



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



CERTIFICATION TEST

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

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.

Correct Answer: C

Community vote distribution

C (100%)

 **QAZdbarhate12345678** Highly Voted 6 months, 2 weeks 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 10 times

 **oluonimole** 4 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

 **lovina_mike** Most Recent 5 days, 10 hours ago

Selected Answer: C

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 itexamsp.com 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

🗳️ 👤 **Manish0427** 1 week, 2 days 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 week, 3 days 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 week, 4 days 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** 2 weeks 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** 2 weeks, 5 days ago

Selected Answer: C

we share the gold layer lakehouse

upvoted 1 times

🗳️ 👤 **rickiebn_binss** 4 weeks ago

Selected Answer: C

Thank you

upvoted 1 times

🗳️ 👤 **ne4** 1 month, 3 weeks 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** 4 months, 2 weeks 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** 4 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** 6 months ago

Selected Answer: C

Answer is C

upvoted 2 times

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

Correct Answer: A

Community vote distribution

A (100%)

  **38578c4** Highly Voted 4 months, 3 weeks ago

Selected Answer: B



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

upvoted 22 times

  **Jayjay5** 2 weeks, 5 days ago

A is still correct because it says "Read"

upvoted 1 times

  **shmmmini** 2 months, 3 weeks 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 Lakehouse, 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 4 times

  **fffssggg** Highly Voted 5 months, 1 week 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 8 times

  **VenkataPhaniPavanKumar** Most Recent 1 week, 1 day 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 1 times

🗳️ 👤 **Manish0427** 1 week, 2 days 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 1 times

🗳️ 👤 **ogbenisho** 1 week, 3 days 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 1 times

🗳️ 👤 **TayaC** 1 week, 4 days ago

Selected Answer: A

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

upvoted 1 times

🗳️ 👤 **Wipag** 1 week, 4 days 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 1 times

🗳️ 👤 **smanzana** 1 week, 5 days ago

Selected Answer: A

The correct is A

upvoted 1 times

🗳️ 👤 **vikramkumar** 2 weeks ago

Selected Answer: A

A is the correct answer

upvoted 1 times

🗳️ 👤 **NNPRN** 2 weeks, 5 days ago

Selected Answer: B

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

upvoted 1 times

🗳️ 👤 **DarioReymago** 2 weeks, 5 days ago

Selected Answer: A

I select A. Request don't talk about stream data

upvoted 1 times

🗳️ 👤 **malik777** 3 weeks ago

Selected Answer: A

event house meant for time series data

upvoted 1 times

🗳️ 👤 **MohanNaidu08** 3 weeks, 2 days 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 2 times

  **Ahmadpbi** 3 weeks, 3 days 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** 3 weeks, 3 days 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** 3 weeks, 5 days ago

Selected Answer: B

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

upvoted 1 times

  **Madhu2023** 3 weeks, 5 days 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

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

Correct Answer: B

Community vote distribution

B (100%)

 **robertlavigne** Highly Voted 5 months, 1 week 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 10 times

 **ogbenisho** Most Recent 1 week, 3 days 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 week, 4 days 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** 2 weeks, 5 days ago

Selected Answer: B

is correct

upvoted 1 times

 **Danielnimi** 3 weeks, 5 days ago

Selected Answer: B

The pipeline in Fabric works same way as pipeline n Data Factory. 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** 2 months, 2 weeks 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

 **CMDev** 3 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 efficiency and minimal effort.

Check this link:

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

upvoted 1 times

🗨️ 👤 **arunmewada** 3 months, 2 weeks 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

🗨️ 👤 **arunmewada** 3 months, 2 weeks ago

answer is A

upvoted 2 times

🗨️ 👤 **mmanrik** 5 months, 2 weeks 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** 6 months, 2 weeks 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** 6 months, 2 weeks ago

Selected Answer: B

Dataflow Gen 1 is deprecated.

upvoted 2 times

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

Correct Answer: B

Community vote distribution

B (100%)

 **robertlavigne** Highly Voted 5 months, 1 week 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 7 times

 **Wipag** Most Recent 1 week, 4 days 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** 2 weeks, 5 days ago

Selected Answer: B

otra vez?

upvoted 1 times

 **sesav4742** 2 months, 2 weeks 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** 2 months, 2 weeks ago

Selected Answer: B

Correct Answer is B

upvoted 1 times

 **MSka** 2 months, 3 weeks ago

Selected Answer: B

It's B

upvoted 2 times

 **QAZdbarhate12345678** 6 months, 2 weeks 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** 6 months, 3 weeks ago

Selected Answer: B

Dataflow Gen 1 is deprecated.
upvoted 1 times

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.

Correct Answer: D

Community vote distribution

C (50%)

D (50%)

 **Tuki93**  6 months, 3 weeks 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 19 times

 **AzurePart** 6 days, 18 hours 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 1 times

 **Tuki93** 3 weeks, 5 days 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 5 times

 **Adriel_1996** 3 months, 1 week ago

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

upvoted 4 times

 **GHill1982** 6 months, 1 week ago

Would V-Order not already be enabled by default?

upvoted 4 times

  **h2o_molecule** 4 months, 3 weeks 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

  **dazero** 4 months ago



V-order is enabled by default.

upvoted 6 times

  **satyaamm** 3 weeks, 3 days ago



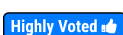
Only when you are saving in optimized format .. otherwise its not enabled by default.

upvoted 1 times

  **Raibeth** 4 months, 3 weeks ago

I still have the same doubt.



upvoted 2 times

  **dazero**  4 months ago

Selected Answer: B

500 million rows is over the F32 table limit.

upvoted 16 times

  **el_massakre** 2 months, 2 weeks 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 3 times

  **VenkataPhaniPavanKumar**  1 week, 1 day 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 week, 2 days 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 week, 3 days 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 week, 4 days ago

Selected Answer: B

Increase the capacity is correct for me.

upvoted 1 times

🗳️ 👤 **Wipag** 1 week, 4 days 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 week, 5 days 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 week, 5 days ago

Selected Answer: C

The correct is C

upvoted 1 times

🗳️ 👤 **Shw7** 1 week, 6 days ago

Selected Answer: D

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

upvoted 2 times

🗳️ 👤 **pramod4lk** 2 weeks 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** 2 weeks, 1 day 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** 2 weeks, 3 days 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 activated 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** 2 weeks, 4 days 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** 3 weeks ago

Selected Answer: C

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



upvoted 1 times

  **Madhu2023** 3 weeks, 4 days ago

Selected Answer: B

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

upvoted 1 times

  **Tuki93** 3 weeks, 5 days ago

Selected Answer: B

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 1 times

HOTSPOT -

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

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

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.

Answer Area

Correct Answer:

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 20 times


  **rnd_27** 3 months, 2 weeks 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 1 times

  **tpositive** 4 months, 1 week ago

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

upvoted 1 times

  **mixonfreddy** Highly Voted 6 months, 2 weeks ago

Selected Answer: A

SELECT YEAR

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

upvoted 20 times

  **Madhu2023** 4 months, 1 week 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



  **rnd_27** 3 months, 2 weeks 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

  **shmmini** 3 months, 1 week ago

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

upvoted 3 times

  **Manish0427** Most Recent 1 week, 2 days 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 1 times

  **Wipag** 1 week, 3 days 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 1 times

  **407a475** 2 weeks, 2 days ago

SELECT YEAR

GROUPING SETS(...)

upvoted 2 times

🗨️ 👤 **renan_moraes37** 3 weeks, 5 days ago

First answer: YEAR

Second answer: ROLLUP...

upvoted 1 times

🗨️ 👤 **Runu** 1 month, 1 week 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 1 times

🗨️ 👤 **Runu** 1 month, 1 week ago

Selected Answer: C

upvoted 1 times

🗨️ 👤 **yourock43mfamily** 2 months, 3 weeks ago

ROLLUP is correct

upvoted 1 times

🗨️ 👤 **hebertorosillo** 4 months ago

SELECT YEAR

GROUPING SETS(...)

upvoted 9 times

🗨️ 👤 **clux** 4 months, 2 weeks ago

1. Year

2. Grouping sets

upvoted 7 times

🗨️ 👤 **sakis213** 5 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** 4 months, 3 weeks ago

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

upvoted 1 times

🗨️ 👤 **realexamguru** 4 months, 3 weeks 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 3 times

🗨️ 👤 **tomaszstaroszczyk** 5 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

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.

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

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.

