensure that combinations of values stored in the Category and Subcategory attributes are unique. What should you do?

A. Create an attribute group that consists of the Category and Subcategory attributes. Publish a business rule for the attribute group. B. Publish a business rule that will be used by the Product entity. C. Create a derived hierarchy based on the Category and Subcategory attributes. Use the Category attribute as the top level for the hierarchy. D. Set the value of the Attribute Type property for the Category and Subcategory attributes to Domain-based. Answer: B Explanation: In Master Data Services, business rule actions are the consequence of business rule condition evaluations. If a condition is true, the action is initiated. The Validation action "must be unique": The selected attribute must be unique independently or in combination with defined attributes. Incorrect: Not A: In Master Data Services, attribute groups help organize attributes in an entity. When an entity has lots of attributes, attribute groups improve the way an entity is displayed in the Master Data Manager web application. Not C: A Master Data Services derived hierarchy is derived from the domain-based attribute relationships that already exist between entities in a model. Not D: In Master Data Services, a domain-based attribute is an attribute with values that are populated by members from another entity. You can think of a domain-based attribute as a constrained list; domain-based attributes prevent users from entering attribute values that are not valid. To select an attribute value, the user must pick from a list.

<https://docs.microsoft.com/en-us/sql/master-data-services/business-rule-actions-master-data-services>

NEW QUESTION 338 You create a Master Data Services (MDS) model that manages the master data for a Product dimension. The Product dimension has the following properties: - All the members of the Product dimension have a product type, a product subtype, and a unique product name. - Each product has a single product type and a single product subtype. - The product type has a one-to-many relationship to the product subtype. You need to ensure that the relationship between the product name, the product type, and the product subtype is maintained when products are added to or updates in the database. What should you add to the model?

A. a subscription view B. a derived hierarchy C. a recursive hierarchy D. an explicit hierarchy Answer: B Explanation: A Master Data Services derived hierarchy is derived from the domain-based attribute relationships that already exist between entities in a model. You can create a derived hierarchy to highlight any of the existing domain-based attribute relationships in the model. Incorrect: Not C: In Master Data Services, a recursive hierarchy is a derived hierarchy that includes a recursive relationship. A recursive relationship exists when an entity has a domain-based attribute based on the entity itself. Not D: In Master Data Services, an explicit hierarchy organizes members from a single entity in any way you specify. The structure can be ragged and unlike derived hierarchies, explicit hierarchies are not based on domain-based attribute relationships.

<https://docs.microsoft.com/en-us/sql/master-data-services/derived-hierarchies-master-data-services>

NEW QUESTION 339 You are loading data from an OLTP database to a data warehouse. The database contains a table named Sales. Sales contains details of records that have a type of refund and records that have a type of sales. The data warehouse design contains a table for sales data and a table for refund data. Which component should you use to load the data to the warehouse? A. the Slowly Changing Dimension transformation B. the Conditional Split transformation C. the Merge transformation D. the Data Conversion transformation E. an Execute SQL task F. the Aggregate transformation G. the Lookup transformation Answer: B Explanation: The Conditional Split transformation can route data rows to different outputs depending on the content of the data. The implementation of the Conditional Split transformation is similar to a CASE decision structure in a programming language. The transformation evaluates expressions, and based on the results, directs the data row to the specified output. This transformation also provides a default output, so that if a row matches no expression it is directed to the default output.

<https://docs.microsoft.com/en-us/sql/integration-services/data-flow/transformations/conditional-split-transformation>

NEW QUESTION 340 You are developing a Microsoft SQL Server Integration Services (SSIS) package. You need to cleanse a data flow source by removing duplicate records based on approximate matches. Which SSIS Toolbox item should you use?

A. CDC Control task B. CDC Splitter C. Union All D. XML task E. Fuzzy Grouping F. Merge G. Merge Join Answer: E Explanation: The Fuzzy Grouping transformation performs data cleaning tasks

by identifying rows of data that are likely to be duplicates and selecting a canonical row of data to use in standardizing the data.

