

Get 1 year Of Updated Microsoft DP-700 Exam Question Dumps

Answer Area

```
SELECT last_run_start_time, last_run_command
FROM
    queryinsights.exec_requests_history
    queryinsights.exec_sessions_history
    queryinsights.frequently_run_queries
    queryinsights.long_running_queries
WHERE last_run_total_elapsed_time_ms > 7200000
AND
    max_run_total_elapsed_time_ms > 7200000
    median_total_elapsed_time_ms > 7200000
    number_of_canceled_runs > 1
    number_of_failed_runs > 1
    number_of_runs > 1
```

P.S. Free 2026 Microsoft DP-700 dumps are available on Google Drive shared by Prep4King: <https://drive.google.com/open?id=1AUqEkomPvWDdwMDKO7s-fXYpvr6R3W-x>

A second format is a DP-700 web-based practice exam that can take for self-assessment. However, it differs from desktop-based DP-700 practice exam software as it can be taken via any browser, including Chrome, Firefox, Safari, and Opera. This Microsoft DP-700 web-based practice exam does not require any other plugins. You can take this DP-700 self-assessment test on Windows, iOS, Linux, Mac, and Android. It also includes all of the functionalities of desktop DP-700 software and will assist you in passing the DP-700 certification test.

The second format is a web-based format that can be accessed from browsers like Firefox, Microsoft Edge, Chrome, and Safari. It means you don't need to download or install any software or plugins to take the Implementing Data Engineering Solutions Using Microsoft Fabric practice test. The web-based format of the Microsoft DP-700 Certification Exams practice test supports all operating systems. The third and last format is desktop software format which can be accessed after installing the software on your Implementing Data Engineering Solutions Using Microsoft Fabric (DP-700) Windows PC or Laptop. These formats are built especially for the students so they don't stop preparing for the Implementing Data Engineering Solutions Using Microsoft Fabric (DP-700) certification.

>> DP-700 Training Tools <<

2026 DP-700 Training Tools - Trustable Microsoft DP-700 Test King: Implementing Data Engineering Solutions Using Microsoft Fabric

If you are troubled with DP-700 exam, you can consider down our free demo. You will find that our latest DP-700 exam torrent are perfect paragon in this industry full of elucidating content for exam candidates of various degree to use. Our results of latest DP-700 exam torrent are startlingly amazing, which is more than 98 percent of exam candidates achieved their goal successfully. The latest DP-700 Exam Torrent covers all the qualification exam simulation questions in recent years, including the corresponding matching materials at the same time.

Microsoft DP-700 Exam Syllabus Topics:

Topic	Details

Topic 1	<ul style="list-style-type: none"> Implement and manage an analytics solution: This section of the exam measures the skills of Microsoft Data Analysts regarding configuring various workspace settings in Microsoft Fabric. It focuses on setting up Microsoft Fabric workspaces, including Spark and domain workspace configurations, as well as implementing lifecycle management and version control. One skill to be measured is creating deployment pipelines for analytics solutions.
Topic 2	<ul style="list-style-type: none"> Monitor and optimize an analytics solution: This section of the exam measures the skills of Data Analysts in monitoring various components of analytics solutions in Microsoft Fabric. It focuses on tracking data ingestion, transformation processes, and semantic model refreshes while configuring alerts for error resolution. One skill to be measured is identifying performance bottlenecks in analytics workflows.
Topic 3	<ul style="list-style-type: none"> Ingest and transform data: This section of the exam measures the skills of Data Engineers that cover designing and implementing data loading patterns. It emphasizes preparing data for loading into dimensional models, handling batch and streaming data ingestion, and transforming data using various methods. A skill to be measured is applying appropriate transformation techniques to ensure data quality.

Microsoft Implementing Data Engineering Solutions Using Microsoft Fabric Sample Questions (Q44-Q49):

NEW QUESTION # 44

You have a Fabric workspace that contains a lakehouse named Lakehousel. Lakehousel contains a table named Status_Target that has the following columns:

- * Key
- * Status
- * LastModified

The data source contains a table named Status.Source that has the same columns as Status_Target. Status.Source is used to populate Status_Target. In a notebook name Notebook!, you load Status_Source to a DataFrame named sourceDF and Status_Target to a DataFrame named targetDF. You need to implement an incremental loading pattern by using Notebook-!. The solution must meet the following requirements:

- * For all the matching records that have the same value of key, update the value of LastModified in Status_Target to the value of LastModified in Status_Source.
- * Insert all the records that exist in Status_Source that do NOT exist in Status_Target.
- * Set the value of Status in Status_Target to inactive for all the records that were last modified more than seven days ago and that do NOT exist in Status.Source.