- A. Date
- B. ProductName
- C. ProductColor
- D. TransactionID
- E. SalesAmount
- F. ProductID

Correct Answer: BCF

Community vote distribution

BCF (100%)

QAZdbarhate12345678 Highly Voted 6 months, 2 weeks ago

Selected Answer: BCF

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 8 times

Wipag Most Recent 1 week, 3 days ago

Selected Answer: BCF

✓ 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 1 times

smanzana 1 week, 5 days ago

Selected Answer: BCF

Correct BCF

upvoted 1 times

DarioReymago 2 weeks, 5 days ago

Selected Answer: BCF

ProductID

ProductName

ProductColor

upvoted 1 times

  **renan_moraes37** 4 weeks ago

Selected Answer: BCE

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

upvoted 1 times

  **oose** 1 month ago

Selected Answer: BCF

The transaction date and the sum should be the other table

upvoted 1 times

  **Kiracloud** 1 month, 3 weeks ago

Selected Answer: BCF



- 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** 3 months, 1 week 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 1 times

  **yue_young** 1 month ago

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

upvoted 1 times

  **shmmi** 3 months, 1 week 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

  **Madhu2023** 4 months, 1 week ago

Selected Answer: BCF

BCF are the attributes of Product. TranID, Date and SalesAmt qualify for fact table instead

upvoted 2 times

  **robertlavigne** 5 months, 1 week ago

Selected Answer: BCF

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** 6 months, 3 weeks ago

Selected Answer: BCF

I believe the provided answer is correct

upvoted 4 times

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.

Correct Answer: A

Community vote distribution

A (100%)

🗳️ 👤 **Tuki93** Highly Voted 6 months, 3 weeks 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 15 times

🗳️ 👤 **QAZdbarhate12345678** Highly Voted 6 months, 2 weeks 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

🗳️ 👤 **Wipag** Most Recent 1 week, 3 days 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** 2 weeks, 5 days ago

Selected Answer: A

Enabling high concurrency

upvoted 1 times

🗳️ 👤 **renan_moraes37** 3 weeks, 5 days ago

Selected Answer: A

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

upvoted 1 times

🗳️ 👤 **renan_moraes37** 4 weeks ago

Selected Answer: A

High concurrency allows multiple notebooks in same Spark session

upvoted 1 times

🗳️ 👤 **turlakDE** 1 month ago

Selected Answer: A

Enabling high concurrency

upvoted 1 times

🗳️ 👤 **oose** 1 month ago

Selected Answer: A

High concurrency Will amake the work
upvoted 1 times

🗨️ 👤 **AbhirupAz900** 2 months, 2 weeks ago

Selected Answer: A

A is correct because without high concurrency multiple notebook can not use same spark session
upvoted 2 times

🗨️ 👤 **sesav4742** 2 months, 2 weeks 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

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.

Correct Answer: DEF

Community vote distribution

DEF (100%)

 **fassil** Highly Voted 4 months ago

Selected Answer: DEF

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 10 times

 **robertlavigne** Highly Voted 5 months, 1 week ago

Selected Answer: DEF

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 8 times

 **VenkataPhaniPavanKumar** Most Recent 1 week, 1 day ago

Selected Answer: DEF

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 1 times

  **Wipag** 1 week, 3 days ago

Selected Answer: DEF

✓ 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 1 times

  **smanzana** 1 week, 5 days ago

Selected Answer: DEF

Correct DEF

upvoted 1 times

  **DarioReymago** 2 weeks, 5 days ago

Selected Answer: DEF

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** 2 weeks, 5 days ago

Selected Answer: DEF

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

upvoted 1 times

  **MohanNaidu08** 3 weeks, 2 days ago

Selected Answer: ADE

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 1 times

  **alirana3** 2 months, 1 week ago

Selected Answer: DEF

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 3 times

  **AbhirupAz900** 2 months, 2 weeks ago

Selected Answer: ADE

Since only 3 answers needed to be selected ADE



upvoted 2 times

  **Etensel** 3 months, 1 week ago

Selected Answer: ADE

belive me..

upvoted 1 times

  **ChenFu** 4 months, 1 week ago

Selected Answer: CEF

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** 4 months, 1 week ago

Selected Answer: DEF

The steps are F, E, D

upvoted 3 times

  **MultiCloudIronMan** 4 months, 2 weeks ago

Selected Answer: DEF

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



  **i_have_a_name** 5 months, 2 weeks ago

Selected Answer: AEF

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** 4 months, 2 weeks 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** 4 months, 3 weeks 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** 6 months, 3 weeks ago

Selected Answer: DEF

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

upvoted 3 times

  **IshtarSQL** 6 months, 3 weeks ago

Selected Answer: DEF

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 3 times

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

Correct Answer: D

Community vote distribution

D (100%)

  **QAZdbarhate12345678** Highly Voted 6 months, 2 weeks 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 9 times

  **Wipag** Most Recent 1 week, 3 days 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 1 times

  **smanzana** 1 week, 5 days ago

Selected Answer: D

Función or Security Policy


upvoted 1 times

  **Evincible** 3 weeks, 4 days ago

Selected Answer: D

In the exam, instead of function, they are given "Security Policy" option

upvoted 4 times

  **Adriel_1996** 3 months, 1 week 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 4 times

  **prabhjot** 4 months, 3 weeks 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

  **Tuki93** 6 months, 3 weeks ago

Selected Answer: D

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

upvoted 3 times

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.	<input type="radio"/>	<input type="radio"/>
The Dataflow Gen1 dataflows will be deployed to the target stage.	<input type="radio"/>	<input type="radio"/>
The scheduled refresh policies will be deployed to the target stage.	<input type="radio"/>	<input type="radio"/>

	Statements	Yes	No
Correct Answer:	Data from the semantic models will be deployed to the target stage.	<input type="radio"/>	<input checked="" type="radio"/>
	The Dataflow Gen1 dataflows will be deployed to the target stage.	<input checked="" type="radio"/>	<input type="radio"/>
	The scheduled refresh policies will be deployed to the target stage.	<input type="radio"/>	<input checked="" type="radio"/>

SamuComqi Highly Voted 5 months, 4 weeks ago

NO: only metadata will be deployed to the target stage (<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 30 times

ChenFu 4 months ago

No

No

Yes

upvoted 6 times

Juanes55 Highly Voted 2 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 5 times

  **Sr18** 1 month, 1 week 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 1 times

  **VenkataPhaniPavanKumar** Most Recent 1 week, 1 day 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 2 times

  **juaidkiani** 1 week, 1 day ago

Statements Yes / No Explanation

Data from the semantic models will be deployed to the target stage. ✗ No Only the model metadata (e.g., schema, measures, relationships) is deployed – not the actual data. You must refresh the model in the target workspace to load data.

The Dataflow Gen1 dataflows will be deployed to the target stage. ✗ No Dataflow Gen1 is not supported in deployment pipelines in Fabric. Only Dataflow Gen2 is supported.

The scheduled refresh policies will be deployed to the target stage. ✗ No Refresh schedules are not deployed. You need to configure scheduling separately in each stage (DEV, TEST, PROD).

upvoted 1 times

  **Wipag** 1 week, 3 days 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 1 times

🗨️ 👤 **CertPeople** 1 week, 6 days 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 1 times

🗨️ 👤 **NNPRN** 2 weeks, 5 days ago

No, Yes, No

upvoted 1 times

🗨️ 👤 **5dcd3a8** 1 month 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 4 times

🗨️ 👤 **Dwarakkrishna** 1 month 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 1 times

  **saikot** 2 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** 4 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 4 times

🗨️ 👤 **mixonfreddy** 6 months, 2 weeks ago

Selected Answer: A

Yes/Yes/No

upvoted 1 times

🗨️ 👤 **GHill1982** 6 months ago

No/Yes/No

upvoted 4 times

🗨️ 👤 **Marco44** 5 months, 3 weeks 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 8 times

🗨️ 👤 **GHill1982** 5 months, 3 weeks 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

🗨️ 👤 **SannaLH** 3 months, 4 weeks 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** 3 months, 4 weeks ago

No sorry, Semantic models are supported as long as they're not PUSH. <https://learn.microsoft.com/en-us/fabric/cicd/deployment-pipelines/intro-to-deployment-pipelines?tabs=new-ui#supported-items>

upvoted 2 times

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

Correct Answer: B

Community vote distribution

B (100%)

🗳️ 👤 **QAZdbarhate12345678** Highly Voted 🏆 6 months, 2 weeks 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 9 times

🗳️ 👤 **Wipag** Most Recent 🕒 1 week, 3 days 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 1 times

