EXAMTOPICS

-Expert Verified, Online, **Free**.



CERTIFICATION TEST

-CertificationTest.net - Cheap & Quality Resources With Best Support

Question #1 *Topic 1*

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.

You need to ensure that the data analysts can access the gold layer lakehouse.

What should you do?

A. Add the DataAnalyst group to the Viewer role for WorkspaceA.

B. Share the lakehouse with the DataAnalysts group and grant the Build reports on the default semantic model permission.

C. Share the lakehouse with the DataAnalysts group and grant the Read all SQL Endpoint data permission.

D. Share the lakehouse with the DataAnalysts group and grant the Read all Apache Spark permission.

Suggested Answer: *C*

Community vote distribution

C (100%)

QAZdbarhate12345678 Highly Voted 1 year, 11 months ago

By granting Read all SQL Endpoint data permission, the analysts get the necessary and sufficient access to query the gold layer data while adhering to the principle of least privilege.Hi QAZdbarhate12345678 I have 2 questions please 1. Which lakehouse is Answer C referring to: Lakehouse1 or Lakehouse2? 2. The case study overview says "The data analysts must NOT have access to the data in the bronze and silver layers." If we share the lakehouse with the DataAnalysts group will this not give the data analysts access to the data in the bronze and silver layers? Thank youAnswering much time after, but it refer to Lakehouse2 or lakehouse3. The Lakehouse will be the Lakehouse with the gold layer.

upvoted 14 times

QAZdbarhate12345678 Highly Voted 1 year, 11 months ago

Selected Answer: C

By granting Read all SQL Endpoint data permission, the analysts get the necessary and sufficient access to query the gold layer data while adhering to the principle of least privilege.

upvoted 14 times

Hanif53 Most Recent 5 months, 3 weeks ago

Selected Answer: C

C. Share the lakehouse with the DataAnalysts group and grant the Read all SQL Endpoint data permission.

upvoted 2 times

oluonimole 1 year, 8 months ago

Hi QAZdbarhate12345678 I have 2 questions please 1. Which lakehouse is Answer C referring to: Lakehouse1 or Lakehouse2? 2. The case study overview says "The data analysts must NOT have access to the data in the bronze and silver layers." If we share the lakehouse with the DataAnalysts group will this not give the data analysts access to the data in the bronze and silver layers?

Thank you

upvoted 2 times

Maggiee 6 months ago

Selected Answer: C

Viewer role provides read-only access to Delta tables.

Analysts won't inherit permissions to bronze/silver if those are isolated.

upvoted 1 times

juamd 6 months, 1 week ago

Selected Answer: C

The system as a whole will follow the medallion architecture, bronze and silver will be in lakehouse1 and gold will be in lakehouse2, so the data analysts must have access only to this one and only though SQL

upvoted 2 times

Josekonoor 8 months, 3 weeks ago

Selected Answer: C

read all sql end point

upvoted 1 times

Josekonoor 8 months, 3 weeks ago

Selected Answer: C

Requirement - "The Data Analysts must have access to Delta Tables" ->

C. Share the lakehouse with the DataAnalysts group and grant the Read all SQL Endpoint data permission.

upvoted 1 times

jjveliz 10 months ago

Selected Answer: C

-allows the data analysts to read the data from the Delta tables in the gold layer using the SQL endpoint

-It also ensures they do not have access to the underlying files or other layers (bronze and silver)

-The SQL endpoint in a lakehouse is designed for querying the Delta tables without needing access to the file system

upvoted 1 times

8ac3742 1 year, 3 months ago

Selected Answer: C

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

upvoted 1 times

Manish0427 1 year, 5 months ago

Selected Answer: C

Option C is the best choice among given options. Though Option A is also correct as all the 3 options of Item level control is valid only when users are viewers (and not admin, members or contributors), but only Option A is not sufficient. When a role is defined as a viewer, it automatically gets the "Read all SQL endpoint data" access as the readonly access to all the lakehouse data using SQL endpoints only. "Read all Apache Spark" provides access to the underlying files also and not only the Gold layer datawarehouse tables. "Build reports on the default semantic model" provides access to semantic model for building power bi reports which is not the requirement in this case. Since the Bronze and Silver layer is in Lakehouse 1 and the Gold Layer is in Lakehouse 2 as a Datawarehouse, the Data Analyst role will be provided to Lakehouse 2 only.

upvoted 1 times

ogbenisho 1 year, 5 months ago

Selected Answer: C

Data Analysts group don't need access to Workspace A since they do not have any business with Bronze and Silver layers. The Gold layer is in Workspace B, but since that is not the only item in the workspace, giving them workspace view access is does not obey the least privilege's rule. So, they only need to access Lakehouse 2 and granting the Read all SQL Endpoint data permission is the answer to that.

upvoted 1 times

Wipag 1 year, 5 months ago

Selected Answer: C

This question centers on controlled access to data layers in Fabric, specifically:

Data analysts must access only the gold layer in the lakehouse.

They must not access bronze or silver layers.

They need to read data in the gold layer (i.e., for analytics).

SQL-based access is the right fit for data analysts (T-SQL and Power BI reports).

The gold layer is typically exposed through the SQL analytics endpoint of the lakehouse.

upvoted 1 times

SilvanoRamalho 1 year, 5 months ago

Selected Answer: C

O acesso ao modelo semântico é muito limitado. O requisito especifica as Tabelas Delta da Camada Dourada upstream (ou seja, no Lakehouse).

upvoted 1 times

DarioReymago 1 year, 5 months ago

Selected Answer: C

we share the gold layer lakehouse

upvoted 1 times

rickiebn_binss 1 year, 5 months ago

Selected Answer: C

Thank you

upvoted 1 times

ne4 1 year, 6 months ago

Selected Answer: C

Semantic model access is too limited. Requirement specifies the upstream Gold Layer Delta Tables (i.e. in the Lakehouse)

upvoted 2 times

LasAnsias 1 year, 9 months ago

Selected Answer: C

Sharing the lakehouse with the DataAnalysts group and granting the Read all SQL Endpoint data permission would allow them to query the Delta tables in the gold layer using the SQL endpoint. This aligns with the requirement that they only have read access to the gold layer.

upvoted 2 times

oluonimole 1 year, 8 months ago

Hi LasAnsias. I have 2 questions please?

1. Which lakehouse is Answer C referring to: Lakehouse1 or Lakehouse2?

2. The case study overview says "The data analysts must NOT have access to the data in the bronze and silver layers." If we share the lakehouse with the DataAnalysts group will this not give the data analysts access to the data in the bronze and silver layers?

Thank you

upvoted 1 times

mixonfreddy 1 year, 11 months ago

Selected Answer: C

Answer is C

upvoted 2 times

LasAnsias 1 year, 9 months ago

Sharing the lakehouse with the DataAnalysts group and granting the Read all SQL Endpoint data permission would allow them to query the Delta tables in the gold layer using the SQL endpoint. This aligns with the requirement that they only have read access to the gold layer.Hi LasAnsias. I have 2 questions please? 1. Which lakehouse is Answer C referring to: Lakehouse1 or Lakehouse2? 2. The case study overview says "The data analysts must NOT have access to the data in the bronze and silver layers." If we share the lakehouse with the DataAnalysts group will this not give the data analysts access to the data in the bronze and silver layers? Thank you

upvoted 2 times

oluonimole 1 year, 8 months ago

Hi QAZdbarhate12345678 I have 2 questions please

1. Which lakehouse is Answer C referring to: Lakehouse1 or Lakehouse2?

2. The case study overview says "The data analysts must NOT have access to the data in the bronze and silver layers." If we share the lakehouse with the DataAnalysts group will this not give the data analysts access to the data in the bronze and silver layers?

Thank you

upvoted 2 times

Gostosao 1 year ago

Answering much time after, but it refer to Lakehouse2 or lakehouse3. The Lakehouse will be the Lakehouse with the gold layer.

upvoted 2 times

Question #2 *Topic 1*

You have a Fabric workspace.
You have semi-structured data.
You need to read the data by using T-SQL, KQL, and Apache Spark. The data will only be written by using Spark.
What should you use to store the data?

A. a lakehouse

B. an eventhouse

C. a datamart

D. a warehouse

Suggested Answer: *A*

Community vote distribution

A (52%) B (48%)

38578c4 Highly Voted 1 year, 9 months ago

KQL is avaialble in eventhouse only https://learn.microsoft.com/en-us/fabric/fundamentals/decision-guide-data-storeA is still correct because it say's "Read"I think both A and B could be a correct answer for this. If my data is in an eventhouse, I can query it using T-SQL, KQL and PySpark If my data is in a Laekehouse, I can query it using SQL and PySpark and Create shortcuts for it in the eventhouse then query it using KQL. I guess this question needs some other precision to only one possible correct answerquestion also said -- data will only be 'written' by spark..

upvoted 37 times

38578c4 Highly Voted 1 year, 9 months ago

Selected Answer: B

KQL is avaialble in eventhouse only https://learn.microsoft.com/en-us/fabric/fundamentals/decision-guide-data-store

upvoted 37 times

alqazqzy.asseel Most Recent 1 month, 1 week ago

Selected Answer: A

Write via Apache Spark, so u must use lakehouse

upvoted 1 times

Bs2009 3 months ago

Selected Answer: A

correct answer is A because data is written "The data will only be written by using Spark" so in event house primary write engine is KQL .

upvoted 3 times

shmmini 1 year, 7 months ago

I think both A and B could be a correct answer for this. If my data is in an eventhouse, I can query it using T-SQL, KQL and PySpark If my data is in a Laekehouse, I can query it using SQL and PySpark and Create shortcuts for it in the eventhouse then query it using KQL. I guess this question needs some other precision to only one possible correct answer

upvoted 10 times

shmmini 1 year, 7 months ago

I think both A and B could be a correct answer for this. If my data is in an eventhouse, I can query it using T-SQL, KQL and PySpark If my data is in a Laekehouse, I can query it using SQL and PySpark and Create shortcuts for it in the eventhouse then query it using KQL. I guess this question needs some other precision to only one possible correct answerquestion also said -- data will only be 'written' by spark..

upvoted 10 times

Maggiee 6 months ago

Selected Answer: A

Lakehouse in Microsoft Fabric supports:

Semi-structured data

Reading data using:

T-SQL

KQL

Apache Spark

Writing data using Spark

Delta Lake storage format

upvoted 2 times

juamd 6 months, 2 weeks ago

Selected Answer: A

The only restriction to the lakehouses while reading the question is about the need to use KQL, but with the use of shortcuts in a KQL database, you can query your data in a lakehouse without moving it. Just create the shortcut and point it to the data.

upvoted 3 times

Mhunity 6 months, 4 weeks ago

Selected Answer: B

B - KQL only available in eventhouse

upvoted 1 times

279312 7 months ago

Selected Answer: A

Lakehouse is correct considering - Due to additional point in the question - The data will only be written by using Spark.

If we consider other points, others are possible in eventhouse too

upvoted 2 times

realuv 7 months ago

Selected Answer: A

semi-structured data

written apache spark

upvoted 1 times

Yasminmv110 7 months ago

Selected Answer: A

https://learn.microsoft.com/en-us/fabric/real-time-intelligence/eventhouse-as-endpoint

The only Fabric storage that satisfies all three requirements is a Lakehouse.

upvoted 1 times

loisistor 7 months, 1 week ago

Selected Answer: B

KQL read operations is only available in Eventhouse

upvoted 1 times

32950cb 7 months, 1 week ago

Selected Answer: B

Eventhouse is the only that can read KQL

upvoted 1 times

zulajariwala 8 months, 1 week ago

Selected Answer: A

Semi structured data.

upvoted 1 times

Josekonoor 8 months, 3 weeks ago

Selected Answer: A

lakehouse

upvoted 1 times

Josekonoor 8 months, 3 weeks ago

Selected Answer: A

semi structured - lakehouse

upvoted 1 times

Josekonoor 8 months, 3 weeks ago

Selected Answer: A

Semi Structured Data

upvoted 1 times

chinhuy 8 months, 4 weeks ago

Selected Answer: A

be simple bro

upvoted 1 times

pa1d 8 months, 4 weeks ago

Selected Answer: B

A lake-house data cannot be read from KQL without mirroring it to Event-House so option B(Event house) is the correct answer.

upvoted 1 times

zulajariwala 9 months ago

Selected Answer: A

A Lakehouse in Microsoft Fabric supports Spark for writing and allows you to read the same data using T-SQL, KQL, and Spark.

It is designed for semi-structured data (JSON, Parquet, etc.).

Eventhouse is mainly for streaming/KQL scenarios.

Warehouse and Datamart are primarily structured, SQL-focused storage and not intended for Spark-written semi-structured data.

upvoted 2 times

Prashi_31 9 months ago

Selected Answer: A

Lakehouse is the right answer

upvoted 1 times

jjveliz 10 months ago

Selected Answer: A

lakehouse is the ideal storage solution for this scenario because it is specifically designed to handle semi-structured data and supports read operations using T-SQL, KQL, and Apache Spark, while also allowing write operations exclusively with Spark.

upvoted 2 times

nelrosell 10 months, 1 week ago

Selected Answer: A

A because, data will only be 'written' by spark

upvoted 1 times

realthiagoluna 11 months, 1 week ago

Selected Answer: A

To read data using T-SQL, KQL, and Apache Spark, while ensuring the data is only written by Apache Spark, the appropriate storage solution is a lakehouse. An eventhouse, which is typically used for real-time streaming data, does not fully support this specific combination of read/write operations and data type flexibility.

upvoted 2 times

RAO_N 1 year, 1 month ago

Selected Answer: A

A lakehouse allows you to store semi-structured data in open formats, write using Spark, and read using T-SQL (through the SQL analytics endpoint), KQL (through KQL database shortcuts), and Apache Spark. It is specifically designed for these multi-engine scenarios.

While eventhouse does support reading with KQL and some T-SQL, it is not designed for writing or reading with Apache Spark.

upvoted 1 times

yo1233 1 year, 2 months ago

Selected Answer: A

Written only with spark

Readable via T-SQL, KQL, and spark

Semi structured data

upvoted 1 times

RG_123 1 year, 2 months ago

Selected Answer: A

Answer is A:

We should consider the fact that it is a semi structured data and not the real time data. We use Eventhouse when we have a streaming/real time data. So for Semi-Structured data, we will store it in a lakehouse and create a short for it in eventhouse. In this way, we can read the data stored in lakehouse using KQL.https://learn.microsoft.com/en-us/fabric/fundamentals/decision-guide-data-store (Here it is clearly mentioned that if you want to store the semi structure data, use Lakehouse and for streaming data, use eventhouse)

upvoted 2 times

Thatoyaka 1 year, 2 months ago

Selected Answer: A

A Lakehouse provides the flexibility to write with Spark and query with T-SQL, KQL, and Spark, making it the best fit for your scenario.

upvoted 2 times

Thatoyaka 1 year, 2 months ago

Selected Answer: A

Eventhouse in Microsoft Fabric is not designed to be written to directly from Apache Spark.

It is optimized for ingesting real-time and high-velocity telemetry data, typically using Kusto Ingest mechanisms or built-in connectors (like Azure Event Hubs, IoT Hub, etc.).

upvoted 2 times

LuluVIP 1 year, 2 months ago

Selected Answer: B

KQL need Eventhouse

upvoted 1 times

d0bdf05 1 year, 2 months ago

Selected Answer: A

Lakehouse:

- Lakehouse supports reading with T-SQL, KQL, and Apache Spark.

- Allows writing with Spark.

- Optimized for semi-structured data like JSON, Parquet, etc.

- It exposes data via SQL and KQL endpoints and supports Spark access.

An eventhouse

- Designed for real-time event streaming.

Datamart

- Only supports T-SQL.

- Not compatible with Spark or KQL.

- Not suitable for semi-structured data.

Datawarehouse

- Designed for structured data and T-SQL.

- Does not support Spark or KQL.

- Not suitable for semi-structured data.

upvoted 1 times

8115575 1 year, 3 months ago

Selected Answer: A

eventhouse is optimized for streaming/event data, mainly used with KQL, not suitable for Spark write or T-SQL.

upvoted 1 times

AndrewsOtis 1 year, 3 months ago

Selected Answer: A

A. a lakehouse

Built on Delta Lake format.

Can be accessed with T-SQL, KQL, and Spark in Fabric.

Perfect for semi-structured + multi-engine scenarios.

B. an eventhouse

Purpose: real-time analytics with event streams.

Not designed for semi-structured data with Spark + T-SQL + KQL.

upvoted 1 times

alex_sql 1 year, 3 months ago

Selected Answer: B

KQL is avaialble in eventhouse only https://learn.microsoft.com/en-us/fabric/fundamentals/decision-guide-data-store

upvoted 1 times

Yvoe 1 year, 3 months ago

Selected Answer: A

Eventhouse is optimized for real-time event data. While it supports both KQL and Apache Spark, it is not designed for Spark-centric scenarios involving semi-structured data written exclusively using Spark.

upvoted 1 times

8ac3742 1 year, 3 months ago

Selected Answer: A

The eventhouse is not designed for semi-structured data, it is designed for columnar data such as log, you can use event stream and KQL pipeline to write data into eventhouse, but you cannot directly use Spark to write data into eventhouse.

In a Lakehouse, you can create a shortcut to its underlying Delta tables, enabling KQL to query the data directly through that shortcut.

upvoted 1 times

nirvana183 1 year, 3 months ago

Selected Answer: A

Lakehouse in Microsoft Fabric is designed to store semi-structured and unstructured data in an open format (Delta/Parquet). It supports being read by T-SQL, KQL, and Apache Spark in the same workspace. You can write data using Spark and still query it from SQL endpoints or KQL without extra ingestion steps.

upvoted 1 times

Kumar_PHBVK 1 year, 3 months ago

Selected Answer: A

For Fabric, storage account is lakehouse.

upvoted 1 times

alex_sql 1 year, 3 months ago

Selected Answer: A

La opción correcta sigue siendo A. a lakehouse, según la documentación oficial de Microsoft Fabric.

upvoted 1 times

Mahlodi 1 year, 3 months ago

Answer A :Why a lakehouse?

Supports multi-engine querying

T-SQL endpoint for SQL queries

KQL endpoint for log and telemetry analysis

Apache Spark for distributed processing

Optimized for semi-structured data

Natively handles JSON, Parquet, Delta formats

Schema-on-read lets you flexibly explore data

Spark-native writes with immediate cross-engine reads

Use Spark to ingest and write as Delta files

T-SQL and KQL can query the same tables without extra ETL

upvoted 1 times

ziggy1117 1 year, 4 months ago

Selected Answer: B

Only Eventhouse allows KQL

upvoted 2 times

YSAu 1 year, 4 months ago

Selected Answer: B

https://learn.microsoft.com/en-us/fabric/fundamentals/decision-guide-data-store

upvoted 1 times

d727c6d 1 year, 4 months ago

Selected Answer: A

A is correct

upvoted 2 times

eskimolight 1 year, 4 months ago

Selected Answer: A

I think it is A

upvoted 2 times

LuisPBI25 1 year, 4 months ago

Selected Answer: B

Lakehouse read operations available are only T-SQL and Spark. However, Eventhouse read operations permitted are T-SQL,Spark and also KQL as it is required in the question.

upvoted 2 times

Shw7 1 year, 5 months ago

Selected Answer: A

a Lakehouse is the correct answer

upvoted 1 times

VenkataPhaniPavanKumar 1 year, 5 months ago

Selected Answer: A

Here's why a lakehouse fits perfectly for your scenario:

It supports semi-structured data, such as JSON or Parquet.

It can be accessed using T-SQL, KQL, and Apache Spark—offering maximum flexibility across personas.

You mentioned data is only written using Spark, and lakehouses are optimized for Spark-based ingestion and processing.

It stores data in Delta Lake format in OneLake, making it ACID-compliant and performant across engines.

Source: Perplexity.ai and Microsoft Copilot (both says Lakehouse)

Quite surprised to understand the gap from microsoft learn and copilot

https://learn.microsoft.com/en-us/fabric/fundamentals/decision-guide-data-store

upvoted 2 times

Manish0427 1 year, 5 months ago

Selected Answer: A

Although data is only in Unstructured format and only Apache Spark is used to write, which is fulfilled by both a Lakehouse and Eventhouse. However, provided that there is no mention of streaming data, I would prefer to use Lakehouse as a data storage. Eventhouse is primarily designed for Streaming data and KQL analytics. Writing data using Apache Spark is not the ideal usecase for Eventhouse. On the other hand, Lakehouse can store unstructured data, supports native write access via Apache Spark, can read using T-SQL (using SQL endpoint) and KQL (using shortcuts to Kusto). In real world scenario, Lakehouse is the ideal choice in such a case.

upvoted 2 times

ogbenisho 1 year, 5 months ago

Selected Answer: B

Obviously, the answer is B

- T-SQL, KQL and Spark can be used on Eventhouse

- T-SQL & Spark can be used on Lakehouse

upvoted 2 times

TayaC 1 year, 5 months ago

Selected Answer: A

A. a lakehouse is correct as the constraint is "only written by using Spark"

upvoted 2 times

Wipag 1 year, 5 months ago

Selected Answer: A

A. a lakehouse

Designed to handle structured, semi-structured, and unstructured data.

Stored in Delta Lake format, making it accessible from:

Spark (native)

T-SQL via the SQL analytics endpoint

KQL via OneLake integration

Supports multi-engine access.

Ideal for big data + analytics scenarios, especially with semi-structured data like JSON, Parquet, etc.

Spark is commonly used to write to lakehouses.

upvoted 2 times

smanzana 1 year, 5 months ago

Selected Answer: A

The correct is A

upvoted 1 times

vikramkumar 1 year, 5 months ago

Selected Answer: A

A is the correct answer

upvoted 1 times

NNPRN 1 year, 5 months ago

Selected Answer: B

Eventhouse is the right answer because KQL is avaialble in eventhouse only https://learn.microsoft.com/en-us/fabric/fundamentals/decision-guide-data-store

upvoted 1 times

DarioReymago 1 year, 5 months ago

Selected Answer: A

I select A. Request dont talk about stream data

upvoted 1 times

malik777 1 year, 5 months ago

Selected Answer: A

event house meant for time series data

upvoted 1 times

MohanNaidu08 1 year, 5 months ago

Selected Answer: A

You need a storage solution that supports:

Semi-structured data

Read access via T-SQL, KQL, and Apache Spark

Write access via Spark

A lakehouse in Microsoft Fabric is designed exactly for this:

Supports semi-structured and structured data (e.g., JSON, Parquet, CSV).

Can be queried using:

T-SQL (via SQL analytics endpoint)

KQL (via Kusto query endpoint)

Apache Spark (via notebooks or jobs)

Optimized for big data processing and analytics.

Allows Spark-based writes and multi-engine reads.

upvoted 3 times

Ahmadpbi 1 year, 5 months ago

Selected Answer: A

 Why Option B (an Eventhouse) is NOT Correct

 What is an Eventhouse in Microsoft Fabric?

An Eventhouse is designed for large-scale, time-series, and event data, such as:

Logs

Telemetry

IoT device data

It is optimized for append-only workloads.

Queries are executed using KQL (Kusto Query Language).

It is NOT built to be written using Apache Spark.

It does not natively support T-SQL or Spark-based analytics.

upvoted 1 times

akmorsh 1 year, 5 months ago

Selected Answer: A

Explanation by chatGPT:

A Lakehouse in Microsoft Fabric is designed to store semi-structured data and supports multiple query languages including T-SQL, KQL, and Apache Spark.

The question specifies that the data will be written only by Spark, and needs to be read by T-SQL, KQL, and Spark, which fits perfectly with the Lakehouse architecture.

Eventhouse is designed mainly for event streaming and ingestion.

Datamart is typically used for relational, structured data with a semantic model and T-SQL querying but does not natively support Spark or KQL.

Warehouse is optimized for structured, relational data and T-SQL querying but does not support KQL or Spark.

upvoted 2 times

Madhu2023 1 year, 5 months ago

Selected Answer: B

Sorry I take that back, Read operations can be done using spark

upvoted 1 times

Madhu2023 1 year, 5 months ago

Selected Answer: A

https://learn.microsoft.com/en-us/fabric/fundamentals/decision-guide-data-store

Apache is only available with Lakehouse

upvoted 1 times

yue_young 1 year, 6 months ago

Selected Answer: B

Lakehouse doesn't support KQL for reading data.

upvoted 1 times

PBridge 1 year, 6 months ago

Selected Answer: A

A lakehouse in Microsoft Fabric supports semi-structured data formats (e.g., JSON, Parquet).

It enables access using:

T-SQL via SQL endpoint

KQL via shortcuts or integration with real-time data

Apache Spark for data engineering tasks

Since the data will only be written by Spark, and needs to be read using multiple languages, the lakehouse is the most flexible and performant storage choice.

upvoted 2 times

Kiracloud 1 year, 6 months ago

Selected Answer: B

https://learn.microsoft.com/en-us/fabric/fundamentals/decision-guide-data-store

upvoted 1 times

ZZIN 1 year, 6 months ago

Selected Answer: A

Isn't the correct answer Lakehouse since it is stated that only Spark is used to write data?

Eventhouse does not write data by Spark.

upvoted 1 times

hazz_hari 1 year, 7 months ago

Selected Answer: B

Eventhouse supports unstructured data, supports read operations using T-SQL, KQL and PySpark, and supports write operations using Spark.

Ref Article: https://learn.microsoft.com/en-us/fabric/fundamentals/decision-guide-data-store

upvoted 1 times

alirana3 1 year, 7 months ago

Selected Answer: A

 Correct Answer: A. Lakehouse

 Here's why:

Even though Eventhouse supports Spark and KQL for writing and reading, it is optimized for telemetry, time series, and log/event data. It is not designed for general-purpose Spark-based data engineering.

Your scenario describes:

Semi-structured data

Write via Spark

Read via T-SQL, Spark, and KQL

According to the guide:

Lakehouse supports:

 Spark (read/write)

 T-SQL (via SQL analytics endpoint)

 KQL (not directly supported)

Eventhouse supports:

 Spark (read/write)

 KQL

 T-SQL (limited)

BUT: It's meant for time series, telemetry/logs, not general-purpose data processing

upvoted 1 times

niteen82 1 year, 7 months ago

Selected Answer: A

Lakehouse Best fit. A lakehouse in Microsoft Fabric is designed to support semi-structured and unstructured data, allows writing via Apache Spark, and supports querying via T-SQL, KQL, and Spark.

upvoted 1 times

suhovet 1 year, 7 months ago

Selected Answer: B

KQL is avaialble in eventhouse only

upvoted 1 times

zxc01 1 year, 8 months ago

Selected Answer: B

I think eventhouse should be correct answer because it can support KQL, Spark and T-SQL. And can support semi-structured data event not good like lakehouse. The problem if we keep semi-structured data in lakehouse, we still need eventhouse to execute KQL on them. However, we just can build shortcut in KQL database to point delta tables in lakehouse(by OneLake), we cannot point to semi-structured files.

upvoted 1 times

DataConsultant 1 year, 8 months ago

Selected Answer: A

A Lakehouse in Microsoft Fabric is the best choice for storing semi-structured data while allowing access using T-SQL, KQL, and Apache Spark

upvoted 2 times

CMDev 1 year, 8 months ago

Selected Answer: B

Check this link and choose an answer, it must be B.

https://learn.microsoft.com/en-us/fabric/fundamentals/decision-guide-data-store

upvoted 2 times

Adriel_1996 1 year, 8 months ago

Selected Answer: B

Based on documentation read on eventhouse allows SQL,T-SQL and Spark.

https://learn.microsoft.com/en-us/fabric/fundamentals/decision-guide-data-store

upvoted 1 times

Pierrejul 1 year, 8 months ago

Selected Answer: B

https://learn.microsoft.com/en-us/fabric/fundamentals/decision-guide-data-store

upvoted 1 times

ayushc 1 year, 8 months ago

Selected Answer: A

By default you cannot query a lakehouse with kql but if you integrate it with Real time analytics you are able to query the lakehouse with kql

upvoted 2 times

pierre314 1 year, 8 months ago

Selected Answer: B

Eventhouse:

Read operations: KQL, Spark and T-SQL

Write operations: KQL, Spark

upvoted 1 times

JensQ 1 year, 8 months ago

Selected Answer: B

Out of the available options, only eventhouse can be read using KQL

upvoted 1 times

Pierrejul 1 year, 8 months ago

Selected Answer: A

lakehouse permit to read Spark data

We are not in KQL

upvoted 1 times

SannaLH 1 year, 8 months ago

Selected Answer: B

Like 38578c4 link says: Eventhouse support read with: KQL, T-SQL, Spark & write with: KQL, Spark, connector ecosystem

upvoted 2 times

568f95c 1 year, 9 months ago

Selected Answer: A

A Lakehouse in Microsoft Fabric is the best choice for storing semi-structured data while allowing access using T-SQL, KQL, and Apache Spark

B. Eventhouse - Designed for real-time event processing.

- Does not support T-SQL or Apache Spark writes.

C. Datamart - Built for self-service BI, supports T-SQL, but not Spark-based writing.

- Not optimized for semi-structured data.

D. Warehouse - Optimized for structured relational data (fact & dimension tables).

- Does not support KQL or Spark-based writes.

upvoted 3 times

GLG 1 year, 9 months ago

Selected Answer: A

Lakehouse: data will only be written bij using Spark. You can use shortcuts to access the data from een Eventhouse for KQL queries. If real-time processing or KQL were requitements, Eventhouse would be more appropiate.

upvoted 3 times

basene 1 year, 9 months ago

Selected Answer: A

You are deal with semi-structured data and we would like to use only spark.

upvoted 1 times

Jayjay5 1 year, 5 months ago

A is still correct because it say's "Read"

upvoted 1 times

certs4pk 11 months ago

question also said -- data will only be 'written' by spark..

upvoted 2 times

fffsssggg 1 year, 10 months ago

Selected Answer: B

Eventhouse:

Read operations: KQL, Spark and T-SQL

Write operations: KQL, Spark

https://learn.microsoft.com/en-us/fabric/get-started/decision-guide-data-store

upvoted 13 times

chimuelo69 1 year, 10 months ago

Selected Answer: A

A lakehouse in Microsoft Fabric is a versatile storage option designed for semi-structured and unstructured data, combining the flexibility of a data lake with the structure of a data warehouse. It supports Apache Spark, T-SQL, and KQL queries, making it ideal for the requirements.

Why the other options are not suitable:

B. Eventhouse: Eventhouse is not a valid Microsoft Fabric storage option.

C. Datamart: A datamart is designed for structured data and primarily supports T-SQL, but not Spark or KQL.

D. Warehouse: A warehouse is optimized for structured data and supports T-SQL, but it is not designed for semi-structured data or Spark.

A lakehouse is the best choice for storing data in a format that allows Spark to write it while enabling T-SQL and KQL for querying.

upvoted 4 times

mmanrik 1 year, 10 months ago

Selected Answer: B

The eventhouse is the only one out of the four options that allows the creation of a KQL Database, therefore it is the only option that fully meets the criteria (being able to read the date via KQL, T-SQL and Spark).

upvoted 3 times

QAZdbarhate12345678 1 year, 11 months ago

Selected Answer: A

lakehouse is the most appropriate choice for this scenario

upvoted 3 times

Tuki93 1 year, 11 months ago

Selected Answer: A

You have semi-structured data.

The data will only be written by using Spark.

upvoted 3 times

kryper_reg 1 year, 11 months ago

Selected Answer: A

Semi structured data should be stored in a lakehouse

upvoted 2 times

Question #3 *Topic 1*

You have a Fabric workspace that contains a warehouse named Warehouse1.
You have an on-premises Microsoft SQL Server database named Database1 that is accessed by using an on-premises data gateway.
You need to copy data from Database1 to Warehouse1.
Which item should you use?

A. a Dataflow Gen1 dataflow

B. a data pipeline

C. a KQL queryset

D. a notebook

Suggested Answer: *B*

Community vote distribution

B (100%)

alqazqzy.asseel Most Recent 1 month, 1 week ago

Selected Answer: B

Data pipeline using copy data activity, u can also ingest data using dataflow gen2 dataflow when you need want to transform, cleaning or shaping data while ingestion

upvoted 1 times

Waloo2_11 5 months, 3 weeks ago

Selected Answer: B

B.data pipeline, using copy data pipeline

upvoted 1 times

juamd 6 months, 2 weeks ago

Selected Answer: B

For the Fabric Warehouse the data is not external, here the copy activity acts as a bridge.

upvoted 1 times

Josekonoor 8 months, 3 weeks ago

Selected Answer: B

"Copy Data from SQL Server to a lakehouse" -> B. a data pipeline

upvoted 2 times

Thatoyaka 1 year, 2 months ago

Selected Answer: B

Why B. a data pipeline is correct:

Data Pipelines in Microsoft Fabric are used for data movement and orchestration, much like Azure Data Factory.

They support:

On-premises SQL Server as a source via the on-premises data gateway

Warehouse as a sink (destination)

Pipelines are designed for copying data between sources and targets — perfect for your scenario.

upvoted 3 times

nirvana183 1 year, 3 months ago

Selected Answer: B

Data pipeline in Microsoft Fabric

upvoted 1 times

ogbenisho 1 year, 5 months ago

Selected Answer: B

Copy data is meant to ingest data from source systems to Landing zone (Bronze layer/ Raw Layer). Copy data is a capacity from Data Pipeline.

upvoted 1 times

Wipag 1 year, 5 months ago

Selected Answer: B

Purpose-built for orchestrating data movement.

Supports Copy Data activity, which works with:

On-prem SQL Server via gateway

Fabric warehouse as a sink

Provides scheduling, retries, logging, etc.

Correct and most robust solution in Fabric for this task.

upvoted 1 times

DarioReymago 1 year, 5 months ago

Selected Answer: B

is correct

upvoted 1 times

Danielnimi 1 year, 5 months ago

Selected Answer: B

The pipeline in Fabric works same way as pipeline n Data Dactory. It supports the use of self-hosted Integrated Runtime (IR) which is used to establish connection between a private network (on-prem) and Azure cloud.

upvoted 1 times

sesav4742 1 year, 7 months ago

Selected Answer: B

A data pipeline, since it has the "Copy data" activity, the most efficient way to copy data into fabric when there are no transformations to be done

upvoted 3 times

arunmewada 1 year, 8 months ago

Selected Answer: B

I attempted to use a **data pipeline** to load data directly from **on-premises** to the **warehouse**, but since the warehouse does not support external sources, I first loaded the data into a **Lakehouse** and then moved it to the **warehouse**.

upvoted 1 times

CMDev 1 year, 8 months ago

Selected Answer: B

It says, you need to copy data, dont specify if we need to do some transformations. In this case, a data pipeline with an activity "copy data" is enough and the best choice in terms of efficency and minimal effort.

Check this link:

https://learn.microsoft.com/en-us/fabric/fundamentals/decision-guide-pipeline-dataflow-spark

upvoted 1 times

arunmewada 1 year, 8 months ago

I attempted to use a **data pipeline** to load data directly from **on-premises** to the **warehouse**, but since the warehouse does not support external sources, I first loaded the data into a **Lakehouse** and then moved it to the **warehouse**.answer is A

upvoted 1 times

arunmewada 1 year, 8 months ago

answer is A

upvoted 2 times

robertlavigne Highly Voted 1 year, 10 months ago

Selected Answer: B

B.

Notebooks can't go through a gateway, KQL isn't appropriate for a sql db. Gen1 dataflows are deprecated. Data pipelines work through a datagateway and are fast

upvoted 18 times

mmanrik 1 year, 10 months ago

Selected Answer: B

A data pipeline, since it has the "Copy data" activity, the most efficient way to copy data into fabric when there are no transformations to be done (if there are a Dataflow Gen 2 might be better, depends on the amount of data to be copied).

upvoted 2 times

QAZdbarhate12345678 1 year, 11 months ago

Selected Answer: B

• In this case, a data pipeline can be used to transfer data from Database1 (the SQL Server database) to Warehouse1 (in the Fabric workspace).

upvoted 3 times

Tuki93 1 year, 11 months ago

Selected Answer: B

Dataflow Gen 1 is deprecated.

upvoted 2 times

Question #4 *Topic 1*

You have a Fabric workspace that contains a warehouse named Warehouse1.
You have an on-premises Microsoft SQL Server database named Database1 that is accessed by using an on-premises data gateway.
You need to copy data from Database1 to Warehouse1.
Which item should you use?

A. an Apache Spark job definition

B. a data pipeline

C. a Dataflow Gen1 dataflow

D. an eventstream

Suggested Answer: *B*

Community vote distribution

B (100%)

Josekonoor Most Recent 8 months, 3 weeks ago

Selected Answer: B

a pipeline

upvoted 2 times

Josekonoor 8 months, 3 weeks ago

Selected Answer: B

Copy Data from Database1 to Warehouse1 -> B. a data pipeline

upvoted 1 times

nirvana183 1 year, 3 months ago

Selected Answer: B

Data pipeline in Microsoft Fabric

upvoted 1 times

Wipag 1 year, 5 months ago

Selected Answer: B

You are copying data from:

Source: On-premises SQL Server (Database1)

Destination: A Microsoft Fabric warehouse (Warehouse1)

Connectivity: Via on-premises data gateway

You need a solution that:

Supports copying structured data

Can connect through a gateway

Can write to a warehouse

Supports scheduling, error handling, and data flow control

upvoted 1 times

DarioReymago 1 year, 5 months ago

Selected Answer: B

otra vez?

upvoted 1 times

sesav4742 1 year, 7 months ago

Selected Answer: B

B.

Notebooks and eventstreams can't go through a gateway, . Gen1 dataflows are deprecated. Data pipelines work through a datagateway and are fast

upvoted 1 times

sesav4742 1 year, 7 months ago

Selected Answer: B

Correct Answer is B

upvoted 1 times

MSka 1 year, 7 months ago

Selected Answer: B

It's B

upvoted 2 times

robertlavigne Highly Voted 1 year, 10 months ago

Selected Answer: B

B.

Notebooks and eventstreams can't go through a gateway, . Gen1 dataflows are deprecated. Data pipelines work through a datagateway and are fast

upvoted 9 times

QAZdbarhate12345678 1 year, 11 months ago

Selected Answer: B

In this case, a data pipeline can be used to transfer data from Database1 (the SQL Server database) to Warehouse1 (in the Fabric workspace).

upvoted 4 times

Tuki93 1 year, 11 months ago

Selected Answer: B

Dataflow Gen 1 is deprecated.

upvoted 1 times

Question #5 *Topic 1*

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. Increase the capacity.