How should you complete the statement? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer:

Explanation:

NEW QUESTION # 45

HOTSPOT

You have a Fabric workspace.

You are debugging a statement and discover the following issues:

Sometimes, the statement fails to return all the expected rows.

The PurchaseDate output column is NOT in the expected format of mmm dd, yy.

You need to resolve the issues. The solution must ensure that the data types of the results are retained. The results can contain blank cells.

How should you complete the statement? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer:

Explanation:

Explanation:

Topic 1, Litware, Inc Case Study

This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other questions in this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next section of the exam. After you begin a new section, you cannot return to this section.

To start the case study

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an All Information tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a question, click the Question button to return to the question.

Overview

Litware, Inc. is a publishing company that has an online bookstore and several retail bookstores worldwide.

Litware also manages an online advertising business for the authors it represents.

Existing Environment. Fabric Environment

Litware has a Fabric workspace named Workspace1. High concurrency is enabled for Workspace1.

The company has a data engineering team that uses Python for data processing.

Existing Environment. Data Processing

The retail bookstores send sales data at the end of each business day, while the online bookstore constantly provides logs and sales data to a central enterprise resource planning (ERP) system.

Litware implements a medallion architecture by using the following three layers: bronze, silver, and gold. The sales data is ingested from the ERP system as Parquet files that land in the Files folder in a lakehouse.

Notebooks are used to transform the files in a Delta table for the bronze and silver layers. The gold layer is in a warehouse that has V-Order disabled.

Litware has image files of book covers in Azure Blob Storage. The files are loaded into the Files folder.

Existing Environment. Sales Data

Month-end sales data is processed on the first calendar day of each month. Data that is older than one month never changes.

In the source system, the sales data refreshes every six hours starting at midnight each day.

The sales data is captured in a Dataflow Gen1 dataflow. When the dataflow runs, new and historical data is captured. The dataflow captures the following fields of the source:

Sales Date

Author

Price

Units

SKU

A table named AuthorSales stores the sales data that relates to each author. The table contains a column named AuthorEmail.

Authors authenticate to a guest Fabric tenant by using their email address.

Existing Environment. Security Groups

Litware has the following security groups:

Sales

Fabric Admins

Streaming Admins

Existing Environment. Performance Issues

Business users perform ad-hoc queries against the warehouse. The business users indicate that reports against the warehouse sometimes run for two hours and fail to load as expected. Upon further investigation, the data engineering team receives the following error message when the reports fail to load: "The SQL query failed while running." The data engineering team wants to debug the issue and find queries that cause more than one failure.

When the authors have new book releases, there is often an increase in sales activity. This increase slows the data ingestion process. The company's sales team reports that during the last month, the sales data has NOT been up-to-date when they arrive at work in the morning.

Requirements. Planned Changes

Litware recently signed a contract to receive book reviews. The provider of the reviews exposes the data in Amazon Simple Storage Service (Amazon S3) buckets.

Litware plans to manage Search Engine Optimization (SEO) for the authors. The SEO data will be streamed from a REST API.

Requirements. Version Control

Litware plans to implement a version control solution in Fabric that will use GitHub integration and follow the principle of least privilege.

Requirements. Governance Requirements

To control data platform costs, the data platform must use only Fabric services and items. Additional Azure resources must NOT be provisioned.

Requirements. Data Requirements

Litware identifies the following data requirements:

Process the SEO data in near-real-time (NRT).

Make the book reviews available in the lakehouse without making a copy of the data.

When a new book cover image arrives in the Files folder, process the image as soon as possible.

NEW QUESTION # 46

You need to resolve the sales data issue. The solution must minimize the amount of data transferred.

What should you do?

- A. Spilt the dataflow into two dataflows.
- B. **Configure incremental refresh for the dataflow. Set Refresh rows from the past to 1 Month.**
- C. Configure scheduled refresh for the dataflow.
- D. Configure incremental refresh for the dataflow. Set Refresh rows from the past to 1 Year.
- E. Configure incremental refresh for the dataflow. Set Store rows from the past to 1 Month.

Answer: B

Explanation:

The sales data issue can be resolved by configuring incremental refresh for the dataflow. Incremental refresh allows for only the new or changed data to be processed, minimizing the amount of data transferred and improving performance.