🗳️ 👤 **smanzana** 1 week, 5 days ago

Selected Answer: B

B- deployment pipeline or event stream

upvoted 1 times

🗳️ 👤 **DarioReymago** 2 weeks, 5 days ago

Selected Answer: B

B las otras opciones son repositorios externos no generan el movimiento

upvoted 1 times

🗳️ 👤 **Evincible** 3 weeks, 4 days ago

Selected Answer: B

In the exam, they replace a deployment pipeline option with "an event stream".

upvoted 4 times

🗳️ 👤 **alirana3** 2 months, 1 week 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 1 times

  **hebertorosillo** 4 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

  **Tuki93** 6 months, 3 weeks 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

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

Correct Answer: B

Community vote distribution

B (100%)

  **QAZdbarhate12345678** Highly Voted 6 months, 2 weeks 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 17 times

  **tomaszstaroszczyk** 5 months ago

What kind of references may a Warehouse have? We don't have shortcuts in a warehouse. It seems to be an independent item, there is no possibility to break a reference to anything external. Therefore 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

  **toniteo** 4 months, 1 week 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

  **CloudKaKida** 2 months, 1 week ago

A

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

upvoted 1 times

  **Sr18** 1 month, 1 week 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 1 times

  **Sr18** 1 month, 1 week 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

  **smanzana** Most Recent 1 week, 5 days ago

Selected Answer: B

B is correct

upvoted 1 times

  **upliftinghut** 1 month, 3 weeks 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** 4 months ago

Selected Answer: A

A. a database project for minimal effort



upvoted 4 times

  **Anushareddysri** 4 months, 4 weeks 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** 4 months, 1 week ago

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

upvoted 1 times

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

Correct Answer: E

Community vote distribution

E (86%)

14%

 **Tuki93** Highly Voted 6 months, 3 weeks 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 17 times

 **Sr18** 1 month, 1 week 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 3 times


 **exnaniantwort** 2 weeks, 5 days 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 1 times

 **toniteo** 4 months, 1 week ago

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


upvoted 1 times

 **IshtarSQL** Highly Voted 6 months, 3 weeks 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 8 times

 **Thanos29** Most Recent 4 days, 6 hours ago

Selected Answer: E

E is Correct.

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

If OneLake Availability 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 1 times

🗨️ 👤 **VenkataPhaniPavanKumar** 1 week, 1 day 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 1 times

🗨️ 👤 **mvg100** 2 weeks, 3 days 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** 2 weeks, 5 days 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 1 times

🗨️ 👤 **AnetaK** 2 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** 3 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 1 times

🗨️ 👤 **CMDev** 3 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** 3 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** 3 months, 1 week 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** 4 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 2 times

  **DarkDerf** 4 months, 1 week 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** 4 months, 1 week ago

Selected Answer: E

Only new data will be loaded to OneLake

upvoted 3 times

  **Anushreddysri** 4 months, 2 weeks 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** 4 months, 1 week ago

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

upvoted 1 times

🗨️ 👤 **Sr18** 1 month, 1 week 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** 4 months, 3 weeks 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** 4 months, 4 weeks ago

Selected Answer: E

I think E is correct:

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

upvoted 2 times

🗨️ 👤 **QAZdbarhate12345678** 6 months, 2 weeks 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

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

Correct Answer: A

Community vote distribution

A (67%)

D (33%)

🗳️ 👤 **VenkataPhaniPavanKumar** 1 week, 1 day 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 1 times

🗳️ 👤 **smanzana** 1 week, 5 days ago

Selected Answer: A

A is correct

upvoted 1 times

🗳️ 👤 **a85becd** 2 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

🗳️ 👤 **hebertorosillo** 4 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

🗳️ 👤 **henryphchan** 4 months, 1 week 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

🗳️ 👤 **shmmmini** 3 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 month ago

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



upvoted 1 times

  **prabhjot** 4 months, 3 weeks 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** 4 months, 3 weeks 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

  **robertlavigne** 4 months, 3 weeks 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** 5 months, 1 week 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

  **Meir** 6 months, 1 week 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

  **GHill1982** 6 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** 6 months, 2 weeks ago

Selected Answer: D

Use of service principal is discouraged

upvoted 2 times

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

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

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.

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


You need to read data from all the shortcuts.

Which shortcuts will retrieve data from the cache?

- A. Stores only
- B. Products only
- C. Stores and Products only
- D. Products, Stores, and Trips
- E. Trips only
- F. Products and Trips only

Correct Answer: C

Community vote distribution



Tuki93 Highly Voted 6 months, 3 weeks 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 16 times

Sr18 1 month, 1 week ago

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

Sr18 1 month, 1 week 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 1 times

Sr18 1 month, 1 week ago

I mean 1gb, stupid typo
upvoted 1 times

Kiket2ride Highly Voted 3 months, 1 week ago

Selected Answer: C

By default is C but now D can also be true. See here:

The retention period for cached files can be set from 1-28 days

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

upvoted 6 times

Adriel_1996 3 months, 1 week ago

I am dealing with same issue. Could be C or D depending on your configuration.

upvoted 2 times

🗄️ 👤 **smanzana** Most Recent 1 week, 5 days ago

Selected Answer: D

D: Cache duration can now be configured between 1 and 28 days, not just the default 24 hours, but files larger than 1 GB are excluded from the cache and will not be cached.

upvoted 1 times

🗄️ 👤 **rehoboth2165** 3 weeks ago

Selected Answer: D

Caching is for 1-28 days on last accessed and the individual files are less than 1GB each

upvoted 2 times

🗄️ 👤 **f4 added67** 3 weeks, 6 days ago

Selected Answer: C

Here the default retention period is 24 hours, but can be configured to 1-28 days. But, they didn't mention about any retention period. Hence, default retention period is applied and option C is the correct

upvoted 2 times

🗄️ 👤 **DarioReymago** 1 month ago

Selected Answer: D

D: he retention period for cached files can be set from 1-28 days Files are less than 1gb all sources type are supported

upvoted 2 times

🗄️ 👤 **vernillen** 1 month ago

Selected Answer: D

Due to recent changes

upvoted 2 times

🗄️ 👤 **5 added3a8** 1 month, 1 week ago

Selected Answer: D

With the new update cached files can be set from 1-28. Also the Individual files greater than 1 GB in size aren't cached.

<https://learn.microsoft.com/en-us/fabric/onelake/onelake-shortcuts#cached>

upvoted 3 times

🗄️ 👤 **upliftinghut** 1 month, 3 weeks ago

Selected Answer: D

less than 1 GB & retention up to 18 days => all are ok

upvoted 2 times

🗄️ 👤 **ts_au** 1 month, 3 weeks ago

Selected Answer: D

caching can now be retained up to 28 days

upvoted 5 times

🗄️ 👤 **henryphchan** 4 months, 1 week ago

Selected Answer: C

Data within 1GB and accessed less than 24 hours will be cached

upvoted 3 times

🗄️ 👤 **Bharat** 6 months, 2 weeks ago

Selected Answer: D

I think that the correct answer is D instead of C because the caching is enabled for ALL types of files in a Lakehouse whether they are Parquet, JSON, or CSV.

upvoted 4 times

🗄️ 👤 **Bharat** 6 months, 2 weeks ago

Sorry, I didn't pay attention to the retention period. Because it exceeds 24 hours, it will not be read from the cache. Therefore, the answer is correct.

upvoted 2 times

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

Correct Answer: B

  **GHill1982** Highly Voted 6 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 5 times

  **realexamguru** 5 months, 1 week ago

correct

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

upvoted 4 times

  **smanzana** Most Recent 1 week, 5 days ago

Selected Answer: B

Correct

upvoted 1 times

  **2d073ed** 2 weeks, 1 day 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 1 times

  **DarioReymago** 2 weeks, 4 days ago

Selected Answer: B

B is correct

upvoted 1 times

  **Tuki93** 5 months, 1 week ago

Selected Answer: B

Provided answer is correct

upvoted 3 times

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.

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

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?

- A. Trips and Stores only
- B. Products and Store only
- C. Stores only
- D. Products only
- E. Products, Stores, and Trips

Correct Answer: B

Community vote distribution

B (100%)

 **Tuki93**  5 months, 1 week 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 16 times

 **DarioReymago** 1 month ago

for this reason is B

upvoted 2 times

 **exnaniantwort** 2 weeks, 5 days 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 2 times

 **AnetaK** 2 months ago

The retention period for cached files can be set from 1-28 days.

(same link as above)

upvoted 3 times

 **prabhjot**  4 months, 3 weeks 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 7 times

 **exnaniantwort** 2 weeks, 5 days 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 1 times

🗨️ **smanzana** Most Recent 1 week, 5 days ago

Selected Answer: C

C is correct.