C. Enable V-Order.

D. Modify the surrogate keys to use a different data type.

E. Create views.

Suggested Answer: *D*

Community vote distribution



Tuki93 Highly Voted 1 year, 11 months ago

Selected Answer: C

Enable V-Order: This should be your first step as it directly targets performance improvements with minimal cost. Consider the following before deciding to disable V-Order: Microsoft Fabric Direct Lake mode depends on V-Order.

https://learn.microsoft.com/en-us/fabric/data-warehouse/v-order#performance-considerations

upvoted 31 times

Tuki93 Highly Voted 1 year, 11 months ago

Enable V-Order: This should be your first step as it directly targets performance improvements with minimal cost. Consider the following before deciding to disable V-Order: Microsoft Fabric Direct Lake mode depends on V-Order. https://learn.microsoft.com/en-us/fabric/data-warehouse/v-order#performance-considerationsAfter reviewing the public documentation again, I understand that V-Order is enabled by default in Data Warehouses, as the documentation only refers to how it can be disabled—and notes that this action is irreversible. https://learn.microsoft.com/en-us/fabric/data-warehouse/v-order#scenarios-where-v-order-might-not-be-beneficial At the same time, I agree with my peers who mentioned that the row limit per table is 300 million in an F32 capacity. Therefore, I'm updating my answer to option B: Increase the capacity. https://learn.microsoft.com/en-us/fabric/fundamentals/direct-lake-overview#fabric-capacity-requirementsDid you read the contents of the url document you wrote? https://learn.microsoft.com/en-us/fabric/data-warehouse/v-order#performance-considerations "Currently, disabling V-Order can only be done at the warehouse level, and it is irreversible: once disabled, it cannot be enabled again. Users must consider the performance if they choose to Disable V-Order in Fabric Warehouse."Would V-Order not already be enabled by default?is V-order automatically enabled on warehouse? No, V-Order is not automatically enabled on a warehouse. You would need to enable it manually to take advantage of its benefits for improved query performance and data compression.V-order is enabled by default.Only when you are saving in optimized format .. otherwise its not enabled by default.I still have the same doubt.What about the limit on the rows limitation (https://learn.microsoft.com/en-us/fabric/fundamentals/direct-lake-overview) ?

upvoted 32 times

Nancy Examtopics Most Recent 2 weeks ago

Selected Answer: C

The expected optimization is often V-Order, because it specifically optimizes data storage for analytical reads without requiring a larger capacity. B. Increase the capacity

Would likely improve performance.

However, it increases operational cost, violating the requirement to minimize costs.

upvoted 1 times

yk2085 1 month, 1 week ago

Selected Answer: C

I will opt in for Enable V-Order

upvoted 1 times

alqazqzy.asseel 1 month, 1 week ago

Selected Answer: D

MD5 hash string length is 32 characters hexadecimal, so when the table grows to 500 million rows, the memory footprint exceeded the F32 capcity guardrail and that will force Power BI reports to fall back to direct query mode. The solution is to use another data type such as numeric data type instead of MD5 hash

upvoted 2 times

GHill1982 1 year, 11 months ago

Would V-Order not already be enabled by default?

upvoted 4 times

GHill1982 1 year, 11 months ago

Would V-Order not already be enabled by default?is V-order automatically enabled on warehouse? No, V-Order is not automatically enabled on a warehouse. You would need to enable it manually to take advantage of its benefits for improved query performance and data compression.V-order is enabled by default.Only when you are saving in optimized format .. otherwise its not enabled by default.I still have the same doubt.

upvoted 4 times

Ce____ 1 month, 4 weeks ago

Selected Answer: C

Given the Operational Cost for the Increase in Capacity, I will opt in for Enable V-Order

upvoted 1 times

h2o_molecule 1 year, 9 months ago

is V-order automatically enabled on warehouse? No, V-Order is not automatically enabled on a warehouse. You would need to enable it manually to take advantage of its benefits for improved query performance and data compression.

upvoted 2 times

h2o_molecule 1 year, 9 months ago

is V-order automatically enabled on warehouse? No, V-Order is not automatically enabled on a warehouse. You would need to enable it manually to take advantage of its benefits for improved query performance and data compression.V-order is enabled by default.Only when you are saving in optimized format .. otherwise its not enabled by default.

upvoted 2 times

priya_rastogi 2 months, 1 week ago

Selected Answer: D

The fact table has grown to 500 million rows, and it uses MD5 hash surrogate keys. Hash keys are typically large string/binary values, which increase storage and memory usage and can hurt Direct Lake query performance.

upvoted 1 times

dazero 1 year, 9 months ago

V-order is enabled by default.

upvoted 8 times

dazero 1 year, 9 months ago

V-order is enabled by default.Only when you are saving in optimized format .. otherwise its not enabled by default.

upvoted 8 times

Bs2009 3 months ago

Selected Answer: B

Microsoft currently documents an F32 guardrail of 300 million rows per table. "https://learn.microsoft.com/is-is/fabric/fundamentals/direct-lake-overview", V-Order can improve Direct Lake scan/compression performance but does not solve the F32 300-million-row guardrail

upvoted 6 times

AmirAhmadabadi 3 months ago

Selected Answer: D

MD5 produces a 32-character hexadecimal string (128-bit value).

For a fact table with 500 million rows, using string-based MD5 keys causes:

Larger storage footprint

Higher memory consumption

Slower joins between fact and dimension tables

Reduced Direct Lake performance

Direct Lake performs best when using:

Integer surrogate keys (BIGINT, INT)

Narrow columns

Efficient joins

Replacing MD5 hash keys with numeric surrogate keys significantly reduces:

Storage

Scan time

Memory usage

This improves query performance without increasing capacity costs.

Since the requirement includes minimizing operational costs, optimizing the model is better than buying more capacity.

Meets performance requirement

Meets cost requirement

upvoted 2 times

guy3hil 5 months ago

Selected Answer: C

Justification

1. Performance Needs

The warehouse DW1 has grown significantly (200M → 500M rows).

Reports use Direct Lake, which queries data directly from the lake without import.

As data volume increases, query scans become heavier.

V-Order optimizes the physical layout of columnar data (Delta tables) by re-ordering and compressing it, which reduces I/O and speeds up scans.

This directly improves query performance for Power BI visuals.

2. Cost Considerations

Increasing capacity (Option B) would improve performance but at a higher operational cost (more compute units).

V-Order achieves performance gains without increasing compute costs, making it the cost-efficient solution.

upvoted 1 times

kzaki 5 months, 2 weeks ago

Selected Answer: C

v-order targets query performance

upvoted 1 times

dazero 1 year, 9 months ago

Selected Answer: B

500 million rows is over the F32 table limit.

upvoted 29 times

dazero 1 year, 9 months ago

500 million rows is over the F32 table limit.He's right. When its over the limit the query performance gets reduces by the fallback.
https://learn.microsoft.com/en-us/fabric/fundamentals/direct-lake-overview#fallback

upvoted 30 times

juamd 6 months, 2 weeks ago

Selected Answer: D

MD5 hash keys are stored as wide string/binary types (16–32 bytes), which are inefficient in Direct Lake. As the fact table grew to 500M rows, these keys trigger DirectQuery fallback, causing slow performance and visual errors. Converting surrogate keys to BIGINT keeps queries within Direct Lake guardrails, enables better compression, and speeds up joins — meeting both performance and cost requirements. The other options either worsen the problem (A), increase costs (B), or don't address the root cause (C, E).

upvoted 1 times

ChandrajitS 7 months ago

Selected Answer: B

The existing capacity limitation reached 300 million rows for F32.

upvoted 3 times

32950cb 7 months, 1 week ago

Selected Answer: D

The root cause of the performance degradation is the use of MD5 hash surrogate keys. MD5 hashes are typically stored as 32-character strings or binary values, which create several problems for Direct Lake and the VertiPaq engine as the table grows

upvoted 2 times

cd93 8 months, 1 week ago

Selected Answer: B

4,500$/month for F32 capacity, and get a Direct Lake limit of only 300 million rows table...

300 million table is actually pretty tiny even for a middle-sized company

This is too costly... onprem Postgres can handle this amount just fine

upvoted 4 times

Josekonoor 8 months, 3 weeks ago

Selected Answer: C

increase capacity

upvoted 1 times

Josekonoor 8 months, 3 weeks ago

Selected Answer: B

Fabric F32 capacity fact table that has grown from 200 million rows to 500 million rows -> increase capacity

upvoted 3 times

Mustafa_Shazly 9 months, 2 weeks ago

Selected Answer: B

after trying I found for warehouse not lakehouse

- V order enabled by default on all warehouses.

- It cannot be enabled/disabled for individual tables.

- You can disable but can't re-enable again.

And on the other hand F32 can handle up to 300 M rows so I think B is true

upvoted 3 times

e0f0ce6 10 months ago

Selected Answer: C

I think it's C. Sure Capacity increase gives more performance. V-order does increase performance as well but at lower costs.

upvoted 1 times

chandu_3812 10 months ago

Selected Answer: C

Enable V-Order

upvoted 1 times

KTechLeg 10 months, 3 weeks ago

Selected Answer: C

V-order is for gold delta tables optimizing them in a way that Power BI (especially NRT) read it better and faster.

upvoted 1 times

Mustafa_Shazly 10 months, 3 weeks ago

Selected Answer: B

I change my anser from D to B

becuase F32 has limit 300M Rows

https://learn.microsoft.com/en-us/fabric/fundamentals/direct-lake-overview#fabric-capacity-guardrails

upvoted 2 times

Mustafa_Shazly 10 months, 3 weeks ago

Selected Answer: D

I think D is true

Problem is performance issues because increase in data mean memory lake

- MD5 and SHA256 Take a lot of memory so SHA256 will raise the problem

-V order used to speed up not solve errors

- increase cap true but we have condition (min ops cost)

upvoted 3 times

TomGray1989 10 months, 4 weeks ago

Selected Answer: C

The thing is, increasing capacity is a potential solution here - but it's really not the first solution to jump to. Yes you'll get more compute power accessibility to potentially resolve performance issue but the cost is significant.

Enable V-Order must be the 1st step (provide best query performance whilst keeping cost under control.)

I think this question is testing your ability to prioritise solution testing. Enabling V Order may resolve your issue quickly, it won't increase your costs. Increasing Capacity may resolve your issue quickly, but it actually maximises operational cost.

The answer here is C

upvoted 1 times

dantat 11 months ago

Selected Answer: B

F32 has a limitation of 300 millions Rows per table

https://learn.microsoft.com/en-us/fabric/fundamentals/direct-lake-overview#fabric-capacity-requirements

upvoted 3 times

a10 11 months, 1 week ago

Selected Answer: B

Consider row limits

upvoted 3 times

Anand_20061987 1 year ago

Selected Answer: D

The best answer is D – modify the surrogate keys to use a different data type.

Using MD5 hash keys (large 16–32 byte values) causes memory and performance issues as the fact table grows. Switching to integer (INT/BIGINT) surrogate keys reduces memory usage, speeds up joins, and improves storage efficiency.

Other options (like SHA256, increasing capacity, or using views) don't fix the root problem — the inefficient data type. Integer keys are the best practice for performance and scalability.

upvoted 4 times

toure1 11 months, 1 week ago

Et qu'en est-il de la limite à 300 millions de la capacité F32?

upvoted 1 times

Patrikkhalaf 1 year, 2 months ago

Selected Answer: C

B. Increase the capacity

Would improve performance but increases costs, violating the requirement to minimize operational costs.

C. Enable V-Order

V-Order is an optimization in Fabric Warehouses/Lakehouses that orders data column-wise for better compression and query performance.

Specifically designed to boost Direct Lake performance.

Cost-neutral (storage optimization, not compute scaling).

Best fit here.

upvoted 3 times

certs4pk 11 months ago

even after enabling v-order, it still falls back 2direct query mode, coz capacity guardrails were exceeded.. hence leaning towards 'expanding capacity' option

upvoted 1 times

d0bdf05 1 year, 2 months ago

Selected Answer: C

Enable V-Order: reshape parquet files (optimizing its). So, performance should be better, with option B (increase capacity) we have more costs which is one of the question requirement.

upvoted 2 times

raweb 1 year, 3 months ago

Selected Answer: B

F2 <> F3 300 milions row.... Increase capacity is needed

upvoted 2 times

8115575 1 year, 3 months ago

Selected Answer: B

An F32 SKU in Microsoft Fabric has a limit of 10 GB of memory per semantic model and can handle 300 million Direct Lake rows per table

upvoted 3 times

jakkuc92 1 year, 3 months ago

Selected Answer: B

F32 supports:

300 million rows per table

40 GB model size on disk

10 GB max memory

Your fact table has 500 million rows, which exceeds the supported limit for F32 in Direct Lake mode.

This mismatch can cause:

Visual errors

Slow performance

Query failures or memory overflows

 Summary:

If you're using Direct Lake and your table exceeds the row or memory limits of your SKU, increasing capacity is necessary to ensure stability and performance.

V-Order is still a great optimization, but it won't help if you're over the SKU limits.

 Final Answer: B. Increase the capacity

upvoted 5 times

127fe02 1 year, 3 months ago

Selected Answer: C

Bullshit question, the V-Order is turn-on by default. But the question still ask about this :(

upvoted 1 times

hergoi 1 year, 3 months ago

Selected Answer: C

On a table with 500 million rows covering data from multiple years, the main cost is not the JOIN operation but scanning the huge fact table for the relevant dates — so V-Order will make more sense than changing the primary key from MD5 to SHA256.

upvoted 2 times

be79ac3 1 year, 4 months ago

Selected Answer: C

With it saying minimize costs I thought it had to be C.

upvoted 2 times

medeirosLF 1 year, 4 months ago

Selected Answer: B

B is the correct answer. V-order is enabled by default; if you disable it, you can't turn it enable again.

upvoted 2 times

medeirosLF 1 year, 4 months ago

Selected Answer: B

V-order is enabled by default.

upvoted 2 times

necktru 1 year, 4 months ago

Selected Answer: B

Vorder cant be enabled

upvoted 2 times

ba2bfdf 1 year, 4 months ago

Selected Answer: B

Direct Lake mode limit 300 Mill. Fallback option mode switch to Direct Query which degrade the performance. Hence Capacity had to be increased.

upvoted 2 times

Shw7 1 year, 5 months ago

Selected Answer: C

Min operational Cost = Enable V Order

upvoted 1 times

VenkataPhaniPavanKumar 1 year, 5 months ago

Selected Answer: C

Why Other Options Are Unsuitable:

Option: A (Change MD5→SHA256)

Reason for Exclusion: Increases key size (128→256 bits), worsening storage/join performance.

Option:B (Increase capacity)

Reason for Exclusion:Raises costs (contradicts "minimize operational costs").

Option:D (Modify surrogate keys)

Reason for Exclusion: Requires rebuilding the entire data model (high effort/cost).

Option:E (Create views)

Reason for Exclusion: Adds abstraction without optimizing underlying data scans.

upvoted 1 times

VenkataPhaniPavanKumar 1 year, 5 months ago

Selected Answer: C

The best option to address performance degradation while minimizing operational costs is C. Enable V-Order.

Here's why:

V-Order is a storage format optimization for Delta Lake tables (used by Direct Lake mode in Power BI). It improves data skipping, which enables faster query performance by reducing the amount of data scanned. Especially with large fact tables like yours (500 million rows), enabling V-Order can have a significant impact on performance without increasing costs.

It's a low-cost, high-impact optimization since it doesn't require more compute resources (unlike increasing capacity).

Why not Improve Capacity?

Effective, but it increases costs. Only do this if V-Order and other optimizations don't solve the problem.

upvoted 1 times

ogbenisho 1 year, 5 months ago

Selected Answer: C

V-Order is a write-time optimization for Parquet files in Microsoft Fabric that significantly enhances read performance within Fabric compute engines like Power BI, SQL, and Spark. It achieves this by applying special sorting, row group distribution, dictionary encoding, and compression to the Parquet files. This optimization is crucial for Fabric's Direct Lake mode, which relies on V-Order for efficient data access

upvoted 1 times

Ariel_Subia 1 year, 5 months ago

Selected Answer: B

Increase the capacity is correct for me.

upvoted 1 times

Wipag 1 year, 5 months ago

Selected Answer: C

V-Order is a performance optimization feature for Delta tables in Fabric.

It reorders data for faster analytical queries by optimizing columnar layout and file structure.

Improves read performance in Direct Lake mode dramatically without extra cost.

No need to reload or refactor the data, so minimal operational impact.

upvoted 1 times

DaniMb91 1 year, 5 months ago

Selected Answer: B

300 million rows is not a hard limit, but rather a performance recommendation for optimal operation under Direct Lake with F32 capacity.

Your table has 500 million rows, which exceeds this recommended limit by 66% — a likely contributor to performance issues.

upvoted 1 times

smanzana 1 year, 5 months ago

Selected Answer: C

The correct is C

upvoted 1 times

Shw7 1 year, 5 months ago

Selected Answer: D

. Modify the surrogate keys to use a different data type.

upvoted 2 times

pramod4lk 1 year, 5 months ago

Selected Answer: C

Rather than buying more capacity or overhauling your surrogate-key design, simply enable V-Order clustering on your most selective columns.

upvoted 1 times

HI1612 1 year, 5 months ago

Selected Answer: C

Enabling V-Order in Microsoft Fabric (specifically for Direct Lake datasets) changes how data is physically stored and accessed, optimizing both read and write operations.

upvoted 1 times

lord_mufi 1 year, 5 months ago

Selected Answer: B

F32 tenant capacity has 300 million rows maximum consumption. The scenario has exhausted the capacity. You must increase the capacity to rectify the problem.

Enable V-Order is actived by default except disabled.

SHA5 has high byte encryption size which will even the performance.

The last option trigger the connection mode to DirectQuery which is even worst.

Option B is the answer

upvoted 1 times

Stef6108 1 year, 5 months ago

Selected Answer: B

Direct lake row limit per table is 300 million in an F32 capacity => Fallback to Direct Query

upvoted 2 times

malik777 1 year, 5 months ago

Selected Answer: C

Improves query performance by optimizing how data is stored and accessed.

upvoted 1 times

Madhu2023 1 year, 5 months ago

Selected Answer: B

https://learn.microsoft.com/en-us/fabric/fundamentals/direct-lake-overview#fallback

upvoted 1 times

ANayakOdisha 1 year, 5 months ago

Selected Answer: C

Enabling V-Order (Vectorized Order-Preserving Encryption for Data Lakes) can significantly improve query performance for analytics workloads, especially in scenarios where you have large datasets like in your case with 500 million rows. V-Order helps optimize query processing and can enhance the efficiency of operations like year-over-year comparisons, which are common in Power BI reports. This optimization can lead to better performance without the need for costly infrastructure upgrades, aligning with the goal of minimizing operational costs.

upvoted 1 times

MP035 1 year, 5 months ago

Selected Answer: B

V-order is enabled by default for warehouse. so Increase capacity to support higher record is the correct answer

upvoted 1 times

turlakDE 1 year, 5 months ago

Selected Answer: B

F32 and limit on the rows - 300M

upvoted 1 times

DarioReymago 1 year, 6 months ago

Selected Answer: B

Option B direct Lake limitation by sku (fallback)

https://learn.microsoft.com/en-us/fabric/fundamentals/direct-lake-overview#fabric-capacity-requirements

upvoted 2 times

yue_young 1 year, 6 months ago

Selected Answer: E

Obviously the answer should be C and E(if there are two answers). Surely to increase the capacity could help improve the performance, but it also demands more money. Considering the report only shows the value year-over-year, to create a more coarse-grained data view would be best.

upvoted 1 times

SebGo 1 year, 6 months ago

Selected Answer: B

Because of Row Limitation

upvoted 1 times

oose 1 year, 6 months ago

Selected Answer: C

Enable V order should be the first step

upvoted 1 times

upliftinghut 1 year, 6 months ago

Selected Answer: B

V-Order is enabled by default. F32 has max number of rows of 300M per table: https://learn.microsoft.com/en-us/fabric/fundamentals/direct-lake-overview

upvoted 1 times

ZZIN 1 year, 6 months ago

Selected Answer: C

Surely you are confused between answer B and C. Since it is stated to minimize operational costs, isn't C the correct answer?

upvoted 1 times

reemprive 1 year, 6 months ago

Selected Answer: B

You can't enable V-Order. It's enabled by default. You can disable V-Order, but you can't turn it back on.

upvoted 1 times

mDp0r 1 year, 6 months ago

Selected Answer: B

Agree to the 300kk row limit on direct lake for an f32

upvoted 1 times

easihene 1 year, 7 months ago

Selected Answer: C

The documentation (https://learn.microsoft.com/en-us/fabric/fundamentals/direct-lake-overview) confirms that the row-count guardrail of an F32 capacity (300 million rows per table) is now being breached by your 500-million-row fact table, and that exceeding a guardrail pushes Direct Lake queries into DirectQuery fallback, which is exactly what users are experiencing as “slow visuals” and intermittent errors.

Because the only reliable way to lift that guardrail is to move to the next SKU, the best fix is to scale the capacity to F64 (or higher), where the limit rises to 1.5 billion rows per table. This leaves plenty of headroom and restores Direct Lake performance without any refactoring.

upvoted 1 times

easihene 1 year, 7 months ago

Admin, please note: I met to select B and not C.

upvoted 1 times

AnetaK 1 year, 7 months ago

Selected Answer: E

As a PBI developer I would prefer to have views created.

Visuals show error because there is two much data to be calculated.

I would prefer to have less data, so the granular data should be aggregated, we need only yearly data. Then there would be less data to be calculated in the visual query.

upvoted 3 times

CloudKaKida 1 year, 7 months ago

Selected Answer: B

https://learn.microsoft.com/en-us/fabric/real-time-intelligence/event-streams/process-events-using-event-processor-editor?pivots=enhanced-capabilities

Fabric SKU Parquet files per table Row groups per table Rows per table (millions) Max model size on disk/OneLake (GB) Max memory (GB) 1

F2 1,000 1,000 300 10 3

F4 1,000 1,000 300 10 3

F8 1,000 1,000 300 10 3

F16 1,000 1,000 300 20 5

F32 1,000 1,000 300 40 10

F64/FT1/P1 5,000 5,000 1,500 Unlimited 25

F128/P2 5,000 5,000 3,000 Unlimited 50

F256/P3 5,000 5,000 6,000 Unlimited 100

F512/P4 10,000 10,000 12,000 Unlimited 200

F1024/P5 10,000 10,000 24,000 Unlimited 400

F2048 10,000 10,000 24,000 Unlimited 400

upvoted 1 times

el_massakre 1 year, 7 months ago

Selected Answer: B

Raise the capacity so Fabric can use DireqtQuery without fallback.

https://learn.microsoft.com/en-us/fabric/fundamentals/direct-lake-overview#fallback

upvoted 3 times

yourock43mfamily 1 year, 7 months ago

Selected Answer: C

V-order improves read performance

upvoted 1 times

SameerL 1 year, 7 months ago

Selected Answer: C

The best option is:

C. Enable V-Order.

Explanation:

Best Query Performance:

V-Order is an optimization feature in Fabric that improves query performance by organizing Parquet files more efficiently.

It enhances compression, speeds up scan performance, and reduces I/O, making it ideal for large datasets like your 500 million-row fact table.

Minimize Operational Costs:

Enabling V-Order does not require scaling up capacity, unlike option B (Increasing the capacity), which adds costs.

It optimizes storage and query performance without additional compute resources.

upvoted 1 times

kumarzzz 1 year, 8 months ago

Selected Answer: D

My suggestion to be both things

To minimize operational costs, you can modify surrogate keys by using a smaller integer data type (e.g., INT instead of BIGINT) and enabling V-order (Virtual Order) for faster data retrieval, if applicable.

upvoted 1 times

MultiCloudIronMan 1 year, 8 months ago

Selected Answer: C

Changing the surrogate key data type would likely involve significant re-engineering efforts without guaranteeing a performance improvement.

upvoted 3 times

zxc01 1 year, 8 months ago

V-Order in fabric warehouse is enabled by default, and you cannot disable it.

upvoted 1 times

Sher_Lock_Alexa 1 year, 8 months ago

Selected Answer: C

V-Order is a highly optimized columnar storage format in Microsoft Fabric that compresses and accelerates queries

upvoted 2 times

hebertorosillo 1 year, 9 months ago

Selected Answer: D

B. Minimize operational costs is false. A. Distraction. C. Warehouse is not possible in this sentence. V-Order is configured in a notebook because it is incorrect. D. It is true because

https://learn.microsoft.com/en-us/fabric/data-engineering/delta-optimization-and-v-order?tabs=sparksqlit reduces and increments performance by utilizing minimal operational costs.

upvoted 2 times

el_massakre 1 year, 7 months ago

He's right. When its over the limit the query performance gets reduces by the fallback.

https://learn.microsoft.com/en-us/fabric/fundamentals/direct-lake-overview#fallback

upvoted 4 times

ChenFu 1 year, 9 months ago

Selected Answer: C

V-Order is a storage optimization format used in Microsoft Fabric for Delta tables. It organizes data in a highly compressed and sorted format, improving query performance while reducing storage costs.

Since your reports are based on Direct Lake, enabling V-Order will directly enhance read performance without increasing operational costs.

Increasing capacity (B) would add cost, while changing hash keys (A) or modifying data types (D) are more complex and don't guarantee significant performance gains.

Creating views (E) does not optimize the underlying storage and, therefore, won't address the root cause effectively.

upvoted 2 times

henryphchan 1 year, 9 months ago

Selected Answer: D

V-Ordering is enabled by default in Fabric Datawarehouse. I have created a new datawarehouse and verify the default v-order status:

SELECT [name], [is_vorder_enabled]

FROM sys.databases;

Disabling V-Order is irreversible: once disabled, it cannot be enabled again. So C is not a valid option.

https://learn.microsoft.com/en-us/fabric/data-warehouse/v-order

upvoted 2 times

clux 1 year, 9 months ago

Selected Answer: C

V-Order improves read performance by applying special optimizations such as sorting, row group distribution, dictionary encoding, and compression on Parquet files. This enhances query performance significantly, especially for large datasets. Additionally, V-Order is cost-effective as it reduces the amount of resources needed for reading data, leading to improved performance without increasing operational costs.

upvoted 1 times

prabhjot 1 year, 9 months ago

Selected Answer: C

i think the answer should be V-order since it will sort data before writing and hence give you better read performance without increasing the capacity which also optimizes on operational costs!

upvoted 1 times

h2o_molecule 1 year, 9 months ago

Selected Answer: C

V-Order is a columnar storage format that optimizes data storage and retrieval. It can significantly improve query performance and reduce storage costs by compressing data and minimizing the amount of data read during queries. This makes it a suitable choice for large fact tables and scenarios where you need to improve performance without increasing operational costs.

upvoted 1 times

robertlavigne 1 year, 10 months ago

Selected Answer: B

B. Although B is expensive I think it is the only viable option.

https://learn.microsoft.com/en-us/fabric/get-started/direct-lake-overview

See the table on Directlake limitations in the above link. For F32 it is 300 million rows before direct query fallback occurs. For F64 it is 1500 million rows.

V Order or Key modification don't minimize the row count and won't avoid these limitations. Views aren't supported in DirectLake and will always fallback to direct query so thats not helpful.

upvoted 7 times

i_have_a_name 1 year, 10 months ago

Selected Answer: D

SHA26 generates a 64 character length key, which will affect the performance further.

V-Order is enabled by default in Fabric.

Modifying the surrogate key should be the cost effective solution.

upvoted 3 times

mixonfreddy 1 year, 11 months ago

Selected Answer: D

You have to modify the surrogate keys

upvoted 4 times

manolet 1 year, 11 months ago

Selected Answer: C

Changing surrogate keys is not a straightforward solution. The key is that you need to make the decision based on minimizing operational cost.

upvoted 2 times

QAZdbarhate12345678 1 year, 11 months ago

Selected Answer: D

The best solution to resolve the performance issues while meeting the requirements of best query performance and minimizing operational costs is:

D. Modify the surrogate keys to use a different data type.

upvoted 4 times

AzurePart 1 year, 5 months ago

Did you read the contents of the url document you wrote?

https://learn.microsoft.com/en-us/fabric/data-warehouse/v-order#performance-considerations

"Currently, disabling V-Order can only be done at the warehouse level, and it is irreversible: once disabled, it cannot be enabled again. Users must consider the performance if they choose to Disable V-Order in Fabric Warehouse."

upvoted 2 times

Tuki93 1 year, 5 months ago

After reviewing the public documentation again, I understand that V-Order is enabled by default in Data Warehouses, as the documentation only refers to how it can be disabled—and notes that this action is irreversible.

https://learn.microsoft.com/en-us/fabric/data-warehouse/v-order#scenarios-where-v-order-might-not-be-beneficial

At the same time, I agree with my peers who mentioned that the row limit per table is 300 million in an F32 capacity. Therefore, I'm updating my answer to option B: Increase the capacity.

https://learn.microsoft.com/en-us/fabric/fundamentals/direct-lake-overview#fabric-capacity-requirements

upvoted 21 times

Adriel_1996 1 year, 8 months ago

What about the limit on the rows limitation (https://learn.microsoft.com/en-us/fabric/fundamentals/direct-lake-overview) ?

upvoted 6 times

satyaammm 1 year, 5 months ago

Only when you are saving in optimized format .. otherwise its not enabled by default.

upvoted 1 times

Raibeth 1 year, 9 months ago

I still have the same doubt.

upvoted 3 times

IshtarSQL 1 year, 11 months ago

Selected Answer: D

While MD5 hashes are deterministic and ensure uniqueness, they can be less efficient for join operations compared to integer-based keys. This inefficiency arises because joining on lengthy string keys demands more computational resources than joining on shorter, integer-based keys.

Recommendation: Modify the surrogate keys to use a different data type, specifically integers.

upvoted 8 times

HOTSPOT -

You have a Fabric workspace that contains a warehouse named DW1. DW1 contains the following tables and columns.

You need to create an output that presents the summarized values of all the order quantities by year and product. The results must include a summary of the order quantities at the year level for all the products.

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

NOTE: Each correct selection is worth one point.

Table name	Column name	Description
SalesOrderDetail	ProductID	Contains the product ID of the ordered product
SalesOrderDetail	ModifiedDate	Contains the date of an order
SalesOrderDetail	OrderQty	Contains the order quantity
Product	ProductID	Contains the unique ID of a product
Product	Name	Contains a product name

Answer Area

```
[ ▼ ] (SO.ModifiedDate) AS OrderDate
```

SELECT CAST
SELECT CONVERT
SELECT YEAR

```
    ,P.Name AS ProductName
    ,SUM(SO.OrderQty) AS OrderQty
FROM [dbo].[SalesOrderDetail] SO
INNER JOIN [dbo].[Product] P
  ON P.ProductID = SO.ProductID
GROUP BY
    [ ▼ ]
ORDER BY OrderDate
```

CUBE(YEAR(SO.ModifiedDate), P.Name)
GROUPING SETS ((YEAR(SO.ModifiedDate), P.Name), (YEAR(SO.ModifiedDate)))
ROLLUP(YEAR(SO.ModifiedDate), P.Name)
YEAR(SO.ModifiedDate), P.Name

Suggested Answer:

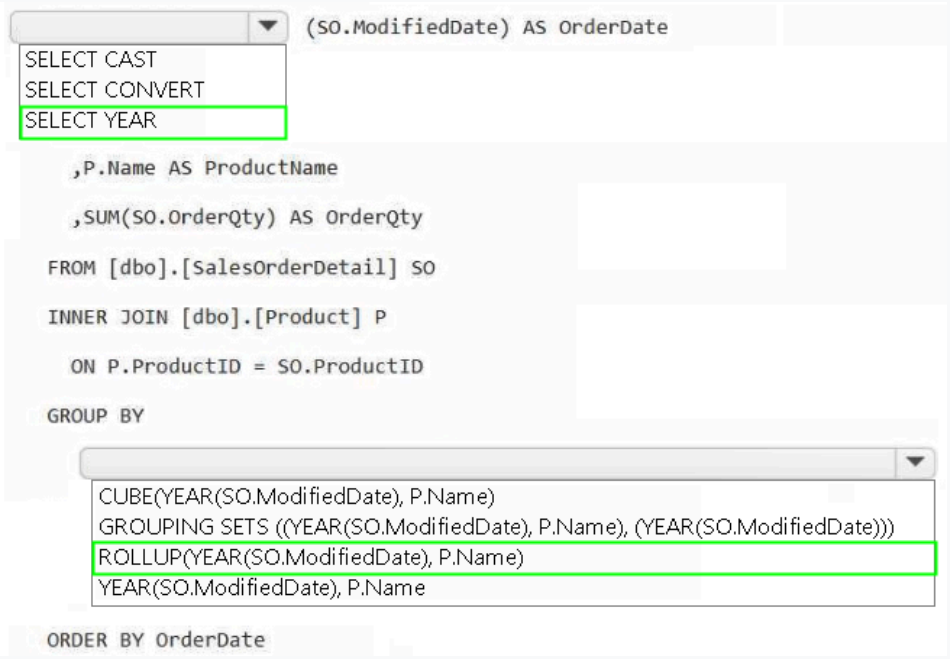

robertlavigne Highly Voted 1 year, 10 months ago

SELECT YEAR

Grouping Sets: Best answer. No extra stuff. Just what we want Cube: Lots of extra combinations. Has a row with the total for all years and all products, and rows with the total for each product for all years. Rollup: Has the unnecessary total of all products and all years Normal Group by: No summary at the year level for all products.

https://learn.microsoft.com/en-us/sql/t-sql/queries/select-group-by-transact-sql?view=sql-server-ver16

upvoted 43 times

robertlavigne Highly Voted 1 year, 10 months ago

SELECT YEAR Grouping Sets: Best answer. No extra stuff. Just what we want Cube: Lots of extra combinations. Has a row with the total for all years and all products, and rows with the total for each product for all years. Rollup: Has the unnecessary total of all products and all years Normal Group by: No summary at the year level for all products. https://learn.microsoft.com/en-us/sql/t-sql/queries/select-group-by-transact-sql?view=sql-server-ver16grouping sets is the best answerFor example, GROUP BY ROLLUP (Country, Region) and GROUP BY GROUPING SETS (ROLLUP (Country, Region)) return the same results.But the requirement says: The results must include a summary of the order quantities at the year level for all the products. Therefore ROLLUP should be correct.

upvoted 44 times

AmirAhmadabadi Most Recent 3 months ago

SELECT YEAR(SO.ModifiedDate) AS OrderDate,

P.Name AS ProductName,

SUM(SO.OrderQty) AS OrderQty

FROM [dbo].[SalesOrderDetail] SO

INNER JOIN [dbo].[Product] P

ON P.ProductID = SO.ProductID

GROUP BY

ROLLUP(YEAR(SO.ModifiedDate), P.Name)

ORDER BY OrderDate;

upvoted 3 times

Josekonoor 7 months, 3 weeks ago

what is meant by "correct Answer A" !?

upvoted 1 times

Josekonoor 8 months, 3 weeks ago

select year(), grouping sets

upvoted 1 times

Josekonoor 8 months, 3 weeks ago

select Year() Grouping Sets

upvoted 1 times

mixonfreddy 1 year, 11 months ago

Selected Answer: A

SELECT YEAR ROLLUP(YEAR(SO.ModifiedDATE), P.Name)

upvoted 31 times

mixonfreddy 1 year, 11 months ago

SELECT YEAR ROLLUP(YEAR(SO.ModifiedDATE), P.Name)Why not plain Group By when it serves the purpose. We can get the summary anyway without adding an additional clause ?Because it does not serve the purpose. Requirements are: 1. Summarize order quantities by year and product. 2. Include a summary of order quantities at the year level for all products. A simple GROUP BY cannot generate subtotals or a grand total. You would need to write multiple queries and combine them with UNION ALL, which is inefficient and cumbersome.Rollup adds extra row for grand total (all years, all products), which is not requested. Grouping Sets works better

upvoted 32 times

Mustafa_Shazly 9 months, 2 weeks ago

If we have 2 years and 3 Products

ROLLUP will give the true result + one row for grand total >> will return 9 Rows

GROUPING SETS will give the exact result >> will return 8 Rows

Our Issue with the last row come with ROLLUP is not required

upvoted 1 times

Madhu2023 1 year, 9 months ago

Why not plain Group By when it serves the purpose. We can get the summary anyway without adding an additional clause ?Because it does not serve the purpose. Requirements are: 1. Summarize order quantities by year and product. 2. Include a summary of order quantities at the year level for all products. A simple GROUP BY cannot generate subtotals or a grand total. You would need to write multiple queries and combine them with UNION ALL, which is inefficient and cumbersome.Rollup adds extra row for grand total (all years, all products), which is not requested. Grouping Sets works better

upvoted 2 times

Madhu2023 1 year, 9 months ago

Why not plain Group By when it serves the purpose. We can get the summary anyway without adding an additional clause ?

upvoted 2 times

digitalqloud 11 months, 2 weeks ago

select year and group sets

upvoted 1 times

rnd_27 1 year, 8 months ago

Because it does not serve the purpose. Requirements are: 1. Summarize order quantities by year and product. 2. Include a summary of order quantities at the year level for all products. A simple GROUP BY cannot generate subtotals or a grand total. You would need to write multiple queries and combine them with UNION ALL, which is inefficient and cumbersome.

upvoted 1 times

rnd_27 1 year, 8 months ago

Because it does not serve the purpose. Requirements are: 1. Summarize order quantities by year and product. 2. Include a summary of order quantities at the year level for all products. A simple GROUP BY cannot generate subtotals or a grand total. You would need to write multiple queries and combine them with UNION ALL, which is inefficient and cumbersome.Rollup adds extra row for grand total (all years, all products), which is not requested. Grouping Sets works better

upvoted 1 times

vkp123 1 year, 1 month ago

ROLLUP is typically the best choice because:

It's concise and readable

Most databases support it (SQL Server, MySQL, PostgreSQL, Oracle)

It automatically generates the hierarchical summary

It's efficient for large datasets

upvoted 1 times

AndrewsOtis 1 year, 3 months ago

SELECT YEAR... To extract the year like a field

ROLLUP ... to achieve the full sales.

upvoted 3 times

jaysun 1 year, 4 months ago

SELECT YEAR

ROLLUP(YEAR(SO.ModifiedDATE), P.Name)

upvoted 2 times

Manish0427 1 year, 5 months ago

Between ROLLUP and GROUPING_SETS in this case, GROUPING_SETS is the correct answer.

Though both options will give us the rows stated in the requirement, ROLLUP will give an extra row with total_qty over all the products and all the years (rolling up the products and years). This is not part of the requirement.

upvoted 3 times

Wipag 1 year, 5 months ago

SELECT YEAR(SO.ModifiedDate) AS OrderDate,

P.Name AS ProductName,

SUM(SO.OrderQty) AS OrderQty

FROM [dbo].[SalesOrderDetail] SO

INNER JOIN [dbo].[Product] P ON P.ProductID = SO.ProductID

GROUP BY ROLLUP(YEAR(SO.ModifiedDate), P.Name)

ORDER BY OrderDate

upvoted 2 times

407a475 1 year, 5 months ago

SELECT YEAR

GROUPING SETS(...)

upvoted 3 times

renan_moraes37 1 year, 5 months ago

First answer: YEAR

Second answer: ROLLUP...

upvoted 1 times

Runu 1 year, 6 months ago

Selected Answer: C

Explanation:

YEAR(SO.ModifiedDate) extracts the year.

ROLLUP(YEAR(SO.ModifiedDate), P.Name) groups the data by year and product, and adds a subtotal at the year level (i.e., all products for a given year).

The result will include rows with NULL in ProductName for the yearly totals.

upvoted 2 times

Runu 1 year, 6 months ago

Selected Answer: C

upvoted 1 times

yourock43mfamily 1 year, 7 months ago

ROLLUP is correct

upvoted 1 times

hebertorosillo 1 year, 9 months ago

SELECT YEAR

GROUPING SETS(...)

upvoted 9 times

clux 1 year, 9 months ago

1.Year

2. Grouping sets

upvoted 7 times

sakis213 1 year, 9 months ago

Question asks for summarized values of all the order quantities by year/product AND year alone , so grouping sets its correct

1.Year

2. Grouping sets

upvoted 8 times

sakis213 1 year, 9 months ago

But ROLLUP performs better for these hierarchical aggregations, it might be better to go with ROLLUP

upvoted 1 times

realexamguru 1 year, 9 months ago

ROLLUP will add one extra Total row that is not requested in the question, so the only correct answer is grouping sets here.

upvoted 4 times

tomaszstaroszczyk 1 year, 10 months ago

SELECT YEAR

GROUPING SETS(...)