Incorrect: Not C: UNION ALL. Incorporates all rows into the results. This includes duplicates. If not specified, duplicate rows are removed. <https://docs.microsoft.com/en-us/sql/integration-services/data-flow/transformations/fuzzy-grouping-transformation>

NEW QUESTION 341 You are developing a Microsoft SQL Server Integration Services (SSIS) package. The package design consists of two differently structured sources in a single data flow. The Sales source retrieves sales transactions from a SQL Server database, and the Product source retrieves product details from an XML file. You need to combine the two data flow sources into a single output dataset. Which SSIS Toolbox item should you use? A. CDC Control task

B. CDC Splitter C. Union All D. XML task

E. Fuzzy Grouping F. Merge G. Merge Join Answer: G

Explanation: The Merge Join transformation provides an output that is generated by joining two sorted datasets using a FULL, LEFT, or INNER join. For example, you can use a LEFT join to join a table that includes product information with a table that lists the country/region in which a product was manufactured. The result is a table that lists all products and their country/region of origin. <https://docs.microsoft.com/en-us/sql/integration-services/data-flow/transformations/merge-join-transformation> NEW

QUESTION 342 You are developing a Microsoft SQL Server Integration Services (SSIS) package. You need to ensure that the package records the current Log Sequence Number (LSN) in the source database before the package begins reading source tables.

Which SSIS Toolbox item should you use? A. CDC Control task B. CDC Splitter

C. Union All D. XML task E. Fuzzy Grouping

F. Merge G. Merge Join Answer: A Explanation: The CDC Control task is used

to control the life cycle of change data capture (CDC) packages. It handles CDC package synchronization with the initial load package, the management of Log Sequence Number (LSN) ranges that are processed in a run of a CDC package.

<https://docs.microsoft.com/en-us/sql/integration-services/control-flow/cdc-control-task> NEW QUESTION 343 Hotspot You are a data warehouse developer. You need to create a Microsoft SQL Server Integration Services (SSIS) catalog on a production SQL Server instance. Which features are needed? (To answer, select the appropriate options in the answer area.)

Answer Area

Fea

CLR

Automatic SSIS

in Memory OL

Answer:

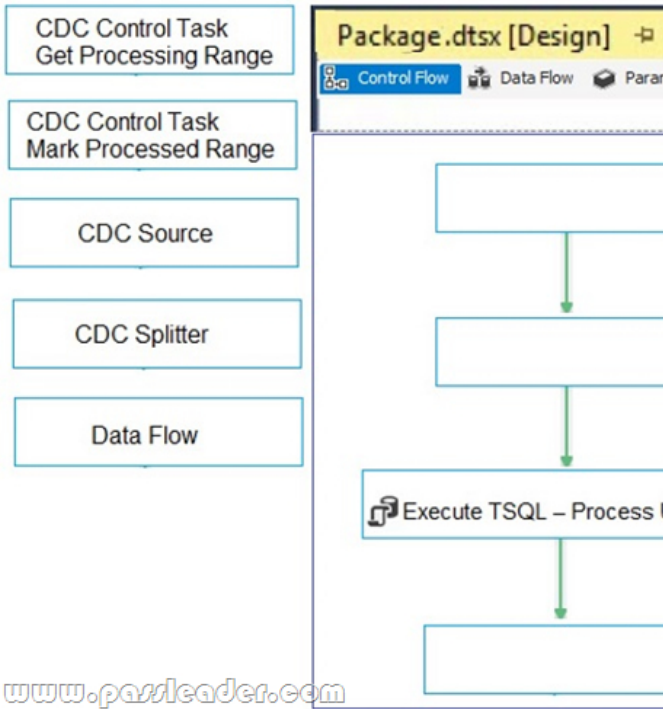
Answer Area

| Feature | Yes | No |
|----------------------------------|----------------------------------|----------------------------------|
| CLR | <input checked="" type="radio"/> | <input type="radio"/> |
| Automatic SSIS Package Execution | <input checked="" type="radio"/> | <input type="radio"/> |
| in Memory OLTP | <input type="radio"/> | <input checked="" type="radio"/> |

Explanation: Box 1: Yes. "Enable CLR Integration" must be selected because the catalog uses CLR stored procedures. Box 2: Yes. Once you have selected the "Enable CLR Integration" option, another checkbox will be enabled named "Enable automatic execution of Integration Services stored procedure at SQL Server startup". Click on this check box to enable the catalog startup stored procedure to run each time the SSIS server instance is restarted.