-Shortcut caching is currently supported for Google Cloud Storage (GCS), Amazon S3, and S3-compatible storage.

-Azure Data Lake Storage Gen2 (ADLS Gen2) does not require or use caching since it integrates natively with OneLake, so it is not cached.

-Files larger than 1 GB are not cached.

upvoted 1 times

🗨️ **gretito_g** 2 weeks, 3 days ago

Selected Answer: C

ProductFile is in storage1 (Azure) → not cached.

TripsFile is in storage2 (Amazon S3), but it's 2 GB, which is large → not cached.

StoreFile is in storage2 (Amazon S3) and only 25 MB → likely to be cached.

upvoted 1 times

🗨️ **exnaniantwort** 2 weeks, 5 days ago

Selected Answer: C

The question is asking "The data from which shortcuts will be retrieved FROM THE CACHE?".

ADLS does NOT need a cache. -> Products is incorrect

Tripfile exceeds 1GB -> Trips is incorrect

Stores is from S3 and supported by shortcut caching. -> Stores is correct

upvoted 1 times

🗨️ **DarioReymago** 1 month ago

Selected Answer: B

Individual files greater than 1 GB in size aren't cached

upvoted 3 times

🗨️ **5dcd3a8** 1 month, 1 week ago

Selected Answer: B

Most Probably the correct answer is B

Shortcut caching is currently supported for GCS, S3, S3 compatible, and on-premises data gateway shortcuts. Individual files greater than 1 GB in size aren't cached.

<https://learn.microsoft.com/en-us/fabric/onelake/onelake-shortcuts#cached>

upvoted 4 times

🗨️ **Ryan2025** 2 months ago

Selected Answer: D

The key detail in this question is that Fabric can cache shortcut data only from Azure Data Lake Storage Gen2, not from Amazon S3.

- Only files stored in Azure Data Lake Gen2 (i.e., storage1) can be cached.

- That means only ProductFile (Products) is eligible for shortcut caching.

upvoted 1 times

🗨️ **Biju1** 2 months, 2 weeks ago

Selected Answer: B

Remember . in documentation, still you cant see ADLS. but ADLS and Dataverse support shortcut now in Fabric. Access one lake house and try. I verified

upvoted 3 times

🗨️ **exnaniantwort** 2 weeks, 5 days ago

ADLS support 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 1 times

🗨️ **Tamele001** 3 months ago

Selected Answer: B

Products and store because they have less than a GB

upvoted 4 times

🗨️ **Adriel_1996** 3 months, 1 week ago

Selected Answer: C

ProductFile is part of ADLS GEN2 which according documentation "Shortcut caching is currently only supported for GCS, S3 and S3 compatible shortcuts." hence is not caching.

TripsFile exceeds 1 GB size

Only StoreFile support caching and has the right size

Based on: <https://learn.microsoft.com/en-us/fabric/onelake/onelake-shortcuts>



Caching and S3 shortcuts

upvoted 4 times

  **moons_123** 3 weeks ago

I saw this link: <https://learn.microsoft.com/en-us/fabric/onelake/create-adls-shortcut>

upvoted 1 times

  **cuomuqi** 3 months, 3 weeks ago

Selected Answer: B

s3 can be saved <https://learn.microsoft.com/en-us/fabric/onelake/create-s3-shortcut>

and items bigger than 1GB can not be cached. "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>

upvoted 2 times

  **vish9** 3 months, 3 weeks ago

Selected Answer: B

In Microsoft Fabric, when the cache for shortcuts is enabled, the caching mechanism applies based on the file size and source location.

ProductFile (50MB) from storage1 (Azure Data Lake Storage Gen2)

Cached ✓ (because it's from Azure Data Lake and is small)

TripsFile (2GB) from storage2 (Amazon S3)

Not cached ✗ (because it's larger than 1GB)

StoreFile (25MB) from storage2 (Amazon S3)

Cached ✓ (because it's smaller than 1GB)

Cache Behavior:

Files ≤ 1GB from Azure Data Lake Storage (ADLS) or Amazon S3 are cached.

Files > 1GB are not cached.

Answer:

The data retrieved from the cache will be Products (ProductFile) and Stores (StoreFile).

upvoted 3 times

  **JensQ** 3 months, 3 weeks ago

Selected Answer: B

S3 and ADLS G2 are both supported.

Only limitation is the size, meaning Trips is too big to be cached.

upvoted 3 times

  **exnaniantwort** 2 weeks, 5 days 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 1 times

  **DarkDerf** 4 months ago

Selected Answer: D

Since caching works only for Azure Data Lake Storage Gen2, the shortcuts pointing to Amazon S3 (storage2) will not be retrieved from the cache.

Evaluating the shortcuts:

Products (shortcut to ProductFile - ADLS Gen2) → ✓ Cached

Stores (shortcut to StoreFile - Amazon S3) → ✗ Not Cached

Trips (shortcut to TripsFile - Amazon S3) → ✗ Not Cached



Answer:

The cached data will be retrieved only for Products.

Thus, the correct answer is:

D. Products only

upvoted 2 times

  **ChenFu** 4 months, 1 week ago

Selected Answer: E



In Microsoft Fabric, when cache for shortcuts is enabled, the data from the shortcuts is automatically cached in OneLake for faster access. The key factors that determine cache availability are:

Shortcut source – Files in Azure Data Lake Storage Gen2 and Amazon S3 are both supported for caching.

Cache enabled – If caching is enabled for the workspace, all accessible shortcuts are cached.

Data access – Cached data is updated when the source changes or when the cache expires (typically after 30 days of inactivity).

upvoted 1 times

  **Adriel_1996** 3 months, 1 week ago

Shortcut source – Files in Azure Data Lake Storage Gen2 where is the documentation that supports ?

upvoted 2 times

  **henryphchan** 4 months, 1 week ago


Selected Answer: C

The answer is C.

Please refer to MS docs (<https://learn.microsoft.com/en-us/fabric/onelake/onelake-shortcuts>)

- Shortcut caching can be used to reduce egress costs associated with cross-cloud data access.
- Shortcut caching is currently only supported for GCS, S3 and S3 compatible shortcuts.

upvoted 5 times

  **henryphchan** 4 months, 1 week ago

- Cached files have a retention period of 24 hours.
- Individual files greater than 1 GB in size aren't cached.

upvoted 2 times

  **AnetaK** 2 months ago

The retention period for cached files can be set from 1-28 days.

<https://learn.microsoft.com/en-us/fabric/onelake/onelake-shortcuts#cached>

upvoted 1 times

HOTSPOT -

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

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

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.

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.

Correct Answer:

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.

  **ChenFu** 4 months, 1 week 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 1 times

  **4371883**  5 months, 1 week 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 5 times

  **DarioReymago**  2 weeks, 4 days 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

  **boonfidelio** 6 months ago

 **Selected Answer: A**

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

upvoted 2 times

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

Correct Answer: B

Community vote distribution

B (67%)

C (33%)

🗳️ 👤 **neo_20** Highly Voted 👍 5 months, 3 weeks ago

Selected Answer: D

Only a Fabric Admin can create Domains.
upvoted 15 times

🗳️ 👤 **smanzana** Most Recent ⌚ 1 week, 5 days ago

Selected Answer: D

D is correct
upvoted 1 times

🗳️ 👤 **DarioReymago** 2 weeks, 4 days ago

Selected Answer: D

just Fabric Adm can create domains
upvoted 1 times

🗳️ 👤 **bettermakeme** 1 month, 2 weeks 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 1 times

🗳️ 👤 **Sreeven** 3 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 1 times

🗳️ 👤 **benni_ale** 3 months, 4 weeks ago

Selected Answer: D

only fabric admin can create and edit domains
upvoted 3 times

🗳️ 👤 **Adriel_1996** 3 months, 1 week ago

But first say that have to create a new domain
upvoted 1 times

🗳️ 👤 **ChenFu** 4 months, 1 week 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 1 times

🗳️ 👤 **8d6881f** 5 months ago

Selected Answer: D

Only a Fabric Admin can create Domains.