If you use ROLLUP then you will get an extra row that sums up the OrderQty total.

This is not requested, ROLLUP returns exactly what we need.

upvoted 9 times

aws123 7 months, 2 weeks ago

grouping sets is the best answer

upvoted 2 times

rnd_27 1 year, 8 months ago

But the requirement says:

The results must include a summary of the order quantities at the year level for all the products.

Therefore ROLLUP should be correct.

upvoted 5 times

tpositive 1 year, 9 months ago

For example, GROUP BY ROLLUP (Country, Region) and GROUP BY GROUPING SETS (ROLLUP (Country, Region)) return the same results.

upvoted 2 times

shmmini 1 year, 8 months ago

Rollup adds extra row for grand total (all years, all products), which is not requested. Grouping Sets works better

upvoted 6 times

Question #7 *Topic 1*

You have a Fabric workspace that contains a lakehouse named Lakehouse1. Data is ingested into Lakehouse1 as one flat table. The table contains the following columns.

You plan to load the data into a dimensional model and implement a star schema. From the original flat table, you create two tables named FactSales and DimProduct. You will track changes in DimProduct.
You need to prepare the data.
Which three columns should you include in the DimProduct table? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

Name	Description
TransactionID	Contains a unique ID for each transaction
Date	Contains the date of a transaction
ProductID	Contains a unique ID for each product
ProductColor	Contains a descriptive attribute that describes the color of each product
ProductName	Contains a unique name for each product
SalesAmount	Contains the sales amount of a transaction

A. Date

B. ProductName

C. ProductColor

D. TransactionID

E. SalesAmount

F. ProductID

Suggested Answer: *B, C, F*

Community vote distribution

BCF (90%) (5%) (5%)

QAZdbarhate12345678 Highly Voted 1 year, 11 months ago

Columns to Include in DimProduct ProductID - This is the unique identifier for each product and acts as the primary key in the DimProduct table. ProductName - Contains the name of the product, which is a descriptive attribute. ProductColor - Contains the color description of the product, which is also a descriptive attribute.

upvoted 13 times

Josekonoor Most Recent 8 months, 3 weeks ago

dimProduct -> Unique Product

upvoted 1 times

balanceBananes 12 months ago

Everything about the products

upvoted 2 times

fdsdfgxcvbdsfhshfg 1 year, 3 months ago

BCF seem to be correct, but it won't allow us to track changes so on a fence with these answer, some additional logic coming from Date column is required

upvoted 3 times

AndrewsOtis 1 year, 3 months ago

Correct BCF It is relation with a Dimension table of Products. The others are Fact items

upvoted 2 times

Wipag 1 year, 5 months ago

✔ F. ProductID Primary/business key of the dimension. Used to join with the fact table (FactSales). ✔ B. ProductName Descriptive attribute of the product. Typically shown in reports. ✔ C. ProductColor Another descriptive attribute. Useful for slicing data and analysis.

upvoted 2 times

smanzana 1 year, 5 months ago

Correct BCF

upvoted 2 times

DarioReymago 1 year, 5 months ago

ProductID ProductName ProductColor

upvoted 2 times

renan_moraes37 1 year, 5 months ago

For dimension needs to be every information about the products like Color and name, also the key which is productID

upvoted 2 times

oose 1 year, 6 months ago

The trabsactio the date and the sumbshould be the other table

upvoted 1 times

Kiracloud 1 year, 6 months ago

- ProductID acts as a unique identifier for each product (primary key). - ProductName provides a descriptive name, useful for reporting and analysis. - ProductColor adds valuable categorical information, potentially useful in filtering and segmentation. Other columns like Date, TransactionID, and SalesAmount are better suited for fact tables (FactSales), since they contain transaction-specific details rather than product attributes.

upvoted 1 times

maham 1 year, 8 months ago

Selected Answer: ABF

The dimension table has to track the changes also. Then date field is important for SCD. So Product color takes second priority in this list

upvoted 2 times

maham 1 year, 8 months ago

The dimension table has to track the changes also. Then date field is important for SCD. So Product color takes second priority in this listin this case, the transaction date is referenced. you cannot use it as it is. If you have to track SCD, you have to create your own Date Column, and populate it differently, not from this tablePlease also note that the Date column may not be the product modified date.

upvoted 2 times

shmmini 1 year, 8 months ago

in this case, the transaction date is referenced. you cannot use it as it is. If you have to track SCD, you have to create your own Date Column, and populate it differently, not from this table

upvoted 3 times

yue_young 1 year, 6 months ago

Please also note that the Date column may not be the product modified date.

upvoted 1 times

Madhu2023 1 year, 9 months ago

BCF are the attributes of Product. TranId , Date and SalesAmt quaify for fact table instead

upvoted 2 times

robertlavigne 1 year, 10 months ago

Date: Should be in the FactSales table ProductName: Attribute of the product and belongs in DimProduct ProductColour: Attribute of the product and belongs in DimProduct TransactionID: Belongs in the FactSales Table Sales Amount: Belongs in the FactSales Table ProductID: Will be the primary key for DimProduct

upvoted 1 times

Tuki93 1 year, 11 months ago

I believe the provided answer is correct

upvoted 4 times

Question #8 *Topic 1*

You have a Fabric workspace named Workspace1 that contains a notebook named Notebook1.
In Workspace1, you create a new notebook named Notebook2.
You need to ensure that you can attach Notebook2 to the same Apache Spark session as Notebook1.
What should you do?

A. Enable high concurrency for notebooks.

B. Enable dynamic allocation for the Spark pool.

C. Change the runtime version.

D. Increase the number of executors.

Suggested Answer: *A*

Community vote distribution

A (100%)

AmirAhmadabadi Most Recent 3 months ago

Selected Answer: A

In Microsoft Fabric, High Concurrency Mode allows multiple notebooks to share the same Spark session. This helps by:

Reusing Spark sessions

Reducing startup time

Sharing compute resources across notebooks

Allowing Notebook2 to attach to the session already used by Notebook1

Therefore, to enable session sharing between notebooks, you must:

Enable high concurrency for notebooks

upvoted 2 times

Josekonoor 8 months, 3 weeks ago

Selected Answer: A

2 notebooks attached to same session -> High Concurrency

upvoted 1 times

92ef42d 1 year, 2 months ago

Selected Answer: A

High Concurrency Mode allows multiple notebooks to share a single Apache Spark session, improving performance, reducing costs, and speeding up execution.

upvoted 1 times

Wipag 1 year, 5 months ago

Selected Answer: A

To allow multiple notebooks (e.g., Notebook1 and Notebook2) in Microsoft Fabric to share the same Apache Spark session, you must enable high concurrency mode. This mode allows for concurrent execution of commands in shared sessions — ideal for collaborative or parallel workloads.

upvoted 1 times

DarioReymago 1 year, 5 months ago

Selected Answer: A

Enabling high concurrency

upvoted 1 times

renan_moraes37 1 year, 5 months ago

Selected Answer: A

Enabling high concurrency gives opportunity for two notebooks work in the same session

upvoted 1 times

renan_moraes37 1 year, 5 months ago

Selected Answer: A

High concurrency allows multiple notebooks in same Spark session

upvoted 1 times

turlakDE 1 year, 5 months ago

Selected Answer: A

Enabling high concurrency

upvoted 1 times

oose 1 year, 6 months ago

Selected Answer: A

High concurrency Will amake the work

upvoted 1 times

AbhirupAz900 1 year, 7 months ago

Selected Answer: A

A is correct because without high concurrency multiple notebook can not use same spark session

upvoted 2 times

sesav4742 1 year, 7 months ago

Selected Answer: A

https://learn.microsoft.com/en-us/fabric/data-engineering/configure-high-concurrency-session-notebooks

In Azure Fabric, when you want to share the same Apache Spark session across multiple notebooks, enabling high concurrency mode is essential. High concurrency mode allows multiple notebooks to share the same Spark session, facilitating collaboration and efficient resource usage.

upvoted 2 times

QAZdbarhate12345678 Highly Voted 1 year, 11 months ago

Selected Answer: A

In Azure Fabric, when you want to share the same Apache Spark session across multiple notebooks, enabling high concurrency mode is essential. High concurrency mode allows multiple notebooks to share the same Spark session, facilitating collaboration and efficient resource usage.

upvoted 7 times

Tuki93 Highly Voted 1 year, 11 months ago

Selected Answer: A

Enabling high concurrency allows multiple notebooks to share the same Spark session.

https://learn.microsoft.com/en-us/fabric/data-engineering/configure-high-concurrency-session-notebooks

https://learn.microsoft.com/en-us/fabric/data-engineering/high-concurrency-overview

upvoted 18 times

Question #9 *Topic 1*

You have a Fabric workspace named Workspace1 that contains a lakehouse named Lakehouse1. Lakehouse1 contains the following tables:

Orders -

Customer -

Employee -
The Employee table contains Personally Identifiable Information (PII).
A data engineer is building a workflow that requires writing data to the Customer table, however, the user does NOT have the elevated permissions required to view the contents of the Employee table.
You need to ensure that the data engineer can write data to the Customer table without reading data from the Employee table.
Which three actions should you perform? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. Share Lakehouse1 with the data engineer.

B. Assign the data engineer the Contributor role for Workspace2.

C. Assign the data engineer the Viewer role for Workspace2.

D. Assign the data engineer the Contributor role for Workspace1.

E. Migrate the Employee table from Lakehouse1 to Lakehouse2.

F. Create a new workspace named Workspace2 that contains a new lakehouse named Lakehouse2.

G. Assign the data engineer the Viewer role for Workspace1.

Suggested Answer: *D, E, F*

Community vote distribution

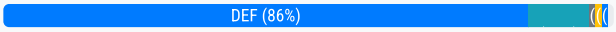

alqazqzy.asseel Most Recent 1 month, 1 week ago

Selected Answer: D,E,F

D, E, F, create new workspace2 , and new lakehouse2. Then migrate the employee table to the new workflow. Finally assign the data engineer the Contributor role for Workspace1 so they can write and do their job.

upvoted 1 times

devpool 5 months, 1 week ago

Selected Answer: B,E,F

data engineer can write data to the Customer table : B

data engineer does not perform migration.

upvoted 1 times

Josekonoor 8 months, 3 weeks ago

Selected Answer: D,E,F

basically you have to move the high security item to another workspace

upvoted 1 times

Mustafa_Shazly 9 months, 2 weeks ago

Selected Answer: D,E,F

the order F >> E >> D

- A and G not true because sharing and viewer not give writing permssion

- B and C not true because user mustn't see Emp table where we will migrate it to WS2

upvoted 1 times

Bin_Hashim 1 year ago

Selected Answer: D,E,F

Separate the PII table from the lakehouse the engineer will use.

F. Create a new workspace named Workspace2 that contains a new lakehouse named Lakehouse2.

E. Migrate the Employee table from Lakehouse1 to Lakehouse2.

Only users who need PII access are given permissions to Workspace2/Lakehouse2.

Give the engineer write access only where Customer lives.

D. Assign the data engineer the Contributor role for Workspace1.

This lets them write to the Customer table (and Orders) in Lakehouse1, while they have no access to Workspace2/Lakehouse2 and therefore cannot read Employee.

So the correct actions are: D, E, and F.

upvoted 4 times

SaswataNad 1 year, 3 months ago

Selected Answer: ADE

As per chatGPT

upvoted 2 times

SaswataNad 1 year, 3 months ago

As per chatGPTI wouldn't rely on that... It is 50/50...

upvoted 2 times

92ef42d 1 year, 2 months ago

I wouldn't rely on that... It is 50/50...

upvoted 2 times

VenkataPhaniPavanKumar 1 year, 5 months ago

Selected Answer: D,E,F

Solution Steps:

Isolate sensitive data:

Create a new workspace (Workspace2) with a new lakehouse (Lakehouse2) to host the PII data separately.

→ F. Create a new workspace named Workspace2 that contains a new lakehouse named Lakehouse2.

Move PII data:

Migrate the Employee table (PII) from Lakehouse1 to Lakehouse2 to remove it from the original environment.

→ E. Migrate the Employee table from Lakehouse1 to Lakehouse2.

Grant minimal write permissions:

Assign the Contributor role for Workspace1 to the data engineer. This role allows writing to objects (like the Customer table) but does not grant access to Workspace2 (where PII resides).

→ D. Assign the data engineer the Contributor role for Workspace1.

upvoted 4 times

Wipag 1 year, 5 months ago

Selected Answer: D,E,F

 D. Assign the data engineer the Contributor role for Workspace1

Contributor role is required to write data to tables in the workspace, including Customer. This role allows adding or editing content but not managing access or seeing hidden data like Employee if it's secured separately.

 E. Migrate the Employee table from Lakehouse1 to Lakehouse2

Since the Employee table contains PII, and the engineer should not have access to it, moving it to a separate lakehouse ensures isolation of sensitive data, while still allowing access to Customer in Lakehouse1.

F. Create a new workspace named Workspace2 that contains a new lakehouse named Lakehouse2

upvoted 3 times

smanzana 1 year, 5 months ago

Selected Answer: D,E,F

Correct DEF

upvoted 1 times

DarioReymago 1 year, 5 months ago

Selected Answer: D,E,F

FED:

Create a new workspace and lakehouse (Workspace2 and Lakehouse2).

Migrate the Employee table to the new lakehouse.

Assign the data engineer the Contributor role in the original workspace.

con esto los ingenieros no pueden ver ni escribir los datos PII

upvoted 1 times

NNPRN 1 year, 5 months ago

Selected Answer: D,E,F

The order is F---E---D.

upvoted 1 times

MohanNaidu08 1 year, 5 months ago

Selected Answer: A,D,E

A. Share Lakehouse1 with the data engineer

This gives the data engineer access to the lakehouse so they can interact with its tables.

Sharing is required before any role-based permissions can be applied.

D. Assign the data engineer the Contributor role for Workspace1

The Contributor role allows users to read and write data to lakehouse tables.

However, table-level security can still restrict access to specific tables like Employee.

E. Migrate the Employee table from Lakehouse1 to Lakehouse2

This isolates the PII data in a separate lakehouse.

Since the data engineer won't have access to Lakehouse2, they won't be able to read the Employee table, even if they have Contributor access to Lakehouse1.

upvoted 2 times

alirana3 1 year, 7 months ago

Selected Answer: D,E,F

The three correct actions to take are:

E. Migrate the Employee table from Lakehouse1 to Lakehouse2 — This isolates sensitive data (Employee table) from the data engineer's access.

D. Assign the data engineer the Contributor role for Workspace1 — This allows the data engineer to write to the Customer table in Lakehouse1.

F. Create a new workspace named Workspace2 that contains a new lakehouse named Lakehouse2 — This action is a result of migrating the Employee table to Lakehouse2 in a separate workspace (Workspace2) to ensure restricted access to sensitive data (PII).

Why Option A is Incorrect:

A. Share Lakehouse1 with the data engineer:

Sharing only gives read access to the lakehouse, which is not sufficient for the data engineer to write data to the Customer table in Lakehouse1. They need write permissions to modify the table, so Contributor role for Workspace1 is the correct option.

upvoted 4 times

AbhirupAz900 1 year, 7 months ago

Selected Answer: A,D,E

Since only 3 answers needed to be selected ADE

upvoted 2 times

Etensel 1 year, 8 months ago

Selected Answer: A,D,E

belive me..

upvoted 1 times

fassil Highly Voted 1 year, 9 months ago

Selected Answer: D,E,F

Assign the data engineer the Contributor role for Workspace1 (D)

- This will provide the necessary permissions to write to the Customer table.

Migrate the Employee table from Lakehouse1 to Lakehouse2 (E)

- This will isolate the table with PII in a different lakehouse.

Create a new workspace named Workspace2 that contains a new lakehouse named Lakehouse2 (Option F)

- This is necessary to separate the Employee table into a different workspace to restrict access.

upvoted 11 times

ChenFu 1 year, 9 months ago

Selected Answer: C,E,F

C. Assign the Viewer role for Workspace2:

This ensures the data engineer can see the structure and metadata of the workspace without accessing sensitive Employee table data. They can still write to the Customer table without elevated permissions.

E. Migrate the Employee table from Lakehouse1 to Lakehouse2:

This separates Personally Identifiable Information (PII) from the original lakehouse, ensuring the data engineer does not have access to sensitive data while working with non-PII data.

F. Create Workspace2 with Lakehouse2:

This creates an isolated environment specifically for sensitive data like the Employee table. It ensures workspace-level access control and adheres to privacy and security best practices.

upvoted 1 times

henryphchan 1 year, 9 months ago

Selected Answer: D,E,F

The steps are F, E, D

upvoted 3 times

MultiCloudIronMan 1 year, 9 months ago

Selected Answer: D,E,F

So, in this case, your answer DEF makes sense as it covers the necessary steps:

Create a new workspace and lakehouse (Workspace2 and Lakehouse2).

Migrate the Employee table to the new lakehouse.

Assign the data engineer the Contributor role in the original workspace.

upvoted 2 times

robertlavigne Highly Voted 1 year, 10 months ago

Selected Answer: D,E,F

First we need to do F and E to separate the Employee table from the other tables. Then they need contributor access to be able to write to the Customer table which is in Workspace1

A is incorrect as sharing only gives read access. https://learn.microsoft.com/en-us/fabric/data-engineering/lakehouse-sharing

B is incorrect as Workspace2 has the Employee table which the data engineer shouldn't access.

C is incorrect as viewer won't allow them to write and the Customer table isn't even in that workspace

E is incorrect as viewer won't allow them to write

upvoted 20 times

i_have_a_name 1 year, 10 months ago

Selected Answer: A,E,F

If we are migrating the Employee table to Lakehouse 2 in the workspace 2 , what is wrong in sharing the Lakehouse 1 with the data engineer ?

My answer would be A, E, F

upvoted 1 times

dcprice 1 year, 9 months ago

Currently uou can't give write access via Lakehouse sharing. https://learn.microsoft.com/en-us/fabric/data-engineering/lakehouse-sharing. So it is DEF

upvoted 1 times

prabhjot 1 year, 9 months ago

but then you ultimately would have to give the data engineer a contributor access to workspace 1 for him to work on the lakehouse as far as i think.

upvoted 1 times

Tuki93 1 year, 11 months ago

Selected Answer: D,E,F

I think D, E, and F are the working solution.

upvoted 4 times

IshtarSQL 1 year, 11 months ago

Selected Answer: D,E,F

Why A is incorrect: A. Share Lakehouse1 with the data engineer.

Sharing Lakehouse1 directly grants access to all its contents, including the Employee table, which violates the requirement.

upvoted 4 times

Question #10 *Topic 1*

You have a Fabric warehouse named DW1. DW1 contains a table that stores sales data and is used by multiple sales representatives.
You plan to implement row-level security (RLS).
You need to ensure that the sales representatives can see only their respective data.
Which warehouse object do you require to implement RLS?

A. STORED PROCEDURE

B. CONSTRAINT

C. SCHEMA

D. FUNCTION

Suggested Answer: *D*

Community vote distribution



alqazqzy.asseel Most Recent 1 month, 1 week ago

Selected Answer: D

To implement RLS in warehouse we should first define a FUCNTION with filtering parameters in this case can be for example the user email, then create SECURITY POLICY that define the defined function as FILTER PREDICATE

upvoted 2 times

Adriel_1996 1 year, 8 months ago

Selected Answer: D

-- Creating schema for Security CREATE SCHEMA Security; GO

-- Creating a function for the SalesRep evaluation CREATE FUNCTION Security.tvf_securitypredicate(@UserName AS varchar(50)) RETURNS TABLE WITH SCHEMABINDING AS RETURN SELECT 1 AS tvf_securitypredicate_result WHERE @UserName = USER_NAME() OR USER_NAME() = 'BatchProcess@contoso.com'; GO

-- Using the function to create a Security Policy CREATE SECURITY POLICY YourSecurityPolicy ADD FILTER PREDICATE Security.tvf_securitypredicate(UserName_column) ON sampleschema.sampletable WITH (STATE = ON); GO

upvoted 13 times

Josekonoor 8 months, 3 weeks ago

Selected Answer: D

for RLS at SQL world, function to be used

upvoted 1 times

Mustafa_Shazly 10 months, 3 weeks ago

Selected Answer: D

D apply security

but we used it when applying RLS schema to organize database objects only.

upvoted 1 times

Wipag 1 year, 5 months ago

Selected Answer: D

To implement Row-Level Security (RLS) in a Microsoft Fabric warehouse (DW1), you need to define a security policy that uses a user-defined function to filter rows based on the current user.

upvoted 3 times

smanzana 1 year, 5 months ago

Selected Answer: D

Función or Security Policy

upvoted 1 times

Adriel_1996 Highly Voted 1 year, 8 months ago

Selected Answer: D

-- Creating schema for Security

CREATE SCHEMA Security;

GO

-- Creating a function for the SalesRep evaluation

CREATE FUNCTION Security.tvf_securitypredicate(@UserName AS varchar(50))

RETURNS TABLE

WITH SCHEMABINDING

AS

RETURN SELECT 1 AS tvf_securitypredicate_result

WHERE @UserName = USER_NAME()

OR USER_NAME() = '';

GO

-- Using the function to create a Security Policy

CREATE SECURITY POLICY YourSecurityPolicy

ADD FILTER PREDICATE Security.tvf_securitypredicate(UserName_column)

ON sampleschema.sampletable

WITH (STATE = ON);

GO

upvoted 15 times

prabhjot 1 year, 9 months ago

Selected Answer: D

you create a predicate function which gets evaluated to filter user access to certain rows in a Yes or No manner

upvoted 2 times

QAZdbarhate12345678 Highly Voted 1 year, 11 months ago

Selected Answer: D

To implement Row-Level Security (RLS) in a Fabric warehouse like DW1, need to use a FUNCTION to define the filtering logic. Specifically, a user-defined function (UDF) is created and associated with the RLS policy to determine which rows each user can access.

upvoted 15 times

Tuki93 1 year, 11 months ago

Selected Answer: D

https://learn.microsoft.com/en-us/fabric/data-warehouse/tutorial-row-level-security#2-define-security-policies

upvoted 3 times

Question #11 *Topic 1*

HOTSPOT -
You have a Fabric workspace named Workspace1_DEV that contains the following items:
10 reports

Four notebooks -

Three lakehouses -

Two data pipelines -

Two Dataflow Gen1 dataflows -

Three Dataflow Gen2 dataflows -
Five semantic models that each has a scheduled refresh policy
You create a deployment pipeline named Pipeline1 to move items from Workspace1_DEV to a new workspace named Workspace1_TEST.
You deploy all the items from Workspace1_DEV to Workspace1_TEST.
For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

Answer Area

Statements	Yes	No
Data from the semantic models will be deployed to the target stage.	○	○
The Dataflow Gen1 dataflows will be deployed to the target stage.	○	○
The scheduled refresh policies will be deployed to the target stage.	○	○

Suggested Answer:

Statements	Yes	No
Data from the semantic models will be deployed to the target stage.	○	◉
The Dataflow Gen1 dataflows will be deployed to the target stage.	◉	○
The scheduled refresh policies will be deployed to the target stage.	○	◉

SamuComqi Highly Voted 1 year, 10 months ago

NO: only metadata will be deployed to the target starge (https://learn.microsoft.com/en-us/fabric/cicd/deployment-pipelines/understand-the-deployment-process?tabs=new#item-properties-copied-during-deployment) YES: Gen1 dataflows will be deployed NO: Refresh schedules will not be deployed (https://learn.microsoft.com/en-us/fabric/cicd/deployment-pipelines/understand-the-deployment-process?tabs=new#item-properties-copied-during-deployment)Gen1 is not in the list of the supported items: https://learn.microsoft.com/en-us/fabric/cicd/deployment-pipelines/intro-to-deployment-pipelines?tabs=new-ui#supported-itemsin list but there some confusion = Get 2 is listed under Data Factory items = Gen 1 some times called dataflow only or legacy is listed under Power BI itemsNo No YesN Y N Gen 1 dataflows will indeed be deployed: https://learn.microsoft.com/en-us/fabric/cicd/deployment-pipelines/create-rules?utm_source=chatgpt.com&tabs=new-ui
upvoted 51 times

SamuComqi Highly Voted 1 year, 10 months ago

NO: only metadata will be deployed to the target starge (https://learn.microsoft.com/en-us/fabric/cicd/deployment-pipelines/understand-the-deployment-process?tabs=new#item-properties-copied-during-deployment)

YES: Gen1 dataflows will be deployed

NO: Refresh schedules will not be deployed (https://learn.microsoft.com/en-us/fabric/cicd/deployment-pipelines/understand-the-deployment-process?tabs=new#item-properties-copied-during-deployment)

upvoted 50 times

92ef42d 1 year, 2 months ago

Gen1 is not in the list of the supported items: https://learn.microsoft.com/en-us/fabric/cicd/deployment-pipelines/intro-to-deployment-pipelines?tabs=new-ui#supported-items

upvoted 4 times

92ef42d 1 year, 2 months ago

Gen1 is not in the list of the supported items: https://learn.microsoft.com/en-us/fabric/cicd/deployment-pipelines/intro-to-deployment-pipelines?tabs=new-ui#supported-itemsin list but there some confusion = Get 2 is listed under Data Factory items = Gen 1 some times called dataflow only or legacy is listed under Power BI items

upvoted 4 times

Mustafa_Shazly 9 months, 2 weeks ago

in list but there some confusion = Get 2 is listed under Data Factory items = Gen 1 some times called dataflow only or legacy is listed under Power BI items

upvoted 1 times

ChenFu 1 year, 9 months ago

No No Yes

upvoted 8 times

ziggy1117 1 year, 3 months ago

N Y N Gen 1 dataflows will indeed be deployed: https://learn.microsoft.com/en-us/fabric/cicd/deployment-pipelines/create-rules?utm_source=chatgpt.com&tabs=new-ui

upvoted 2 times

Juanes55 1 year, 6 months ago

Data from the semantic models will be deployed to the target stage. No – The data itself is not deployed. Only the structure/definition (metadata) of the semantic model is deployed. You'll need to refresh the model in the target stage to populate it with data. The Dataflow Gen1 dataflows will be deployed to the target stage. No – Only Dataflow Gen2 is supported in Microsoft Fabric deployment pipelines. Gen1 dataflows are not deployable through Fabric pipelines. The scheduled refresh policies will be deployed to the target stage. No – Scheduled refresh settings are not included in deployments. You must configure them separately in the target workspace.No, Yes, No I have pipelines and they do deploy Gen1 Dataflows :) jokes apart, Indeed Dataflows Gen 1 are supported with deployment pipelines. Only problem you need to configure the gateways and migrate connections and schedules and creds also.

upvoted 16 times

Juanes55 1 year, 6 months ago

Data from the semantic models will be deployed to the target stage. No – The data itself is not deployed. Only the structure/definition (metadata) of the semantic model is deployed. You'll need to refresh the model in the target stage to populate it with data.

The Dataflow Gen1 dataflows will be deployed to the target stage. No – Only Dataflow Gen2 is supported in Microsoft Fabric deployment pipelines. Gen1 dataflows are not deployable through Fabric pipelines.

The scheduled refresh policies will be deployed to the target stage. No – Scheduled refresh settings are not included in deployments. You must configure them separately in the target workspace.

upvoted 13 times

Sr18 1 year, 6 months ago

No, Yes, No I have pipelines and they do deploy Gen1 Dataflows :) jokes apart, Indeed Dataflows Gen 1 are supported with deployment pipelines. Only problem you need to configure the gateways and migrate connections and schedules and creds also.

upvoted 3 times

Josekonoor Most Recent 8 months, 3 weeks ago

Model alone get deployed but no data; dataflow gen 1 can be deployed; Schedules do not get along with pipeline deployemnt

upvoted 1 times

Mustafa_Shazly 10 months, 3 weeks ago

No https://learn.microsoft.com/en-us/fabric/cicd/deployment-pipelines/understand-the-deployment-process?tabs=new-ui#assign-a-workspace-to-an-empty-stage Yes (ncremental refresh settings aren't copied in Gen 1.) mean Gen 1 will go but with limit https://learn.microsoft.com/en-us/fabric/cicd/deployment-pipelines/understand-the-deployment-process?tabs=new-ui#item-properties-that-are-not-copied No https://learn.microsoft.com/en-us/fabric/cicd/deployment-pipelines/understand-the-deployment-process?tabs=new-ui#item-properties-that-are-not-copied

upvoted 2 times

certs4pk 11 months ago

y is this question still asking on gen1? isn't it deprecated as of dec'25? should this question b updated?

upvoted 1 times

Psycho 11 months, 2 weeks ago

No, no, no

upvoted 2 times

mhal 1 year, 1 month ago

A. No. Data will not be transported B/C. Somehow a trick question in the sense, that (what I get from the documentation) B: Dataflow Gen1 will be cloned, but don't have really deployment support, so I would say No C: Refresh schedules are copied, but not activated, so I would again say No

upvoted 1 times

Acrious 1 year, 4 months ago

N, Y, N Gen1 flows are supported, just have to fix potential other things (connections etc)

upvoted 4 times

VenkataPhaniPavanKumar 1 year, 5 months ago

Answer 1) No 2) No 3) Yes Here's how each statement holds up in the context of Microsoft Fabric deployment pipelines: 1) Data from the semantic models will be deployed to the target stage → No Only the metadata (definition) of the semantic models is deployed. The actual data is not copied—you'll need to refresh the models in the target workspace after deployment. 2) The Dataflow Gen1 dataflows will be deployed to the target stage → No Dataflow Gen1 is not supported in Fabric deployment pipelines. Only Dataflow Gen2 items are deployable. 3) The scheduled refresh policies will be deployed to the target stage → Yes Scheduled refresh settings for semantic models are included in the deployment and will be carried over to the target stage

upvoted 4 times

Wipag 1 year, 5 months ago

Statement 1: Data from the semantic models will be deployed to the target stage. □ No • Explanation: Deployment pipelines do not deploy data from semantic models. Only the metadata (like tables, measures, and relationships) is deployed. Data must be refreshed in the target workspace after deployment. —— Statement 2: The Dataflow Gen1 dataflows will be deployed to the target stage. □ No • Explanation: Dataflow Gen1 items are not supported by Fabric deployment pipelines. Only Dataflow Gen2 is supported. Gen1 must be recreated or migrated manually. —— Statement 3: The scheduled refresh policies will be deployed to the target stage. □ No • Explanation: Scheduled refresh policies are not carried over in deployment pipelines. You must manually configure the refresh schedules in the target workspace after deploying.

upvoted 4 times

CertPeople 1 year, 5 months ago

1 -> NO, only metadata are deployed 2 -> NO, dataflow GEN1 are not supported by deployment pipelines 3 -> NO, Refresh schedules are semantic model object attributes that are not deployed by deployment pipelines

upvoted 4 times

NNPRN 1 year, 5 months ago

No, Yes, No

upvoted 1 times

5dcd3a8 1 year, 5 months ago

No/No/No 1.The following item properties aren't copied during deployment: Data - Data isn't copied. Only metadata is copied 2.When you deploy content from one pipeline stage to another, the copied content can contain the following items: Data Engineering items: Environment GraphQL (preview) Lakehouse (preview) Notebook Spark Job Definitions (preview) User Data Functions (preview) Data Factory items: Copy Job (preview) Dataflows gen2 3.The following semantic model properties are also not copied during deployment: Role assignment Refresh schedule https://learn.microsoft.com/en-us/fabric/cicd/deployment-pipelines/understand-the-deployment-process?tabs=new-ui#item-properties-that-are-not-copied https://learn.microsoft.com/en-us/fabric/cicd/deployment-pipelines/intro-to-deployment-pipelines?tabs=new-ui#supported-items

upvoted 6 times

Dwarakkrishna 1 year, 6 months ago

No, Yes, No Item properties copied during deployment 1.Data sources (deployment rules are supported) 2.Parameters (deployment rules are supported) 3.Report visuals 4.Report pages 5.Dashboard tiles 6.Model metadata 7.Item relationships 8.Sensitivity labels Item properties that are not copied 1.Data - Data isn't copied. Only metadata is copied 2.URL 3.ID 4.Permissions - For a workspace or a specific item 5.Workspace settings - Each stage has its own workspace 6.App content and settings - To update your apps, see Update content to Power BI apps 7.Personal bookmarks 8.The following semantic model properties are also not copied during deployment: 9.Role assignment 10.Refresh schedule 11.Data source credentials 12.Query caching settings (can be inherited from the capacity) 13.Endorsement settings Link-https://learn.microsoft.com/en-us/fabric/cicd/deployment-pipelines/understand-the-deployment-process?tabs=new-ui

upvoted 2 times

saikot 1 year, 7 months ago

No Yes Yes data from semantic models will be deployed in target: No. Deployment pipelines deploy the metadata of semantic models (datasets), not the actual data. The data in the target workspace will depend on the data sources configured for those semantic models and will need to be refreshed in the target environment. data flow gen1 will be deployed in target: Yes. Dataflow Gen1 dataflows are deployed as part of the pipeline. Their definition and query logic will be moved to the target workspace. However, similar to semantic models, the data they process will be refreshed based on the target environment's data sources and refresh configurations. scheduled refresh policy will be deployed in target: Yes. The scheduled refresh policy configured for the semantic models in the source workspace is typically deployed along with the semantic model metadata to the target workspace. You should verify the schedule in the target environment after deployment.

upvoted 1 times

fassil 1 year, 9 months ago

No/NO/NO Item properties that are not copied The following item properties aren't copied during deployment: Data - Data isn't copied. Only metadata is copied URL ID Permissions - For a workspace or a specific item Workspace settings - Each stage has its own workspace App content and settings - To update your apps, see Update content to Power BI apps Personal bookmarks The following semantic model properties are also not copied during deployment: Role assignment Refresh schedule Data source credentials Query caching settings (can be inherited from the capacity) Endorsement settings Incremental refresh settings aren't copied in Gen 1.

upvoted 5 times

mixonfreddy 1 year, 11 months ago

Selected Answer: A

Yes/Yes/No

upvoted 1 times

mixonfreddy 1 year, 11 months ago

Yes/Yes/NoNo/Yes/NoNo/No/No Deployment pipelines do not support Dataflow Gen1 dataflows. These need to be manually migrated or recreated in the target stage. Only Dataflow Gen2 is supported in deployment pipelines.I believe Dataflow Gen1 is supported (aka Power BI Dataflows) , see https://learn.microsoft.com/en-us/fabric/cicd/deployment-pipelines/intro-to-deployment-pipelines?tabs=new#supported-itemsGHill is correct. Power BI dataflows are listed in the first link. And here it is called Dataflow Gen1: https://learn.microsoft.com/en-us/fabric/data-factory/dataflows-gen2-overviewNO sorry again. I misread. The first question is - is DATA from semantic model copied. It is not. GHill is correct after all: https://learn.microsoft.com/en-us/fabric/cicd/deployment-pipelines/understand-the-deployment-process?tabs=new-ui#assign-a-workspace-to-an-empty-stage

upvoted 1 times

GHill1982 1 year, 11 months ago

No/Yes/No

upvoted 4 times

GHill1982 1 year, 11 months ago

No/Yes/NoNo/No/No Deployment pipelines do not support Dataflow Gen1 dataflows. These need to be manually migrated or recreated in the target stage. Only Dataflow Gen2 is supported in deployment pipelines.I believe Dataflow Gen1 is supported (aka Power BI Dataflows) , see https://learn.microsoft.com/en-us/fabric/cicd/deployment-pipelines/intro-to-deployment-pipelines?tabs=new#supported-itemsGHill is correct. Power BI dataflows are listed in the first link. And here it is called Dataflow Gen1: https://learn.microsoft.com/en-us/fabric/data-factory/dataflows-gen2-overviewNO sorry again. I misread. The first question is - is DATA from semantic model copied. It is not. GHill is correct after all: https://learn.microsoft.com/en-us/fabric/cicd/deployment-pipelines/understand-the-deployment-process?tabs=new-ui#assign-a-workspace-to-an-empty-stage

upvoted 4 times

Marco44 1 year, 10 months ago

No/No/No Deployment pipelines do not support Dataflow Gen1 dataflows. These need to be manually migrated or recreated in the target stage. Only Dataflow Gen2 is supported in deployment pipelines.I believe Dataflow Gen1 is supported (aka Power BI Dataflows) , see https://learn.microsoft.com/en-us/fabric/cicd/deployment-pipelines/intro-to-deployment-pipelines?tabs=new#supported-itemsGHill is correct. Power BI dataflows are listed in the first link. And here it is called Dataflow Gen1: https://learn.microsoft.com/en-us/fabric/data-factory/dataflows-gen2-overviewNO sorry again. I misread. The first question is - is DATA from semantic model copied. It is not. GHill is correct after all: https://learn.microsoft.com/en-us/fabric/cicd/deployment-pipelines/understand-the-deployment-process?tabs=new-ui#assign-a-workspace-to-an-empty-stage

upvoted 9 times

Marco44 1 year, 10 months ago

No/No/No

Deployment pipelines do not support Dataflow Gen1 dataflows. These need to be manually migrated or recreated in the target stage. Only Dataflow Gen2 is supported in deployment pipelines.

upvoted 9 times

GHill1982 1 year, 10 months ago

I believe Dataflow Gen1 is supported (aka Power BI Dataflows) , see https://learn.microsoft.com/en-us/fabric/cicd/deployment-pipelines/intro-to-deployment-pipelines?tabs=new#supported-items

upvoted 9 times

GHill1982 1 year, 10 months ago

I believe Dataflow Gen1 is supported (aka Power BI Dataflows) , see https://learn.microsoft.com/en-us/fabric/cicd/deployment-pipelines/intro-to-deployment-pipelines?tabs=new#supported-itemsGHill is correct. Power BI dataflows are listed in the first link. And here it is called Dataflow Gen1: https://learn.microsoft.com/en-us/fabric/data-factory/dataflows-gen2-overviewNO sorry again. I misread. The first question is - is DATA from semantic model copied. It is not. GHill is correct after all: https://learn.microsoft.com/en-us/fabric/cicd/deployment-pipelines/understand-the-deployment-process?tabs=new-ui#assign-a-workspace-to-an-empty-stage

upvoted 9 times

SannaLH 1 year, 8 months ago

GHill is correct. Power BI dataflows are listed in the first link. And here it is called Dataflow Gen1: https://learn.microsoft.com/en-us/fabric/data-factory/dataflows-gen2-overviewNO sorry again. I misread. The first question is - is DATA from semantic model copied. It is not. GHill is correct after all: https://learn.microsoft.com/en-us/fabric/cicd/deployment-pipelines/understand-the-deployment-process?tabs=new-ui#assign-a-workspace-to-an-empty-stage

upvoted 2 times

SannaLH 1 year, 8 months ago

GHill is correct. Power BI dataflows are listed in the first link. And here it is called Dataflow Gen1: https://learn.microsoft.com/en-us/fabric/data-factory/dataflows-gen2-overview

upvoted 2 times

SannaLH 1 year, 8 months ago

NO sorry again. I misread. The first question is - is DATA from semantic model copied. It is not. GHill is correct after all: https://learn.microsoft.com/en-us/fabric/cicd/deployment-pipelines/understand-the-deployment-process?tabs=new-ui#assign-a-workspace-to-an-empty-stage

upvoted 1 times

Question #12 *Topic 1*

You have a Fabric deployment pipeline that uses three workspaces named Dev, Test, and Prod.
You need to deploy an eventhouse as part of the deployment process.
What should you use to add the eventhouse to the deployment process?