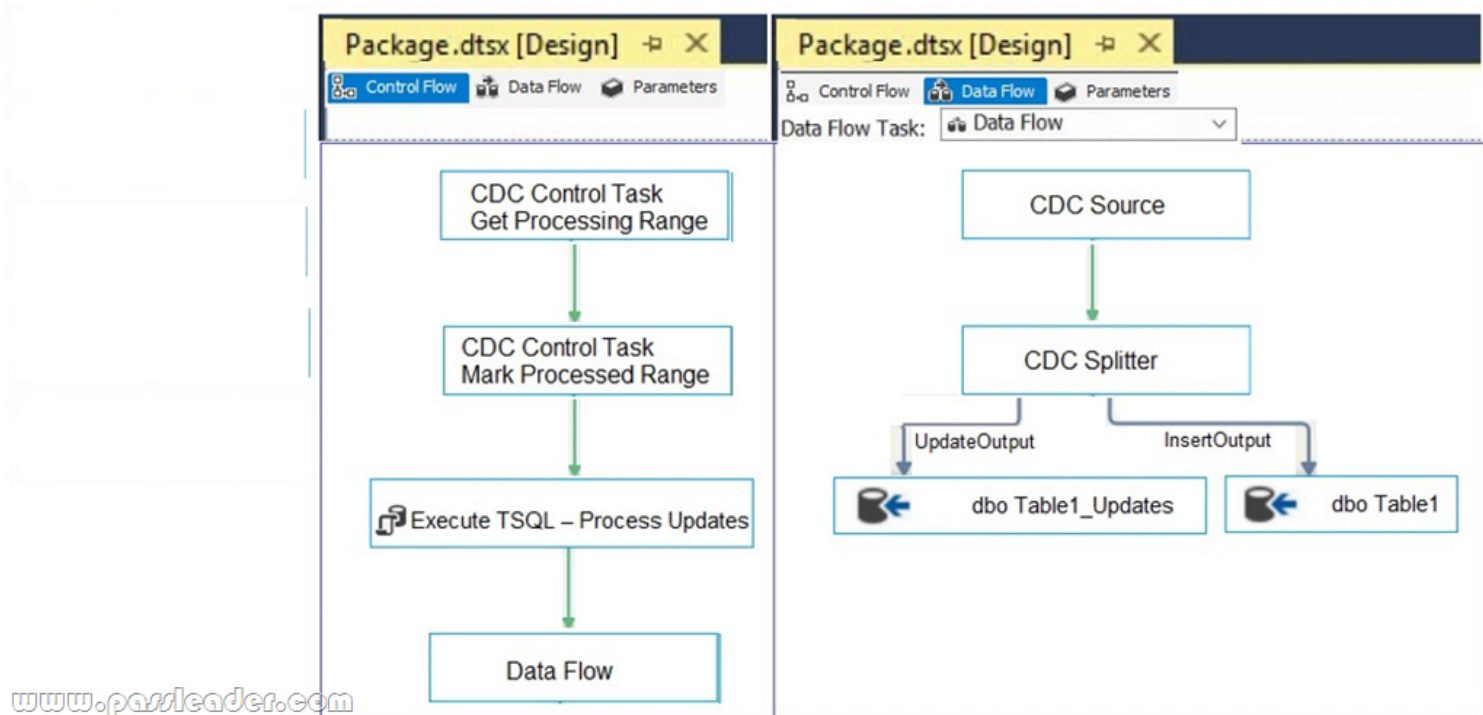
<https://www.mssqltips.com/sqlservertip/4097/understanding-the-sql-server-integration-services-catalog-and-creating-the-ssisdb-catalog/> NEW QUESTION 344 Drag and Drop You are developing a Microsoft SQL Server Integration Services (SSIS) package to incrementally load new and changed records from a data source. The SSIS package must load new records into Table1 and updated records into Table1_Updates. After loading records, the package must call a Transact-SQL statement to process updated rows according to existing business logic. You need to complete the design of the SSIS package. Which tasks should you use? (To answer, drag the appropriate SSIS objects to the correct targets. Each SSIS object 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.)