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

upvoted 2 times

🗳️ 👤 **MDWPartners** 5 months, 2 weeks ago

Selected Answer: D

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

upvoted 4 times

🗳️ 👤 **harshalt10** 5 months, 3 weeks 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 4 times

🗳️ 👤 **Zekes** 6 months, 1 week 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** 3 months, 3 weeks 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 2 times

🗳️ 👤 **Bharat** 6 months, 2 weeks 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** 6 months, 2 weeks 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** 6 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** 5 months, 4 weeks 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

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

Correct Answer: D

Community vote distribution

D (50%)

B (50%)

 **Bovote** Highly Voted 5 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 15 times

 **bettermakeme** 1 month, 2 weeks ago

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

upvoted 2 times

 **01ceacf** Highly Voted 6 months, 2 weeks 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

 **zxc01** 2 months, 2 weeks 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 1 times

 **smanzana** Most Recent 1 week, 5 days ago

Selected Answer: D

D is correct

upvoted 1 times

 **DarioReymago** 2 weeks, 4 days 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 higher option]

Viewer + object-level is enough

upvoted 1 times

🗨️ 👤 **Jay63002** 2 weeks, 5 days 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 1 times

🗨️ 👤 **PBridge** 1 month, 1 week 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 2 times

🗨️ 👤 **Sr18** 1 month, 1 week ago

Selected Answer: D

Contributor role

upvoted 1 times

🗨️ 👤 **5dcd3a8** 1 month, 3 weeks 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 month, 3 weeks ago

Selected Answer: C

they already have access to DW1

upvoted 2 times

  **d98a23f** 3 months ago

Selected Answer: D

contributor 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** 3 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** 3 months, 1 week 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 4 times

  **GadAugust** 3 months, 1 week 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 4 times

  **zxc01** 3 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 2 times


  **Goye** 3 months, 4 weeks 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** 3 months, 4 weeks 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** 3 months, 4 weeks ago

please correct me if i am wrong otherwise is D

upvoted 1 times

  **prabhjot** 4 months, 3 weeks 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** 5 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

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.

Correct Answer: B

Community vote distribution

A (100%)

01ceacf Highly Voted 6 months, 2 weeks ago

Selected Answer: A

Read Access to Table Data.

B/D - Viewer/Member can access other items

C - cannot access SQL DB

upvoted 8 times

IshtarSQL Highly Voted 6 months, 2 weeks 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

smanzana Most Recent 1 week, 5 days ago

Selected Answer: A

A is correct

upvoted 1 times

DarioReymago 2 weeks, 4 days ago

Selected Answer: A

just read to LH

upvoted 1 times

vernillen 1 month ago

Selected Answer: A

The most logical one

upvoted 1 times

benni_ale 4 months ago

Selected Answer: A

Agree with other solutions

upvoted 3 times

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 to 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.

Badges	Answer Area
<input type="checkbox"/> Certified	Entity1: <input type="text"/>
<input type="checkbox"/> Master data	Entity2: <input type="text"/>
<input type="checkbox"/> Promoted	Entity3: <input type="text"/>
<input type="checkbox"/> Cannot be endorsed	Entity4: <input type="text"/>

Correct Answer:

Badges	Answer Area
<input checked="" type="checkbox"/> Certified	Entity1: <input type="text"/>
<input checked="" type="checkbox"/> Master data	Entity2: <input type="text"/>
<input checked="" type="checkbox"/> Promoted	Entity3: <input type="text"/>
<input checked="" type="checkbox"/> Cannot be endorsed	Entity4: <input type="text"/>

fassil Highly Voted 4 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 21 times

DarioReymago 2 weeks, 4 days ago

this is correct

upvoted 3 times

Kiket2ride Most Recent 3 months, 1 week ago

These are drag and drop so it makes no sense that two of those are of the same type

upvoted 1 times

Clode 1 month, 1 week ago

Each badge may be used once, more than once, or not at all

upvoted 1 times

HAARS 3 months, 2 weeks ago

4 IS CANNOT endorse

upvoted 2 times

🗨️ 👤 **MultiCloudIronMan** 4 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 1 times

🗨️ 👤 **henryphchan** 4 months, 1 week 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 2 times

🗨️ 👤 **8d6881f** 5 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

🗨️ 👤 **Tuki93** 5 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 3 times

🗨️ 👤 **doctordodge** 5 months, 1 week 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 4 times

🗨️ 👤 **4371883** 5 months, 1 week 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

HOTSPOT -

You have three users named User1, User2, and User3.

You have the Fabric workspaces shown in the following table.

Name	Workspace admin
Workspace1	User1
Workspace2	User2

You have a security group named Group1 that contains User1 and User3.

The Fabric admin creates the domains shown in the following table.

Name	Domain admin
Domain1	User1
Domain2	User2

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.

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.

☐
☐

Correct Answer:

Statements

Yes

No

User3 has Viewer role access to Workspace3.

☒
☐

User3 has Domain contributor access to Domain1.

☐
☒

User2 has Contributor role access to Workspace3.

☐
☒

 **giraf** Highly Voted 4 months, 2 weeks 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 18 times

 **38578c4** Highly Voted 4 months, 1 week ago

NYN

<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 10 times

 **smanzana** Most Recent 1 week, 5 days ago

NYN

Users or groups assigned to a domain's default group automatically receive the Domain Contributor role.

upvoted 1 times

 **Ricauduro** 3 weeks 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 1 times

  **bettermakeme** 1 month, 2 weeks 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 3 times

  **5dcd3a8** 1 month, 3 weeks 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 month, 3 weeks ago

N/Y/N

<https://learn.microsoft.com/en-us/fabric/governance/domains#default-domain>

upvoted 1 times

🗨️ 👤 **vish9** 3 months, 1 week 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 7 times

🗨️ 👤 **MultiCloudIronMan** 3 months, 3 weeks ago

NYN is correct

upvoted 1 times

🗨️ 👤 **henryphchan** 4 months, 1 week 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

🗨️ 👤 **8d6881f** 5 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** 5 months ago

I correct myself. The answer is Correct. Y N N.

Mind the order of things

upvoted 1 times

🗨️ 👤 **doctordodge** 5 months ago

Y Y N

<https://learn.microsoft.com/en-us/fabric/governance/domains>

upvoted 3 times

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.

Workspace	Items
Workspace1	Model1 Notebook1 Report1 Lakehouse1 Pipeline1
Workspace2	Model1 Notebook2 Report1 Lakehouse2

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?

- 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.

Correct Answer: D

 **smanzana** 1 week, 5 days ago

Selected Answer: D

D is correct

upvoted 1 times

 **DarioReymago** 2 weeks, 4 days ago

Selected Answer: D

D is correct

upvoted 1 times

 **exnaniantwort** 2 weeks, 5 days ago

Selected Answer: D

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 1 times

 **bettermakeme** 1 month, 2 weeks ago

Selected Answer: D

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

🗨️ 👤 **benni_ale** 4 months ago

Selected Answer: D

Given minimal effort is required deploying the deployment pipeline is enough

upvoted 4 times

🗨️ 👤 **Zoran_S** 4 months, 1 week ago

Selected Answer: D

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** 4 months, 1 week ago

Selected Answer: D

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

upvoted 1 times

🗨️ 👤 **giraf** 4 months, 2 weeks ago

Selected Answer: A

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 1 times

🗨️ 👤 **Tuki93** 5 months ago

Selected Answer: D

I believe the provided answer 'D' is correct.

upvoted 2 times

🗨️ 👤 **GHill1982** 6 months ago

Selected Answer: D

This requires minimal effort as it has no additional steps such as renaming or deleting items.

upvoted 4 times

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.

Correct Answer: A

🗲️ 👤 **smanzana** 1 week, 5 days ago

Selected Answer: A

A is correct

upvoted 1 times

🗲️ 👤 **ce6fb55** 3 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 1 times

🗲️ 👤 **MultiCloudIronMan** 3 months, 3 weeks ago

Selected Answer: A

Correct

upvoted 1 times

🗲️ 👤 **Goye** 3 months, 4 weeks ago

Selected Answer: A

A is correct, I tested it in fabric

upvoted 2 times

🗲️ 👤 **8d6881f** 5 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 3 times

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.

Actions

⋮ Change the runtime version.

⋮ Install the libraries.

⋮ Create a pool.

⋮ Create an environment.

⋮ Set the default environment.

Answer Area

Correct Answer:

Actions	Answer Area
⋮ Change the runtime version.	
⋮ Install the libraries.	
⋮ Create a pool.	
⋮ Create an environment.	
⋮ Set the default environment.	