A. GitHub Actions

B. a deployment pipeline

C. an Azure DevOps pipeline

Suggested Answer: *B*

Community vote distribution

B (97%) (3

xDeadEnd123x Most Recent 3 weeks ago

Selected Answer: C

deployment pipeline

upvoted 1 times

DasyGY 1 month, 2 weeks ago

Selected Answer: B

deployement pipeline

upvoted 1 times

Josekonoor 8 months, 3 weeks ago

Selected Answer: B

deployment pipeline

upvoted 1 times

Wipag Highly Voted 1 year, 5 months ago

Selected Answer: B

Microsoft Fabric deployment pipelines now support eventhouses along with other Fabric items like:

• Lakehouses

• Warehouses

• Notebooks

• Dataflows Gen2

• Data pipelines

• Semantic models

• Reports

• Eventhouses

upvoted 7 times

smanzana 1 year, 5 months ago

Selected Answer: B

B- deployment pipeline or event stream

upvoted 1 times

DarioReymago 1 year, 5 months ago

Selected Answer: B

B las otras opciones son repositorios externos no generan el movimiento

upvoted 1 times

alirana3 1 year, 7 months ago

Selected Answer: B

B. a deployment pipeline

Here's why:

Microsoft Fabric deployment pipelines are designed for managing the deployment of resources (including eventhouses) across different environments (e.g., Dev, Test, Prod). You use a deployment pipeline to automate and control the deployment process across workspaces, ensuring that the eventhouse (along with other resources) is deployed correctly through the stages of your workflow.

GitHub Actions and Azure DevOps pipelines could also be used for deployment, but the Fabric deployment pipeline is specifically built for this use case within Microsoft Fabric and integrates directly with the workspace and stage management.

For details, check the official documentation:

Get started with deployment pipelines

https://learn.microsoft.com/en-us/fabric/cicd/deployment-pipelines/get-started-with-deployment-pipelines?tabs=from-fabric%2Cnew-ui

upvoted 2 times

hebertorosillo 1 year, 9 months ago

Selected Answer: C

This is a use case question. The company Contoso already has DevOps.for my is letter C.

upvoted 1 times

QAZdbarhate12345678 Highly Voted 1 year, 11 months ago

Selected Answer: B

In Microsoft Fabric, deployment pipelines are the recommended and built-in method for managing and deploying artifacts (like eventhouses, lakehouses, datasets, and reports) across environments such as Dev, Test, and Prod.

upvoted 17 times

Tuki93 1 year, 11 months ago

Selected Answer: B

https://learn.microsoft.com/en-us/fabric/cicd/deployment-pipelines/get-started-with-deployment-pipelines?tabs=from-fabric%2Cnew%2Cstage-settings-new

https://learn.microsoft.com/en-us/fabric/cicd/deployment-pipelines/understand-the-deployment-process?tabs=new

upvoted 2 times

Question #13 *Topic 1*

You have a Fabric workspace named Workspace1 that contains a warehouse named Warehouse1.
You plan to deploy Warehouse1 to a new workspace named Workspace2.
As part of the deployment process, you need to verify whether Warehouse1 contains invalid references. The solution must minimize development effort.
What should you use?

A. a database project

B. a deployment pipeline

C. a Python script

D. a T-SQL script

Suggested Answer: *B*

Community vote distribution

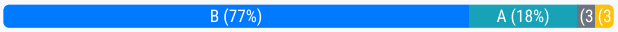

QAZdbarhate12345678 Highly Voted 1 year, 11 months ago

Microsoft Fabric's deployment pipelines provide a built-in mechanism to manage and validate the deployment of artifacts like warehouses. When you use a deployment pipeline to move Warehouse1 from one workspace (Workspace1) to another (Workspace2), the pipeline automatically checks for issues such as invalid references or missing dependencies during the deployment process.What kind of references may a Warehouse have? We don't have shortcuts in a warehouse. It seems to be an indepentent item, there is no possiblity to break a reference to anything external. Therfor I think this relates to internal references, i.e. stored procedure operating on non-existing table. In this case a database project (answer A) would be necessary, but I don't find it related to a Deployment Pipeline in any way.agreed with answer A. Here is the hint. https://community.fabric.microsoft.com/t5/Data-Engineering/Microsoft-fabric-WareHouse-migration-issue-via-deployment/td-p/4281976A https://learn.microsoft.com/en-us/fabric/cicd/deployment-pipelines/intro-to-deployment-pipelines?tabs=new-ui#supported-itemsB: Its preview like most of the things in Fabric but it does work. Read deployment with pipelines. Don't get distracted with git. https://learn.microsoft.com/en-us/fabric/data-warehouse/source-control#deployment-pipelinesAnd believe me MS is still in Marketing mode, so even if it breaks sometimes. The feature is there and you must use it specially deployment pipeline :)

upvoted 21 times

QAZdbarhate12345678 Highly Voted 1 year, 11 months ago

Selected Answer: B

Microsoft Fabric's deployment pipelines provide a built-in mechanism to manage and validate the deployment of artifacts like warehouses. When you use a deployment pipeline to move Warehouse1 from one workspace (Workspace1) to another (Workspace2), the pipeline automatically checks for issues such as invalid references or missing dependencies during the deployment process.

upvoted 21 times

alqazqzy.asseel Most Recent 1 month, 1 week ago

Selected Answer: A

deployment pipeline moves items from Workspace1 to Workspace2, it does not inspect the code for invalid database reference bugs prior to deployment.

upvoted 5 times

tomaszstaroszczyk 1 year, 10 months ago

What kind of references may a Warehouse have? We don't have shortcuts in a warehouse. It seems to be an indepentent item, there is no possiblity to break a reference to anything external. Therfor I think this relates to internal references, i.e. stored procedure operating on non-existing table. In this case a database project (answer A) would be necessary, but I don't find it related to a Deployment Pipeline in any way.

upvoted 4 times

tomaszstaroszczyk 1 year, 10 months ago

What kind of references may a Warehouse have? We don't have shortcuts in a warehouse. It seems to be an indepentent item, there is no possiblity to break a reference to anything external. Therfor I think this relates to internal references, i.e. stored procedure operating on non-existing table. In this case a database project (answer A) would be necessary, but I don't find it related to a Deployment Pipeline in any way.agreed with answer A. Here is the hint. https://community.fabric.microsoft.com/t5/Data-Engineering/Microsoft-fabric-WareHouse-migration-issue-via-deployment/td-p/4281976A https://learn.microsoft.com/en-us/fabric/cicd/deployment-pipelines/intro-to-deployment-pipelines?tabs=new-ui#supported-itemsB: Its preview like most of the things in Fabric but it does work. Read deployment with pipelines. Don't get distracted with git. https://learn.microsoft.com/en-us/fabric/data-warehouse/source-control#deployment-pipelinesAnd believe me MS is still in Marketing mode, so even if it breaks sometimes. The feature is there and you must use it specially deployment pipeline :)

upvoted 5 times

juamd 6 months, 2 weeks ago

Selected Answer: A

Deployment pipelines are not capable of detecting invalid references in T-SQL code.

upvoted 5 times

toniteo 1 year, 9 months ago

agreed with answer A. Here is the hint. https://community.fabric.microsoft.com/t5/Data-Engineering/Microsoft-fabric-WareHouse-migration-issue-via-deployment/td-p/4281976A https://learn.microsoft.com/en-us/fabric/cicd/deployment-pipelines/intro-to-deployment-pipelines?tabs=new-ui#supported-itemsB: Its preview like most of the things in Fabric but it does work. Read deployment with pipelines. Don't get distracted with git. https://learn.microsoft.com/en-us/fabric/data-warehouse/source-control#deployment-pipelinesAnd believe me MS is still in Marketing mode, so even if it breaks sometimes. The feature is there and you must use it specially deployment pipeline :)

upvoted 4 times

toniteo 1 year, 9 months ago

agreed with answer A. Here is the hint. https://community.fabric.microsoft.com/t5/Data-Engineering/Microsoft-fabric-WareHouse-migration-issue-via-deployment/td-p/4281976

upvoted 4 times

32950cb 7 months, 1 week ago

Selected Answer: A

To verify whether a Fabric Warehouse contains invalid references (such as views or stored procedures referencing tables or columns that have been deleted or renamed) with minimal effort, a database project is the most effective tool.

upvoted 6 times

CloudKaKida 1 year, 7 months ago

A https://learn.microsoft.com/en-us/fabric/cicd/deployment-pipelines/intro-to-deployment-pipelines?tabs=new-ui#supported-items

upvoted 1 times

CloudKaKida 1 year, 7 months ago

A https://learn.microsoft.com/en-us/fabric/cicd/deployment-pipelines/intro-to-deployment-pipelines?tabs=new-ui#supported-itemsB: Its preview like most of the things in Fabric but it does work. Read deployment with pipelines. Don't get distracted with git. https://learn.microsoft.com/en-us/fabric/data-warehouse/source-control#deployment-pipelinesAnd believe me MS is still in Marketing mode, so even if it breaks sometimes. The feature is there and you must use it specially deployment pipeline :)

upvoted 1 times

Josekonoor 8 months, 3 weeks ago

Selected Answer: B

deployment pipeline

upvoted 1 times

Sr18 1 year, 6 months ago

B: Its preview like most of the things in Fabric but it does work. Read deployment with pipelines. Don't get distracted with git. https://learn.microsoft.com/en-us/fabric/data-warehouse/source-control#deployment-pipelinesAnd believe me MS is still in Marketing mode, so even if it breaks sometimes. The feature is there and you must use it specially deployment pipeline :)

upvoted 2 times

Sr18 1 year, 6 months ago

B: Its preview like most of the things in Fabric but it does work. Read deployment with pipelines. Don't get distracted with git. https://learn.microsoft.com/en-us/fabric/data-warehouse/source-control#deployment-pipelines

upvoted 2 times

EddEe 10 months ago

Selected Answer: B

https://blog.fabric.microsoft.com/en-us/blog/bridging-the-gap-automate-warehouse-sql-endpoint-deployment-in-microsoft-fabric?ft=Lakehouse:category#:~:text=Warehouse%20with%20Warehouse%20dependencies,warehouse%20connections%20are%20functioning%20properly.

upvoted 1 times

Rjb_073086 10 months ago

Selected Answer: B

while deployment pipelines do this (hence i picked B), just clicking view dependencies from warehouse does the job with zero dev effort

upvoted 2 times

8115575 1 year, 3 months ago

Selected Answer: B

a database project is more relevant for traditional SQL Server development using tools like Visual Studio; not integrated with Fabric's deployment model.

Scripts are not minimizing effort.

upvoted 3 times

bcbc59b 1 year, 4 months ago

Selected Answer: B

Since the question does not define any constraints regarding how the verification result should be presented or treated, according to the documentation, I'd go with deployment pipelines: They check dependencies https://learn.microsoft.com/en-us/fabric/cicd/deployment-pipelines/understand-the-deployment-process?utm_source=chatgpt.com&tabs=new-ui#autobinding-in-the-same-workspace

upvoted 1 times

jaysun 1 year, 4 months ago

Selected Answer: B

Detects invalid references automatically during deployment.

Minimal configuration needed — aligns directly with your goal to "minimize development effort."

upvoted 1 times

smanzana 1 year, 5 months ago

Selected Answer: B

B is correct

upvoted 1 times

upliftinghut 1 year, 6 months ago

Selected Answer: D

T-SQL below:

SELECT

referencing_object_name = OBJECT_NAME(referencing_id),

referenced_entity_name,

is_ambiguous,

is_missing_reference

FROM

sys.sql_expression_dependencies

WHERE

is_missing_reference = 1;

upvoted 1 times

hebertorosillo 1 year, 9 months ago

Selected Answer: A

A. a database project for minimal effor

upvoted 6 times

Anushareddysri 1 year, 9 months ago

Selected Answer: C

Question is focusing on "As part of the deployment process, you need to verify whether Warehouse1 contains invalid references" for this python script can be written to identify invalid references and can add it as a notebook activity in pipeline. So answer is python script.

upvoted 1 times

Anushareddysri 1 year, 9 months ago

Selected Answer: C

Question is focusing on "As part of the deployment process, you need to verify whether Warehouse1 contains invalid references" for this python script can be written to identify invalid references and can add it as a notebook activity in pipeline. So answer is python script.

upvoted 1 times

toniteo 1 year, 9 months ago

Python script could be a lengthy solution, while the question asks for minimal effort.

upvoted 2 times

Sr18 1 year, 6 months ago

And believe me MS is still in Marketing mode, so even if it breaks sometimes. The feature is there and you must use it specially deployment pipeline :)

upvoted 1 times

Question #14 *Topic 1*

You have a Fabric workspace that contains a Real-Time Intelligence solution and an eventhouse.
Users report that from OneLake file explorer, they cannot see the data from the eventhouse.
You enable OneLake availability for the eventhouse.
What will be copied to OneLake?

A. only data added to new databases that are added to the eventhouse

B. only the existing data in the eventhouse

C. no data

D. both new data and existing data in the eventhouse

E. only new data added to the eventhouse

Suggested Answer: *E*

Community vote distribution

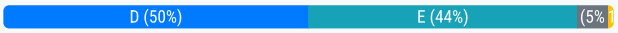

Tuki93 Highly Voted 1 year, 11 months ago

Selected Answer: E

- Existing tables aren't affected. New tables are made available in OneLake.

https://learn.microsoft.com/en-us/fabric/real-time-intelligence/event-house-onelake-availability#how-it-works

upvoted 24 times

Tuki93 Highly Voted 1 year, 11 months ago

- Existing tables aren't affected. New tables are made available in OneLake. https://learn.microsoft.com/en-us/fabric/real-time-intelligence/event-house-onelake-availability#how-it-worksDoc say new table not new dataagreed with answer E after reading and evaluating the article carefully.As per documentation, now you have option to select historical backfill. So I think question's verbiage is incomplete now. Answer D : Seems more right answer now, because you have this option and may be it was selected while making the One Lake Availability.It is still E. You can turn on OneLake availability at the database or table level. When enabled at the database level, all new tables and their data are made available in OneLake. When turning on the feature, >>. Turning on at the table level makes only that table and its data available in OneLake Only if the question mentions explicitly together the "Apply to existing tables option", otherwise, only "enable OneLake availability" does not do the backfill.

upvoted 26 times

soonapaana001 Most Recent 3 weeks ago

Selected Answer: D

Why D is Correct

In Microsoft Fabric Real-Time Intelligence, enabling OneLake availability at the Eventhouse/database level includes historic data backfill:

According to the official documentation ("Turn on OneLake availability for an eventhouse"), when turning on the feature, it allows including historic backfill so that existing tables and historical data already resident in the Eventhouse are converted and mirrored into Delta Parquet format in OneLake.

Going forward, any new data streamed or ingested into the Eventhouse is continuously synchronized.

Where the Confusion Came From

During earlier preview releases of Microsoft Fabric Real-Time Analytics, enabling mirroring/OneLake availability on a KQL database only streamed incoming data prospectively (requiring manual re-ingestion for legacy data). Microsoft subsequently updated the behavior to support historical backfill upon enablement, making D the intended answer for the current Fabric DP-600 / DP-700 exams.

upvoted 2 times

pmg314 3 weeks ago

Selected Answer: D

Enabling OneLake availability in a Microsoft Fabric KQL database or eventhouse includes both new data and existing historic data

upvoted 1 times

alqazqzy.asseel 1 month, 1 week ago

Selected Answer: D

D is the correct answer if existing tables are included when enabling availability

upvoted 1 times

Sr18 1 year, 6 months ago

As per documentation, now you have option to select historical backfill. So I think question's verbiage is incomplete now. Answer D : Seems more right answer now, because you have this option and may be it was selected while making the One Lake Availability.

upvoted 7 times

Sr18 1 year, 6 months ago

As per documentation, now you have option to select historical backfill. So I think question's verbiage is incomplete now. Answer D : Seems more right answer now, because you have this option and may be it was selected while making the One Lake Availability.It is still E. You can turn on OneLake availability at the database or table level. When enabled at the database level, all new tables and their data are made available in OneLake. When turning on the feature, >>. Turning on at the table level makes only that table and its data available in OneLake Only if the question mentions explicitly together the "Apply to existing tables option", otherwise, only "enable OneLake availability" does not do the backfill.

upvoted 7 times

gpasini 1 month, 1 week ago

Selected Answer: E

Copilot seem to suggest that this is a common DP-700 trap : if you don't specify "Apply to existing tables" it just goes with the new data

upvoted 2 times

DasyGY 1 month, 2 weeks ago

Selected Answer: E

Onelake availability at database level implies only new tables and new data are set to be available in onelake. Backfill is done when option enabled but it's not automatic. So the answer is E

upvoted 2 times

dila123 1 month, 3 weeks ago

Selected Answer: D

OneLake availability = ongoing new data + existing data through backfill

upvoted 1 times

CamiU 5 months ago

Selected Answer: D

You can turn on OneLake availability at the database or table level. When enabled at the database level, all new tables and their data are made available in OneLake. When turning on the feature, you can also choose to apply this option to existing tables by selecting the Apply to existing tables option, to include historic backfill. Turning on at the table level makes only that table and its data available in

upvoted 4 times

montoya207 5 months, 2 weeks ago

Selected Answer: E

The correct answer is E. only new data added to the eventhouse.

Here is the technical justification based on the sources:

No Retroactive Syncing: When you enable OneLake integration for a KQL database or eventhouse, it does not "backfill" or copy data that already exists in the tables

.

New Data Only: The integration is designed to work only for new tables created after the setting is enabled or for new data ingested from that point forward

.

Default Storage: By default, KQL data is stored in the eventhouse engine and is not visible in OneLake unless this specific setting is toggled, and even then, it only applies to subsequent data movements

upvoted 2 times

nicolerita 7 months, 1 week ago

Selected Answer: C

Enabling OneLake availability for an eventhouse in Microsoft Fabric does not copy any data into OneLake. Instead, it:

Exposes the eventhouse data in place through OneLake so it can be accessed via the OneLake file explorer and APIs.

Makes the eventhouse appear as a logical item in the OneLake namespace, but no existing or new data is physically copied as part of simply enabling availability.

upvoted 1 times

garrity0007 7 months, 3 weeks ago

Selected Answer: E

Solo los datos nuevos. No se por que hay tantas dudas en esta pregunta.

upvoted 2 times

MedMed009 7 months, 4 weeks ago

Selected Answer: C

When you enable OneLake availability for an eventhouse, it does not physically copy data; instead, it creates a shortcut in OneLake that points to the eventhouse storage. This shortcut makes all existing data and any new data added later accessible through OneLake file explorer.

upvoted 1 times

Rahalcer 9 months, 1 week ago

Selected Answer: E

Only new data : “When enabled at the database level, all new tables and their data are made available in OneLake. When turning on the feature, you can also choose to apply this option to existing tables by selecting the Apply to existing tables option, to include historic backfill.”

Link : https://learn.microsoft.com/en-us/fabric/real-time-intelligence/event-house-onelake-availability

upvoted 2 times

Mustafa_Shazly 9 months, 2 weeks ago

Selected Answer: C

turn on OneLake availability at the database or table level

- database level >> new tables and their data are made available in OneLake. (existing tables need to anther opttion >> Apply to existing tables option)

=So he didn't say (select Apply to existing tables option) and say that DB level and no new table >>>>>> So >> no data

- A wrong because new database need to enable feature first (database level)

- B Wrong (true if we make it on table level or DB level with Apply to existing tables option)

- D and E Wrong because no new table (DB level need new table after enable option)

In the last users need to see old data in exist table so we need to enable OneLake availability with Apply to existing tables option

upvoted 1 times

EddEe 10 months ago

Selected Answer: B

I read this as the option having been turned off and on in a sequence

Turning off OneLake availability soft deletes your data from OneLake. When you turn availability back on, all data is made available in OneLake, including historic backfill.

upvoted 1 times

Rjb_073086 10 months ago

Selected Answer: D

should be D as it copies all of the data from the delta tables

upvoted 1 times

Psycho 11 months, 2 weeks ago

Selected Answer: D

When you enable OneLake availability for an eventhouse, it does not physically copy data; instead, it creates a shortcut in OneLake that points to the eventhouse storage. This shortcut makes all existing data and any new data added later accessible through OneLake file explorer

upvoted 2 times

Psycho 11 months, 2 weeks ago

Selected Answer: E

Only for new data

upvoted 2 times

larinate 11 months, 2 weeks ago

Selected Answer: C

When enable OneLake availability for an Eventhouse, NOTHING is copied automatically.

Simply enabling OneLake availability does not copy any data at all. A new real-time connection is created that streams future ingestions into OneLake and only after new data would be visible

upvoted 1 times

Ana_L 11 months, 2 weeks ago

Selected Answer: C

Action What Happens?

Enable OneLake availability Shortcut created

Existing data copied? □ No

New data copied? □ No

Data remains in eventhouse? □ Yes

upvoted 1 times

Ron8900 11 months, 3 weeks ago

Selected Answer: E

The question should be clear.

upvoted 2 times

johnny78 11 months, 3 weeks ago

Selected Answer: C

it only creates a shortcut. no acutal data copied

upvoted 1 times

ageorgieva 11 months, 3 weeks ago

Selected Answer: D

I did that in real time and it is everything - past and new data.

upvoted 2 times

learnazureportal 11 months, 3 weeks ago

Selected Answer: D

D - enabling OneLake availability for an eventhouse copies both the existing data and continuously ingested new data to OneLake, making the full dataset accessible.

upvoted 2 times

Bin_Hashim 1 year ago

Selected Answer: D

Microsoft's current docs say when you turn on OneLake availability for an eventhouse (KQL database), you can apply it to existing tables to include historic backfill, and once enabled it continues to sync new data as it arrives. So users will see both the historical data and ongoing new data in OneLake. Microsoft Learn

"When enabled at the database level… you can also choose to apply this option to existing tables, to include historic backfill." Microsoft Learn

"When you turn availability back on, all data is made available in OneLake, including historic backfill." Microsoft Learn

Note: Earlier guidance (pre-2025) often said "only new data" with no backfill. The behavior has since changed/been clarified in the official docs.

upvoted 7 times

tall 1 year ago

Selected Answer: D

When you enable OneLake availability for an eventhouse, both new and existing data in the eventhouse will be copied to OneLake. This feature ensures that data, whether newly ingested or already present, becomes available for access through OneLake, making it easier for users to interact with and explore the data directly from OneLake file explorer.

upvoted 3 times

ShazzyHawk 1 year, 1 month ago

Selected Answer: D

Earlier behavior (Preview phase, mid-2024)

When OneLake availability for Eventhouse was first introduced,

only new data ingested after enabling the feature was written to OneLake.

Existing data remained in Kusto storage unless re-ingested or exported manually.

That's why older documentation (and many early exam questions) listed

E. only new data added to the eventhouse as the correct answer. Current behavior (GA – 2025 update)

Per the Microsoft Fabric Real-Time Intelligence documentation, when you enable OneLake availability on an Eventhouse now:

"When OneLake availability is turned on, existing data in the Eventhouse and any new data ingested afterward is made available in OneLake in Delta format."

So now:

Existing data is copied and made visible in OneLake.

New data continues to sync automatically going forward.

upvoted 6 times

wellingtonluis 1 year, 2 months ago

Selected Answer: E

You have the possibility to choose existing data, but only by enabling the option to include this, the data is not presented automatically and as the statement does not mention "possibility", the final answer is E.

upvoted 2 times

kaytee_05 1 year, 4 months ago

Selected Answer: D

When you enable OneLake availability for an eventhouse, both the existing data withini the eventhouse and any new data subsequently added to it will be copied to OneLake. This ensures that the eventhouse data is acceble through OneLake file explorer for users who previously couldn't see it.

upvoted 3 times

Acrious 1 year, 4 months ago

Selected Answer: D

I am going with D, as many have mentioned you can enable it to copy the existing data. The question is... is the question old and the question text just not updated OR are they truly trying to trick you... for anyone who has taken the test recently is the content updated to say you selected the option to include existing data?

upvoted 2 times

Thanos29 1 year, 5 months ago

Selected Answer: E

E is Correct.

Question says : User cannot see the data from the eventhouse.

If OneLake Availablilty is enavled, only New Tables are made available. Also no mention of whether the option : Apply to exisiting Table is checked, and also Users cannot see data, meaning Users cannot see Old data, as only new Tables appear

https://learn.microsoft.com/en-us/fabric/real-time-intelligence/event-house-onelake-availability#how-it-works

upvoted 2 times

VenkataPhaniPavanKumar 1 year, 5 months ago

Selected Answer: D

Best thing to check the answers is to use Microsoft Copilot

Copilot says

When you enable OneLake availability for an eventhouse in Microsoft Fabric, the platform copies both existing and newly ingested data into the OneLake-backed file storage. This ensures that users can access all relevant telemetry data—past and present—through tools like the OneLake file explorer or other storage-compatible interfaces.

This is particularly helpful when enabling lake-centric analytics or blending real-time and historical data into broader lakehouse or warehouse solutions.

upvoted 3 times

wellingtonluis 1 year, 2 months ago

Didn't you realize yet. Copilot and ChatGPT will always give you a "Yes" answer. Even if it is not a full answer. Because yes it is possible, but not just enabling availability.

upvoted 1 times

mvg100 1 year, 5 months ago

Selected Answer: D

This explains why both new and existing data will be copied to OneLake.

Check this link:

https://learn.microsoft.com/en-us/fabric/real-time-intelligence/event-house-onelake-availability#how-it-works

Also chatbot says:

When you enable OneLake availability for an eventhouse in Microsoft Fabric:

Existing data in the eventhouse becomes immediately available in OneLake.

New data added to the eventhouse is also automatically streamed and accessible via OneLake moving forward.

This allows users to access both historical and real-time streaming data through OneLake interfaces, including the OneLake file explorer.

upvoted 1 times

Jay63002 1 year, 5 months ago

Selected Answer: D

You can turn on OneLake availability at the database or table level. When enabled at the database level, all new tables and their data are made available in OneLake. When turning on the feature, you can also choose to apply this option to existing tables by selecting the Apply to existing tables option, to include historic backfill. Turning on at the table level makes only that table and its data available in OneLake. The Data retention policy of your KQL database is also applied to the data in OneLake. Data removed from your KQL database at the end of the retention period is also removed from OneLake. If you turn off OneLake availability, data is soft deleted from OneLake.

upvoted 2 times

AnetaK 1 year, 7 months ago

Selected Answer: D

This is what chatbot said: "When you enable OneLake availability for the eventhouse, option D is correct: both new data and existing data in the eventhouse will be copied to OneLake. This ensures that users can access all the data from the eventhouse through OneLake file explorer."

upvoted 1 times

zxc01 1 year, 8 months ago

Selected Answer: E

We can enable it on both database and table levels. I think Microsoft just want to ask for default settings with them. on database level, default option is "Apply to existing tables" which will map all tables to Onelake but don't include any old data. same issue when we try it on table level though UI.

upvoted 2 times

CMDev 1 year, 8 months ago

Selected Answer: D

The correct answer is:

D. both new data and existing data in the eventhouse

When you enable OneLake availability for an eventhouse, you have the option to apply this setting to existing tables by selecting the "Apply to existing tables" option. This means that both new data and existing data will be available in OneLake.

As stated in the article: "When turning on the feature, you can also choose to apply this option to existing tables by selecting the Apply to existing tables option, to include historic backfill."

This explains why both new and existing data will be copied to OneLake.

Check this link:

https://learn.microsoft.com/en-us/fabric/real-time-intelligence/event-house-onelake-availability#how-it-works

upvoted 5 times

CMDev 1 year, 8 months ago

When you enable OneLake availability for an eventhouse, you have the option to apply this setting to existing tables, which includes the data already present in those tables. This means that both new data (data added after enabling OneLake availability) and existing data (data that was already in the eventhouse at the time of enabling OneLake availability) will be available in OneLake.

As the article states: "When turning on the feature, you can also choose to apply this option to existing tables by selecting the Apply to existing tables option, to include historic backfill."

This indicates that enabling OneLake availability can include the historical data, meaning the existing data, in addition to any new data added afterwards.

upvoted 1 times

Kiket2ride 1 year, 8 months ago

Selected Answer: D

I think it used to be that only new data was available but documentation now says otherwise:

https://learn.microsoft.com/en-us/fabric/real-time-intelligence/event-house-onelake-availability#how-it-works

When you turn availability back on, all data is made available in OneLake, including historic backfill.

I vote for D then

upvoted 4 times

DarkDerf 1 year, 9 months ago

Selected Answer: E

If we strictly follow the documentation's explanation that only new tables are made available in OneLake, then your choice E. only new data added to the eventhouse is the most precise answer.

 Final Answer: E. only new data added to the eventhouse

upvoted 3 times

DarkDerf 1 year, 9 months ago

Selected Answer: D

The correct answer is:

D. both new data and existing data in the eventhouse

Explanation:

When you enable OneLake availability for an eventhouse in Microsoft Fabric, all existing data within the eventhouse, along with any new data added afterward, is copied to OneLake. This ensures that users can access both past and future data through OneLake file explorer.

This means that enabling OneLake availability makes the entire dataset in the eventhouse visible in OneLake, not just new data moving forward.

upvoted 2 times

henryphchan 1 year, 9 months ago

Selected Answer: E

Only new data will be loaded to OneLake

upvoted 5 times

Anushareddysri 1 year, 9 months ago

Selected Answer: D

As per the Microsoft document you can enable onelake along with an option of existing tables as well. Here is the reference document.
https://learn.microsoft.com/en-us/fabric/real-time-intelligence/event-house-onelake-availability

upvoted 2 times

henryphchan 1 year, 9 months ago

It applies to existing tables but not exist data, so I vote for E

upvoted 1 times

Sr18 1 year, 6 months ago

Historial Backfill means data + tables : Read properly "You can turn on OneLake availability at the database or table level. When enabled at the database level, all new tables and their data are made available in OneLake. When turning on the feature, you can also choose to apply this option to existing tables by selecting the Apply to existing tables option, to include historic backfill. Turning on at the table level makes only that table and its data available in OneLake."

upvoted 1 times

h2o_molecule 1 year, 9 months ago

Selected Answer: D

You can turn on OneLake availability at the database or table level. When enabled at the database level, all new tables and their data are made available in OneLake. When turning on the feature, you can also choose to apply this option to existing tables by selecting the Apply to existing tables option.

upvoted 3 times

e0f0ce6 1 year, 9 months ago

Selected Answer: E

I think E is correct:

https://learn.microsoft.com/en-us/fabric/real-time-intelligence/event-house-onelake-availability

upvoted 3 times

QAZdbarhate12345678 1 year, 11 months ago

Selected Answer: D

The correct answer is:

D. both new data and existing data in the eventhouse

When you enable OneLake availability for the eventhouse, both the existing data and any new data added to the eventhouse will be copied to OneLake. This allows users to access the eventhouse data from OneLake file explorer.

upvoted 5 times

IshtarSQL 1 year, 11 months ago

Selected Answer: E

When you enable OneLake availability for an eventhouse in Microsoft Fabric, only new data ingested into the eventhouse will be synced and made available in OneLake.

upvoted 11 times

Mustafa_Shazly 9 months, 2 weeks ago

Doc say new table not new data

upvoted 1 times

exnaniantwort 1 year, 5 months ago

It is still E.

You can turn on OneLake availability at the database or table level. When enabled at the database level, all new tables and their data are made available in OneLake. When turning on the feature, <<<you can also choose to apply this option to existing tables by selecting the Apply to existing tables option, to include historic backfill>>>. Turning on at the table level makes only that table and its data available in OneLake

Only if the question mentions explicitly together the "Apply to existing tables option", otherwise, only "enable OneLake availability" does not do the backfill.

upvoted 9 times

toniteo 1 year, 9 months ago

agreed with answer E after reading and evaluating the article carefully.

upvoted 2 times

Question #15 *Topic 1*

You have a Fabric workspace named Workspace1.
You plan to integrate Workspace1 with Azure DevOps.
You will use a Fabric deployment pipeline named deployPipeline1 to deploy items from Workspace1 to higher environment workspaces as part of a medallion architecture. You will run deployPipeline1 by using an API call from an Azure DevOps pipeline.
You need to configure API authentication between Azure DevOps and Fabric.
Which type of authentication should you use?

A. service principal

B. Microsoft Entra username and password

C. managed private endpoint

D. workspace identity

Suggested Answer: *A*

Community vote distribution



gpasini Most Recent 1 month, 1 week ago

Selected Answer: A

This is a common DP-600/DP-700 style question: whenever an Azure DevOps pipeline or another external automation tool needs to call Fabric REST APIs (including deployment pipeline APIs), the expected answer is usually Service Principal rather than a user account or managed identity.
https://learn.microsoft.com/en-us/fabric/cicd/git-integration/automate-git-integration-with-service-principal

upvoted 1 times

jfranzen 1 year, 3 months ago

Selected Answer: A

The relevant documentation I found to support Service Principals is this: https://learn.microsoft.com/en-us/rest/api/fabric/articles/get-started/fabric-api-quickstart

Specifically calling out integrating with the API by creating a Service Principal (App Registration) in the Entra tenant.

upvoted 2 times

asitp9 1 year, 4 months ago

Selected Answer: A

A is correct

upvoted 1 times

VenkataPhaniPavanKumar Highly Voted 1 year, 5 months ago

Selected Answer: A

A service principal is the recommended authentication method when automating deployments and integrations between Azure DevOps and Microsoft Fabric. It provides a secure, scalable, and manageable way to authenticate API calls—especially from CI/CD systems like Azure DevOps pipelines—without requiring user credentials.

upvoted 6 times

smanzana 1 year, 5 months ago

Selected Answer: A

A is correct

upvoted 1 times

a85becd 1 year, 7 months ago

Selected Answer: A

Announcing Service Principal support for Fabric APIs:

https://blog.fabric.microsoft.com/en-au/blog/announcing-service-principal-support-for-fabric-apis/#:~:text=Service%20principal%20is%20a%20security,integrate%20them%20with%20other%20services.

upvoted 2 times

henryphchan 1 year, 9 months ago

A service principal is an identity created for use with applications, hosted services, and automated tools to access Azure resources. It provides a secure, non-interactive way for Azure DevOps pipelines to authenticate and interact with Fabric APIs without needing user credentials.Service principal is only supported for Power BI Items deployment. In this question no details is given so assume warehouses and laekhouses are there as well. So the right qnswer is UserI don't understand this. I use Services Principal for some tools that connect to warehouse

upvoted 3 times

henryphchan 1 year, 9 months ago

Selected Answer: A

A service principal is an identity created for use with applications, hosted services, and automated tools to access Azure resources. It provides a secure, non-interactive way for Azure DevOps pipelines to authenticate and interact with Fabric APIs without needing user credentials.

upvoted 3 times

hebertorosillo 1 year, 9 months ago

Selected Answer: B

Microsoft Entra supported identities

This API supports the Microsoft identities listed in this section.

Identity Support

User Yes

Service principal and Managed identities Only Power BI Items are supported . ref https://learn.microsoft.com/en-us/rest/api/fabric/core/deployment-pipelines/deploy-stage-content?tabs=HTTP

upvoted 3 times

shmmini 1 year, 8 months ago

Service principal is only supported for Power BI Items deployment. In this question no details is given so assume warehouses and laekhouses are there as well. So the right qnswer is UserI don't understand this. I use Services Principal for some tools that connect to warehouse

upvoted 1 times

shmmini 1 year, 8 months ago

Service principal is only supported for Power BI Items deployment. In this question no details is given so assume warehouses and laekhouses are there as well. So the right qnswer is User

upvoted 1 times

DarioReymago 1 year, 6 months ago

I don't understand this. I use Services Principal for some tools that connect to warehouse

upvoted 1 times

prabhjot 1 year, 9 months ago

Selected Answer: A

to interact with azure services and devops to implement automation, service principal is most commonly used

upvoted 2 times

h2o_molecule 1 year, 9 months ago

Selected Answer: A

Here's a brief overview of the options:

Service principal: Provides a secure and scalable method for authenticating and authorizing access between Azure DevOps and other services. Ideal for automation and integration scenarios.

Microsoft Entra username and password: Less secure and not recommended for automated processes.

Managed private endpoint: Used for secure, private connectivity between services, but not specifically for authentication in this context.

Workspace identity: Not typically used for cross-service API authentication in Azure DevOps integration scenarios.

upvoted 2 times

Meir 1 year, 11 months ago

See https://learn.microsoft.com/en-us/answers/questions/2074603/run-pipeline-via-azure-function-by-rest-apiAgree. Workspace identity is primarily designed for connecting Fabric items to storage accounts and is not the most suitable option for API authentication between Azure DevOps and Fabric.

upvoted 4 times

Meir 1 year, 11 months ago

Selected Answer: A

See https://learn.microsoft.com/en-us/answers/questions/2074603/run-pipeline-via-azure-function-by-rest-api

upvoted 4 times

robertlavigne 1 year, 9 months ago

Selected Answer: B

https://learn.microsoft.com/en-us/rest/api/fabric/core/deployment-pipelines/deploy-stage-content?tabs=HTTP

The above documentation shows that the deploy pipeline API can only use a service principal to deploy PowerBI content. Since we are deploying content related to a medallion architecture we will need to use user based authentication.

Workspace ID is for fabric to reach out to other things which isn't what we want. And a managed private-endpoint is not relevant.

upvoted 3 times

viskas 1 year, 10 months ago

Selected Answer: A

To configure API authentication between Azure DevOps and Fabric, a service principal is the recommended approach.

Workspace identity: Workspace identities are a feature within Fabric itself. They are used to control access to resources within the Fabric workspace. While relevant for managing access within Fabric, they don't directly facilitate authentication with external services like Azure DevOps.

upvoted 3 times

GHill1982 1 year, 11 months ago

Agree. Workspace identity is primarily designed for connecting Fabric items to storage accounts and is not the most suitable option for API authentication between Azure DevOps and Fabric.

upvoted 2 times

mixonfreddy 1 year, 11 months ago

Selected Answer: D

Use of service principal is discouraged

upvoted 2 times

Question #16 *Topic 1*

You have a Google Cloud Storage (GCS) container named storage1 that contains the files shown in the following table.

You have a Fabric workspace named Workspace1 that has the cache for shortcuts enabled. Workspace1 contains a lakehouse named Lakehouse1. Lakehouse1 has the shortcuts shown in the following table.

You need to read data from all the shortcuts.
Which shortcuts will retrieve data from the cache?