The solution specifies that data older than one month never changes, so setting the refresh period to 1 Month is appropriate. This ensures that only the most recent month of data will be refreshed, reducing unnecessary data transfers.

Topic 2, Contoso, Ltd Case Study

This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other questions in this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next section of the exam. After you begin a new section, you cannot return to this section.

To start the case study

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an All Information tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a question, click the Question button to return to the question.

Overview. Company Overview

Contoso, Ltd. is an online retail company that wants to modernize its analytics platform by moving to Fabric.

The company plans to begin using Fabric for marketing analytics.

Overview. IT Structure

The company's IT department has a team of data analysts and a team of data engineers that use analytics systems.

The data engineers perform the ingestion, transformation, and loading of data. They prefer to use Python or SQL to transform the data.

The data analysts query data and create semantic models and reports. They are qualified to write queries in Power Query and T-SQL.

Existing Environment. Fabric

Contoso has an F64 capacity named Cap1. All Fabric users are allowed to create items.

Contoso has two workspaces named WorkspaceA and WorkspaceB that currently use Pro license mode.

Existing Environment. Source Systems

Contoso has a point of sale (POS) system named POS1 that uses an instance of SQL Server on Azure Virtual Machines in the same Microsoft Entra tenant as Fabric. The host virtual machine is on a private virtual network that has public access blocked. POS1 contains all the sales transactions that were processed on the company's website.

The company has a software as a service (SaaS) online marketing app named MAR1. MAR1 has seven entities. The entities contain data that relates to email open rates and interaction rates, as well as website interactions. The data can be exported from MAR1 by calling REST APIs. Each entity has a different endpoint.

Contoso has been using MAR1 for one year. Data from prior years is stored in Parquet files in an Amazon Simple Storage Service (Amazon S3) bucket. There are 12 files that range in size from 300 MB to 900 MB and relate to email interactions.

Existing Environment. Product Data

POS1 contains a product list and related data. The data comes from the following three tables:

Products

ProductCategories

ProductSubcategories

In the data, products are related to product subcategories, and subcategories are related to product categories.

Existing Environment. Azure

Contoso has a Microsoft Entra tenant that has the following mail-enabled security groups:

DataAnalysts: Contains the data analysts

DataEngineers: Contains the data engineers

Contoso has an Azure subscription.

The company has an existing Azure DevOps organization and creates a new project for repositories that relate to Fabric.

Existing Environment. User Problems

The VP of marketing at Contoso requires analysis on the effectiveness of different types of email content. It typically takes a week to manually compile and analyze the data. Contoso wants to reduce the time to less than one day by using Fabric.

The data engineering team has successfully exported data from MAR1. The team experiences transient connectivity errors, which causes the data exports to fail.

Requirements. Planned Changes

Contoso plans to create the following two lakehouses:

Lakehouse1: Will store both raw and cleansed data from the sources

Lakehouse2: Will serve data in a dimensional model to users for analytical queries. Additional items will be added to facilitate data ingestion and transformation.

Contoso plans to use Azure Repos for source control in Fabric.

Requirements. Technical Requirements

The new lakehouses must follow a medallion architecture by using the following three layers: bronze, silver, and gold. There will be extensive data cleansing required to populate the MAR1 data in the silver layer, including deduplication, the handling of missing values, and the standardizing of capitalization.

Each layer must be fully populated before moving on to the next layer. If any step in populating the lakehouses fails, an email must be sent to the data engineers.

Data imports must run simultaneously, when possible.

The use of email data from the Amazon S3 bucket must meet the following requirements:

Minimize egress costs associated with cross-cloud data access.

Prevent saving a copy of the raw data in the lakehouses.

Items that relate to data ingestion must meet the following requirements:

The items must be source controlled alongside other workspace items.

Ingested data must land in the bronze layer of Lakehouse1 in the Delta format.

No changes other than changes to the file formats must be implemented before the data lands in the bronze layer.

Development effort must be minimized and a built-in connection must be used to import the source data.

In the event of a connectivity error, the ingestion processes must attempt the connection again.

Lakehouses, data pipelines, and notebooks must be stored in WorkspaceA. Semantic models, reports, and dataflows must be stored in WorkspaceB.

Once a week, old files that are no longer referenced by a Delta table log must be removed.

Requirements. Data Transformation

In the POS1 product data, ProductID values are unique. The product dimension in the gold layer must include only active products from product list. Active products are identified by an IsActive value of 1.

Some product categories and subcategories are NOT assigned to any product. They are NOT analytically relevant and must be omitted from the product dimension in the gold layer.