🗲️ **Tuki93** Highly Voted 5 months ago

Seems correct.

upvoted 6 times

🗲️ **VenkataPhaniPavanKumar** Most Recent 1 week, 1 day ago

In Microsoft Fabric warehouses, primary keys must be defined as NONCLUSTERED and NOT ENFORCED

upvoted 1 times

🗲️ **smanzana** 1 week, 5 days ago

Correct

upvoted 1 times

🗲️ **DarioReymago** 2 weeks, 4 days ago

correct

upvoted 1 times

🗲️ **Sunnyb** 2 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 4 times

  **benni_ale** 4 months ago

Seems corret: Create (virtual) environment, install libraries, set the newly created environment as default from workspace settings

upvoted 2 times

  **henryphchan** 4 months, 1 week ago

The answer is correct

upvoted 3 times

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

Correct Answer: C

🗲️ 👤 **smanzana** 1 week, 5 days ago

Selected Answer: C

C is correct

upvoted 1 times

🗲️ 👤 **benni_ale** 4 months ago

Selected Answer: C

Optimise!

upvoted 1 times

🗲️ 👤 **henryphchan** 4 months, 1 week 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 3 times

🗲️ 👤 **prabhjot** 4 months, 3 weeks ago

Selected Answer: C

Optimize solves the "small file problem" by consolidating multiple small files into larger parquet files

upvoted 2 times

🗲️ 👤 **Tuki93** 5 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 1 times

🗲️ 👤 **doctordodge** 5 months ago

Selected Answer: C

Link: <https://learn.microsoft.com/en-us/fabric/data-engineering/lakehouse-table-maintenance>

upvoted 1 times

🗲️ 👤 **Ghill1982** 6 months ago

Selected Answer: C

To consolidate the underlying Parquet files in Table1, you should run the OPTIMIZE command.

upvoted 4 times

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

Correct Answer: G

 **GHill1982** Highly Voted 6 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 7 times

 **smanzana** Most Recent 1 week, 5 days ago

Selected Answer: G

Location

upvoted 1 times

 **akmorsh** 2 weeks, 6 days 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

 **Besmarlyn** 2 months, 3 weeks ago

Selected Answer: G

Answer is correct

upvoted 1 times

 **benni_ale** 4 months ago

Selected Answer: G

Location

upvoted 1 times

 **henryphchan** 4 months, 1 week ago

Selected Answer: G

In Monitoring Hub, Location = Workspace

upvoted 2 times

 **Tuki93** 5 months ago

Selected Answer: G

Provided answer "Location" is correct

upvoted 1 times

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

Correct Answer: C

  **GHill1982** Highly Voted  6 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 12 times

  **smanzana** Most Recent  1 week, 5 days ago

Selected Answer: D

Monitor
upvoted 1 times

  **benni_ale** 4 months ago

Selected Answer: D

Monitor
upvoted 1 times

  **Tuki93** 5 months ago

Selected Answer: D

We can find this in Monitor. example: Runtime information
Runtime 1.2 (Spark 3.4, Delta 2.4)
upvoted 2 times

DRAG DROP -

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

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

```
CREATE TABLE dbo.DimCustomer (  
    CustomerKey VARCHAR(255) NOT NULL,  
    Name VARCHAR(255) NOT NULL,  
    Email VARCHAR(255) NOT NULL  
);
```

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.

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

Correct Answer:

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

 **amli123**  5 months, 1 week 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 41 times

 **mvg100** 2 weeks ago

This is correct,

Primary Key

Non Clustered

Not Enforced

<https://learn.microsoft.com/en-us/fabric/data-warehouse/table-constraints>

upvoted 1 times

🗳️ 👤 **bettermakeme** 1 month, 2 weeks ago

This is the correct option, I have tried it in my warehouse

```
ALTER TABLE [dbo].[dim_cust]
```

```
ADD CONSTRAINT [PK_dim_cust_CustomerKey] PRIMARY KEY NONCLUSTERED (CustomerKey) NOT ENFORCED
```

upvoted 4 times

🗳️ 👤 **ChenFu** 4 months, 1 week ago

```
ALTER TABLE dbo.DimCustomer
```

```
ADD CONSTRAINT PK_DimCustomer PRIMARY KEY CLUSTERED (CustomerKey)
```

```
ENFORCED
```

upvoted 3 times

🗳️ 👤 **Othman2110** Highly Voted 👍 4 months ago

Primary Key

Non Clustered

Not Enforced

upvoted 10 times

🗳️ 👤 **smanzana** Most Recent 🕒 1 week, 5 days ago

```
ALTER TABLE
```

```
NONCLUSTERED
```

```
NOT ENFORCED
```

upvoted 1 times

🗳️ 👤 **DataSturdy** 2 weeks, 3 days ago

```
ALTER TABLE dbo.dimcustomer
```

```
ADD CONSTRAINT PK_DimCustomer PRIMARY KEY NONCLUSTERED (customerkey)
```

```
not ENFORCED
```

this is the right answer

upvoted 1 times

🗳️ 👤 **turlakDE** 2 weeks, 4 days 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 1 times

🗳️ 👤 **Rajib** 2 weeks, 4 days 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 1 times

🗳️ 👤 **PBridge** 1 month 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 month, 2 weeks 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 3 times

  **fassil** 4 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 5 times

  **Tuki93** 5 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

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



Correct Answer: D

  **GHill1982** Highly Voted 6 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 12 times

  **smanzana** Most Recent 1 week, 5 days ago

Selected Answer: B

B is correct

upvoted 1 times



  **smanzana** 1 week, 5 days 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 week, 6 days 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 1 times

  **kaushikbs** 3 weeks 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 1 times

  **Ricauduro** 3 weeks, 1 day ago

Selected Answer: B

Not clear the meaning of dynamic....checking a failed execution using the monitor hub we can rerun and follow the execution

upvoted 1 times

  **DarioReymago** 1 month 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 manually check the result the option is B


upvoted 1 times

  **Thaan** 1 month, 1 week ago

Selected Answer: B

This is correct



upvoted 1 times

  **d98a23f** 3 months ago

Selected Answer: B

i think that both of B or D allow to achive this monitoring 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** 3 months, 4 weeks 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** 4 months, 1 week 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** 4 months, 1 week 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** 4 months, 2 weeks 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** 2 months ago

In Monitor Hub you cannot execute refresh.

upvoted 2 times

  **Tuki93** 5 months ago

Selected Answer: D

I think provided answer is correct: a semantic link in a notebook

upvoted 2 times

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:

```
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
```

Does this meet the goal?

A. Yes

B. No

Correct Answer: B

 **SamuComqi** Highly Voted 6 months ago

Selected Answer: B

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 14 times

 **smanzana** Most Recent 1 week, 5 days ago

Selected Answer: B

B is correct

upvoted 1 times

 **smanzana** 1 week, 6 days ago

Selected Answer: B

B-No, in KQL by default "sort by" is desc order

upvoted 1 times

 **turlakDE** 2 weeks, 4 days ago

Selected Answer: B

sort in kql by default is descending

upvoted 1 times

🗋️ 👤 **8d1e8fe** 2 months ago

Selected Answer: B

<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 1 times

🗋️ 👤 **benni_ale** 4 months ago

Selected Answer: B

sort by in kql is by default in descending order

upvoted 1 times

🗋️ 👤 **MuffiSan** 5 months, 1 week ago

Selected Answer: B

sort is desc by default in kql

upvoted 4 times

🗋️ 👤 **GHill1982** 6 months ago

Selected Answer: B

In KQL the where keyword should be used instead of filter.

upvoted 3 times

🗋️ 👤 **GHill1982** 6 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** 6 months ago

Also consider that the default for order by/sort by in KQL 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** 6 months ago

The duplicated project, even though not required, does however work and does not throw any errors.

upvoted 1 times

🗋️ 👤 **QAZdbarhate12345678** 6 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

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:

```
bike_location
| filter Neighbourhood == "Sands End" and No_Bikes >= 15
| order by No_Bikes
| project BikepointID, Street, Neighbourhood, No_Bikes, No_Empty_Docks, Timestamp
```

Does this meet the goal?