Name	Size
ProductFile.parquet	8 MB
StoreFile.json	500 MB
TripsFile.csv	99 MB

Name	Source	Last accessed
Products	ProductFile	12 hours ago
Stores	StoreFile	4 hours ago
Trips	TripsFile	48 hours ago

A. Stores only

B. Products only

C. Stores and Products only

D. Products, Stores, and Trips

E. Trips only

F. Products and Trips only

Suggested Answer: *C*

Community vote distribution

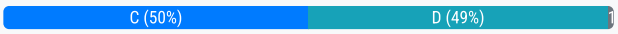

Tuki93 Highly Voted 1 year, 11 months ago

Selected Answer: C

Shortcut caching can be used to reduce egress costs associated with cross-cloud data access. As files are read through an external shortcut, the files are stored in a cache for the Fabric workspace. Subsequent read requests are served from cache rather than the remote storage provider. Cached files have a retention period of 24 hours.

https://learn.microsoft.com/en-us/fabric/onelake/onelake-shortcuts#caching

upvoted 21 times

Tuki93 Highly Voted 1 year, 11 months ago

Shortcut caching can be used to reduce egress costs associated with cross-cloud data access. As files are read through an external shortcut, the files are stored in a cache for the Fabric workspace. Subsequent read requests are served from cache rather than the remote storage provider. Cached files have a retention period of 24 hours. https://learn.microsoft.com/en-us/fabric/onelake/onelake-shortcuts#cachingAnswer should be B, as ADLS shortcuts are supported as well as gcb, s3. But cap is 1gb.My bad this is answer for 18. Here it should be all 3. As all of them are less than 2 gb and 28 days is now max retention.I mean 1gb, stupid typo

upvoted 21 times

Sr18 1 year, 6 months ago

Answer should be B, as ADLS shortcuts are supported as well as gcb, s3. But cap is 1gb.

upvoted 1 times

Sr18 1 year, 6 months ago

Answer should be B, as ADLS shortcuts are supported as well as gcb, s3. But cap is 1gb.My bad this is answer for 18. Here it should be all 3. As all of them are less than 2 gb and 28 days is now max retention.I mean 1gb, stupid typo

upvoted 1 times

Sr18 1 year, 6 months ago

My bad this is answer for 18. Here it should be all 3. As all of them are less than 2 gb and 28 days is now max retention.I mean 1gb, stupid typo

upvoted 5 times

Sr18 1 year, 6 months ago

My bad this is answer for 18. Here it should be all 3. As all of them are less than 2 gb and 28 days is now max retention.

upvoted 5 times

Sr18 1 year, 6 months ago

I mean 1gb, stupid typo

upvoted 1 times

ts_au 1 year, 6 months ago

caching can now be retained up to 28 days

upvoted 19 times

luckybalaji Most Recent 5 months, 4 weeks ago

Based on information from this page https://learn.microsoft.com/en-us/fabric/onelake/onelake-shortcuts#caching. It is upto 28 days and also missing of the details of “Default” setting or “default” setting in the reference page.

upvoted 1 times

MedMed009 7 months, 4 weeks ago

->Shortcut : a logical pointer to external data (like GCS, ADLS, S3). With cache enabled, Fabric keeps a local copy in OneLake for faster queries. ->Cache enabled : When cache for shortcuts is enabled in a Fabric workspace, data that is read from the shortcut is cached in OneLake. The cache is refreshed when the shortcut is accessed. Maximum retention period for cached shortcut data is 28 days. -> OneLake responds to subsequent read requests from the cache rather than the remote storage provider. You can set the retention period for cached files between 1-28 days.

upvoted 2 times

zulajariwala 8 months, 1 week ago

Given file sizes ProductFile.parquet → 8 MB StoreFile.json → 500 MB TripsFile.csv → 99 MB   All are < 1 GB, so size is NOT a limiting factor Apply caching rules (from doc) Caching works when: Source = GCS (supported) Cache enabled = Yes File size < 1 GB File format = Supported formats (Parquet, CSV, JSON all supported)

upvoted 1 times

Josekonoor 8 months, 3 weeks ago

all that the question say is "Cache is enabled". When it is simply enabled without specifically changing the retention , it goes with the default retention of 1 day. So any data that was accessed within recent 24 hours will be there in the cache untill next 24 hours

upvoted 3 times

reactortyz 9 months, 1 week ago

According to current documentation answer should be D: OneLake responds to subsequent read requests from the cache rather than the remote storage provider. You can set the retention period for cached files between 1-28 days. Individual files greater than 1 GB in size aren't cached.

upvoted 1 times

Rahalcer 9 months, 1 week ago

When shortcut caching is enabled, files are cached after access and remain in cache for up to 24 hours since last access. Trips 48 hours so C

upvoted 2 times

jennsun 9 months, 2 weeks ago

Shortcut cache keeps files ≤1 GB for 24 hours since last access. Trip.csv would be expired.

upvoted 2 times

friik 9 months, 3 weeks ago

Caching can now be up to 28 days. The question is hopefully updated by Microsoft during the test, this question was more relevant back when Cached files had a retention period of 24 hours.

upvoted 1 times

a74cb2b 11 months ago

Each file is "The retention period for cached files can be set from 1-28 days. Each time the file is accessed, the retention period is reset." [...] "If a file hasn't been accessed for more than the selected retention period, it's purged from the cache. Individual files greater than 1 GB in size aren't cached."

upvoted 4 times

9a69b12 11 months, 1 week ago

This question is outdated due to increased retention limit, but C was right when max retention was 24H.

upvoted 3 times

Psycho 11 months, 2 weeks ago

When cache for shortcuts is enabled, recently accessed files are cached. The cache duration is 24 hours (default in Fabric). If a shortcut was accessed within the last 24 hours, it will read from the cache. If accessed more than 24 hours ago, it will read from the original source.

upvoted 3 times

digitalqloud 11 months, 2 weeks ago

Fabric only supports shortcut caching when Fabric can convert the source into Delta.

upvoted 1 times

Bin_Hashim 1 year ago

Answer: C — Stores and Products only Why: With Cache for shortcuts enabled, Fabric serves data from cache only if the shortcut was accessed within the last 24 hours. Products → last accessed 12 hours ago → cached Stores → last accessed 4 hours ago → cached Trips → last accessed 48 hours ago → cache expired → reads from source, not cache Therefore, only Stores and Products will be retrieved from the cache.

upvoted 1 times

YeloDeo01_ 1 year, 1 month ago

The retention period for cached files can be set from 1-28 days. https://learn.microsoft.com/en-us/fabric/onelake/onelake-shortcuts

upvoted 5 times

8115575 1 year, 2 months ago

As nothing has been explicitly stated regarding cache retention period, one can assume it is still the default value of 1 day. Hence, cache retention of 24 hours has been expired for Trips and it will be read from source again. From doc page: Shortcut caching can reduce egress costs associated with cross-cloud data access. As files are read through an external shortcut, the files are stored in a cache for the Fabric workspace. Subsequent read requests are served from cache rather than the remote storage provider. The retention period for cached files can be set from 1-28 days. Each time the file is accessed, the retention period is reset. If the file in remote storage provider is more recent than the file in the cache, the request is served from remote storage provider and the updated file will then be stored in cache. If a file hasn't been accessed for more than the selected retention period, it's purged from the cache. Individual files greater than 1 GB in size aren't cached.

upvoted 4 times

Question #17 *Topic 1*

You have a Fabric workspace named Workspace1 that contains an Apache Spark job definition named Job1.
You have an Azure SQL database named Source1 that has public internet access disabled.
You need to ensure that Job1 can access the data in Source1.
What should you create?

A. an on-premises data gateway

B. a managed private endpoint

C. an integration runtime

D. a data management gateway

Suggested Answer: *B*

Community vote distribution

B (100%)

GHill1982 Highly Voted 1 year, 10 months ago

To ensure that Job1 can access the data in Source1, you need to create a managed private endpoint. This will allow the Spark job to securely connect to the Azure SQL database without requiring public internet access.correct https://learn.microsoft.com/en-us/fabric/security/security-managed-private-endpoints-overview#supported-item-types

upvoted 7 times

GHill1982 Highly Voted 1 year, 10 months ago

Selected Answer: B

To ensure that Job1 can access the data in Source1, you need to create a managed private endpoint. This will allow the Spark job to securely connect to the Azure SQL database without requiring public internet access.

upvoted 7 times

lacuna quid Most Recent 1 month, 3 weeks ago

Selected Answer: B

The key clue is public internet access disabled.

upvoted 2 times

Hanif53 5 months, 3 weeks ago

Selected Answer: B

In Fabric, when accessing Azure resources that are behind private networking:

Azure SQL Database

Azure Storage

Azure Key Vault

You create a Managed Private Endpoint (MPE) from Fabric.

This provides secure private connectivity from Fabric compute (Spark) to the Azure resource.

upvoted 3 times

smanzana 1 year, 5 months ago

Selected Answer: B

Correct

upvoted 1 times

2d073ed 1 year, 5 months ago

Selected Answer: B

B correct answer.

Managed private endpoints are connections that workspace admins can create to access data sources that are behind a firewall or that are blocked from public internet access.

upvoted 2 times

DarioReymago 1 year, 5 months ago

Selected Answer: B

B is correct

upvoted 1 times

Tuki93 1 year, 10 months ago

Selected Answer: B

Provided answer is correct

upvoted 3 times

realexamguru 1 year, 10 months ago

correct

https://learn.microsoft.com/en-us/fabric/security/security-managed-private-endpoints-overview#supported-item-types

upvoted 5 times

Question #18 *Topic 1*

You have an Azure Data Lake Storage Gen2 account named storage1 and an Amazon S3 bucket named storage2.
You have the Delta Parquet files shown in the following table.

You have a Fabric workspace named Workspace1 that has the cache for shortcuts enabled. Workspace1 contains a lakehouse named Lakehouse1. Lakehouse1 has the following shortcuts:
A shortcut to ProductFile aliased as Products
A shortcut to StoreFile aliased as Stores
A shortcut to TripsFile aliased as Trips
The data from which shortcuts will be retrieved from the cache?

Name	Stored in	Size	Description
ProductFile	storage1	50 MB	Contains a list of products and their details
TripsFile	storage2	2 GB	Contains one month's worth of taxi trip data
StoreFile	storage2	25 MB	Contains a list of stores and their addresses

A. Trips and Stores only

B. Products and Store only

C. Stores only

D. Products only

E. Products, Stores, and Trips

Suggested Answer: *B*

Community vote distribution

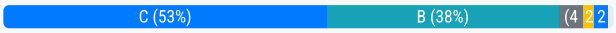

Tuki93 Highly Voted 1 year, 10 months ago
Shortcut caching is currently only supported for GCS, S3 and S3 compatible shortcuts. If a file hasn't been accessed for more than 24 hrs it's purged from the cache. Individual files greater than 1 GB in size aren't cached https://learn.microsoft.com/en-us/fabric/onelake/onelake-shortcuts#cachingThe retention period for cached files can be set from 1-28 days. (same link as above)for this reason is BADLS supports shortcuts, BUT the question is asking "The data from which shortcuts will be retrieved FROM THE CACHE?". ADLS does NOT need a cache. -> Products is incorrect
upvoted 32 times

Tuki93 Highly Voted 1 year, 10 months ago
Selected Answer: C
Shortcut caching is currently only supported for GCS, S3 and S3 compatible shortcuts.

If a file hasn't been accessed for more than 24 hrs it's purged from the cache. Individual files greater than 1 GB in size aren't cached

https://learn.microsoft.com/en-us/fabric/onelake/onelake-shortcuts#caching
upvoted 30 times

AnetaK 1 year, 7 months ago
The retention period for cached files can be set from 1-28 days. (same link as above)
upvoted 6 times

DarioReymago 1 year, 6 months ago
for this reason is B
upvoted 8 times

DarioReymago 1 year, 6 months ago
for this reason is BADLS supports shortcuts, BUT the question is asking "The data from which shortcuts will be retrieved FROM THE CACHE?". ADLS does NOT need a cache. -> Products is incorrect
upvoted 8 times

exnaniantwort 1 year, 5 months ago

ADLS supports shortcuts, BUT the question is asking "The data from which shortcuts will be retrieved FROM THE CACHE?". ADLS does NOT need a cache. -> Products is incorrect

upvoted 8 times

prabhjot 1 year, 9 months ago

Shortcuts are accessible from both ADLS and Amazon S3 as of the time of writing this answer. Trips is too large to be cachedADLS supports shortcuts, BUT the question is asking "The data from which shortcuts will be retrieved FROM THE CACHE?". ADLS does NOT need a cache. -> Products is incorrect02/09/2025.- Shortcut caching is currently supported for GCS, S3, S3 compatible, and on-premises data gateway shortcuts.

upvoted 12 times

prabhjot 1 year, 9 months ago

Selected Answer: B

Shortcuts are accessible from both ADLS and Amazon S3 as of the time of writing this answer. Trips is too large to be cached

upvoted 11 times

AndrewsOtis 1 year, 2 months ago

02/09/2025.- Shortcut caching is currently supported for GCS, S3, S3 compatible, and on-premises data gateway shortcuts.

upvoted 1 times

Hanif53 Most Recent 5 months, 3 weeks ago

ProductFile – 50 MB → ≤ 1 GB → □ Cached StoreFile – 25 MB → ≤ 1 GB → □ Cached TripsFile – 2 GB → > 1 GB → □ Not cached B. Products and Store only

upvoted 1 times

MedMed009 7 months, 4 weeks ago

How caching works in Fabric shortcuts -When caching is enabled, Fabric will cache small external files locally in the workspace for faster access. -Large files (like the 2 GB TripsFile) are not cached — they are read directly from the external source. -Files in Azure Data Lake Gen2 are cached when shortcuts are enabled. -Files in Amazon S3 are cached only if they are small enough (like StoreFile at 25 MB).

upvoted 1 times

MedMed009 8 months, 2 weeks ago

Selected Answer: E

-Caching is supported in ADSL Gen2, here is a documentation that shows how to create shortcut in ADSL that will be used in caching : https://learn.microsoft.com/en-us/fabric/onelake/create-adls-shortcut?utm_source=copilot.com -As for file limitation, there is no official documentation that states file size should be under 1GB. With that being said, I would go with answer E. Feel free to add your comments or usefull links.

upvoted 1 times

zulajariwala 8 months, 1 week ago

Based on the current Microsoft Learn article, shortcut caching currently supports GCS, S3, S3-compatible, and on-premises gateway shortcuts—not ADLS Gen2. The same article also says individual files larger than 1 GB are not cached. So the rule is: S3-backed shortcut + file ≤ 1 GB → cached ADLS Gen2-backed shortcut → not retrieved from shortcut cache Any file > 1 GB → not cached Using that rule on your question, the cached shortcuts are the ones pointing to the Amazon S3 bucket and under 1 GB. In this question, that is: B. Products and Store only

upvoted 2 times

MedMed009 8 months, 2 weeks ago

-Caching is supported in ADSL Gen2, here is a documentation that shows how to create shortcut in ADSL that will be used in caching : https://learn.microsoft.com/en-us/fabric/onelake/create-adls-shortcut?utm_source=copilot.com -As for file limitation, there is no official documentation that states file size should be under 1GB. With that being said, I would go with answer E. Feel free to add your comments or usefull links.https://learn.microsoft.com/en-us/fabric/onelake/onelake-shortcuts#caching "Individual files greater than 1 GB in size aren't cached."

upvoted 1 times

Attollo 8 months, 1 week ago

https://learn.microsoft.com/en-us/fabric/onelake/onelake-shortcuts#caching "Individual files greater than 1 GB in size aren't cached."

upvoted 2 times

Rahalcer 9 months, 1 week ago

From Microsoft Learn : OneLake shortcuts : "Shortcut caching is supported for shortcuts that reference Amazon S3 and Google Cloud Storage (GCS). It isn't supported for Azure Data Lake Storage Gen2." Source: https://learn.microsoft.com/en-us/fabric/onelake/onelake-shortcuts And we have also : “Individual files greater than 1 GB in size aren’t cached.”

upvoted 2 times

jennsun 9 months, 2 weeks ago

Trip file is 2GB which is over 1BG cache data from shortcut. So only Products and Store are available.

upvoted 1 times

TomGray1989 10 months, 4 weeks ago

Selected Answer: A

Azure storage is not necessary to cache and so isn't cached, to reduce the Egress costs. As Fabric evolves, sizes and durations will increase or become more efficient. At this moment in time, Retention can be set from 1-28 days. Trips and Stores are both in Amazon S3 bucket which is cached. 2GB file size supported. The answer is A, Trips and Stores. I would not be surprised if the retention period increases in the near future.

upvoted 1 times

friik 9 months, 3 weeks ago

Stores Only. The Trips dataset is too large to be cached, and the Products are in an Azure Data Lake, which are not cached, always retrieved directly, since it's "within" the organization.

upvoted 1 times

TomGray1989 10 months, 4 weeks ago

Azure storage is not necessary to cache and so isn't cached, to reduce the Egress costs. As Fabric evolves, sizes and durations will increase or become more efficient. At this moment in time, Retention can be set from 1-28 days. Trips and Stores are both in Amazon S3 bucket which is cached. 2GB file size supported. The answer is A, Trips and Stores. I would not be surprised if the retention period increases in the near future.I think reviewing this I believe the answer to be B. Products and Stores only.

upvoted 1 times

TomGray1989 10 months, 4 weeks ago

I think reviewing this I believe the answer to be B. Products and Stores only.

upvoted 1 times

Psycho 11 months, 2 weeks ago

When cache for shortcuts is enabled, Fabric caches small files for performance. Large files (multi-GB) are typically not cached because caching is optimized for smaller datasets. Based on Microsoft Fabric behavior: Files under 1 GB are cached. Files larger than 1 GB are not cached.

upvoted 1 times

digitalqloud 11 months, 2 weeks ago

Shortcut caching is currently supported for GCS, S3, S3 compatible, and on-premises data gateway shortcuts.

upvoted 1 times

Bin_Hashim 1 year ago

Answer: A — Trips and Stores only Why: Fabric's Cache for shortcuts currently caches data only for shortcuts that point to external cloud storage outside Azure (e.g., Amazon S3 and Google Cloud Storage). Shortcuts to Azure Data Lake Storage Gen2 (ADLS Gen2) are not cached. ProductFile → storage1 (ADLS Gen2) → not cached TripsFile → storage2 (Amazon S3) → cached StoreFile → storage2 (Amazon S3) → cached So reads from the cache will come from Trips and Stores only.

upvoted 2 times

2bc006e 1 year, 1 month ago

It is B because "Shortcut caching is currently supported for GCS, S3, S3 compatible, and on-premises data gateway shortcuts." please, read the notes in the documentation in the section Caching https://learn.microsoft.com/en-us/fabric/onelake/onelake-shortcuts#cache-shortcuts

upvoted 3 times

Deano99 1 year, 1 month ago

Only the Products shortcut points to a file in ADLS Gen2, which is the only supported source for shortcut caching. Therefore, only Products will be retrieved from the cache. Stores points to a file in Amazon S3, which does not support shortcut caching. So even though the shortcut exists, data will not be retrieved from the cache.

upvoted 1 times

9a1b660 1 year, 1 month ago

- Files must be under 1 GB to be eligible for caching. - Cached data is refreshed when accessed and expires after 24 hours of inactivity. - Caching applies to shortcuts in lakehouses with cache enabled.

upvoted 1 times

92ef42d 1 year, 2 months ago

Selected Answer: E

Both Amazon S3 and ADLS Gen2 are supported: https://community.fabric.microsoft.com/t5/Data-Engineering/Fabric-Shortcuts-and-Cache/m-p/4766667 There is no exact strict limit for the file size. So, the correct answer includes all the tables, i.e. E

upvoted 1 times

92ef42d 1 year, 2 months ago

It is C. ADLS Gen2 is not supported: https://learn.microsoft.com/en-us/fabric/onelake/onelake-shortcuts#caching Files > 1GB are not supported: https://learn.microsoft.com/en-us/fabric/onelake/onelake-shortcuts#caching

upvoted 2 times

92ef42d 1 year, 2 months ago

It is B. There is a limit for individual file size: https://learn.microsoft.com/en-us/fabric/onelake/onelake-shortcuts#caching

upvoted 1 times

92ef42d 1 year, 2 months ago

It is C. ADLS Gen2 is not supported: https://learn.microsoft.com/en-us/fabric/onelake/onelake-shortcuts#caching

upvoted 1 times

HOTSPOT -

You have a Fabric workspace named Workspace1 that contains the items shown in the following table.

For Model1, the Keep your Direct Lake data up to date option is disabled.

You need to configure the execution of the items to meet the following requirements:

Notebook1 must execute every weekday at 8:00 AM.

Notebook2 must execute when a file is saved to an Azure Blob Storage container.

Model1 must refresh when Notebook1 has executed successfully.

How should you orchestrate each item? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Name	Type
Notebook1	Notebook
Notebook2	Notebook
Lakehouse1	Lakehouse
Pipeline1	Data pipeline
Model1	Semantic model

Answer Area

Notebook1:
- Add Notebook1 to an Apache Spark job definition.
- Add Notebook1 to Pipeline1.
- From Real-Time hub, configure the execution of Notebook1.

Notebook2:
- Add Notebook2 to an Apache Spark job definition.
- Add Notebook2 to Pipeline1.
- From Real-Time hub, configure the execution of Notebook2.

Pipeline1:
- Add Pipeline1 to an Apache Spark job definition.
- Configure the execution of Pipeline1 by using a schedule.
- From Real-Time hub, configure the execution of Pipeline1.

Model1:
- Add Model1 to Pipeline1.
- From Real-Time hub, configure Model1 to refresh.
- Set Keep your Direct Lake data up to date to On.

Suggested Answer:

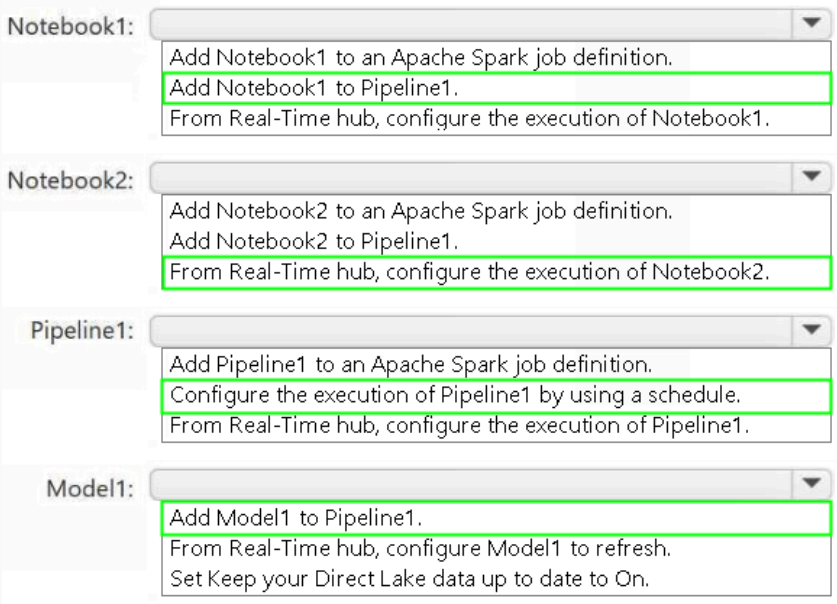

GHill1982 Highly Voted 1 year, 10 months ago

All four answers are correct.Notebook1: From Real-Time hub, configure the execution of Notebook1. Notebook2: Add Notebook2 to Pipeline1. Pipeline1: Configure the execution of Pipeline1 by using a schedule. Model1: Add Model1 to Pipeline1.

upvoted 20 times

GHill1982 Highly Voted 1 year, 10 months ago

All four answers are correct.

upvoted 20 times

hanifabu Most Recent 5 months, 3 weeks ago

Notebook1: Add Notebook1 to Pipeline1

Notebook2: From Real-Time hub, configure execution

Pipeline1: Configure execution using schedule

Model1: Add Model1 to Pipeline1

upvoted 2 times

wellingtonluis 1 year, 2 months ago

Totally agree with the answer.

N1 - Pipeline

N2 - Real-time Hub

Pipeline - Schedule

Model - Add to pipeline

upvoted 6 times

8ac3742 1 year, 3 months ago

Notebook1

Must run every weekday at 8:00 AM → that is a scheduled job.

Best option: Add Notebook1 to an Apache Spark job definition (because Spark jobs can be scheduled).

Notebook2

Must run when a file arrives in Azure Blob Storage.

That's an event-based trigger from Real-Time hub.

Best option: From Real-Time hub, configure the execution of Notebook2.

Pipeline1

No special condition given, but it is listed in the items table.

Since it's a pipeline, its execution must be tied to something. The natural match is scheduled pipelines (not Spark job).

Best option: Configure the execution of Pipeline1 by using a schedule.

Model1

Must refresh when Notebook1 finishes successfully.

That means the refresh should be tied to Notebook1's job completion.

Best option: Add Model1 to Pipeline1 (since pipelines can orchestrate dependent tasks in order).

upvoted 1 times

e71e79b 1 year, 3 months ago

Notebook1 - 2) Add Notebook1 to Pipeline1 (Allows you to schedule execution at 8:00 AM every weekday. Pipelines offer precise scheduling.)

Notebook2 - 3) From Real-Time hub, configure the execution of Notebook2 (Enables event-driven triggers, such as reacting to file uploads in Azure Blob Storage.)

Pipeline1 - 2) Configure the execution of pipeline1 by using a schedule (Pipelines can be scheduled directly without embedding them in Spark jobs.)

Model1 - 1) Add Model1 to Pipeline1 (Since Direct Lake refresh is disabled, pipeline-based execution triggered post-Notebook1 ensures Model1 refreshes on success.)

upvoted 3 times

DarioReymago 1 year, 5 months ago

NB1 add to Pipeline

NB2 realtime hub

Pipeline: Schedule

M1: add to Pipeline

:

In Microsoft Fabric, the Real-Time hub allows you to ingest and manage streaming data from various sources, including Azure Blob Storage events. You can create event streams in Fabric that connect to your Azure Blob Storage account and trigger actions based on events like file creation, deletion, or modification. These event streams can then be used to process and analyze the real-time data for various applications like data analysis, machine learning, or creating alerts

upvoted 2 times

4371883 1 year, 10 months ago

Notebook 1 runs at 8AM, this can be scheduled via Pipeline (Pipeline1), so the configuration of scheduling is in Pipeline 1. Because Model1 refreshes on completion of Notebook1, it should be added to Pipeline 1 after Notebook 1. Notebook 2 is separate to the Pipeline 1 sequence, it requires a trigger (file saved in Blob) from real-time hub.

upvoted 10 times

ChenFu 1 year, 9 months ago

Notebook1: From Real-Time hub, configure the execution of Notebook1.

Notebook2: Add Notebook2 to Pipeline1.

Pipeline1: Configure the execution of Pipeline1 by using a schedule.

Model1: Add Model1 to Pipeline1.

upvoted 2 times

boonfidelio 1 year, 11 months ago

Selected Answer: A

Based on the information provided, I believe A is correct.

upvoted 2 times

Question #20 *Topic 1*

Your company has a sales department that uses two Fabric workspaces named Workspace1 and Workspace2.
The company decides to implement a domain strategy to organize the workspaces.
You need to ensure that a user can perform the following tasks:
Create a new domain for the sales department.
Create two subdomains: one for the east region and one for the west region.
Assign Workspace1 to the east region subdomain.
Assign Workspace2 to the west region subdomain.
The solution must follow the principle of least privilege.
Which role should you assign to the user?

A. workspace Admin

B. domain admin

C. domain contributor

D. Fabric admin

Suggested Answer: *B*

Community vote distribution

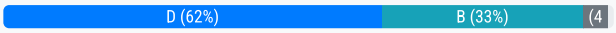

neo_20 Highly Voted 1 year, 10 months ago

Selected Answer: D

Only a Fabric Admin can create Domains.

upvoted 23 times

neo_20 Highly Voted 1 year, 10 months ago

Only a Fabric Admin can create Domains.Domain admin can create domains too. Fabric admin violates the principle of least privileges.

upvoted 25 times

alqazqzy.asseel Most Recent 1 month, 1 week ago

Selected Answer: D

As per Official Fabric docs "Fabric admin (or higher): Fabric admins can create and edit domains, specify domain admins and domain contributors, and associate workspaces with domains. Fabric admins see all the defined domains on the Domains tab in the admin portal, and they can edit and delete domains."

upvoted 1 times

AmirAhmadabadi 3 months ago

Selected Answer: B

Create and manage domains

Create subdomains

Assign workspaces to domains and subdomains

Manage the domain hierarchy

upvoted 1 times

Hanif53 5 months, 3 weeks ago

Selected Answer: B

Role Scope

Fabric Admin Tenant-wide (too broad)

Domain Admin Domain governance

Domain Contributor Limited domain actions

Workspace Admin Workspace only

upvoted 1 times

luckybalaji 5 months, 4 weeks ago

Selected Answer: B

Fabric admin give full access, question is to get least privilege access

upvoted 1 times

zulajariwala 8 months, 1 week ago

Selected Answer: D

To create domain, you must be a Fabric adminhttps://learn.microsoft.com/en-us/fabric/governance/domains#create-a-domain

upvoted 2 times

Mustafa_Shazly 9 months, 2 weeks ago

Selected Answer: D

https://learn.microsoft.com/en-us/fabric/governance/domains#create-a-domain

=To create domain, you must be a Fabric admin.

=To create subdomains for a domain, you must be Fabric admin or domain admin

upvoted 4 times

TorresW 9 months, 3 weeks ago

Selected Answer: B

Why domain contributor

Creating a new domain and subdomains, and assigning workspaces to those subdomains, are domain-level actions, not workspace-level.

Microsoft's domains documentation + DP-600 discussion show that domain contributor is the least-privilege role that can:

Work with domains/subdomains as delegated by a domain admin

Assign workspaces they administer to those domains/subdomains.

Why not the others

Workspace Admin: Can manage Workspace1/2, but cannot create domains or subdomains.

Domain Admin: Can create/edit domains and manage settings, but is more powerful than needed (violates least privilege).

Fabric Admin: Tenant-wide powerful role; far exceeds what's required.

upvoted 2 times

friik 9 months, 3 weeks ago

Selected Answer: B

Domain admin. They can create Domains and solves the requirement.

Fabric Admin has too many privileges and is disqualified from "Principle of Least Privilege"

upvoted 1 times

Rjb_073086 10 months, 1 week ago

Selected Answer: B

answer should've been B. domain admin can create domains too. fabric admin is for entire tenancy

upvoted 1 times

TomGray1989 10 months, 4 weeks ago

Selected Answer: D

A Fabric Admin can manage domains themselves, a Domain Admin manages within a domain.

upvoted 1 times

eegdo 11 months ago

Selected Answer: D

Only Fabric Admins can create domains: https://learn.microsoft.com/en-us/fabric/governance/domains

upvoted 1 times

Psycho 11 months, 2 weeks ago

Selected Answer: B

Workspace Admin → Can manage items within a workspace, but cannot create domains or subdomains.

Domain Contributor → Can add workspaces to an existing domain but cannot create domains or subdomains.

Fabric Admin → Full administrative rights across the entire Fabric tenant (too much privilege for this scenario).

Domain Admin → Can create domains, subdomains, and assign workspaces. This satisfies all requirements while following the principle of least privilege.

upvoted 4 times

Psycho 11 months, 2 weeks ago

Workspace Admin → Can manage items within a workspace, but cannot create domains or subdomains. Domain Contributor → Can add workspaces to an existing domain but cannot create domains or subdomains. Fabric Admin → Full administrative rights across the entire Fabric tenant (too much privilege for this scenario). Domain Admin → Can create domains, subdomains, and assign workspaces. This satisfies all requirements while following the principle of least privilege.Wrong as in Jan 2026 Domain admins can access to the Domains tab in the admin portal, but they can only see and edit the domains they're admins of. Domain admins can update the domain description, define/update domain contributors, and associate workspaces with the domain.

upvoted 4 times

8d82b67 10 months, 1 week ago

Wrong as in Jan 2026

Domain admins can access to the Domains tab in the admin portal, but they can only see and edit the domains they're admins of. Domain admins can update the domain description, define/update domain contributors, and associate workspaces with the domain.

upvoted 1 times

9a69b12 11 months, 2 weeks ago

Selected Answer: B

Domain Admin can create a new domain with sub-domains. “The solution must follow the principle of least privilege”. Would be weird if the Fabric Admin was the right answer.

upvoted 3 times

larinate 11 months, 2 weeks ago

Selected Answer: B

Fabric Admin oversees domain governance at the tenant level, but the question is not about tenant governance, it’s about: 1. creating one new domain for the sales department 2. creating its subdomains 3. assigning two specific workspaces and most important we must follow least privilege. All of this is contained within a single domain, not global. governance. "Domain Admin" can create domains/subdomain/assign worspaces/permissions only for theirs domain -> least privilege

upvoted 4 times

AzureGeek79 12 months ago

Selected Answer: B

B is correct answer since it follows the principle of least privilege, not D.

upvoted 3 times

wellingtonluis 1 year, 2 months ago

Selected Answer: D

Only a Fabric Admin can create Domains.

upvoted 3 times

d0bdf05 1 year, 3 months ago

Selected Answer: B

I´m agree with B, because "Domain Admin" can create and manage domains and is covering the principle of less privilege.

upvoted 4 times

d0bdf05 1 year, 3 months ago

Selected Answer: B

I´m agree with B, because "Domain Admin" can create and manage domains and is covering the principle of less privilege.

upvoted 4 times

fdsdfgxcvbdsfhshfg 1 year, 3 months ago

Domain Admin can work on already created domain, not create it and appoint himself as the admin out of thin air.

Fabric Admin is required

upvoted 5 times

e71e79b 1 year, 3 months ago

Selected Answer: B

A domain admin has the necessary permissions to:

Create a new domain

Create subdomains

Assign workspaces to domains or subdomains

upvoted 3 times

smanzana 1 year, 5 months ago

Selected Answer: D

D is correct

upvoted 1 times

DarioReymago 1 year, 5 months ago

Selected Answer: D

just Fabric Adm can create domains

upvoted 2 times

bettermakeme 1 year, 6 months ago

Selected Answer: D

https://learn.microsoft.com/en-us/rest/api/fabric/admin/domains/create-domain?tabs=HTTP

Domains - Create Domain - The caller must be a Fabric administrator.

upvoted 3 times

Sreeven 1 year, 8 months ago

Selected Answer: D

User can => Create a new domain for the sales department needs Fabric Admin role. for sub domains Domain Admin should do.

upvoted 2 times

benni_ale 1 year, 8 months ago

Selected Answer: D

only fabric admin can create and edit domains

upvoted 4 times

Adriel_1996 1 year, 8 months ago

But first say that have to create a new domain

upvoted 1 times

ChenFu 1 year, 9 months ago

Selected Answer: C

The Domain Contributor role provides the exact permissions needed to create subdomains and assign workspaces, without granting excessive privileges. Therefore, C is the best choices

upvoted 3 times

8d6881f 1 year, 9 months ago

Selected Answer: D

Only a Fabric Admin can create Domains.

https://learn.microsoft.com/en-us/fabric/governance/domains?source=recommendations

upvoted 4 times

MDWPartners 1 year, 10 months ago

Selected Answer: D

Fabric admin (or higher) are the only ones who can create domains

upvoted 6 times

AzureGeek79 11 months, 4 weeks ago

Domain admin can create domains too. Fabric admin violates the principle of least privileges.

upvoted 4 times

harshalt10 1 year, 10 months ago

Selected Answer: D

The question states that user should be able to create Domains. Hence, answer should be D.

https://learn.microsoft.com/en-us/fabric/onelake/onelake-shortcuts

upvoted 5 times

Zekes 1 year, 11 months ago

Selected Answer: B

Domain Admin

Domain contributors are workspace Admins (will only be relevant to one workspace)

https://learn.microsoft.com/en-us/fabric/governance/domains#domain-roles

upvoted 2 times

Stretchy 1 year, 8 months ago

Domain admins can only manage domains they are assigned to. Fabric Admin: Fabric admins can create and edit domains, specify domain admins and contributors, and associate workspaces with domains. They have full visibility and control over all defined domains.

Domain Admin: Domain admins can manage the domains they are assigned to, including creating subdomains and managing domain-specific settings

upvoted 3 times

Bharat 1 year, 11 months ago

Selected Answer: B

Given answer is correct since the domain admin can create a new subdomain but a domain contributor cannot.

upvoted 3 times

IshtarSQL 1 year, 11 months ago

Selected Answer: C

Not B because: Domain Admin:

While this role provides higher-level control over domains and subdomains, it exceeds the principle of least privilege for this scenario.

upvoted 1 times

GHill1982 1 year, 10 months ago

Only a Domain Admin can create domains. Domain Contributor cannot access the domain tab.

https://learn.microsoft.com/en-us/fabric/governance/domains#domain-roles

upvoted 1 times

GHill1982 1 year, 10 months ago

After further research it appears only a Fabric Admin can create Domains.

https://learn.microsoft.com/en-us/fabric/governance/domains#create-a-domain

upvoted 5 times

Question #21 *Topic 1*

You have a Fabric workspace named Workspace1 that contains a warehouse named DW1 and a data pipeline named Pipeline1.
You plan to add a user named User3 to Workspace1.
You need to ensure that User3 can perform the following actions:
View all the items in Workspace1.
Update the tables in DW1.
The solution must follow the principle of least privilege.
You already assigned the appropriate object-level permissions to DW1.
Which workspace role should you assign to User3?

A. Admin

B. Member

C. Viewer

D. Contributor

Suggested Answer: *D*

Community vote distribution



Bovote Highly Voted 1 year, 10 months ago

Viewer - Can view all content in the workspace, but can't modify it. Contributor - Can view and modify all content in the workspace. Member - Can view, modify, and share all content in the workspace. Can add Members Admin - Can view, modify, share, and manage all content in the workspace, including managing permissions. - Can add Admins, Members and can delete workspace. So Contributor is the least role who can view and update the tables (modify the content).To clear everyone's doubt , even if user have warehouse object level access only contributor role can provide modify permission in warehousel think you don't need any access to the workspace. Connection string and SSMS is enough, if r/w access is granted in the DW.

upvoted 25 times

Bovote Highly Voted 1 year, 10 months ago

Selected Answer: D

Viewer - Can view all content in the workspace, but can't modify it. Contributor - Can view and modify all content in the workspace. Member - Can view, modify, and share all content in the workspace. Can add Members Admin - Can view, modify, share, and manage all content in the workspace, including managing permissions. - Can add Admins, Members and can delete workspace.

So Contributor is the least role who can view and update the tables (modify the content).

upvoted 23 times

gpasini Most Recent 1 month, 1 week ago

Selected Answer: C

Copilot said:

The correct answer is Viewer. □

Reasoning:

The requirement is that User3 can view all items in Workspace1.

User3 must also be able to update tables in DW1.

The question explicitly states that you already assigned the appropriate object-level permissions to DW1.

In Fabric, a Viewer workspace role can:

See all items in the workspace.

Connect to and read warehouse content.

Not modify workspace items.

Since the necessary object-level permissions on the warehouse have already been granted, User3 can perform the required table updates through those SQL permissions without needing elevated workspace permissions.

Assigning Contributor, Member, or Admin would also work, but they would grant unnecessary permissions (such as modifying workspace items, pipelines, warehouses, etc.), violating the principle of least privilege.

upvoted 1 times

bettermakeme 1 year, 6 months ago

To clear everyone's doubt , even if user have warehouse object level access only contributor role can provide modify permission in warehouse

upvoted 3 times

bettermakeme 1 year, 6 months ago

To clear everyone's doubt , even if user have warehouse object level access only contributor role can provide modify permission in warehouseI think you don't need any access to the workspace. Connection string and SSMS is enough, if r/w access is granted in the DW.

upvoted 3 times

c41d525 1 month, 2 weeks ago

Selected Answer: D

Viewer Can view workspace items but cannot modify them.

upvoted 1 times

Ce____ 1 month, 4 weeks ago

Selected Answer: D

D Remains the answer, a viewer can't modify the contents of a workspace, just to view, a most acceptable with least privilege is Contributor

upvoted 2 times

01ceacf 1 year, 11 months ago

Selected Answer: B

Only member or above can modify warehouse items

https://learn.microsoft.com/en-us/fabric/get-started/roles-workspaces

upvoted 8 times

01ceacf 1 year, 11 months ago

Only member or above can modify warehouse items https://learn.microsoft.com/en-us/fabric/get-started/roles-workspacesAs I understand, the explain "Create or modify warehouse items." in this document means member role can create and modify warehouse. If you check Microsoft documents for item level permission, they all said item on high level, such as lakehouse, warehouse. Item doesn't mean objects inside warehouse. Contributor role can update table data inside warehouse.

upvoted 8 times

alinavit 2 months ago

Selected Answer: C

It states: "You already assigned the appropriate object-level permissions to DW1."