Answer Area



Answer:

Answer Area



Explanation: Step 3: The Data Flow task encapsulates the data flow engine that moves data between sources and destinations, and lets the user transform, clean, and modify data as it is moved. Addition of a Data Flow task to a package control flow makes it possible for the package to extract, transform, and load data. Step 4: The CDC source reads a range of change data from SQL Server 2017 change tables and delivers the changes downstream to other SSIS component. Step 5: The CDC splitter splits a single flow of change rows from a CDC source data flow into different data flows for Insert, Update and Delete operations.

<https://docs.microsoft.com/en-us/sql/integration-services/control-flow/cdc-control-task>

<https://docs.microsoft.com/en-us/sql/integration-services/control-flow/data-flow-task>

<https://docs.microsoft.com/en-us/sql/integration-services/data-flow/cdc-splitter?view=sql-server-2017> NEW QUESTION 345

A database has tables named Table1, Table2, and Table3: - Table1 has a foreign key relationship with Table2. - Table2 has a foreign key relationship with Table3. - Table1 does not have a direct relationship with Table3. You need to recommend an appropriate dimension usage relationship. What should you recommend? A. many-to-one relationship B. referenced relationship C. regular dimension relationship D. fact relationship

Answer: B Explanation: A reference dimension relationship between a cube dimension and a measure group exists when the key column for the dimension is joined indirectly to the fact table through a key in another dimension table. Incorrect: Not C: A regular dimension relationship between a cube dimension and a measure group exists when the key column for the dimension is joined directly to the fact table.

<https://docs.microsoft.com/en-us/sql/analysis-services/multidimensional-models-olap-logical-cube-objects/dimension-relationships>

NEW QUESTION 346 You manage Master Data Services (MDS). You plan to create entities and attributes and load them with the data. You also plan to match data before loading it into Data Quality Services (DQS). You need to recommend a solution to perform the actions. What should you recommend? A. MDS Add-in for Microsoft Excel B. MDS Configuration Manager C. Data Quality Matching D. MDS Repository Answer: A Explanation: In the Master Data Services Add-in for Excel, matching functionality is provided by Data Quality Services (DQS). This functionality must be enabled to be used: 1. To enable Data Quality Services integration. 2. Open Master Data Services Configuration Manager. 3. In the left pane, click Web Configuration. 4. On the Web Configuration page, select the website and web application. 5. In the Enable DQS Integration section, click Enable integration with Data Quality Services. 6. On the confirmation dialog box, click OK.

<https://docs.microsoft.com/en-us/sql/master-data-services/install-windows/enable-data-quality-services-integration-with-master-data-services> NEW QUESTION 347

..... Download the newest PassLeader 70-767 dumps from passleader.com now! 100% Pass

Guarantee! 70-767 PDF dumps & 70-767 VCE dumps: <https://www.passleader.com/70-767.html> (350 Q&As) (New Questions Are 100% Available and Wrong Answers Have Been Corrected! Free VCE simulator!) P.S. New 70-767 dumps PDF:
https://drive.google.com/open?id=0B-ob6L_QjGLpcXZXWU14dHhIUVk >> New 70-761 dumps PDF:
https://drive.google.com/open?id=0B-ob6L_QjGLpaEZzRVFnOE9OenM >> New 70-762 dumps PDF:
https://drive.google.com/open?id=0B-ob6L_QjGLpN3RVQ25sVUM5dkU >> New 70-764 dumps PDF:
https://drive.google.com/open?id=0B-ob6L_QjGLpN3N6eHJ6Z2EzZWc >> New 70-765 dumps PDF:
https://drive.google.com/open?id=0B-ob6L_QjGLpZHIHSG5KM09xUms >> New 70-768 dumps PDF:
https://drive.google.com/open?id=0B-ob6L_QjGLpeXAxaUJkWEZnVIU