A. Yes

B. No

Correct Answer: B

Community vote distribution


A (100%)

 **benni_ale** Highly Voted 4 months ago

Selected Answer: B

no, order by by default is desc in kql

upvoted 8 times

 **smanzana** Most Recent 1 week, 5 days ago

Selected Answer: B

B is correct

upvoted 1 times

 **smanzana** 1 week, 6 days ago

Selected Answer: B


B-No, in KQL by default "order by" is desc order

upvoted 1 times

 **akmorsh** 2 weeks, 6 days ago

Selected Answer: B

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 2 times

🗋️ 👤 **568f95c** 4 months, 1 week ago

Selected Answer: A

They are asking for the order of No_Bikes in ascending order in this question

upvoted 2 times

🗋️ 👤 **Tuki93** 5 months ago

Selected Answer: B

default is desc in Kusto

upvoted 4 times

🗋️ 👤 **mauriciorosa8** 5 months ago

Selected Answer: B

The default sorting order in KQL is ****descending (desc)****; ****ascending (asc)**** must be explicitly specified when needed.

upvoted 4 times

🗋️ 👤 **8d6881f** 5 months ago

Selected Answer: B

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 3 times

🗋️ 👤 **amli123** 5 months, 1 week ago

Selected Answer: B

default is desc

upvoted 2 times

🗋️ 👤 **MuffiSan** 5 months, 1 week 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** 5 months, 1 week ago

"The sort and order operators are equivalent"

upvoted 1 times

🗋️ 👤 **QAZdbarhate12345678** 6 months ago

Selected Answer: A

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** 6 months, 2 weeks 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

  **GHill1982** 6 months ago

The answer is no. The default for order by in KQL is 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

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:

```
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
```

Does this meet the goal?

A. Yes

B. No

Correct Answer: A

Community vote distribution

B (100%)

 **GHill1982** Highly Voted 6 months ago

Selected Answer: A

The sort and order operators are equivalent.

<https://learn.microsoft.com/en-us/kusto/query/sort-operator?view=microsoft-fabric>
upvoted 9 times

 **smanzana** Most Recent 1 week, 5 days ago

Selected Answer: A

A is correct

upvoted 1 times

 **smanzana** 1 week, 6 days ago

Selected Answer: A

A (Adding ASC after ORDER BY no_bikes explicitly enforces ascending order)

upvoted 1 times

 **a823528** 4 weeks 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** 2 weeks, 4 days ago

Don't you think answer is "No", as the query should use "where" instead of "filter"

upvoted 1 times

  **benni_ale** 4 months ago

Selected Answer: A

correct!



upvoted 1 times

  **henryphchan** 4 months, 1 week ago

Selected Answer: A

sort by and order by are synonyms


upvoted 1 times

  **Bharat** 6 months, 2 weeks 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

  **henryphchan** 4 months, 1 week ago

"sort by" is legal in KQL, it is the same as "order by"

upvoted 1 times

  **MuffiSan** 5 months, 1 week ago

no, sort by and order by both work and are equal

A is correct.

upvoted 2 times

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:

```
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
```

Does this meet the goal?

A. Yes


B. No

Correct Answer: B

Community vote distribution

A (50%)

B (50%)

 **Mourya959595** Highly Voted 6 months, 1 week ago

Selected Answer: A

SQL By default shows ASC order
upvoted 19 times

 **JensQ** Highly Voted 3 months, 3 weeks 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 8 times

 **smanzana** Most Recent 1 week, 6 days ago

Selected Answer: B

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** 2 weeks, 2 days ago

Selected Answer: A

SQL By default shows ASC order
upvoted 1 times

🗨️ **turlakDE** 2 weeks, 4 days ago

Selected Answer: A

YES, this is SQL and order by is asc by default
upvoted 2 times

🗨️ **rohitbinnani** 2 weeks, 5 days ago

Selected Answer: A

SQL is by-default in asc order while using order by
upvoted 2 times

🗨️ **DarioReymago** 4 weeks ago

Selected Answer: B

you have to add "--" in the beginning
upvoted 1 times

🗨️ **5dcd3a8** 1 month, 1 week ago

Selected Answer: A

Most Probably is A
the -- is not universally required for T-SQL in Fabric.
It's only required when you're: Mixing T-SQL inside a KQL context
upvoted 1 times

🗨️ **bettermakeme** 1 month, 2 weeks ago

Selected Answer: A

SQL By default shows ASC order
upvoted 1 times

🗨️ **xmh5025** 1 month, 3 weeks ago

Selected Answer: B

Remember to preface T-SQL queries with a T-SQL comment line, --, to tell the query editor to interpret the following query as T-SQL and not KQL.
<https://learn.microsoft.com/en-us/kusto/query/t-sql?view=microsoft-fabric#query-with-t-sql>
upvoted 4 times

🗨️ **Sher_Lock_Alexa** 3 months, 4 weeks ago

Selected Answer: A

sql ASC is the default sort order
upvoted 4 times

🗨️ **fassil** 4 months ago

Selected Answer: B

To run a T-SQL query, you need to begin the query with an empty T-SQL comment line (--). This tells the query editor to interpret the following query as T-SQL and not KQL
upvoted 4 times

🗨️ **JensQ** 4 months ago

Selected Answer: B

"To run a T-SQL query, begin the query with an empty T-SQL comment line: --. The -- syntax tells the query editor to interpret the following query as T-SQL and not KQL."
<https://learn.microsoft.com/en-us/kusto/query/t-sql?view=microsoft-fabric#query-with-t-sql>
upvoted 2 times

🗨️ **kostadiv** 4 months, 1 week ago

Selected Answer: B

missing "--" in the beginning
upvoted 2 times

🗨️ **Zoran_S** 4 months, 3 weeks ago

Selected Answer: B

"Remember to preface T-SQL queries with a T-SQL comment line, --, to tell the query editor to interpret the following query as T-SQL and not KQL."
While the code syntax is correct it is missing the '--' in the beginning to tell KQL to run it as T-SQL. Therefore it should be incorrect
upvoted 2 times

🗨️ 👤 **Bharat** 6 months, 2 weeks ago

Selected Answer: B

the answer is correct since there is no mention of 'Enable SQL Endpoint'. If it is enabled then the correct answer would have been A
upvoted 4 times

🗨️ 👤 **GHill1982** 6 months ago

In Fabric you can query a KQL database with SQL. I tested creating a new Eventhouse and ingesting an eventstream and this SQL querying functionality is available by default.

upvoted 6 times

🗨️ 👤 **doctordodge** 5 months ago

While this is true that you can query SQL in a KQL database, don't you need some sort of signifier to let the KQL database know you want to query using SQL (using "--" before your SELECT statement)?

<https://learn.microsoft.com/en-us/kusto/query/t-sql?view=microsoft-fabric>

upvoted 2 times

🗨️ 👤 **zxc01** 3 months ago

for default setting on KQL database, you can execute sql code directly and don't need using "--".

upvoted 1 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.

Correct Answer: B

🗳️ 👤 **12a2ecc** Highly Voted 2 months, 2 weeks 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 5 times

🗳️ 👤 **smanzana** Most Recent 1 week, 6 days 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

🗳️ 👤 **Biju1** 2 months, 2 weeks ago

Selected Answer: B

B is correct
upvoted 1 times

🗳️ 👤 **5e89616** 2 months, 3 weeks 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 3 times

🗳️ 👤 **12a2ecc** 2 months, 2 weeks 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 3 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 to 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)
```

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)
```

Correct Answer:

5e89616 Highly Voted 2 months, 3 weeks 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 6 times

Azure-JL 2 months, 1 week ago

Source: <https://learn.microsoft.com/en-us/fabric/data-warehouse/row-level-security>

upvoted 3 times

smanzana Most Recent 1 week, 5 days ago

The answer is correct

upvoted 1 times

DarioReymago 4 weeks ago

```
-- 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** 2 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
```


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

Correct Answer: A

  **smanzana** 1 week, 6 days ago

Selected Answer: A

A- Yes

upvoted 1 times

  **DarioReymago** 4 weeks ago

Selected Answer: A

<https://learn.microsoft.com/en-us/fabric/data-engineering/notebook-utilities#get-secret-using-user-credentials>

upvoted 1 times

  **lapiscineviden** 1 month, 3 weeks ago

Selected Answer: A

Get secret using user credentials:

```
notebookutils.credentials.getSecret('https://<name>.vault.azure.net/', 'secret name')
```

upvoted 1 times

  **Maryam3018** 2 months, 1 week 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 2 times

  **zxc01** 2 months, 2 weeks 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 4 times

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

Correct Answer: B


 **5e89616** Highly Voted 2 months, 3 weeks 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 5 times

 **smanzana** Most Recent 1 week, 6 days ago

Selected Answer: B

B - No

upvoted 1 times

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

Correct Answer: B

 **smanzana** 1 week, 6 days 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 month, 4 weeks ago


Selected Answer: B

It's B without the name of the secret
upvoted 1 times

 **LasAnsias** 2 months, 2 weeks 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 4 times

 **12a2ecc** 2 months, 2 weeks 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** 2 months, 3 weeks 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** 2 months, 2 weeks 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** 2 months, 2 weeks 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 to 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)
```

Answer Area

Correct Answer:

```
DAG = {
  "activities": [
    {
      "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)
```

5e89616 Highly Voted 2 months, 3 weeks 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 9 times

Rull 1 week, 6 days ago

Agree. <https://learn.microsoft.com/en-us/fabric/data-engineering/microsoft-spark-utilities#reference-run-multiple-notebooks-in-parallel>

upvoted 1 times

indomanish 2 weeks, 6 days 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

smanzana Most Recent 1 week, 5 days ago

The answer is correct:

activities

dependencies

runMultiple

upvoted 1 times

indomanish 2 weeks, 6 days 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

  **DarioReymago** 4 weeks ago


I'm not sure about Notebook or Activity. In Fabric Pipeline use Activity and type Notebook. Others ok

upvoted 1 times

  **kim32** 1 month, 2 weeks ago

<https://learn.microsoft.com/en-us/fabric/data-engineering/microsoft-spark-utilities#reference-run-multiple-notebooks-in-parallel>

upvoted 3 times

  **Sunnyb** 2 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 1 times

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.