So it is enough to select viewer role

Verified in Fabric

upvoted 2 times

90fb958 2 months ago

Selected Answer: C

“You already assigned the appropriate object-level permissions to DW1” is key.

Appropriate Object level permission could allow to UPDATE existing tables (create is not required in this scenario). Thus leaving answer: viewer

upvoted 1 times

95704b2 5 months ago

Selected Answer: C

C is the correct answer based on principle of least privilege Viewer can view all the content of the Workspace and since you already assigned the appropriate object-level permissions to DW1

upvoted 3 times

montoya207 5 months, 2 weeks ago

Selected Answer: C

The correct answer is C. Viewer.

Here is the technical justification based on the sources:

Viewing items: The Viewer role allows the user to view all items in the workspace (Workspace1), fulfilling the first requirement.

Updating data: In Fabric, T-SQL permissions (such as UPDATE or INSERT) are managed at the database engine level. If you have already assigned the appropriate object-level permissions, the user only needs “Read” access to the Warehouse element to connect to its SQL connection point. The Viewer role grants this basic access.

upvoted 3 times

luckybalaji 5 months, 4 weeks ago

Selected Answer: C

Viewer only gives the least privilege, giving contributor role makes permission to amend pipeline1 it is note required as per question.

upvoted 4 times

Iconsulting 6 months, 1 week ago

Selected Answer: C

If you assign the Contributor role to User3, they can also modify Pipeline1, thus violating the principle of least privilege, which is a requirement.

upvoted 3 times

Soteras 6 months, 3 weeks ago

Selected Answer: C

 Por que Colaborador NÃO é a resposta correta

O usuário não precisa criar nem modificar itens do workspace

Já possui permissões no nível do warehouse

A role Colaborador concede privilégios adicionais desnecessários

Portanto, viola o princípio do menor privilégio

upvoted 1 times

aws123 7 months, 2 weeks ago

Selected Answer: C

"You already assigned the appropriate object-level permissions to DW1" so viewer is the best answer with least privileges

upvoted 2 times

MedMed009 8 months, 2 weeks ago

Selected Answer: B

Requirement: User3 must view all items in Workspace1 and update tables in DW1.

Object-level permissions for DW1 are already granted.

-Viewer → too restrictive (view only).

-Contributor → can edit, but does not guarantee visibility of all items.

-Admin → excessive privilege.

-Member → correct fit: can view all items and update DW1 tables.

So the answer is B

Link : https://learn.microsoft.com/en-us/fabric/fundamentals/roles-workspaces

upvoted 1 times

e0f0ce6 9 months, 1 week ago

Selected Answer: C

"You already assigned the appropriate object-level permissions to DW1." This ensures the ability to read write in the DW. For the rest Viewer is enough.

upvoted 2 times

Mustafa_Shazly 9 months, 2 weeks ago

Selected Answer: D

Update any item begin from Contributor role

anything else viewer or sharing will read only

object-level permissions restrict access object to some users Even if viewer or Contributor

upvoted 1 times

emboutchamani 9 months, 2 weeks ago

Selected Answer: C

In Microsoft Fabric, the Viewer role provides "Read" permissions on the Warehouse item, which allows the user to connect to the Warehouse via the SQL endpoint.

The actual ability to modify data (INSERT/UPDATE/DELETE) or schema (DDL) is controlled by SQL object-level permissions inside the Warehouse.

Since the prompt states you have "already assigned the appropriate object-level permissions to DW1" (e.g., GRANT UPDATE, db_datawriter), the User only needs the connectivity provided by the Viewer role to utilize those SQL permissions.

upvoted 2 times

emboutchamani 9 months, 2 weeks ago

Selected Answer: D

the minimum right to modify content is contributor

upvoted 1 times

eoicp 9 months, 3 weeks ago

Selected Answer: C

The Object level permissions have already been set in the DW1, meaning user3 can already access and modify tables without even being added to the workspace. The text says he needs view permissions on the workspace. Due to the principle of least privilige the answwer should be C. Viewer.

upvoted 1 times

Anchalgupta 10 months ago

Selected Answer: C

It satisfies both requirements while following the principle of least privilege

upvoted 1 times

a74cb2b 11 months ago

Selected Answer: C

"You already assigned the appropriate object-level permissions to DW1." => this overwrites the general workspace level restrictions on Viewer, allowing User3 to update the the tables in DW1.

Admin, Member, and Contributor all grant Control access to warehouses, while Viewer grants Connect and ReadData permissions. When combined with object-level SQL permissions, Viewer satisfies the requirements while following least privilege.

upvoted 3 times

9a69b12 11 months, 1 week ago

Selected Answer: C

"You already assigned the appropriate object-level permissions to DW1." meaning that the user, if granted VIEW, has all necessary permissions.

upvoted 2 times

larinate 11 months, 2 weeks ago

Selected Answer: B

Only Member or Admin can modify shared workspace items such as warehouses.

Contributor is limited to their own items. Admin is too much, so Member -> B

upvoted 1 times

rc012345 12 months ago

Selected Answer: C

Viewer access at workspace level.

In the warehouse (OLS) - sql query: Granted UPDATE, INSERT access on table level to "user-email"

Result: User able to update (insert) values to respective table.

upvoted 2 times

Bin_Hashim 1 year ago

Selected Answer: C

According to these words from Question

"You already assigned the appropriate object-level permissions to DW1."

(this seems tricky part of the question)

C. Viewer

upvoted 2 times

ShazzyHawk 1 year, 1 month ago

Selected Answer: C

“You already assigned the appropriate object-level permissions to DW1”

changes everything.

Object-level permissions on DW1 override workspace-level edit restrictions.

A Viewer role on the workspace allows the user to see all workspace items,

while the object-level permissions on DW1 allow them to update data in the warehouse.

Therefore, granting Contributor or higher would violate least privilege, since those roles allow editing other workspace items.

upvoted 4 times

wellingtonluis 1 year, 2 months ago

Selected Answer: C

The Key phrase here is: "You already assigned the appropriate object-level permissions to DW1". So Viewer in Workspace is enought.

upvoted 2 times

fdsdfgxcvbdsfhshfg 1 year, 2 months ago

Selected Answer: D

You would think that SQL Granular Permissions would work even for a viewer workspace role, like in the case of Synpase I believe, but here Fabric won't allow an UPDATE statement for a viewer.

upvoted 2 times

e71e79b 1 year, 3 months ago

Selected Answer: D

A Contributor can view all items in a Fabric workspace and interact with specific objects—like updating tables—as long as they already have object-level permissions.

Since you've already given User3 the necessary object-level permissions on DW1, this role grants just enough access without overstepping boundaries. Perfect alignment with the principle of least privilege.

upvoted 2 times

alex_sql 1 year, 4 months ago

Selected Answer: D

Contributor es el mínimo rol que permite ver y modificar (si ya tiene los permisos necesarios a nivel de objeto)

upvoted 1 times

Acrious 1 year, 4 months ago

Selected Answer: D

It is not clear, that they actually were given the Object level permissions at the Workspace level, again another unclear detail in the question imho. So I am going to go with contributor.

upvoted 1 times

smanzana 1 year, 5 months ago

Selected Answer: D

D is correct

upvoted 1 times

DarioReymago 1 year, 5 months ago

Selected Answer: C

View all the items in Workspace1.

Update the tables in DW1.

You already assigned the appropriate object-level permissions to DW1.

Contributor Can view and modify all content in the workspace. So it is higer option]

Viewer + object-level is enough

upvoted 1 times

alex_sql 1 year, 4 months ago

Viewer + permisos a nivel objeto: solo puede consultar, no puede modificar.

upvoted 2 times

Jay63002 1 year, 5 months ago

Selected Answer: D

as per this doc, contributor can create and modify warehouse item : https://learn.microsoft.com/en-us/fabric/fundamentals/roles-workspaces

upvoted 2 times

PBridge 1 year, 6 months ago

Selected Answer: D

Explanation:

You need to ensure that User3 can:

View all items in the workspace

Update tables in DW1

Follow the principle of least privilege

Roles breakdown:

Admin: Too much access (can delete, manage roles, settings, etc.)

Member: Also more permissions than needed (can publish and manage content broadly)

Viewer: Read-only — cannot update DW1

 Contributor:

Can view all items

Can edit datasets, dataflows, pipelines, and warehouses (if object-level permissions are granted)

Fits the least privilege principle

upvoted 3 times

Sr18 1 year, 6 months ago

Selected Answer: D

Contributor role

upvoted 1 times

5dcd3a8 1 year, 6 months ago

Selected Answer: D

The Viewer role allows a user to:

See all items in the workspace

Cannot open notebooks or pipelines

Cannot run queries or update data, even if object-level permissions are granted

Viewer role is strictly read-only at the workspace level, regardless of object-level permissions.

Why Contributor is correct:

Can view all items in the workspace

Can interact with items (like run pipelines, notebooks)

With object-level permissions, can update data in DW1

Does not have management privileges (unlike Member or Admin)

Contributor + object-level permissions = just enough rights to view and update without overreaching — perfect for least privilege.

upvoted 2 times

ts_au 1 year, 6 months ago

Selected Answer: C

they already have access to DW1

upvoted 2 times

d98a23f 1 year, 8 months ago

Selected Answer: D

contibutor have read/write access to a warehouse:https://learn.microsoft.com/en-us/fabric/data-warehouse/share-warehouse-manage-permissions

upvoted 1 times

Adriel_1996 1 year, 8 months ago

Selected Answer: C

If DW (which need modify permission) already has assigned the right permission only with viewer the User3 could view all items

upvoted 3 times

vish9 1 year, 8 months ago

Selected Answer: C

Assign the Viewer role to User3 to allow viewing all workspace items.

Object-level permissions already cover updates to DW1, ensuring least privilege is maintained.

upvoted 5 times

GadAugust 1 year, 8 months ago

Selected Answer: C

This is a really tricky question and the answer does not just lay in the coarse grained workspace roles

https://learn.microsoft.com/en-us/fabric/get-started/roles-workspaces

The user should be able to view the items in the workspace - VIEWER

The user can already update the tables in the DW1 database through object-level permissions

So if you apply the Principle of Least privilege, VIEWER is all that is required. All other roles will elevate the user privilege over and above the requirements.

upvoted 5 times

zxc01 1 year, 8 months ago

Yes, tricky question again. We just can guess grant user3 has two steps if we choose role viewer on workspace level. the next step should grant user permission on warehouse roles. However, question doesn't give enough details for just can add user on workspace level or both workspace and item level. It will be changed to clear if Microsoft change this question to "You already assigned the appropriate object-level permissions to User3."

upvoted 3 times

Goye 1 year, 8 months ago

Selected Answer: B

member can create or modify warehouse items according to doc, contributors can't

https://learn.microsoft.com/en-us/fabric/fundamentals/roles-workspaces#-workspace-roles

upvoted 1 times

benni_ale 1 year, 8 months ago

Selected Answer: C

you could assign viewer role in the workspace which won't overwrite the the appropriate-already-configured object-level permission on dw1 so the user will be able to view all objects in ws1 but will also be able to run update statements within the dw1

upvoted 1 times

benni_ale 1 year, 8 months ago

please correct me if i am wrong otherwise is D

upvoted 1 times

prabhjot 1 year, 9 months ago

Selected Answer: D

Contributor fits the best as being a member would grant the user the power to add other users to the workspace, thus defeating the whole purpose of the question - "Least Privileged Access"

upvoted 1 times

18e18d0 1 year, 9 months ago

Selected Answer: D

It is not specified that the user should be able to add other users to Workspace. Contributors can modify data, but not add others to Workspace. According to the documentation - Contributor is the least privileged role: https://learn.microsoft.com/en-us/fabric/security/permission-model#workspace-roles

upvoted 1 times

eoicp 9 months, 3 weeks ago

I think you don't need any access to the workspace. Connection string and SSMS is enough, if r/w access is granted in the DW.

upvoted 1 times

Bovote 1 year, 10 months ago

Viewer Can view all content in the workspace, but can't modify it.

Contributor Can view and "modify" all content in the workspace.

Member Can view, modify, and share all content in the workspace. Can add Members

Admin Can view, modify, share, and manage all content in the workspace, including managing permissions.

- Can add Admins, Members and can delete workspace.

So Contributor is the least role who can view and update the tables (modify the content).

upvoted 1 times

4371883 1 year, 10 months ago

Selected Answer: D

Contributor is the least privilege role according to this. https://learn.microsoft.com/en-us/fabric/security/permission-model

upvoted 1 times

gfors 1 year, 11 months ago

Selected Answer: D

This is wrong in the docs. Read this https://community.fabric.microsoft.com/t5/Service/Create-or-modify-warehouse-items-can-do-with-Contributor-role/m-p/4318708

upvoted 2 times

doctordodge 1 year, 10 months ago

This is only a custom-setting that can be given to a Contributor role by an Admin. It is not a default. Default is that contributors cannot modify warehouse items. Only Members and Admins can do so by default.

upvoted 1 times

GHill1982 1 year, 10 months ago

Thanks for the link, very useful

upvoted 1 times

zxc01 1 year, 7 months ago

As I understand, the explain "Create or modify warehouse items." in this document means member role can create and modify warehouse. If you check Microsoft documents for item level permission, they all said item on high level, such as lakehouse, warehouse. Item doesn't mean objects inside warehouse. Contributor role can update table data inside warehouse.

upvoted 2 times

Question #22 *Topic 1*

You have a Fabric capacity that contains a workspace named Workspace1. Workspace1 contains a lakehouse named Lakehouse1, a data pipeline, a notebook, and several Microsoft Power BI reports.
A user named User1 wants to use SQL to analyze the data in Lakehouse1.
You need to configure access for User1. The solution must meet the following requirements:
Provide User1 with read access to the table data in Lakehouse1.
Prevent User1 from using Apache Spark to query the underlying files in Lakehouse1.
Prevent User1 from accessing other items in Workspace1.
What should you do?

A. Share Lakehouse1 with User1 directly and select Read all SQL endpoint data.

B. Assign User1 the Viewer role for Workspace1. Share Lakehouse1 with User1 and select Read all SQL endpoint data.

C. Share Lakehouse1 with User1 directly and select Build reports on the default semantic model.

D. Assign User1 the Member role for Workspace1. Share Lakehouse1 with User1 and select Read all SQL endpoint data.

Suggested Answer: *B*

Community vote distribution

A (97%) (3)

e71e79b Most Recent 1 year, 3 months ago

Selected Answer: A

Read table data using SQL:

The SQL endpoint provides a structured interface for querying lakehouse tables without invoking Apache Spark.

Selecting “Read all SQL endpoint data” gives User1 SQL read access to the tables in Lakehouse1.

Prevent Apache Spark access:

By not assigning workspace-level roles (like Viewer, Member, or Admin), User1 won't see or access notebooks or Spark-powered tools.

Limit visibility to only Lakehouse1:

Sharing Lakehouse1 directly (instead of the entire workspace) ensures User1 has visibility only into what you explicitly share.

upvoted 3 times

asitp9 1 year, 4 months ago

Selected Answer: B

in that case workspace also needs some access, workspace should be having viewer access

upvoted 1 times

smanzana 1 year, 5 months ago

Selected Answer: A

A is correct

upvoted 2 times

DarioReymago 1 year, 5 months ago

Selected Answer: A

just read to LH

upvoted 1 times

vernillen 1 year, 6 months ago

Selected Answer: A

The most logical one

upvoted 1 times

benni_ale 1 year, 9 months ago

Selected Answer: A

Agree with other solutions

upvoted 3 times

01ceacf Highly Voted 1 year, 11 months ago

Selected Answer: A

Read Access to Table Data.

B/D - Viewer/Member can access other items

C - cannot access SQL DB

upvoted 12 times

IshtarSQL Highly Voted 1 year, 11 months ago

Selected Answer: A

Assigning the Viewer role would grant User1 access to all items in the workspace, violating the requirement to restrict access to Lakehouse1 only.

upvoted 7 times

Question #23 *Topic 1*

DRAG DROP -

You are implementing the following data entities in a Fabric environment:

Entity1: Available in a lakehouse and contains data that will be used as a core organization entity

Entity2: Available in a semantic model and contains data that meets organizational standards

Entity3: Available in a Microsoft Power BI report and contains data that is ready for sharing and reuse

Entity4: Available in a Power BI dashboard and contains approved data for executive-level decision making

Your company requires that specific governance processes be implemented for the data.

You need to apply endorsement badges to the entities based on each entity's use case.

Which badge should you apply to each entity? To answer, drag the appropriate badges the correct entities. Each badge 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.

NOTE: Each correct selection is worth one point.

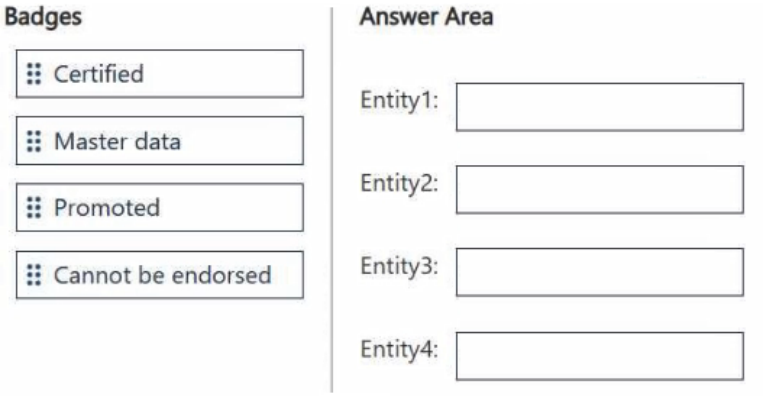

Suggested Answer:

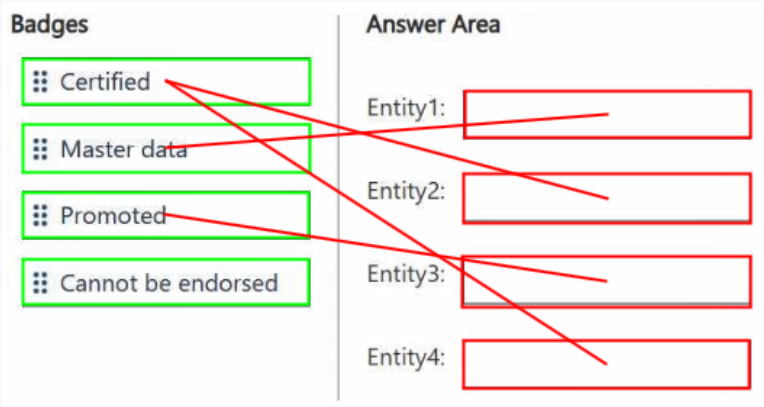

fassil Highly Voted 1 year, 9 months ago

Entity1: core organization entity --> Master Data Entity2: meets organizational standards --> Certified Entity3: ready for sharing and reuse --> Promoted Entity4: Power BI dashboard, executive-level decision making --> Cannot be endorsed, The reason is Types of items that can be endorsed All Fabric items and Power BI items except Power BI dashboards can be promoted or certified.

upvoted 55 times

tejas29 Most Recent 1 month, 4 weeks ago

Entity1 → Master data □

Entity2 → Certified □

Entity3 → Promoted □

Entity4 → Cannot be endorsed □ (Power BI dashboards cannot be promoted or certified)

upvoted 5 times

zulajariwala 8 months, 1 week ago

Entity Badge

Entity1 Master data

Entity2 Certified

Entity3 Promoted

Entity4 Cannot be endorsed

upvoted 3 times

Rjb_073086 10 months, 1 week ago

entity1--> promoted: cant be master data as it isnt governed or enterprise approved

entity2-->certified

entity3-->promoted

entity4-->certified --> exec level dashboards should be labeled as certified

upvoted 1 times

larinate 11 months, 2 weeks ago

Entity 3 “sharing and reuse” and its not equals to “enterprise-wide authoritative data.” so its rather Promoted not Certified.

upvoted 1 times

fdsdfgxcvbdsfhshfg 1 year, 2 months ago

Wtf is a 'data entity' anyway? Dashboard doesn't hold any data and cannot be endorsed

upvoted 2 times

a74cb2b 11 months ago

Wrong -> Dashboards can't be endorsed at all; "The master data label can only be applied to items that contain data, such as lakehouses, and semantic models." https://learn.microsoft.com/en-us/fabric/governance/endorsement-overview

upvoted 2 times

Delafupra 1 year, 4 months ago

Entity1: Master Data Entity2: Certified Entity3: Promoted Entity4: Master Data(Dashboards cannot be promoted or certified, but they can be endorsed as master data), so it is the only option for Entity4. You can find this information in the microsoft documentationWrong -> Dashboards can't be endorsed at all; "The master data label can only be applied to items that contain data, such as lakehouses, and semantic models." https://learn.microsoft.com/en-us/fabric/governance/endorsement-overview

upvoted 1 times

Delafupra 1 year, 4 months ago

Entity1: Master Data Entity2: Certified Entity3: Promoted Entity4: Master Data(Dashboards cannot be promoted or certified, but they can be endorsed as master data), so it is the only option for Entity4. You can find this information in the microsoft documentation

upvoted 1 times

Clode 1 year, 6 months ago

Each badge may be used once, more than once, or not at all

upvoted 2 times

HAARS 1 year, 8 months ago

4 IS CANNOT endorse

upvoted 4 times

Kiket2ride 1 year, 8 months ago

THese are drag and drop so it makes no sense that two of those are of the same type

upvoted 1 times

Kiket2ride 1 year, 8 months ago

THese are drag and drop so it makes no sense that two of those are of the same typeEach badge may be used once, more than once, or not at all

upvoted 1 times

fassil Highly Voted 1 year, 9 months ago

Entity1: core organization entity --> Master Data

Entity2: meets organizational standards --> Certified

Entity3: ready for sharing and reuse --> Promoted

Entity4: Power BI dashboard, executive-level decision making --> Cannot be endorsed, The reason is Types of items that can be endorsed

All Fabric items and Power BI items except Power BI dashboards can be promoted or certified.

upvoted 64 times

DarioReymago 1 year, 5 months ago

this is correct

upvoted 7 times

MultiCloudIronMan 1 year, 9 months ago

Entity1: core organization entity --> Master Data

Entity2: meets organizational standards --> Certified

Entity3: ready for sharing and reuse --> Promoted

Entity4: Power BI dashboard, executive-level decision making --> certified

upvoted 2 times

henryphchan 1 year, 9 months ago

Entity1: core organization entity --> Master Data

Entity2: meets organizational standards --> Certified

Entity3: ready for sharing and reuse --> Promoted

Entity4: Power BI dashboard, executive-level decision making --> Cannot be endorsed, The reason is Master data can only be applied to lakehouse and models.

upvoted 8 times

8d6881f 1 year, 9 months ago

Tricky one: on first sight entity 4 cannot be endorsed: https://learn.microsoft.com/en-us/fabric/governance/endorsement-overview : "Any Fabric or Power BI item except Power BI dashboards can be promoted"

However: is entity 4 in fact referring to the semantic model underlying the report? Since that model is available through the report? In that case it should be 'Certified'.

Very unclear/tricky question

upvoted 1 times

8d6881f 1 year, 9 months ago

Tricky one: on first sight entity 4 cannot be endorsed: https://learn.microsoft.com/en-us/fabric/governance/endorsement-overview : "Any Fabric or Power BI item except Power BI dashboards can be promoted"

However: is entity 4 in fact referring to the semantic model underlying the report? Since that model is available through the report? In that case it should be 'Certified'. Very unclear/tricky question

upvoted 1 times

AdventureChick 11 months, 1 week ago

No, they are not referring to the semantic model. I think it is very clear. Power BI dashboards cannot have endorsements. They are clearly testing whether or not we know that.

upvoted 2 times

Tuki93 1 year, 10 months ago

Entity 4 should be 'Cannot be endorsed.' In my opinion,

"The master data label can only be applied to items that contain data, such as lakehouses and semantic models."

https://learn.microsoft.com/en-us/fabric/governance/endorsement-overview

upvoted 5 times

doctordodge 1 year, 10 months ago

All are correct EXCEPT Entity4: Power BI dashboards cannot be endorsed, according to Microsoft official documentation.

https://learn.microsoft.com/en-us/fabric/governance/endorsement-overview#types-of-items-that-can-be-endorsed

upvoted 7 times

4371883 1 year, 10 months ago

Answer is correct according to this. https://learn.microsoft.com/en-us/fabric/governance/endorsement-overview

Since there is no other suitable badge for Entity 3, it should be labelled as Certified.

upvoted 1 times

Question #24 *Topic 1*

You have three users named User1, User2, and User3.
You have the Fabric workspaces shown in the following table.

You have a security group named Group1 that contains User1 and User3.
The Fabric admin creates the domains shown in the following table.

User1 creates a new workspace named Workspace3.
You add Group1 to the default domain of Domain1.
For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.

Name	Workspace admin
Workspace1	User1
Workspace2	User2

Name	Domain admin
Domain1	User1
Domain2	User2

Answer Area

Statements	Yes	No
User3 has Viewer role access to Workspace3.	○	○
User3 has Domain contributor access to Domain1.	○	○
User2 has Contributor role access to Workspace3.	○	○

Suggested Answer:
Answer: No, Yes, No. Workspace3 is created by User1 (a member of Group1) and is unassigned; when Group1 is added to the default domain of Domain1, unassigned workspaces whose admin is a group member are assigned to Domain1, and Microsoft documents that the specified users/group members generally automatically become domain contributors of workspaces assigned this way, so User3 becomes a Domain1 contributor. Domain assignment never grants workspace access: item visibility and access depend on workspace roles and item permissions, so User3 gets no Viewer role on Workspace3. Nobody added User2 to Workspace3 (User1 created it and is its only admin), so User2 has no Contributor role there. Reference: https://learn.microsoft.com/en-us/fabric/governance/domains

gpasini Most Recent 1 month, 1 week ago

Answer: No / Yes / No

This question is about Fabric domains and default domain assignment.

Let's analyze step by step.

Given

Workspace1 admin = User1

Workspace2 admin = User2

Group1 contains User1 and User3

Domains:

Domain1 admin = User1

Domain2 admin = User2

Then:

User1 creates Workspace3

Group1 is added to the default domain of Domain1

Important Fabric rule

When a user or security group is added to a domain's default domain, any new workspaces they create are automatically associated with that domain.

Since:

Group1 contains User1

Group1 is assigned to Domain1's default domain

User1 creates Workspace3

⇒ Workspace3 automatically belongs to Domain1.

Statement 1

"User3 has Viewer role access to Workspace3."

Being a member of a group assigned to a domain does not grant access to workspaces in that domain.

Workspace permissions and domain membership are separate concepts.

 Answer: No

Statement 2

"User3 has Domain contributor access to Domain1."

Members of a security group assigned to a domain's default domain become Domain contributors for that domain.

User3 belongs to Group1.

 Answer: Yes

Statement 3

"User2 has Contributor role access to Workspace3."

User2 is the admin of Domain2, but Workspace3 belongs to Domain1.

Domain administrators do not automatically become workspace contributors for workspaces in another domain.

No permission was granted to User2 on Workspace3.

 Answer: No

Final Answers

Statement AnswerUser3 has Viewer role access to Workspace3 No

User3 has Domain contributor access to Domain1 Yes

User2 has Contributor role access to Workspace3 No

Answer: No / Yes / No
upvoted 6 times

38578c4 Highly Voted 1 year, 9 months ago

N Y N https://learn.microsoft.com/en-us/fabric/governance/domains Being a domain contributor doesn't give you any permissions on the workspace itself, only on the domain. Workspace roles are set differently. This excludes the first and last option. User 3 will be domain contributor on the new workspace. Even though without domain admin permissions, there's nothing user 3 can do with Workspace 3.
upvoted 17 times

38578c4 1 year, 9 months ago

N Y N https://learn.microsoft.com/en-us/fabric/governance/domains Being a domain contributor doesn't give you any permissions on the workspace itself, only on the domain. Workspace roles are set differently. This excludes the first and last option. User 3 will be domain contributor on the new workspace. Even though without domain admin permissions, there's nothing user 3 can do with Workspace 3.Why "User 3 will be domain contributor on the new workspace"? In https://learn.microsoft.com/en-us/fabric/governance/domains#default-domain, it is written that "The specified users and/or members of the specified security groups GENERALLY automatically become domain contributors of workspaces that are assigned in this manner." but we have a word GENERALLY here...
upvoted 17 times

Waloo2_11 4 months, 3 weeks ago

No

YES

NO

Select Default domain and specify users and/or security groups. When you add people to the default domain list, unassigned workspaces they're admins of, and new workspaces they create, will automatically be assigned to the domain.

-The users and/or members of the security groups specified in the default domain definition generally automatically become domain contributors of the workspaces that get assigned to the domain via the default domain mechanism.

https://learn.microsoft.com/en-us/fabric/governance/domains#define-the-domain-as-a-default-domain

upvoted 5 times

zulajariwala 8 months, 1 week ago

Statement Answer

User3 has Viewer access □ No

User3 has Domain contributor access □ Yes

User2 has Contributor access □ No

upvoted 4 times

larinate 11 months, 2 weeks ago

Domain permissions = metadata + catalog governance

Workspace permissions = item-level access

These are completely separate.

answer is NO / YES / NO

upvoted 7 times

Rozamunde 1 year, 3 months ago

NYN

User1 creates a new workspace named Workspace3 - at this moment the Workspace3 is unassigned, i.e. no domain.

When Group1 (User1 and User3) is added to the default domain of Domain1:

1. "When you add people to the default domain list, unassigned workspaces they're admins of, and new workspaces they create, will automatically be assigned to the domain".

I.e. Workspace3 is assigned to Domain1

2. "The users and/or members of the security groups specified in the default domain definition generally automatically become domain contributors of the workspaces that get assigned to the domain via the default domain mechanism."

I.e User3 becomes Domain1 contributor, not Viewer.

From <https://learn.microsoft.com/en-us/fabric/governance/domains#create-subdomains>

upvoted 6 times

ziggy1117 1 year, 4 months ago

Y-N-N

Y - whenever a specified user or member of a specified security group creates a new workspace, it is assigned to the default domain.

https://learn.microsoft.com/en-us/fabric/governance/domains

N - Domain admins must be explicitly set by the domain admin or Fabric admin

N - Contributor role access to a workspace must be set by the Admin or Member of the workspace

upvoted 3 times

smanzana 1 year, 5 months ago

NYN

Users or groups assigned to a domain's default group automatically receive the Domain Contributor role.

upvoted 4 times

Ricauduro 1 year, 5 months ago

https://learn.microsoft.com/en-us/fabric/governance/domains#define-the-domain-as-a-default-domain

"""

To define a domain as a default domain, you must be a Fabric admin or a domain admin of the domain.

Select Default domain and specify users and/or security groups. When you add people to the default domain list, unassigned workspaces they're admins of, and new workspaces they create, will automatically be assigned to the domain.

"""

User1 creates a new workspace named Workspace3. So WS3 is assigned to the domain 1

"""

Note

The users and/or members of the security groups specified in the default domain definition generally automatically become domain contributors of the workspaces that get assigned to the domain via the default domain mechanism.

"""

User3 is on security Group1, which is assigned to Domain1. So User3 has Domain contributor access to Domain1.

Statement 2 -> Y

upvoted 3 times

bettermakeme 1 year, 6 months ago

N Y N

https://learn.microsoft.com/en-us/fabric/governance/domains

Default domain

A default domain is a domain that has been defined as the default domain for specified users and/or security groups. When you define a domain as the default domain for specified users and/or security groups, the following happens:

The system scans the organization's workspaces. When it finds a workspace whose admin is a specified user or member of a specified security group:

If the workspace already has a domain assignment, it is preserved. The default domain doesn't override the current assignment.

If the workspace is unassigned, it is assigned to the default domain.

After this, whenever a specified user or member of a specified security group creates a new workspace, it is assigned to the default domain.

The specified users and/or members of the specified security groups generally automatically become domain contributors of workspaces that are assigned in this manner.

upvoted 5 times

5dcd3a8 1 year, 6 months ago

When a workspace is created by a user who is a member of a domain, that workspace is automatically associated with that domain.

You can configure a default domain group (e.g., Group1) for a domain, and:

That group gains the Domain Contributor role for all workspaces in the domain.

This does not grant workspace-level Viewer or Contributor roles.

Let's apply that to your scenario:

User3 and Group1:

Group1 contains User3

Group1 is the default domain group for Domain1

So, User3 has the Domain Contributor role in Domain1

Statement 2: "User3 has Domain contributor access to Domain1" YES

Statement 1: "User3 has Viewer role access to Workspace3" NO

Because domain roles don't give workspace Viewer permissions.

User2 and Workspace3:

User2 is not a member of Group1

User2 did not create Workspace3

No indication they were added to it manually

Statement 3: "User2 has Contributor access to Workspace3" NO

upvoted 2 times

5dcd3a8 1 year, 6 months ago

N/Y/N

https://learn.microsoft.com/en-us/fabric/governance/domains#default-domain

upvoted 2 times

vish9 1 year, 8 months ago

N,N,N User3 doesn't have Viewer or Domain contributor access because Group1 was only added to the default domain, not explicitly granted permissions. User2 has no access to Workspace3 since User1 created it and didn't add User2. Therefore, all three statements are false.

upvoted 11 times

vish9 1 year, 8 months ago

N,N,N User3 doesn't have Viewer or Domain contributor access because Group1 was only added to the default domain, not explicitly granted permissions. User2 has no access to Workspace3 since User1 created it and didn't add User2. Therefore, all three statements are false.Default domain When you add people to this list their new and unassigned workspaces will be automatically assigned to the domain. It also makes them domain contributors. This is from MS Fabric itself. So, NYN.finally a voice of sanity in this pile of garbagemy bad, it's NYN actually " The specified users and/or members of the specified security groups generally automatically become domain contributors of workspaces that are assigned in this manner. " https://learn.microsoft.com/en-us/fabric/governance/domains#default-domain

upvoted 11 times

92ef42d 1 year, 2 months ago
Default domain

When you add people to this list their new and unassigned workspaces will be automatically assigned to the domain. It also makes them domain contributors.

This is from MS Fabric itself. So, NYN.
upvoted 1 times

fdsdfgxcvbdsfhshfg 1 year, 2 months ago
finally a voice of sanity in this pile of garbage
upvoted 1 times

fdsdfgxcvbdsfhshfg 1 year, 2 months ago
finally a voice of sanity in this pile of garbagemy bad, it's NYN actually " The specified users and/or members of the specified security groups generally automatically become domain contributors of workspaces that are assigned in this manner. " https://learn.microsoft.com/en-us/fabric/governance/domains#default-domain
upvoted 1 times

fdsdfgxcvbdsfhshfg 1 year, 2 months ago
my bad, it's NYN actually

"

The specified users and/or members of the specified security groups generally automatically become domain contributors of workspaces that are assigned in this manner.

"

https://learn.microsoft.com/en-us/fabric/governance/domains#default-domain
upvoted 2 times

MultiCloudIronMan 1 year, 8 months ago
NYN is correct
upvoted 2 times

8d6881f 1 year, 9 months ago
Should be Y Y N https://learn.microsoft.com/en-us/fabric/governance/domains "After this, whenever a specified user or member of a specified security group creates a new workspace, it is assigned to the default domain. The specified users and/or members of the specified security groups generally automatically become domain contributors of workspaces that are assigned in this manner."
upvoted 2 times

inter00 1 year ago
Why "User 3 will be domain contributor on the new workspace"? In https://learn.microsoft.com/en-us/fabric/governance/domains#default-domain, it is written that "The specified users and/or members of the specified security groups GENERALLY automatically become domain contributors of workspaces that are assigned in this manner." but we have a word GENERALLY here...
upvoted 1 times

henryphchan 1 year, 9 months ago
Y/Y/N

I think User 3 has Domain contributor access based on the document: https://learn.microsoft.com/en-us/fabric/governance/domains

When you define a domain (Domain1) as the default domain for specified users and/or security groups (Group1), the system scans the organization's workspaces. When it finds a workspace whose admin (User1) is a specified user or member of a specified security group (Group1). If the workspace is unassigned (Workspace3), it is assigned to the default domain (Domain1). The specified users and/or members of the specified security groups (User 1 and User3) generally automatically become domain contributors of workspaces (Workspace3) that are assigned in this manner.
upvoted 3 times

giraf Highly Voted 1 year, 9 months ago

NYN

Being a domain contributor doesn't give you any permissions on the workspace itself, only on the domain. Workspace roles are set differently. This excludes the first and last option.

User 3 will be domain contributor on the new workspace. Even though without domain admin permissions, there's nothing user 3 can do with Workspace 3.

upvoted 44 times

8d6881f 1 year, 9 months ago

Should be Y Y N

https://learn.microsoft.com/en-us/fabric/governance/domains

"After this, whenever a specified user or member of a specified security group creates a new workspace, it is assigned to the default domain.

The specified users and/or members of the specified security groups generally automatically become domain contributors of workspaces that are assigned in this manner."

upvoted 2 times

8d6881f 1 year, 9 months ago

I correct myself. The answer is Correct. Y N N.

Mind the order of things

upvoted 1 times

doctordodge 1 year, 10 months ago

Y Y N

https://learn.microsoft.com/en-us/fabric/governance/domains

upvoted 4 times

Question #25 *Topic 1*

You have two Fabric workspaces named Workspace1 and Workspace2.
You have a Fabric deployment pipeline named deployPipeline1 that deploys items from Workspace1 to Workspace2. DeployPipeline1 contains all the items in Workspace1.
You recently modified the items in Workspaces1.
The workspaces currently contain the items shown in the following table.

Items in Workspace1 that have the same name as items in Workspace2 are currently paired.
You need to ensure that the items in Workspace1 overwrite the corresponding items in Workspace2. The solution must minimize effort.
What should you do?

