

DP-700試験問題解説集: Implementing Data Engineering Solutions Using Microsoft Fabric一発合格-簡単DP-700的中関連問題



さらに、GoShiken DP-700ダンプの一部が現在無料で提供されています: https://drive.google.com/open?id=1hr5O-OtF1oL_bhqYZl5bC9EOuCEiS00

GoShikenは、魅力的なキャラクターで世界中の試験受験者を招きます。当社の専門家は彼らの卓越性に大きく貢献しました。したがって、試験をシミュレートするDP-700が最良であると率直に言うことができます。DP-700学習教材のコンテンツを作成する取り組みは、学習ガイドの開発につながり、完成度を高めます。そのため、模擬試験は間違いなくレビューの耐久性を高めています。関心を集め、いくつかの難しい点を簡素化するために、当社の専門家は、DP-700試験の合格に役立つように、DP-700学習教材の設計に最善を尽くしています。

Microsoft DP-700 認定試験の出題範囲:

トピック	出題範囲
トピック 1	<ul style="list-style-type: none">分析ソリューションの実装と管理: このセクションでは、Microsoft Fabric における様々なワークスペース設定の構成に関する Microsoft データアナリストのスキルを評価します。Spark やドメインワークスペースの構成を含む Microsoft Fabric ワークスペースの設定、ライフサイクル管理とバージョン管理の実装に重点が置かれます。評価対象となるスキルの一つに、分析ソリューションのデプロイメントパイプラインの作成があります。
トピック 2	<ul style="list-style-type: none">分析ソリューションの監視と最適化: このセクションでは、Microsoft Fabric の分析ソリューションのさまざまなコンポーネントを監視するデータアナリストのスキルを評価します。データの取り込み、変換プロセス、セマンティックモデルの更新を追跡し、エラー解決のためのアラートを設定することに重点が置かれます。評価対象となるスキルの一つは、分析ワークフローにおけるパフォーマンスのボトルネックを特定することです。

トピック 3

- データの取り込みと変換: このセクションでは、データエンジニアのスキル、特にデータ読み込みパターンの設計と実装能力を評価します。特に、多次元モデルへの読み込みのためのデータ準備、バッチおよびストリーミングデータの取り込み処理、そして様々な手法を用いたデータ変換能力に重点が置かれます。評価対象となるスキルの一つは、データ品質を確保するために適切な変換手法を適用することです。

>> DP-700試験問題解説集 <<

試験の準備方法-正確的なDP-700試験問題解説集試験-認定するDP-700の中関連問題

GoShikenは多くの受験生を助けて彼らにMicrosoftのDP-700試験に合格させることができるのは我々専門的なチームがMicrosoftのDP-700試験を研究して解答を詳しく分析しますから。試験が更新されているうちに、我々はMicrosoftのDP-700試験の資料を更新し続けています。できるだけ100%の通過率を保証使用しています。

Microsoft Implementing Data Engineering Solutions Using Microsoft Fabric 認定 DP-700 試験問題 (Q104-Q109):

質問 # 104

You have an Azure SQL database named DB1.

In a Fabric workspace, you deploy an eventstream named EventStreamDB1 to stream record changes from DB1 into a lakehouse. You discover that events are NOT being propagated to EventStreamDB1.

You need to ensure that the events are propagated to EventStreamDB1.

What should you do?

- A. Enable change data capture (CDC) for DB1.
- B. Create an Azure Stream Analytics job.
- C. Enable Extended Events for DB1.
- D. Create a read-only replica of DB1.

正解: A

質問 # 105

You have a Fabric workspace that contains a lakehouse named Lakehouse1. Lakehouse1 contains a Delta table named Table1.

You analyze Table1 and discover that Table1 contains 2,000 Parquet files of 1 MB each.

You need to minimize how long it takes to query Table1.

What should you do?

- A. Run the OPTIMIZE and VACUUM commands.
- B. Disable V-Order and run the OPTIMIZE command.
- C. Disable V-Order and run the VACUUM command.

正解: A

解説:

Problem Overview:

Solution:

Commands and Their Roles:

- Compacts small Parquet files into larger files to improve query performance.
- It supports optional features like V-Order, which organizes data for efficient scanning.
- Removes old, unreferenced data files and metadata from the Delta table.
- Running VACUUM after OPTIMIZE ensures unnecessary files are cleaned up, reducing storage overhead and improving performance.

質問 # 106

You need to recommend a solution for handling old files. The solution must meet the technical requirements. What should you include in the recommendation?

- A. a data pipeline that includes a Delete data activity
- B. a data pipeline that includes a Copy data activity
- C. a notebook that runs the OPTIMIZE command
- D. a notebook that runs the VACUUM command

正解: D

解説:

Topic 1, 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.

質問 # 107

You are building a data orchestration pattern by using a Fabric data pipeline named Dynamic Data Copy as shown in the exhibit. (Click the Exhibit tab.)

□ Dynamic Data Copy does NOT use parametrization.

You need to configure the ForEach activity to receive the list of tables to be copied.

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

NOTE: Each correct selection is worth one point.

正解:

解説:

Explanation:

質問 # 108

You have a Fabric workspace that contains a semantic model named Model1. You need to monitor the refresh history of Model1 and visualize the refresh history in a chart. What should you use?

- A. the refresh history from the settings of Modell.
- **B. a notebook**
- C. a Dataflow Gen2 dataflow
- D. a data pipeline

正解： B

質問 #109

弊社GoShikenでのMicrosoftのDP-700問題集を購入する予定のあるお客様は何の質問があれば、ライブチャットといい、メールといい、我々の社員は待っていて質問を回復します。当然、購入した後、あなたはどんな疑問があると、次々に丁寧に返答できます。あなたの送信を歓迎しております。あなたに行き届いたサービスを提供できるのは我々の幸いです。

DP-700的中関連問題: <https://www.goshiken.com/Microsoft/DP-700-mondaishu.html>

2026年GoShikenの最新DP-700 PDFダンプおよびDP-700試験エンジンの無料共有: https://drive.google.com/open?id=1hr5O-OtF1oL_bhqYZl5bC9EOuCEiS00