Requirements. Data Security

Security in Fabric must meet the following requirements:

The data engineers must have read and write access to all the lakehouses, including the underlying files.

The data analysts must only have read access to the Delta tables in the gold layer.

The data analysts must NOT have access to the data in the bronze and silver layers.

The data engineers must be able to commit changes to source control in WorkspaceA.

NEW QUESTION # 47

You have a Fabric workspace that contains a warehouse named Warehouse1.

In Warehouse1, you create a table named DimCustomer by running the following statement.

□

You need to set the Customerkey column as a primary key of the DimCustomer table.

Which three code segments should you run in sequence? To answer, move the appropriate code segments from the list of code segments to the answer area and arrange them in the correct order.

Answer:

Explanation:

NEW QUESTION # 48

You have a Fabric F32 capacity that contains a workspace. The workspace contains a warehouse named DW1 that is modelled by using MD5 hash surrogate keys.

DW1 contains a single fact table that has grown from 200 million rows to 500 million rows during the past year.

You have Microsoft Power BI reports that are based on Direct Lake. The reports show year-over-year values.

Users report that the performance of some of the reports has degraded over time and some visuals show errors.

You need to resolve the performance issues. The solution must meet the following requirements:

Provide the best query performance.

Minimize operational costs.

Which should you do?

- A. Change the MD5 hash to SHA256.
- B. **Modify the surrogate keys to use a different data type.**
- C. Increase the capacity.
C Enable V-Order
- D. Create views.

Answer: B

Explanation:

In this case, the key issue causing performance degradation likely stems from the use of MD5 hash surrogate keys. MD5 hashes are 128-bit values, which can be inefficient for large datasets like the 500 million rows in your fact table. Using a more efficient data type for surrogate keys (such as integer or bigint) would reduce the storage and processing overhead, leading to better query performance. This approach will improve performance while minimizing operational costs because it reduces the complexity of querying and indexing, as smaller data types are generally faster and more efficient to process.

NEW QUESTION # 49

.....

Being the most competitive and advantageous company in the market, our DP-700 practice quiz have help tens of millions of exam candidates realize their dreams all these years. If you are the dream-catcher, we are willing to offer help with our DP-700 Study Guide like always. And if you buy our DP-700 exam materials, then you will find that passing the exam is just a piece of cake in front of you.

DP-700 Test King: <https://www.prep4king.com/DP-700-exam-prep-material.html>

- DP-700 Exam Labs □ DP-700 Exam Labs □ DP-700 Accurate Test □ Enter ➤ www.dumpsmaterials.com □ and search for “DP-700” to download for free □ Detailed DP-700 Study Dumps
- DP-700 New Questions □ DP-700 Certification Test Answers □ DP-700 Reliable Exam Pass4sure □ Search for □ DP-700 □ and download it for free on ➤ www.pdfvce.com □ website □ Test DP-700 Guide
- 2026 Microsoft DP-700 Realistic Training Tools Free PDF Quiz □ Search for “DP-700” and download it for free on (www.practicevce.com) website □ Latest DP-700 Exam Cost
- Latest DP-700 Exam Cost □ DP-700 New APP Simulations □ Exam DP-700 Success □ Search on ➤ www.pdfvce.com ↳ for [DP-700] to obtain exam materials for free download □ DP-700 Reliable Exam Pass4sure
- Test DP-700 Guide □ Test DP-700 Cram Pdf □ DP-700 New Questions □ Simply search for ➡ DP-700 □ □ □ for free download on [www.examcollectionpass.com] □ DP-700 Reliable Test Notes
- Avail Latest DP-700 Training Tools to Pass DP-700 on the First Attempt □ The page for free download of ➤ DP-700 □ on [www.pdfvce.com] will open immediately □ DP-700 Certification Test Answers
- DP-700 Certification Exam Dumps □ Exam DP-700 Success □ Vce DP-700 Free □ Open □ www.examcollectionpass.com □ enter { DP-700 } and obtain a free download □ Latest DP-700 Exam Cost
- Test DP-700 Cram Pdf □ Vce DP-700 Free □ DP-700 New APP Simulations □ Search for □ DP-700 □ and

download exam materials for free through ➡ www.pdfvce.com ☐ ☐ DP-700 Exam Labs

DOWNLOAD the newest Prep4King DP-700 PDF dumps from Cloud Storage for free: <https://drive.google.com/open?id=1AUqEkomPvWDdwMDKO7s-fXYpvr6R3W-x>