Workspace	Items
Workspace1	Model1 Notebook1 Report1 Lakehouse1 Pipeline1
Workspace2	Model1 Notebook2 Report1 Lakehouse2

A. Delete all the items in Workspace2, and then run deployPipeline1.

B. Rename each item in Workspace2 to have the same name as the items in Workspace1.

C. Back up the items in Workspace2, and then run deployPipeline1.

D. Run deployPipeline1 without modifying the items in Workspace2.

Suggested Answer: *D*

Community vote distribution

D (93%)

benni_ale Highly Voted 1 year, 9 months ago
Given minimal effort is required deploying the deploymentpipeline is enough
upvoted 6 times

exnaniantwort Highly Voted 1 year, 5 months ago
Item pairing Pairing is the process by which an item in one stage of the deployment pipeline is associated with the same item in the adjacent stage. If items aren't paired, even if they have the same name and type, the item in the target stage isn't overwritten during a deploy. A deploy of an unpaired item is known as a clean deploy and creates a copy of that item in the adjacent stage. Simply deploying will overwrite the corresponding items in Workspace2 by those in Workspace 1 https://learn.microsoft.com/en-us/fabric/cicd/deployment-pipelines/assign-pipeline?tabs=new-ui#item-pairing
upvoted 6 times

fdsdfgxcvbdsfhshfg Most Recent 1 year, 2 months ago
There is no information on whether the new items from Workspace1 should be deployed or not, but the only answer that does make sense is D, although some new objects will be created in Workspace2
upvoted 2 times

smanzana 1 year, 5 months ago
D is correct
upvoted 2 times

DarioReymago 1 year, 5 months ago
D is correct
upvoted 2 times

bettermakeme 1 year, 6 months ago

Deploy content from one stage to another There are several ways to deploy content from one stage to another. You can deploy all the content, or you can select which items to deploy. You can deploy content to any adjacent stage, in either direction. Deploying content from a working production pipeline to a stage that has an existing workspace, includes the following steps: Deploying new content as an addition to the content already there. Deploying updated content to replace some of the content already there.

upvoted 1 times

Zoran_S 1 year, 9 months ago

I interpret it as, the items in Workspace1 that are paired with Workspace2 need to be overwritten. Those that are not paired remain the same. Therefore just run the deployPipeline1 and everything should be good. If I delete everything in Workspace2 and then run deployPipeline1 then it would be the same as to sync Workspace1 to Workspace2

upvoted 1 times

galacaw 1 year, 9 months ago

https://learn.microsoft.com/en-us/fabric/cicd/deployment-pipelines/understand-the-deployment-process?tabs=new-ui#deployment-process

upvoted 1 times

giraf 1 year, 9 months ago

During deployment, notenook2 will not be deleted from Workspace2. So, you need to delete all items (or at least those not in Workspace1) and then run the pipeline.

upvoted 2 times

Tuki93 1 year, 9 months ago

I believe the provided answer 'D' is correct.

upvoted 2 times

GHill1982 1 year, 10 months ago

This requires minimal effort as it has no additional steps such as renaming of deleting items.

upvoted 4 times

Question #26 *Topic 1*

You have a Fabric workspace named Workspace1 that contains a data pipeline named Pipeline1 and a lakehouse named Lakehouse1.
You have a deployment pipeline named deployPipeline1 that deploys Workspace1 to Workspace2.
You restructure Workspace1 by adding a folder named Folder1 and moving Pipeline1 to Folder1.
You use deployPipeline1 to deploy Workspace1 to Workspace2.
What occurs to Workspace2?

A. Folder1 is created, Pipeline1 moves to Folder1, and Lakehouse1 is deployed.

B. Only Pipeline1 and Lakehouse1 are deployed.

C. Folder1 is created, and Pipeline1 and Lakehouse1 move to Folder1.

D. Only Folder1 is created and Pipeline1 moves to Folder1.

Suggested Answer: *A*

Community vote distribution

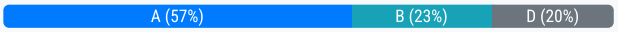

rohit_15 Most Recent 1 month, 2 weeks ago

Selected Answer: A

entire structire will move to Workspace2

upvoted 1 times

57f5e84 11 months, 2 weeks ago

Selected Answer: A

i make test and A is corret.

Folder1 will be create in workspace2 and contain Pipeline1 deploye, Lakehouse will be deploye in workspace2

upvoted 5 times

8115575 1 year, 2 months ago

Selected Answer: D

No reason for Lakehouse1 to be deployed when question only mentions a restructuring, meaning Lakehouse1 is already in workspace2

upvoted 3 times

8115575 1 year, 2 months ago

No reason for Lakehouse1 to be deployed when question only mentions a restructuring, meaning Lakehouse1 is already in workspace2We've selected a deployment of Workspace A to B, it means that all the items are processed even if they had no changes - deployment pipelines rules are reapplied, connections are re-established etc.

upvoted 3 times

learnazureportal Highly Voted 11 months, 4 weeks ago

Selected Answer: B

Only Pipeline1 and Lakehouse1 are deployed. Only the artifacts themselves (Pipeline1, Lakehouse1) are deployed, not the folder hierarchy.

upvoted 8 times

041d753 1 month, 3 weeks ago

Folders in deployment pipelines (preview)

Folders enable users to efficiently organize and manage workspace items in a familiar way. When you deploy content that contains folders to a different stage, the folder hierarchy of the applied items is automatically applied.

https://learn.microsoft.com/en-us/fabric/cicd/deployment-pipelines/understand-the-deployment-process?tabs=new-ui#folders-in-deployment-pipelines-preview

upvoted 1 times

fdsdfgxcvbdsfhshfg 1 year, 2 months ago

We've selected a deployment of Workspace A to B, it means that all the items are processed even if they had no changes - deployment pipelines rules are reapplied, connections are re-established etc.

upvoted 3 times

smanzana 1 year, 5 months ago

Selected Answer: A

A is correct

upvoted 5 times

ce6fb55 1 year, 7 months ago

Selected Answer: D

Wouldn't it be D if Lakehouse1 is already in Workspace2 and wasn't changed?

Usecase says, that only a restructure takes place.

upvoted 3 times

MultiCloudIronMan 1 year, 8 months ago

Selected Answer: A

Correct

upvoted 4 times

Goye 1 year, 8 months ago

Selected Answer: A

A is correct, I tested it in fabric

upvoted 4 times

8d6881f Highly Voted 1 year, 9 months ago

Selected Answer: A

https://learn.microsoft.com/en-us/fabric/cicd/deployment-pipelines/understand-the-deployment-process?tabs=new-ui#deploy-content-from-one-stage-to-another

"When content from the source stage is copied to the target stage, Fabric identifies existing content in the target stage and overwrites it. "

upvoted 6 times

Question #27 *Topic 1*

DRAG DROP -

Your company has a team of developers. The team creates Python libraries of reusable code that is used to transform data.

You create a Fabric workspace name Workspace1 that will be used to develop extract, transform, and load (ETL) solutions by using notebooks.

You need to ensure that the libraries are available by default to new notebooks in Workspace1.

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.

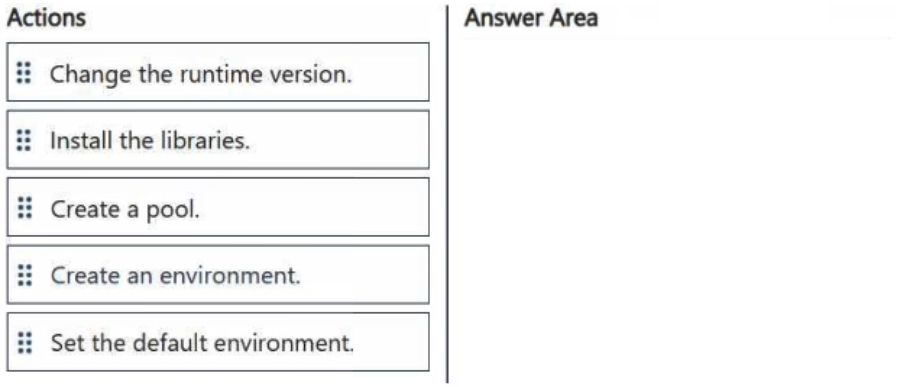

Suggested Answer:

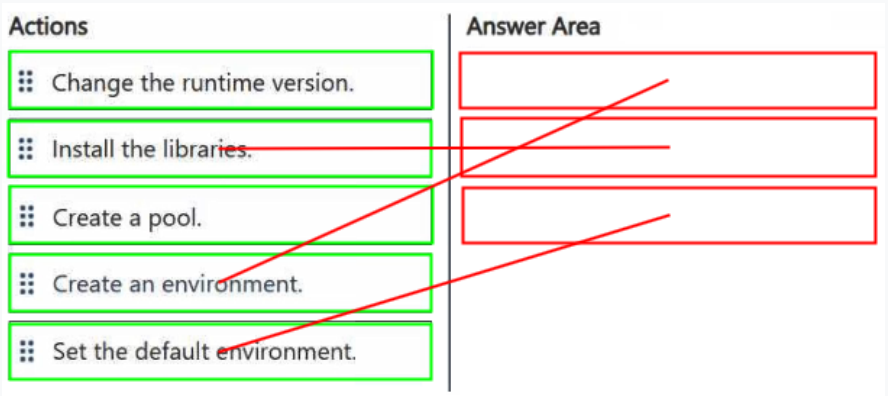

Mustafa_Shazly Most Recent 9 months, 2 weeks ago

I try to create new Env I found the true answer is

- Create an environment (Because we need to install custom lib)

- You can't publish new Env without any change (add new lib or change runtime version) so in our case we will add lib

- creating new Env can does with defult pool so we didn't need to create one

- set default Env true because users run it directry and if make it in pipeline attach new Env directly

upvoted 3 times

learnazureportal 11 months, 4 weeks ago

I go for the following steps - Also, Setting the default environment is optional and typically done after creating the environment.

https://learn.microsoft.com/en-us/fabric/data-engineering/create-and-use-environment

1) Create a pool → 2) Create an environment → 3) Install the libraries

upvoted 1 times

92ef42d 1 year, 2 months ago

I tried this in the Fabric.

First you should toggle ON "Set the default environment".

Then you can create a new environment and set it as default.

And only after that you can install additional libraries (if needed).

So, the answer is correct.

upvoted 3 times

5a51eb2 1 year ago

Well, to be fair first you need to create the environment, then install libraries and then when you toggle you have the already made environment to choose form the drop-down. So the answer is correct, but the toggling i.e., the "Set the defaulat environment" is the last step not the first step.

upvoted 2 times

e71e79b 1 year, 3 months ago

1) Create a pool

2) Create an environment

3) Set the default environment

upvoted 1 times

92ef42d 1 year, 2 months ago

I tried this in the Fabric. First you should toggle ON "Set the default environment". Then you can create a new environment and set it as default. And only after that you can install additional libraries (if needed). So, the answer is correct.

upvoted 3 times

VenkataPhaniPavanKumar 1 year, 5 months ago

In Microsoft Fabric warehouses, primary keys must be defined as NONCLUSTERED and NOT ENFORCED

upvoted 1 times

smanzana 1 year, 5 months ago

Correct

upvoted 1 times

DarioReymago 1 year, 5 months ago

correct

upvoted 1 times

Sunnyb Highly Voted 1 year, 7 months ago

Create an environment

You define a Spark environment that can hold runtime and library settings.

Install the libraries

Upload or reference your custom Python libraries (e.g., .whl, .egg, or .py files) in this environment.

Set the default environment

Assign this environment as the default for the workspace so that every new notebook automatically uses it.

upvoted 6 times

benni_ale 1 year, 9 months ago

Seems corret: Create (virtual) environment, install libraries, set the newly created environment as default from workspace settings

upvoted 4 times

henryphchan 1 year, 9 months ago

The answer is correct

upvoted 3 times

Tuki93 Highly Voted 1 year, 9 months ago

Seems correct.

upvoted 14 times

Question #28 *Topic 1*

You have a Fabric workspace that contains a lakehouse and a notebook named Notebook1. Notebook1 reads data into a DataFrame from a table named Table1 and applies transformation logic. The data from the DataFrame is then written to a new Delta table named Table2 by using a merge operation.
You need to consolidate the underlying Parquet files in Table1.
Which command should you run?

A. VACUUM

B. BROADCAST

C. OPTIMIZE

D. CACHE

Suggested Answer: *C*

Community vote distribution

C (100%)

rohit_15 Most Recent 1 month, 2 weeks ago

Selected Answer: C

Optimise will consolidate the small parquet files into larger files.

upvoted 1 times

smanzana 1 year, 5 months ago

Selected Answer: C

C is correct

upvoted 3 times

benni_ale 1 year, 9 months ago

Selected Answer: C

Optimise!

upvoted 2 times

henryphchan Highly Voted 1 year, 9 months ago

Selected Answer: C

In Delta Lake, the OPTIMIZE command is used to consolidate small Parquet files into larger ones. This improves query performance by reducing the overhead of managing many small files.

VACUUM command is used to clean up and remove files that are no longer needed (e.g., old versions of data files, deleted files)

upvoted 7 times

prabhjot 1 year, 9 months ago

Selected Answer: C

Optimize solves the "small file problem" by consolidating multiple small files into larger parquet files

upvoted 3 times

Tuki93 1 year, 9 months ago

Selected Answer: C

Optimize: Consolidates multiple small Parquet files into large file. Big Data processing engines, and all Fabric engines, benefit from having larger files sizes. Having files of size above 128 MB, and optimally close to 1 GB, improves compression and data distribution, across the cluster nodes. It reduces the need to scan numerous small files for efficient read operations. It's a general best practice to run optimization strategies after loading large tables.

upvoted 2 times

doctordodge 1 year, 10 months ago

Selected Answer: C

Link: https://learn.microsoft.com/en-us/fabric/data-engineering/lakehouse-table-maintenance

upvoted 2 times

GHill1982 Highly Voted 1 year, 10 months ago

Selected Answer: C

To consolidate the underlying Parquet files in Table1, you should run the OPTIMIZE command.

upvoted 6 times

Question #29 *Topic 1*

You have five Fabric workspaces.
You are monitoring the execution of items by using Monitoring hub.
You need to identify in which workspace a specific item runs.
Which column should you view in Monitoring hub?

A. Start time

B. Capacity

C. Activity name

D. Submitter

E. Item type

F. Job type

G. Location

Suggested Answer: *G*

Community vote distribution

G (95%) (5%)

92ef42d Most Recent 1 year, 2 months ago

Selected Answer: G

I've checked Fabric. It is location.

upvoted 2 times

NickNnack 1 year, 2 months ago

Selected Answer: G

Looking at Fabric, the answer is Location = Workspace

upvoted 1 times

ExamWiner 1 year, 3 months ago

Selected Answer: G

Ubicación. Confirmado!

upvoted 1 times

smanzana 1 year, 5 months ago

Selected Answer: G

Location

upvoted 1 times

akmorsh 1 year, 5 months ago

 Explanation: In Microsoft Fabric's Monitoring Hub, when you're tracking execution details (e.g., pipelines, notebooks, dataflows), each execution is associated with a Capacity — and each Fabric Capacity is linked to a specific workspace. The Capacity column in the Monitoring Hub reflects the name of the workspace where the item ran, because each workspace operates under a specific capacity, and the Monitoring Hub shows this to help distinguish execution sources.I have gone into my Monitor in Fabric, it is Location, not Capacity. Capacity provides a GUID that is not unique to a Workspace.

upvoted 1 times

akmorsh 1 year, 5 months ago

Selected Answer: B

 Explanation: In Microsoft Fabric's Monitoring Hub, when you're tracking execution details (e.g., pipelines, notebooks, dataflows), each execution is associated with a Capacity — and each Fabric Capacity is linked to a specific workspace. The Capacity column in the Monitoring Hub reflects the name of the workspace where the item ran, because each workspace operates under a specific capacity, and the Monitoring Hub shows this to help distinguish execution sources.

upvoted 1 times

NickNnack 1 year, 2 months ago

I have gone into my Monitor in Fabric, it is Location, not Capacity.

Capacity provides a GUID that is not unique to a Workspace.

upvoted 1 times

Besmarlyn 1 year, 7 months ago

Selected Answer: G

Answer is correct

upvoted 1 times

benni_ale 1 year, 9 months ago

Selected Answer: G

Location

upvoted 1 times

henryphchan 1 year, 9 months ago

Selected Answer: G

In Monitoring Hub, Location = Workspace

upvoted 2 times

Tuki93 1 year, 9 months ago

Selected Answer: G

Provided answer "Location" is correct

upvoted 1 times

GHill1982 Highly Voted 1 year, 10 months ago

Selected Answer: G

Correct. The Location shows the Workspace.

https://learn.microsoft.com/en-us/training/modules/monitor-fabric-items/3-use-monitor-hub

upvoted 10 times

Question #30 *Topic 1*

You have a Fabric workspace that contains a warehouse named DW1. DW1 is loaded by using a notebook named Notebook1.
You need to identify which version of Delta was used when Notebook1 was executed.
What should you use?

A. Real-Time hub

B. OneLake data hub

C. the Admin monitoring workspace

D. Fabric Monitor

E. the Microsoft Fabric Capacity Metrics app

Suggested Answer: *C*

Community vote distribution

D (100%)

smanzana Most Recent 1 year, 5 months ago

Selected Answer: D

Monitor

upvoted 2 times

benni_ale 1 year, 9 months ago

Selected Answer: D

Monitor

upvoted 2 times

Tuki93 1 year, 9 months ago

Selected Answer: D

We can find this in Monitor. example: Runtime information

Runtime 1.2 (Spark 3.4, Delta 2.4)

upvoted 4 times

GHill1982 Highly Voted 1 year, 10 months ago

Selected Answer: D

You can see this in the Monitor. On the Details section for a Notebook there is Runtime information which gives the version i.e. Runtime 1.3 (Spark 3.5, Delta 3.2)

upvoted 18 times

Question #31 *Topic 1*

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.

```
CREATE TABLE dbo.DimCustomer (
  CustomerKey VARCHAR(255) NOT NULL,
  Name VARCHAR(255) NOT NULL,
  Email VARCHAR(255) NOT NULL
  );
```

Code Segments

- `DROP CONSTRAINT PK_DimCustomer`
- `ADD CONSTRAINT PK_DimCustomer PRIMARY KEY NONCLUSTERED (CustomerKey)`
- `NOT ENFORCED`
- `ALTER TABLE dbo.DimCustomer`
- `ADD CONSTRAINT PK_DimCustomer PRIMARY KEY CLUSTERED (CustomerKey)`
- `ENFORCED`

Answer Area

Suggested Answer:
Answer: ALTER TABLE dbo.DimCustomer, then ADD CONSTRAINT PK_DimCustomer PRIMARY KEY NONCLUSTERED (CustomerKey), then NOT ENFORCED. Fabric Warehouse only lets you add a primary key to an existing table with ALTER TABLE, and a PRIMARY KEY is supported only when both NONCLUSTERED and NOT ENFORCED are specified. The CLUSTERED and ENFORCED segments are therefore invalid, and DROP CONSTRAINT is not needed because no constraint exists yet. Reference: https://learn.microsoft.com/en-us/fabric/data-warehouse/table-constraints

amli123 Highly Voted 1 year, 10 months ago
Alter table dbo.dimcustomer Add constraint PK_dimcustomer primary key nonclustered (customerkey) not enforced

PRIMARY KEY is only supported when NONCLUSTERED and NOT ENFORCED are both used. https://learn.microsoft.com/en-us/fabric/data-warehouse/table-constraints
upvoted 81 times

amli123 Highly Voted 1 year, 10 months ago
Alter table dbo.dimcustomer Add constraint PK_dimcustomer primary key nonclustered (customerkey) not enforced PRIMARY KEY is only supported when NONCLUSTERED and NOT ENFORCED are both used. https://learn.microsoft.com/en-us/fabric/data-warehouse/table-constraintsThis is the correct option, I have tired it in my warehouse ALTER TABLE [dbo].[dim_cust] ADD CONSTRAINT [PK_dim_cust_CustomerKey] PRIMARY KEY NONCLUSTERED (CustomerKey) NOT ENFORCEDALTER TABLE dbo.DimCustomer ADD CONSTRAINT PK_DimCustomer PRIMARY KEY CLUSTERED (CustomerKey) ENFORCEDThis is correct, Primary Key Non Clustered Not Enforced https://learn.microsoft.com/en-us/fabric/data-warehouse/table-constraints
upvoted 85 times

fogwills Most Recent 1 week ago

Non Clustered, Not Enforced.

From Microsoft doco:

SQL analytics endpoint and Warehouse in Microsoft Fabric support adding these table constraints to existing tables:

PRIMARY KEY is only supported when you use NONCLUSTERED and NOT ENFORCED.

FOREIGN KEY is only supported when you use NOT ENFORCED.

UNIQUE constraint is only supported when you use NONCLUSTERED and NOT ENFORCED.

upvoted 1 times

luckybalaji 5 months, 4 weeks ago

please correct the provided answer is incorrect one

upvoted 3 times

RickyK 11 months, 2 weeks ago

Primary Key

Non Clustered

Not enforced

As a SQL developer I would use CLUSTERED and ENFORCED in SQL Server. But this appears to be a particular constraint within Fabric. The Details are as follows.

"SQL analytics endpoint and Warehouse in Microsoft Fabric support these table constraints:

PRIMARY KEY is only supported when NONCLUSTERED and NOT ENFORCED are both used"

Source: https://learn.microsoft.com/en-us/fabric/data-warehouse/table-constraints

upvoted 5 times

92ef42d 1 year, 2 months ago

ALTER TABLE myTable

ADD CONSTRAINT PK_myTable

PRIMARY KEY CLUSTERED (c1)

ENFORCED;

PRIMARY KEY must be enforced to guarantee uniqueness and non-null values.

NOT ENFORCED is only valid for FOREIGN KEY or CHECK constraints, not for PRIMARY KEY.

upvoted 1 times

ziggy1117 1 year, 4 months ago

so answer is

ALTER TABLE

NONCLUSTERED

NOT ENFORCED

upvoted 4 times

smanzana 1 year, 5 months ago

ALTER TABLE

NONCLUSTERED

NOT ENFORCED

upvoted 3 times

DataSturdy 1 year, 5 months ago

ALTER TABLE dbo.dimcustomer

ADD CONSTRAINT PK_DimCustomer PRIMARY KEY NONCLUSTERED (customerkey)

not ENFORCED

this is the right answer

upvoted 3 times

turlakDE 1 year, 5 months ago

Alter table dbo.dimcustomer

Add constraint PK_dimcustomer primary key nonclustered (customerkey)

not enforced

Alter table dbo.dimcustomer

Add constraint PK_dimcustomer primary key nonclustered (customerkey)

not enforced - This is from Microsoft documentation

upvoted 3 times

ziggy1117 1 year, 4 months ago

https://learn.microsoft.com/en-us/fabric/data-warehouse/table-constraints The given example is: CREATE TABLE PrimaryKeyTable (c1 INT NOT NULL, c2 INT); To add PK: ALTER TABLE PrimaryKeyTable ADD CONSTRAINT PK_PrimaryKeyTable PRIMARY KEY NONCLUSTERED (c1) NOT ENFORCED;so answer is ALTER TABLE NONCLUSTERED NOT ENFORCED

upvoted 3 times

ziggy1117 1 year, 4 months ago

https://learn.microsoft.com/en-us/fabric/data-warehouse/table-constraints The given example is: CREATE TABLE PrimaryKeyTable (c1 INT NOT NULL, c2 INT); To add PK: ALTER TABLE PrimaryKeyTable ADD CONSTRAINT PK_PrimaryKeyTable PRIMARY KEY NONCLUSTERED (c1) NOT ENFORCED;

upvoted 3 times

Rajib 1 year, 5 months ago

The provided answer is wrong Alter table dbo.dimcustomer

Add constraint PK_dimcustomer primary key nonclustered (customerkey)

not enforced https://learn.microsoft.com/en-us/fabric/data-warehouse/table-constraints

upvoted 3 times

PBridge 1 year, 6 months ago

Correct Sequence:

ALTER TABLE dbo.DimCustomer

→ Prepares the table to modify its structure.

ADD CONSTRAINT PK_DimCustomer PRIMARY KEY CLUSTERED (CustomerKey)

→ Adds the primary key constraint to the column. Clustered is preferred by default unless specified otherwise.

ENFORCED

→ Ensures referential integrity is applied.

upvoted 1 times

kim32 1 year, 6 months ago

https://learn.microsoft.com/en-us/fabric/data-warehouse/table-constraints

CREATE TABLE PrimaryKeyTable (c1 INT NOT NULL, c2 INT);

ALTER TABLE PrimaryKeyTable

ADD CONSTRAINT PK_PrimaryKeyTable PRIMARY KEY NONCLUSTERED (c1)

NOT ENFORCED;

upvoted 4 times

Othman2110 1 year, 9 months ago

Primary Key

Non Clustered

Not Enforced

upvoted 14 times

fassil 1 year, 9 months ago

Alter table dbo.dimcustomer

Add constraint PK_dimcustomer primary key nonclustered (customerkey)

not enforced

Create a Microsoft Fabric Warehouse table with a primary key:

CREATE TABLE PrimaryKeyTable (c1 INT NOT NULL, c2 INT);

ALTER TABLE PrimaryKeyTable ADD CONSTRAINT PK_PrimaryKeyTable PRIMARY KEY NONCLUSTERED (c1) NOT ENFORCED;

upvoted 6 times

Tuki93 1 year, 9 months ago

Provided answer is WRONG.

https://learn.microsoft.com/en-us/fabric/data-warehouse/table-constraints

SQL analytics endpoint and Warehouse in Microsoft Fabric support these table constraints:

PRIMARY KEY is only supported when NONCLUSTERED and NOT ENFORCED are both used.

FOREIGN KEY is only supported when NOT ENFORCED is used.

UNIQUE constraint is only supported when NONCLUSTERED and NOT ENFORCED are both used.

SQL analytics endpoint and Warehouse don't support default constraints at this time.

upvoted 7 times

mvg100 1 year, 5 months ago

This is correct,

Primary Key

Non Clustered

Not Enforced

https://learn.microsoft.com/en-us/fabric/data-warehouse/table-constraints

upvoted 7 times

bettermakeme 1 year, 6 months ago

This is the correct option, I have tired it in my warehouse

ALTER TABLE [dbo].[dim_cust]

ADD CONSTRAINT [PK_dim_cust_CustomerKey] PRIMARY KEY NONCLUSTERED (CustomerKey) NOT ENFORCED

upvoted 10 times

ChenFu 1 year, 9 months ago

ALTER TABLE dbo.DimCustomer

ADD CONSTRAINT PK_DimCustomer PRIMARY KEY CLUSTERED (CustomerKey)

ENFORCED

upvoted 4 times

Question #32 *Topic 1*

You have a Fabric workspace that contains a semantic model named Model1.
You need to dynamically execute and monitor the refresh progress of Model1.
What should you use?

A. dynamic management views in Microsoft SQL Server Management Studio (SSMS)

B. Monitoring hub

C. dynamic management views in Azure Data Studio

D. a semantic link in a notebook

Suggested Answer: *D*

Community vote distribution

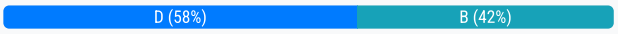

Rjb_073086 Most Recent 10 months, 1 week ago

Selected Answer: D

monitoring hub only tracks historical refreshes, so D

upvoted 1 times

peacegrace 11 months, 1 week ago

Selected Answer: D

Semantic link (via the sempy.fabric library in a Fabric notebook) allows you to dynamically execute a refresh using functions like refresh_dataset (which triggers the refresh and returns a request ID) and monitor progress by polling get_refresh_execution_details or list_refresh_requests for real-time status updates (e.g., "Running", "Succeeded", "Failed", duration, errors).

learn.microsoft.com

This supports programmatic, on-demand control directly in code.

The Monitoring hub (option B) is great for UI-based historical tracking, alerts, and viewing in-progress activities across workspaces, but it doesn't enable "dynamic execution" (programmatic triggering) of the refresh itself—it's more passive observation.

upvoted 3 times

learnazureportal 11 months, 4 weeks ago

Selected Answer: B

B. Monitoring hub

upvoted 2 times

ShazzyHawk Highly Voted 1 year, 1 month ago

Selected Answer: D

A. dynamic management views in SSMS DMVs in SSMS are for SQL databases or warehouses — not for Fabric semantic models.

B. Monitoring hub The Monitoring hub shows execution history and status after the refresh starts, but it does not allow dynamic execution or programmatic monitoring.

C. dynamic management views in Azure Data Studio Same limitation as SSMS — works for SQL engines, not semantic models in Fabric.

D. a semantic link in a notebook Correct — allows dynamic execution (start refresh) and programmatic monitoring of refresh status

upvoted 7 times

92ef42d 1 year, 2 months ago

Selected Answer: B

B. Monitoring hub

 Explanation:

The Monitoring hub in Microsoft Fabric is designed to track and monitor refresh progress of semantic models (like Model1) in real time.

It provides status updates, duration, success/failure logs, and refresh history.

Why not the others?

A & C (DMVs in SSMS/Azure Data Studio): These are used for querying metadata but don't support dynamic execution or real-time monitoring in Fabric.

D (Semantic link in a notebook): Used for querying semantic models, not for refreshing or monitoring them.

upvoted 2 times

fdsdfgxcvbdsfhshfg 1 year, 2 months ago

Selected Answer: B

Maybe the idea behind a dynamic execution it so take action based on the failed refresh so like a retry or something? It makes sense to be semantic link then.

upvoted 2 times

8115575 1 year, 2 months ago

Selected Answer: B

Semantic links are used to query semantic models from notebooks, not for monitoring refresh progress. In Monitoring hub its possible to monitor the refresh status and execute whenever you wish

upvoted 4 times

MarceloManhaes 1 year, 3 months ago

Selected Answer: B

The correct answer is: B. Monitoring hub

In Microsoft Fabric, the Monitoring hub is the appropriate tool to dynamically execute and monitor the refresh progress of semantic models like Model1. It provides visibility into refresh operations, including status, duration, and any errors encountered.

Here's why the other options are not suitable:

• A. Dynamic management views in SSMS: These are used for SQL Server databases, not for monitoring semantic model refreshes in Fabric.

• C. Dynamic management views in Azure Data Studio: Similar to SSMS, these are for database-level diagnostics, not for semantic model refresh tracking in Fabric.

• D. A semantic link in a notebook: Semantic links allow querying semantic models from notebooks but do not provide monitoring capabilities for refresh operations.

Ref: https://learn.microsoft.com/en-us/fabric/data-factory/semantic-model-refresh-activity

upvoted 3 times

bcbc59b 1 year, 4 months ago

Selected Answer: D

I agree with D. Monitoring hub only allows monitoring the refresh progress but not executing the refresh.

upvoted 3 times

LuisPBI25 1 year, 4 months ago

Selected Answer: D

B is not right answer because monitoring hub cannot provide with dinamyc execution

upvoted 1 times

Shw7 1 year, 5 months ago

Selected Answer: B

Monitoring HUB

Real-time visibility into refresh operations.

Execution status, duration, and error details.

upvoted 2 times

smanzana 1 year, 5 months ago

Selected Answer: B

B is correct

upvoted 2 times

smanzana 1 year, 5 months ago

B is correctLet me correct that: Correct is option D (Semantic link in a notebook) because allows both actions: -Dynamically triggering the refresh (using Python, Spark, or KQL code) -Monitoring the status and progress from the same notebook, even reacting to errors

upvoted 2 times

smanzana 1 year, 5 months ago

Let me correct that:

Correct is option D (Semantic link in a notebook) because allows both actions:

-Dynamically triggering the refresh (using Python, Spark, or KQL code)

-Monitoring the status and progress from the same notebook, even reacting to errors

upvoted 1 times

Runu 1 year, 5 months ago

Selected Answer: B

So, Why Is the Official Answer Still B. Monitoring hub?

Scenario Best Option

Manual / UI-based monitoring Monitoring hub

Automated / programmatic monitoring Notebook + semantic link (your method)

So your code is valid and powerful, especially for automation or scheduled monitoring, but if the exam or question context is asking for standard UI tools, then Monitoring hub (B) is the expected answer.

upvoted 3 times

kaushikcbs 1 year, 5 months ago

Selected Answer: B

Question not clear. If it is monitoring the refresh then Option b, if it is dynamically triggering the refresh then Option D

upvoted 2 times

Ricauduro 1 year, 5 months ago

Selected Answer: B

Not clear the meaning of dynamic....checking a failed exection using the monitor hub we can rerun and follow the execution

upvoted 1 times

DarioReymago 1 year, 5 months ago

Selected Answer: D

question is not clear. what they need with dynamic. If I need to react to a specific message I can work with semantic link. but if i need to manualy check the result the option is B

upvoted 1 times

Thaan 1 year, 6 months ago

Selected Answer: B

This is correct

upvoted 1 times

d98a23f 1 year, 8 months ago

Selected Answer: B

i think that both of B or D allow to achive this monitoriung refresh of model 1, however, it was specified that oui need to have the refresh progress (refresh status) and not historical duration for instance. therefore, Monitor hub is the suitable way to see it.

upvoted 1 times

prabhjot 1 year, 8 months ago

Selected Answer: B

I think Monitor Hub- A semantic link to a notebook would not be the best option for dynamically executing and monitoring the refresh progress of Model1 because notebooks are typically used for interactive data exploration and analysis. While you can write custom scripts and queries within a notebook to interact with your data, it does not provide the built-in monitoring capabilities that the Monitoring hub offers. The Monitoring hub is specifically designed to centralize and streamline the monitoring process, giving you real-time insights into the health and status of your refresh operations.

upvoted 2 times

toniteo 1 year, 9 months ago

Selected Answer: D

D is correct. Question partly says "dynamically execute" refresh, which is clearly illustrated below in way #4:

https://www.kevinrchant.com/2025/01/28/four-ways-to-monitor-a-semantic-model-refresh-in-microsoft-fabric/

upvoted 4 times

henryphchan 1 year, 9 months ago

Selected Answer: B

I think Monitor Hub is correct because you need to monitor the refresh progress. The Symantic link is designed to provide a reference or connection to your semantic model in a notebook context

upvoted 4 times

MultiCloudIronMan 1 year, 9 months ago

Selected Answer: B

Therefore, the correct answer is B. Monitoring hub. This tool allows you to dynamically monitor the progress of refresh operations and manage other activities related to your semantic model.

upvoted 3 times

AnetaK 1 year, 7 months ago

In Monitor Hub you cannot execute refresh.

upvoted 2 times

Tuki93 1 year, 9 months ago

Selected Answer: D

I think provided answer is correct: a semantic link in a notebook

upvoted 2 times

GHill1982 Highly Voted 1 year, 10 months ago

Selected Answer: D

If the definition of "dynamic" is the ability to execute and monitor in real-time, a semantic link in a notebook could be more suitable than the monitoring hub.

upvoted 25 times

Question #33 *Topic 1*

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You have a Fabric eventstream that loads data into a table named Bike_Location in a KQL database. The table contains the following columns:

BikepointID -

Street -

Neighbourhood -

No_Bikes -

No_Empty_Docks -

Timestamp -

You need to apply transformation and filter logic to prepare the data for consumption. The solution must return data for a neighbourhood named Sands End when No_Bikes is at least 15. The results must be ordered by No_Bikes in ascending order.

Solution: You use the following code segment:

Does this meet the goal?

```
bike_location
| filter Neighbourhood == "Sands End" and No_Bikes >= 15
| sort by No_Bikes
| project BikepointID, Street, Neighbourhood, No_Bikes, No_Empty_Docks, Timestamp
| project BikepointID, Street, Neighbourhood, No_Bikes, No_Empty_Docks, Timestamp
```

A. Yes

B. No

Suggested Answer: *B*

Community vote distribution

B (98%) 2

SamuComqi Highly Voted 1 year, 10 months ago

The answer is B. No because the "sort by" is sorting values in descending order (default behavior --> https://learn.microsoft.com/en-us/kusto/query/sort-operator?view=microsoft-fabric). One should add "asc" to sort values as required. The double "project" at the end does not affect the final result

upvoted 24 times

MuffiSan Highly Voted 1 year, 10 months ago

sort is desc by default in kql

upvoted 5 times

9114bc3 Most Recent 6 months, 1 week ago

| Requirement | Query implementation | | ---------- | ---------- | | Return only Sands End | `Neighbourhood == "Sands End"` | | No_Bikes at least 15 | `No_Bikes >= 15` | | Sort ascending by No_Bikes | `sort by No_Bikes` (ascending is default) | | Return selected columns | `project ...` |

upvoted 1 times

evest 11 months ago

The table in the requirements is written with first capital letters of the two words Bike_Location while in the KQL script in the question itself it is bike_location. According to the Kusto documentation "Identifiers are case-sensitive. For example, you can't refer to a table called ThisTable as thisTABLE." https://learn.microsoft.com/en-us/kusto/query/schema-entities/entity-names?view=azure-data-explorer&preserve-view=true#reference-identifiers-in-queries

upvoted 1 times

92ef42d 1 year, 2 months ago

B. Default KQL sorting is descending

upvoted 3 times

smanzana 1 year, 5 months ago

B is correct

upvoted 2 times

smanzana 1 year, 5 months ago

B-No, in KQL by default “sort by” is desc order

upvoted 1 times

turlakDE 1 year, 5 months ago

sort in kql by default is descending

upvoted 1 times

8d1e8fe 1 year, 6 months ago

https://learn.microsoft.com/en-us/kusto/query/?view=microsoft-fabric Kusto is case sensitive on almost anything. The table being referenced is not capitalized like the name. Could be a typo in the test or an attention to detail question.

upvoted 3 times

benni_ale 1 year, 9 months ago

sort by in kql is by default in descending order

upvoted 2 times

GHill1982 1 year, 10 months ago

In KQL the where keyword should be used instead of filter.Apologies, I have now learned the where and filter operators are equivalent in KQL. So the answer is correct depending on whether the duplicated project is a mistake in the question.Also consider that the default for order by/sort by in KQL is is desc (high to low). The opposite of SQL. https://learn.microsoft.com/en-gb/kusto/query/sort-operator?view=azure-data-explorer#parametersThe duplicated project, even though not required, does however work and does not throw any errors.

upvoted 3 times

GHill1982 1 year, 10 months ago

Selected Answer: B

In KQL the where keyword should be used instead of filter.

upvoted 3 times

GHill1982 1 year, 10 months ago

Apologies, I have now learned the where and filter operators are equivalent in KQL. So the answer is correct depending on whether the duplicated project is a mistake in the question.

upvoted 1 times

GHill1982 1 year, 10 months ago

Apologies, I have now learned the where and filter operators are equivalent in KQL. So the answer is correct depending on whether the duplicated project is a mistake in the question.Also consider that the default for order by/sort by in KQL is is desc (high to low). The opposite of SQL. https://learn.microsoft.com/en-gb/kusto/query/sort-operator?view=azure-data-explorer#parametersThe duplicated project, even though not required, does however work and does not throw any errors.

upvoted 1 times

GHill1982 1 year, 10 months ago

Also consider that the default for order by/sort by in KQL is is desc (high to low). The opposite of SQL. https://learn.microsoft.com/en-gb/kusto/query/sort-operator?view=azure-data-explorer#parametersThe duplicated project, even though not required, does however work and does not throw any errors.

upvoted 1 times

GHill1982 1 year, 10 months ago

Also consider that the default for order by/sort by in KQL is is desc (high to low). The opposite of SQL.

https://learn.microsoft.com/en-gb/kusto/query/sort-operator?view=azure-data-explorer#parameters

upvoted 1 times

GHill1982 1 year, 10 months ago

The duplicated project, even though not required, does however work and does not throw any errors.

upvoted 1 times

QAZdbarhate12345678 1 year, 11 months ago

The provided code segment meets the goal. Breakdown: Filter Condition: The code uses filter Neighbourhood == "Sands End" and No_Bikes >= 15 to meet the requirement of filtering the data for the "Sands End" neighborhood and where No_Bikes is at least 15. Sorting: The sort by No_Bikes clause ensures that the results are ordered by No_Bikes in ascending order, which aligns with the requirement. Projection: The fields BikepointID, Street, Neighbourhood, No_Bikes, No_Empty_Docks, and Timestamp are projected, which ensures the required fields are returned for the query. Hence, the solution is correct and achieves the stated goal. The correct answer is A. Yes.

upvoted 1 times

Question #34 *Topic 1*

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You have a Fabric eventstream that loads data into a table named Bike_Location in a KQL database. The table contains the following columns:

BikepointID -

Street -

Neighbourhood -

No_Bikes -

No_Empty_Docks -

Timestamp -
You need to apply transformation and filter logic to prepare the data for consumption. The solution must return data for a neighbourhood named Sands End when No_Bikes is at least 15. The results must be ordered by No_Bikes in ascending order.

Solution: You use the following code segment:

Does this meet the goal?

```
bike_location
| filter Neighbourhood == "Sands End" and No_Bikes >= 15
| order by No_Bikes
| project BikepointID, Street, Neighbourhood, No_Bikes, No_Empty_Docks, Timestamp
```

A. Yes

B. No

Suggested Answer: *B*

Community vote distribution

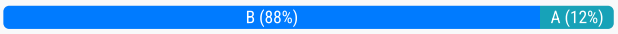

benni_ale Highly Voted 1 year, 9 months ago
no, order by by default is desc in kql
upvoted 11 times

mauriciorosa8 1 year, 9 months ago
The default sorting order in KQL is **descending (desc)**; **ascending (asc) must be explicitly specified when needed.**
upvoted 5 times

Rozamunde Most Recent 1 year, 3 months ago
Duplicate of previous question
upvoted 4 times

smanzana 1 year, 5 months ago
B is correct
upvoted 3 times

smanzana 1 year, 5 months ago
B-No, in KQL by default “order by” is desc order
upvoted 3 times

akmorsh 1 year, 5 months ago

Correct Answer: B. No Explanation: The code is almost identical to the previous correct example, but this time it uses order by instead of sort by, and that's the issue.

upvoted 3 times

568f95c 1 year, 9 months ago

They are asking for the order of No_Bikes in ascending order in this question

upvoted 2 times

Tuki93 1 year, 9 months ago

default is desc in Kusto

upvoted 5 times

8d6881f 1 year, 9 months ago

no. In kql desc is default for sort/order by: "asc sorts into ascending order, low to high. Default is desc, high to low." https://learn.microsoft.com/en-us/kusto/query/sort-operator?view=microsoft-fabric

upvoted 4 times

amli123 1 year, 10 months ago

default is desc

upvoted 2 times

MuffiSan 1 year, 10 months ago

no. In kql desc is default for sort by: "asc sorts into ascending order, low to high. Default is desc, high to low." https://learn.microsoft.com/en-us/kusto/query/sort-operator?view=microsoft-fabric"The sort and order operators are equivalent"

upvoted 2 times

MuffiSan 1 year, 10 months ago

Selected Answer: B

no. In kql desc is default for sort by:

"asc sorts into ascending order, low to high. Default is desc, high to low." https://learn.microsoft.com/en-us/kusto/query/sort-operator?view=microsoft-fabric

upvoted 2 times

MuffiSan 1 year, 10 months ago

"The sort and order operators are equivalent"

upvoted 1 times

QAZdbarhate12345678 1 year, 11 months ago

The provided solution meets the goal. Analysis: Filtering Condition: The code includes filter Neighbourhood == "Sands End" and No_Bikes >= 15, which satisfies the requirement to only include rows where the neighborhood is "Sands End" and No_Bikes is at least 15. Ordering: The order by No_Bikes clause ensures that the results are sorted by No_Bikes in ascending order, meeting the stated requirement. Projection: The project BikepointID, Street, Neighbourhood, No_Bikes, No_Empty_Docks, Timestamp ensures that the necessary columns are returned. Conclusion: The solution fulfills all the requirements. Correct Answer: A. Yes.

upvoted 2 times

Bharat 1 year, 11 months ago

Selected Answer: A

the answer is yes since you don't have to specify the "asc" option. It is the default.

upvoted 2 times

Bharat 1 year, 11 months ago

the answer is yes since you don't have to specify the "asc" option. It is the default.The answer is no. The default for order by in KQL in desc therefore asc must be specified. Confusing at first as it is the opposite of SQL. https://learn.microsoft.com/en-gb/kusto/query/sort-operator?view=azure-data-explorer#parameters

upvoted 2 times

GHill1982 Highly Voted 1 year, 10 months ago

The answer is no. The default for order by in KQL in desc therefore asc must be specified. Confusing at first as it is the opposite of SQL. https://learn.microsoft.com/en-gb/kusto/query/sort-operator?view=azure-data-explorer#parameters

upvoted 8 times

Question #35 *Topic 1*

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You have a Fabric eventstream that loads data into a table named Bike_Location in a KQL database. The table contains the following columns:

BikepointID -

Street -

Neighbourhood -

No_Bikes -

No_Empty_Docks -

Timestamp -
You need to apply transformation and filter logic to prepare the data for consumption. The solution must return data for a neighbourhood named Sands End when No_Bikes is at least 15. The results must be ordered by No_Bikes in ascending order.

Solution: You use the following code segment:

Does this meet the goal?

```
bike_location
| filter Neighbourhood == "Sands End" and No_Bikes >= 15
| sort by No_Bikes asc
| project BikepointID, Street, Neighbourhood, No_Bikes, No_Empty_Docks, Timestamp
```

A. Yes

B. No

Suggested Answer: *A*

Community vote distribution

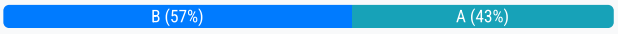

GHill1982 Highly Voted 1 year, 10 months ago
The sort and order operators are equivalent. https://learn.microsoft.com/en-us/kusto/query/sort-operator?view=microsoft-fabric
upvoted 15 times

rc012345 Highly Voted 12 months ago
in KQL, both table and column names are case sensitive. Hence, the table name doesn't match with description. Hence, this query will fail. NO issues with filter as both "where" and "filter" keywords work in KQL.
upvoted 6 times

evest Most Recent 11 months ago
KQL is case-sensitive and the table name is written in different way in the requirements and in the question with the script.
upvoted 5 times

Bin_Hashim 1 year ago
why we are not Considering Table Name **Table name is incorrect. It is case-sensitive. In the scenario the table is named Bike_Location. The query uses bike_location (all lowercase). In KQL, identifiers (table/column names) are case-sensitive. Using the wrong case for the table name will cause the query to fail to resolve the table, even though the rest of the query logic is correct:
upvoted 4 times

Sangeetha112 1 year, 1 month ago

'filter' is wrong, instead 'where' is correct, so I think the answer is 'B'

upvoted 2 times

92ef42d 1 year, 2 months ago

Table name is incorrect. It is case-sensitive. https://learn.microsoft.com/en-us/kusto/query/schema-entities/entity-names?view=azure-data-explorer&preserve-view=true#reference-identifiers-in-queries

upvoted 3 times

wellingtonluis 1 year, 2 months ago

Don't you think answer is "No" , as the query should use "where" instead of "filter"?

upvoted 3 times

smanzana 1 year, 5 months ago

A is correct

upvoted 1 times

smanzana 1 year, 5 months ago

A (Adding ASC after ORDER BY no_bikes explicitly enforces ascending order)

upvoted 1 times

a823528 1 year, 5 months ago

"Identifiers are case-sensitive. For example, you can't refer to a table called ThisTable as thisTABLE." I wonder how many times I have to add a comment before it is shown in the discussion. https://learn.microsoft.com/en-us/kusto/query/schema-entities/entity-names?view=azure-data-explorer&preserve-view=true#reference-identifiers-in-queriesDon't you think answer is "No" , as the query should use "where" instead of "filter"The where and filter operators are equivalent

upvoted 1 times

a823528 1 year, 5 months ago

Selected Answer: B

"Identifiers are case-sensitive. For example, you can't refer to a table called ThisTable as thisTABLE." I wonder how many times I have to add a comment before it is shown in the discussion. https://learn.microsoft.com/en-us/kusto/query/schema-entities/entity-names?view=azure-data-explorer&preserve-view=true#reference-identifiers-in-queries

upvoted 1 times

Jay63002 1 year, 5 months ago

Don't you think answer is "No" , as the query should use "where" instead of "filter"

upvoted 1 times

Jay63002 1 year, 5 months ago

Don't you think answer is "No" , as the query should use "where" instead of "filter"The where and filter operators are equivalent

upvoted 1 times

Massy 11 months, 3 weeks ago

The where and filter operators are equivalent

upvoted 1 times

benni_ale 1 year, 9 months ago

correct!

upvoted 1 times

henryphchan 1 year, 9 months ago

sort by and order by are synonyms

upvoted 1 times

Bharat 1 year, 11 months ago

Selected Answer: B

Actually, the correct answer is No here because 'sort by' is illegal in KQL. It should be 'order by' instead.

upvoted 1 times

Bharat 1 year, 11 months ago

Actually, the correct answer is No here because 'sort by' is illegal in KQL. It should be 'order by' instead."sort by" is legal in KQL, it is the same as "order by"no, sort by and order by both work and are equal A is correct.

upvoted 1 times

henryphchan 1 year, 9 months ago

"sort by" is legal in KQL, it is the same as "order by"

upvoted 1 times

MuffiSan 1 year, 10 months ago

no, sort by and order by both work and are equal A is correct.

upvoted 2 times

Question #36 *Topic 1*

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You have a Fabric eventstream that loads data into a table named Bike_Location in a KQL database. The table contains the following columns:

BikepointID -

Street -

Neighbourhood -

No_Bikes -

No_Empty_Docks -

Timestamp -

You need to apply transformation and filter logic to prepare the data for consumption. The solution must return data for a neighbourhood named Sands End when No_Bikes is at least 15. The results must be ordered by No_Bikes in ascending order.

Solution: You use the following code segment:

Does this meet the goal?

```
SELECT BikepointID, Street, Neighbourhood, No_Bikes, No_Empty_Docks, Timestamp
FROM bike_location
WHERE neighbourhood = 'Sands End'
AND no_bikes >= 15
ORDER BY no_bikes
```

A. Yes

B. No

Suggested Answer: *B*

Community vote distribution

A (51%) B (49%)

Mourya959595 Highly Voted 1 year, 11 months ago

SQL By default shows ASC order

upvoted 29 times

JensQ 1 year, 8 months ago

Selected Answer: B

You need to use "--" to tell the KQL DB to read code as SQL. Otherwise it tries to read it as KQL

upvoted 15 times

JensQ Highly Voted 1 year, 8 months ago

You need to use "--" to tell the KQL DB to read code as SQL. Otherwise it tries to read it as KQLJust run a SQL query in KQL database without any indication that is SQL so without using "" or -- anything likr that and it run just fine. Also I confirmed that ASC is the defaultNow you don't need. The code works properly in Fabric

upvoted 16 times

5a51eb2 1 year ago

Just run a SQL query in KQL database without any indication that is SQL so without using "" or -- anything likr that and it run just fine. Also I confirmed that ASC is the default

upvoted 4 times

92ef42d 1 year, 2 months ago

Now you don't need. The code works properly in Fabric

upvoted 5 times

sooss Most Recent 7 months ago

All KQL options have Filter instead of Where

upvoted 1 times

EddEe 10 months ago

Case sensitivity on the table name - should be Bike_Location

upvoted 2 times

evest 11 months ago

I tested in a KQL database with a table named Taxivendor1 in a KQL Queryset and ran the following statement: select top 100 * from taxivendor1 and it gave error: Semantic error: There is no table '...taxivendor1'. But when I run it with capital letter select top 100 * from Taxivendor1 then it is alright.

upvoted 4 times

peacegrace 11 months ago

Because the given solution is not valid KQL syntax and would fail to run (or return incorrect/no results), it does not meet the goal.

upvoted 2 times

Bin_Hashim 1 year ago

Query should not be Case Sensitive ??

upvoted 3 times

wellingtonluis 1 year, 1 month ago

Need to add "--" on the begining

upvoted 1 times

92ef42d 1 year, 2 months ago

You can use SQL code directly without any additional "--" now (checked in Fabric). Also, it is case insensitive to the column and table names, order by sorts ascending by default. So, everything is correct.

upvoted 2 times

wellingtonluis 1 year, 2 months ago

A Because the -- is not required for T-SQL in Fabric.

upvoted 2 times

e392998 1 year, 3 months ago

General syntax is correct but there are typo in neighbourhood and no_bikes

upvoted 1 times

ziggy1117 1 year, 4 months ago

I tried it in Fabric and it works. Order By is ascending order by default in SQL. (It is desc order in KQL) No need to use "--" in the Eventhouse. Also it is case insensitive so even if Neighbourhood is the column, WHERE neihborhood >=15 works

upvoted 7 times

norkita 1 year, 4 months ago

A , is SQL

upvoted 2 times

necktru 1 year, 4 months ago

Missing- - characters

upvoted 1 times

Delafupra 1 year, 4 months ago

The query it is right, but I believe it a tricky question. The solution is not correct just for the fact that is a SQL query I believe the solution must be a KQL query

upvoted 2 times

smanzana 1 year, 5 months ago

It is not clear but the question states that the table is in a KQL database, so the query targets a KQL DB. Since the syntax is T-SQL, it must be prefixed with -- to work correctly.

upvoted 1 times

Stef6108 1 year, 5 months ago

SQL By default shows ASC order

upvoted 2 times

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.

You need to ensure that processes for the bronze and silver layers run in isolation.

How should you configure the Apache Spark settings?

A. Disable high concurrency.

B. Create a custom pool.

C. Modify the number of executors.

D. Set the default environment.

Suggested Answer: *B*

Community vote distribution

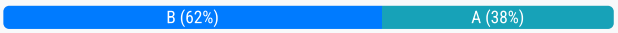

zulajariwala Most Recent 8 months, 1 week ago

Selected Answer: B

B. Create a custom pool.

Why: the workspace already has high concurrency enabled, which is designed to share Spark sessions across workloads for better startup time and cost efficiency. That is the opposite of running bronze and silver processes in isolation. Microsoft's docs describe high concurrency as one Spark application hosting multiple workloads in a shared session model, while custom Spark pools provide deeper control over compute for specific workloads.

So, to ensure the bronze and silver processes run in isolation, use a custom Spark pool for those workloads.

upvoted 1 times

5e89616 1 year, 7 months ago

Selected Answer: A

The requirement is to run notebook sessions for Bronze and Silver in isolation -> disable high concurrency: High concurrency mode allows users to share the same Spark sessions in Apache Spark for Fabric data engineering and data science workloads. An item like a notebook uses a Spark session for its execution and when enabled allows users to share a single Spark session across multiple notebooks. https://learn.microsoft.com/en-us/fabric/data-engineering/workspace-admin-settings#high-concurrency

upvoted 5 times

Mustafa_Shazly 9 months, 2 weeks ago

Selected Answer: A

When H Concurrency is enabled

- the standard and custom pool will shared and no run Isolation

"For custom pools with high concurrency mode, users get 36X faster session start experience compared to a standard Spark session."

https://learn.microsoft.com/en-us/fabric/data-engineering/high-concurrency-overview

upvoted 4 times

EddEe 10 months ago

Selected Answer: B

You use spark pool when you build medallion style transformations or process large files

upvoted 1 times

92ef42d 1 year, 2 months ago

Selected Answer: B

Answer is B.

A would make things worse.

C improves the performance, but would not isolate

D has nothing about performance or isolation

upvoted 1 times

smanzana 1 year, 5 months ago

Selected Answer: B

B- Creating a custom Apache Spark pool allows you to isolate workloads such as the bronze and silver layer processes. This ensures dedicated resources for each process, improves performance, and prevents contention between workloads.

upvoted 1 times

12a2ecc Highly Voted 1 year, 7 months ago

Selected Answer: B

In Microsoft Fabric, when you want to isolate different workloads, such as data processing jobs for different medallion layers (bronze, silver, gold), the best practice is to use custom Spark pools. This allows:

Isolation of resources: Ensuring the bronze and silver layers don't interfere with each other in terms of compute or memory.

Dedicated execution environments: Each pool can be optimized/configured differently based on the workload characteristics.

Better performance and reliability: Avoid contention in high-load scenarios like month-end processing or spikes during new book releases.

upvoted 11 times

Biju1 1 year, 7 months ago

Selected Answer: B

B is correct

upvoted 1 times

5e89616 Highly Voted 1 year, 7 months ago

Selected Answer: A

The requirement is to run notebook sessions for Bronze and Silver in isolation -> disable high concurrency:

High concurrency mode allows users to share the same Spark sessions in Apache Spark for Fabric data engineering and data science workloads. An item like a notebook uses a Spark session for its execution and when enabled allows users to share a single Spark session across multiple notebooks.

https://learn.microsoft.com/en-us/fabric/data-engineering/workspace-admin-settings#high-concurrency

upvoted 8 times

12a2ecc 1 year, 7 months ago

I think it is (B), custom pools:

Why Not the Other Options?

A. Disable high concurrency:

Not suitable here. High concurrency mode allows multiple lightweight queries to run simultaneously. Disabling it doesn't help with workload isolation and might reduce performance for concurrent jobs.

C. Modify the number of executors:

This affects how Spark parallelizes tasks within a single job but doesn't isolate workloads or prevent resource contention between different jobs.

D. Set the default environment:

This option just sets where notebooks run by default. It doesn't provide isolation between different layers.

upvoted 6 times

DRAG DROP
-

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.

You need to ensure that the authors can see only their respective sales data.

How should you complete the statement? To answer, drag the appropriate values the correct targets. Each value 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.

NOTE: Each correct selection is worth one point.

Values

- AuthorSales
- AuthorEmail
- AuthorSales.AuthorEmail
- BLOCK
- FILTER
- INLINE
- SCHEMABINDING
- USER_NAME()

Answer Area

```
CREATE FUNCTION dbo.tvf_rlspredicate(@Author AS varchar(50))
    RETURNS TABLE
WITH [                    ]
AS
    RETURN SELECT 1 AS tvf_rlspredicate_result
WHERE @Author = [                    ]
GO

CREATE SECURITY POLICY RLSFilter
ADD FILTER PREDICATE Security.tvf_rlspredicate(AuthorEmail)
ON [                    ]
WITH (STATE = ON)
```

Suggested Answer:

Answer Area

```
CREATE FUNCTION dbo.tvf_rlspredicate(@Author AS varchar(50))
    RETURNS TABLE
WITH [SCHEMABINDING]
AS
    RETURN SELECT 1 AS tvf_rlspredicate_result
WHERE @Author = [USER_NAME()]
GO

CREATE SECURITY POLICY RLSFilter
ADD FILTER PREDICATE Security.tvf_rlspredicate(AuthorEmail)
ON [AuthorSales]
WITH (STATE = ON)
```

5e89616 Highly Voted 1 year, 7 months ago

The answer is correct. - USER_NAME is compared the value of a column to the user currently running the query. If they match, the user can see the row; if not, they can't. - SCHEMABINDING - used in security policy function to prevent structural changes on the tables that are bound to the policy - TableName on which the security policy function should be applied

upvoted 13 times

5e89616 Highly Voted 1 year, 7 months ago

The answer is correct. - USER_NAME is compared the value of a column to the user currently running the query. If they match, the user can see the row; if not, they can't. - SCHEMABINDING - used in security policy function to prevent structural changes on the tables that are bound to the policy - TableName on which the security policy function should be appliedSource: https://learn.microsoft.com/en-us/fabric/data-warehouse/row-level-security

upvoted 13 times

Azure-JL 1 year, 7 months ago

Source: https://learn.microsoft.com/en-us/fabric/data-warehouse/row-level-security

upvoted 4 times

92ef42d Most Recent 1 year, 2 months ago

- INLINE - USER_NAME() - AuthorSales INLINE functions are more performant and required for RLS predicates. SCHEMABINDING ensures the function is tightly bound to the schema, preventing changes that could break RLS logic.

upvoted 1 times

hergoi 1 year, 3 months ago

- INLINE - USER_NAME - AuthorSales When implementing Row-Level Security (RLS) in Fabric, the goal is to ensure that each author only sees their own sales data. 1. INLINE vs SCHEMABINDING INLINE is the recommended option for RLS because it integrates the predicate directly into the query execution plan, resulting in better performance. Microsoft documentation and best practices specifically recommend INLINE for RLS predicates CREATE FUNCTION dbo.tvf_rlspredicate(@Author AS varchar(50)) RETURNS TABLE WITH INLINE AS RETURN SELECT 1 AS tvf_rlspredicate_result WHERE @Author = USER_NAME() GO CREATE SECURITY POLICY RLSFilter ADD FILTER PREDICATE Security.tvf_rlspredicate(AuthorEmail) ON AuthorSales WITH (STATE = ON);

upvoted 1 times

smanzana 1 year, 5 months ago

The answer is correct

upvoted 1 times

DarioReymago 1 year, 5 months ago

https://learn.microsoft.com/en-us/fabric/data-warehouse/row-level-security#examples -- Creating a function for the SalesRep evaluation CREATE FUNCTION Security.tvf_securitypredicate(@Author AS nvarchar(50)) RETURNS TABLE WITH SCHEMABINDING AS RETURN SELECT 1 AS tvf_securitypredicate_result WHERE @SalesRep = USER_NAME(); GO -- Using the function to create a Security Policy CREATE SECURITY POLICY RLSFilter ADD FILTER PREDICATE Security.tvf_securitypredicate(AuthorEmail) ON Authorsales.AuthorEmail WITH (STATE = ON); GO

upvoted 1 times

mujicalj 1 year, 7 months ago

Create a Security schema, a function Security.tvf_securitypredicate, and a security policy SalesFilter.

SQL

Copy -- Creating schema for Security CREATE SCHEMA Security; GO

-- Creating a function for the SalesRep evaluation CREATE FUNCTION Security.tvf_securitypredicate(@SalesRep AS nvarchar(50)) RETURNS TABLE WITH SCHEMABINDING AS RETURN SELECT 1 AS tvf_securitypredicate_result WHERE @SalesRep = USER_NAME() OR USER_NAME() = 'manager@contoso.com'; GO

-- Using the function to create a Security Policy CREATE SECURITY POLICY SalesFilter ADD FILTER PREDICATE Security.tvf_securitypredicate(SalesRep) ON sales.Orders WITH (STATE = ON); GO

upvoted 2 times

mujicalj 1 year, 7 months ago

Create a Security schema, a function Security.tvf_securitypredicate, and a security policy SalesFilter. SQL Copy -- Creating schema for Security CREATE SCHEMA Security; GO -- Creating a function for the SalesRep evaluation CREATE FUNCTION Security.tvf_securitypredicate(@SalesRep AS nvarchar(50)) RETURNS TABLE WITH SCHEMABINDING AS RETURN SELECT 1 AS tvf_securitypredicate_result WHERE @SalesRep = USER_NAME() OR USER_NAME() = ''; GO -- Using the function to create a Security Policy CREATE SECURITY POLICY SalesFilter ADD FILTER PREDICATE Security.tvf_securitypredicate(SalesRep) ON sales.Orders WITH (STATE = ON); GO

upvoted 2 times

Question #39 *Topic 1*

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You have an Azure key vault named KeyVault1 that contains secrets.

You have a Fabric workspace named Workspace1. Workspace contains a notebook named Notebook1 that performs the following tasks:

• Loads stage data to the target tables in a lakehouse
• Triggers the refresh of a semantic model

You plan to add functionality to Notebook1 that will use the Fabric API to monitor the semantic model refreshes.

You need to retrieve the registered application ID and secret from KeyVault1 to generate the authentication token.

Solution: You use the following code segment:

Use notebookutils.credentials.getSecret and specify the key vault URL and key vault secret.

Does this meet the goal?

A. Yes

B. No

Suggested Answer: *A*

Community vote distribution

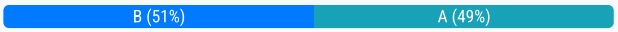

juamd Most Recent 6 months ago

Selected Answer: A

I would say A, since key vault secret i think is related to the secret name instead of the secret value

upvoted 4 times

Rjb_073086 10 months, 1 week ago

Selected Answer: B

the getsecret's input should be secret name, not the URL

upvoted 1 times

larinate 11 months, 2 weeks ago

Selected Answer: B

No, because notebookutils.credentials.getSecret(...) is a Fabric workspace credential store, but credentials are at Azure Key Value, which is different item.

upvoted 2 times

riccardodeana93 12 months ago

Selected Answer: A

According to the documentation, the name of the secret is simply called secret. So I think that's correct.

Help on module notebookutils.credentials in notebookutils:

NAME

notebookutils.credentials - Utility for credentials operations in Fabric

FUNCTIONS

getSecret(akvName, secret) -> str

Gets a secret from the given Azure Key Vault.

:param akvName: The name of the Azure Key Vault.

:param secret: The name of the secret.

:return: The secret value.

upvoted 4 times

92ef42d 1 year, 2 months ago

Selected Answer: B

No.

notebookutils.credentials.getSecret("https://<vault-name>.vault.azure.net/", "mySecretName")

URL works

Secret name should be provided instead of secret itself (why do you need a kv, if you show the secret in your code?)

upvoted 5 times

Jennyfly Highly Voted 1 year, 2 months ago

Selected Answer: B

Answer: B

The use cased mentioned “and specify the key vault URL and key vault SECRET”.

NOT secret name.

Key vault secret is different from secret name

Reference: https://learn.microsoft.com/en-us/fabric/data-engineering/notebook-utilities#get-secret

notebookutils.credentials.getSecret('https://<name>.vault.azure.net/', 'secret name')

upvoted 11 times

wellingtonluis 1 year, 2 months ago

Selected Answer: A

A - No doubts https://learn.microsoft.com/en-us/fabric/data-engineering/notebook-utilities#get-secret

upvoted 1 times

wellingtonluis 1 year, 2 months ago

A - No doubts https://learn.microsoft.com/en-us/fabric/data-engineering/notebook-utilities#get-secretB. By your link, the secret NAME is required, not secret itself

upvoted 1 times

92ef42d 1 year, 2 months ago

B.

By your link, the secret NAME is required, not secret itself

upvoted 2 times

bcbc59b 1 year, 4 months ago

Selected Answer: A

Leaving aside that we do not know with whom the secret is shared in the Key Vault and that this approach is not best practice, with matching Key Vault setup it would work.

upvoted 2 times

smanzana 1 year, 5 months ago

Selected Answer: A

A- Yes

upvoted 1 times

DarioReymago 1 year, 5 months ago

Selected Answer: A

https://learn.microsoft.com/en-us/fabric/data-engineering/notebook-utilities#get-secret-using-user-credentials

upvoted 2 times

lapiscineeviden 1 year, 6 months ago

Selected Answer: A

Get secret using user credentials:

notebookutils.credentials.getSecret('https://<name>.vault.azure.net/', 'secret name')

upvoted 1 times

Maryam3018 1 year, 7 months ago

Selected Answer: B

You do not specify the Key Vault URL. You do specify the Linked Service Name that represents the Key Vault (as configured in Microsoft Fabric). Then you specify the Secret Name within that Key Vault.

upvoted 3 times

Maryam3018 1 year, 7 months ago

You do not specify the Key Vault URL. You do specify the Linked Service Name that represents the Key Vault (as configured in Microsoft Fabric). Then you specify the Secret Name within that Key Vault.You can use Key Vault references which is still in preview and then it would work, but otherwise you specify a link to a key vault as shown there https://learn.microsoft.com/en-us/fabric/data-engineering/notebook-utilities#get-secret

upvoted 3 times

fdsdfgxcvbdsfhshfg 1 year, 2 months ago

You can use Key Vault references which is still in preview and then it would work, but otherwise you specify a link to a key vault as shown there

https://learn.microsoft.com/en-us/fabric/data-engineering/notebook-utilities#get-secret

upvoted 1 times

zxc01 Highly Voted 1 year, 7 months ago

Selected Answer: A

Answer should be Yes. We normally use Key Vault name and linked service, but we still can use Key Vault URL.

notebookutils.credentials.getSecret('https://<name>.vault.azure.net/', 'secret name')

ref link: https://learn.microsoft.com/en-us/fabric/data-engineering/notebook-utilities

upvoted 7 times

Question #40 *Topic 1*

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You have an Azure key vault named KeyVault1 that contains secrets.

You have a Fabric workspace named Workspace1. Workspace contains a notebook named Notebook1 that performs the following tasks:

• Loads stage data to the target tables in a lakehouse
• Triggers the refresh of a semantic model

You plan to add functionality to Notebook1 that will use the Fabric API to monitor the semantic model refreshes.

You need to retrieve the registered application ID and secret from KeyVault1 to generate the authentication token.

Solution: You use the following code segment:

Use notebookutils.credentials.putSecret and specify the key vault URL and key vault secret.

Does this meet the goal?

A. Yes

B. No

Suggested Answer: *B*

Community vote distribution

B (100%)

92ef42d Most Recent 1 year, 2 months ago

Selected Answer: B

B

putSecret writes the secret, while you need get it

upvoted 3 times

smanzana 1 year, 5 months ago

Selected Answer: B

B - No

upvoted 1 times

5e89616 Highly Voted 1 year, 7 months ago

Selected Answer: B

You need to retrieve the registered application ID and secret from KeyVault1 to generate the authentication token.

The function notebookutils.credentials.putSecret is used to store a secret into a secret scope — not retrieve it.

upvoted 12 times

Question #41 *Topic 1*

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You have an Azure key vault named KeyVault1 that contains secrets.

You have a Fabric workspace named Workspace1. Workspace contains a notebook named Notebook1 that performs the following tasks:

• Loads stage data to the target tables in a lakehouse
• Triggers the refresh of a semantic model

You plan to add functionality to Notebook1 that will use the Fabric API to monitor the semantic model refreshes.

You need to retrieve the registered application ID and secret from KeyVault1 to generate the authentication token.

Solution: You use the following code segment:

Use notebookutils.credentials.getSecret and specify the key vault URL and the name of a linked service.

Does this meet the goal?

A. Yes

B. No

Suggested Answer: *B*

Community vote distribution

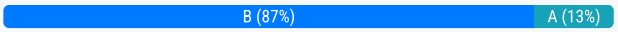

92ef42d Most Recent 1 year, 2 months ago

Selected Answer: B

B

secret name is needed instead of linked service

upvoted 4 times

aaef3f4 1 year, 3 months ago

Selected Answer: B

Correct answer is B

https://learn.microsoft.com/en-us/azure/synapse-analytics/spark/microsoft-spark-utilities?pivots=programming-language-python is mssparkutils

https://learn.microsoft.com/en-us/fabric/data-engineering/notebook-utilities#get-secret-using-user-credentials is NotebookUtils (former MSSparkUtils) for Fabric

Get secret

getSecret returns an Azure Key Vault secret for a given Azure Key Vault endpoint and secret name using user credentials.

notebookutils.credentials.getSecret('https://<name>.vault.azure.net/', 'secret name')

upvoted 4 times

smanzana 1 year, 5 months ago

Selected Answer: B

B - No (You can use either the Key Vault URL or the Linked Service name, but not both at the same time)

upvoted 2 times

minami1998 1 year, 6 months ago

Selected Answer: B

It's B without the name of the secret

upvoted 3 times

LasAnsias Highly Voted 1 year, 7 months ago

Selected Answer: B

No, this solution does not meet the goal. The notebookutils.credentials.getSecret method requires the key vault URL and the name of the secret, not the name of a linked service.

upvoted 6 times

12a2ecc 1 year, 7 months ago

Selected Answer: B

I don't believe the proposed solution works as all forms of getSecret() requires 'secret name', even with 'linked service name' specified:

mssparkutils.credentials.getSecret('azure key vault name','secret name','linked service name')

https://learn.microsoft.com/en-us/azure/synapse-analytics/spark/microsoft-spark-utilities?pivots=programming-language-python

upvoted 2 times

5e89616 1 year, 7 months ago

Selected Answer: A

In my understating (A) is the correct answer:

notebookutils.credentials.getSecret(" ", " ")

is used to retrieve secrets (such as client IDs and secrets) from Azure Key Vault using a linked service.

upvoted 3 times

5e89616 1 year, 7 months ago

Selected Answer: A

In my understating (A) is the correct answer:

notebookutils.credentials.getSecret("<KeyVault URL>", "<Linked Service Name>")

is used to retrieve secrets (such as client IDs and secrets) from Azure Key Vault using a linked service.

upvoted 3 times

zxc01 1 year, 7 months ago

No, you need put secret name as second value if you use Key Vault URL method, you need put Azure Key Vault name if you want to use linked service.

notebookutils.credentials.getSecret('https://<name>.vault.azure.net/', 'secret name')

upvoted 1 times

12a2ecc 1 year, 7 months ago

I believe KY URL and Secret are always required params, while LinkedServiceName is optional

upvoted 1 times

DRAG DROP

-

You have two Fabric notebooks named Load_Salesperson and Load_Orders that read data from Parquet files in a lakehouse. Load_Salesperson writes to a Delta table named dim_salesperson. Load_Orders writes to a Delta table named fact_orders and is dependent on the successful execution of Load_Salesperson.

You need to implement a pattern to dynamically execute Load_Salesperson and Load_Orders in the appropriate order by using a notebook.

How should you complete the code? To answer, drag the appropriate values the correct targets. Each value 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.

NOTE: Each correct selection is worth one point.

Values

- activities
- broadcast
- dependencies
- execute
- notebooks
- runMultiple

Answer Area

```
DAG = {
    "  [          ]  ": [
        {
            "name": "Load_Salesperson",
            "path": "Load_Salesperson",
            "timeoutPerCellInSeconds": 300,
        },
        {
            "name": "Load_Orders",
            "path": "Load_Orders",
            "timeoutPerCellInSeconds": 600,
            "  [          ]  ": ["Load_Salesperson"]
        }
    ],
    "timeoutInSeconds": 43200
}
mssparkutils.notebook.  [          ]  (DAG)
```

Suggested Answer:

Answer Area



5e89616 Highly Voted 1 year, 7 months ago

Correct activities: to list the activities dependencies: to make sure the order and dependency of running runMultiple: there are 2 notebooks to be executed

upvoted 17 times

5e89616 Highly Voted 1 year, 7 months ago

Correct activities: to list the activities dependencies: to make sure the order and dependency of running runMultiple: there are 2 notebooks to be executedAgree. https://learn.microsoft.com/en-us/fabric/data-engineering/microsoft-spark-utilities#reference-run-multiple-notebooks-in-parallelAlthough nothing will run in parallel here because of the dependency, the code is correct.activities: □ This is not a valid key for mssparkutils.notebook.runMultiple(). If you use "activities" instead of "runMultiple", the function will not recognize the execution plan and will fail.

upvoted 17 times

92ef42d Most Recent 1 year, 2 months ago

-activities

-dependencies

-runMultiple()

Executes notebooks based on a DAG (Directed Acyclic Graph) structure.

Supports sequential and parallel execution.

Allows defining dependencies between notebooks.

DAG = {

"activities": [

{"name": "step1", "path": "Notebook1"},

{"name": "step2", "path": "Notebook2", "dependencies": ["step1"]}

],

"timeoutInSeconds": 3600,

"concurrency": 2

}

mssparkutils.notebook.runMultiple(DAG)

https://fabric.guru/using-runmultiple-to-orchastrate-notebook-execution-in-microsoft-fabric

upvoted 1 times

Rull 1 year, 5 months ago

Agree. https://learn.microsoft.com/en-us/fabric/data-engineering/microsoft-spark-utilities#reference-run-multiple-notebooks-in-parallelAlthough nothing will run in parallel here because of the dependency, the code is correct.

upvoted 4 times

Rull 1 year, 5 months ago

Agree. https://learn.microsoft.com/en-us/fabric/data-engineering/microsoft-spark-utilities#reference-run-multiple-notebooks-in-parallel

upvoted 4 times

smanzana 1 year, 5 months ago

The answer is correct:

activities

dependencies

runMultiple

upvoted 1 times

indomanish 1 year, 5 months ago

DAG = {

"runMultiple": [

{

"name": "Load_Salesperson",

"path": "Load_Salesperson",

"timeoutPerCellInSeconds": 300,

},

{

"name": "Load_Orders",

"path": "Load_Orders",

"timeoutPerCellInSeconds": 600,

"dependencies": ["Load_Salesperson"]

}

],

"timeoutInSeconds": 43200

}

mssparkutils.notebook.runMultiple(DAG)

upvoted 1 times

indomanish 1 year, 5 months ago

DAG = { "runMultiple": [{ "name": "Load_Salesperson", "path": "Load_Salesperson", "timeoutPerCellInSeconds": 300, }, { "name": "Load_Orders", "path": "Load_Orders", "timeoutPerCellInSeconds": 600, "dependencies": ["Load_Salesperson"] }], "timeoutInSeconds": 43200 }

mssparkutils.notebook.runMultiple(DAG)

upvoted 1 times

92ef42d 1 year, 2 months ago

The requirement:

execute Load_Salesperson and Load_Orders in the appropriate order

How is it related with "run-multiple-notebooks-in-parallel"?

upvoted 1 times

DarioReymago 1 year, 5 months ago

I'm not sure about Notebook or Activity. In Fabric Pipeline use Activity and type Notebook. Others ok

upvoted 1 times

Acrious 1 year, 4 months ago

No, take a look here https://learn.microsoft.com/en-us/fabric/data-engineering/microsoft-spark-utilities#reference-run-multiple-notebooks-in-parallelThe requirement: execute Load_Salesperson and Load_Orders in the appropriate order How is it related with "run-multiple-notebooks-in-parallel"?

upvoted 1 times

Acrious 1 year, 4 months ago

No, take a look here https://learn.microsoft.com/en-us/fabric/data-engineering/microsoft-spark-utilities#reference-run-multiple-notebooks-in-parallel

upvoted 1 times

Sunnyb 1 year, 7 months ago

Correct answer should be:

"notebooks" defines the list of notebook execution steps.

"dependencies" specifies that Load_Orders depends on Load_Salesperson.

mssparkutils.notebook.runMultiple() is the correct API call to execute a DAG of notebooks in Fabric.

In Azure Data Factory and Synapse Pipelines, workflows use activities ...not in Fabric

upvoted 3 times

kim32 1 year, 6 months ago

https://learn.microsoft.com/en-us/fabric/data-engineering/microsoft-spark-utilities#reference-run-multiple-notebooks-in-parallel

upvoted 3 times

kim32 1 year, 6 months ago

https://learn.microsoft.com/en-us/fabric/data-engineering/microsoft-spark-utilities#reference-run-multiple-notebooks-in-parallelThe requirement: execute Load_Salesperson and Load_Orders in the appropriate order How is it related with "run-multiple-notebooks-in-parallel"?

upvoted 3 times

bcbc59b 1 year, 4 months ago

Although nothing will run in parallel here because of the dependency, the code is correct.

upvoted 1 times

indomanish 1 year, 5 months ago

activities:

 This is not a valid key for mssparkutils.notebook.runMultiple(). If you use "activities" instead of "runMultiple", the function will not recognize the execution plan and will fail.

upvoted 1 times

Question #43 *Topic 1*

HOTSPOT

-

You have a Fabric workspace named Workspace1 that contains a warehouse named Warehouse2.

A team of data analysts has Viewer role access to Workspace1.

You create a table by running the following statement.

You need to ensure that the team can view only the first two characters and the last four characters of the CreditCard attribute.

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

NOTE: Each correct selection is worth one point.

```
CREATE TABLE [warehouse2].[dbo].[CreditCard]
(
  CreditCard varchar(20) NOT NULL
  ,CreditCardType varchar(10) NOT NULL)
GO
```

Answer Area

[ ] TABLE dbo.CreditCard

- ALTER
- CREATE
- DEFAULT
- DROP
- EMAIL
- PARTIAL
- REPLACE
- UPDATE

[ ] COLUMN [CreditCard]

- ALTER
- CREATE
- DEFAULT
- DROP
- EMAIL
- PARTIAL
- REPLACE
- UPDATE

ADD MASKED

WITH (FUNCTION = ' [ ] (2,"XXXXXXXXXX",4)')

- ALTER
- CREATE
- DEFAULT
- DROP
- EMAIL
- PARTIAL
- REPLACE
- UPDATE