```
CREATE TABLE [warehouse2].[dbo].[CreditCard]
(
    CreditCard varchar(20) NOT NULL
    ,CreditCardType varchar(10) NOT NULL)
GO
```

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.

Answer Area

ALTER
CREATE
DEFAULT
DROP
EMAIL
PARTIAL
REPLACE
UPDATE

TABLE dbo.CreditCard

ALTER
CREATE
DEFAULT
DROP
EMAIL
PARTIAL
REPLACE
UPDATE

COLUMN [CreditCard]

ADD MASKED

WITH (FUNCTION = '

ALTER
CREATE
DEFAULT
DROP
EMAIL
PARTIAL
REPLACE
UPDATE

(2, "XXXXXXXXXX", 4) ')

Correct Answer:

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

5e89616 Highly Voted 2 months, 3 weeks ago

Correct

ALTER TABLE- to modify the tabel

ALTER COLUMN - to modify the column

PARTIAL(prefix padding, padding string, suffix padding) - to expose first and last n characters, adding custom padding 'xxx' of a text field

upvoted 10 times

smanzana Most Recent 1 week, 5 days ago

The answer is correct: ALTER - ALTER - PARTIAL

upvoted 1 times

13d2a97 2 months, 1 week ago

ALTER / ALTER / PARTIAL

upvoted 3 times

HOTSPOT

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

Dynamic Data Copy does NOT use parametrization.

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

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

NOTE: Each correct selection is worth one point.

Answer Area

@activity(' Batch Object Copy
Dynamic Data Copy
Extraction Loop
Lookup Schema and Table '). output
output.count
output.pipelineReturnValue
output.value

Correct Answer:

Answer Area

@activity(' Batch Object Copy
Dynamic Data Copy
Extraction Loop
Lookup Schema and Table '). output
output.count
output.pipelineReturnValue
output.value

smanzana 1 week, 5 days ago

The answer is correct:

Lookup Schema and Table

output.value

upvoted 1 times

vernillen 3 weeks, 5 days ago

Your activity output always has a .value attribute. Use that.

upvoted 2 times

  **DarioReymago** 4 weeks ago



correct

upvoted 1 times

  **MasterOfData** 1 month ago

@activity('Lookup Schema and Table').output

upvoted 1 times

  **LasAnsias** 2 months, 2 weeks ago

i agree.

Lookup Schema and Table

output.value

upvoted 4 times

  **acote** 2 months, 2 weeks ago

Correct

upvoted 2 times

HOTSPOT

-

You have a Fabric workspace that contains a warehouse named Warehouse1. Warehouse1 contains a table named DimCustomers. DimCustomers contains the following columns:

- CustomerName
- CustomerID
- BirthDate
- EmailAddress

You need to configure security to meet the following requirements:

- BirthDate in DimCustomer must be masked and display 1900-01-01.
- EmailAddress in DimCustomer must be masked and display only the first leading character and the last five characters.

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

NOTE: Each correct selection is worth one point.

Answer Area

```
ALTER TABLE DimCustomer
```

```
ALTER COLUMN BirthDate
```

```
ADD MASKED WITH (FUNCTION =
```

'default()'
'partial(1900-01-01)'
'random(1900-01-01, 1900-01-01)'

```
ALTER TABLE DimCustomer
```

```
ALTER COLUMN EmailAddress
```

```
ADD MASKED WITH (FUNCTION =
```

'default()'
'email()'
'partial(1,"@",5)'
'random(1,"@",5)'

Answer Area

```
ALTER TABLE DimCustomer
```

```
ALTER COLUMN BirthDate
```

```
ADD MASKED WITH (FUNCTION =
```

'default()'
'partial(1900-01-01)'
'random(1900-01-01, 1900-01-01)'

Correct Answer:

```
ALTER TABLE DimCustomer
```

```
ALTER COLUMN EmailAddress
```

```
ADD MASKED WITH (FUNCTION =
```

'default()'
'email()'
'partial(1,"@",5)'
'random(1,"@",5)'

default()

Partial()

are correct. verified. you can try it in azure sql too

upvoted 7 times

  **PJLB** Highly Voted 2 months, 3 weeks ago

BirthDate: default() <- each datatype has a default

EmailAddress: partial()

upvoted 6 times

  **smanzana** Most Recent 1 week, 5 days ago

Default

Partial

upvoted 1 times

  **DarioReymago** 4 weeks ago

default and partial

upvoted 2 times

  **bettermakeme** 1 month, 2 weeks ago

BirthDate: 'default()'

EmailAddress: 'partial(1,"@",5)'

upvoted 2 times

  **Nilakhi** 2 months ago

1. default() -

For date and time data types, use 1900-01-01 00:00:00.0000000

2. email() -

Masking method that exposes the first letter of an email address and the constant suffix ".com", in the form of an email address. aXXX@XXXX.com.

upvoted 1 times

  **Nilakhi** 2 months ago



sorry, 2nd one should be partial()

upvoted 2 times

  **zxc01** 2 months, 3 weeks ago

the question is not very clear for datatype. However, DOB column should be date data type, then default() can match requirement which will be showed be '1900-01-01'

upvoted 3 times

  **tarmapas** 2 months, 3 weeks ago

first answer is default() for two reason: random is only for numbers and default for a date is 1900-01-01

upvoted 2 times

You have a Fabric workspace named Workspace1 that contains the following items:

- A Microsoft Power BI report named Report1
- A Power BI dashboard named Dashboard1
- A semantic model named Model1
- A lakehouse name Lakehouse1

Your company requires that specific governance processes be implemented for the items.

Which items can you endorse in Fabric?

- A. Lakehouse1, Model1, and Dashboard1 only
- B. Lakehouse1, Model1, Report1 and Dashboard1
- C. Report1 and Dashboard1 only
- D. Model1, Report1, and Dashboard1 only
- E. Lakehouse1, Model1, and Report1 only

Correct Answer: E

🗨️ **smanzana** 1 week, 5 days ago

Selected Answer: E

E is correct

upvoted 1 times

🗨️ **13d2a97** 2 months, 1 week ago

Selected Answer: E

All Fabric items and Power BI items except Power BI dashboards can be promoted or certified.

upvoted 4 times

🗨️ **Biju1** 2 months, 2 weeks ago

Selected Answer: E

Power BI dashboards cannot be promoted or certified.

upvoted 2 times

🗨️ **albertozgz** 2 months, 2 weeks ago

Selected Answer: B

All can be endorsed as a Master data

<https://learn.microsoft.com/en-us/fabric/governance/endorsement-overview#types-of-items-that-can-be-endorsed>

upvoted 1 times

🗨️ **5e89616** 2 months, 3 weeks ago

Selected Answer: E

Dashboard cannot be endorsed

upvoted 4 times

You have a Fabric workspace named Workspace1.

Your company acquires GitHub licenses.

You need to configure source control for Workspace1 to use GitHub. The solution must follow the principle of least privilege.

Which permissions do you require to ensure that you can commit code to GitHub?

- A. Actions (Read and write) and Contents (Read and write)
- B. Actions (Read and write) only
- C. Contents (Read and write) only
- D. Contents (Read) and Commit statuses (Read and write)

Correct Answer: C

🗨️ 👤 **5e89616** Highly Voted 👍 2 months, 3 weeks ago

Selected Answer: C

Contents (Read and write):

Allows reading repository content.

Allows pushing commits, creating branches, updating files – everything you need to commit code.

upvoted 6 times

🗨️ 👤 **smanzana** Most Recent ⌚ 1 week, 6 days ago

Selected Answer: C

C- Contents (Read and write) only

upvoted 1 times

🗨️ 👤 **13d2a97** 2 months, 1 week ago

Selected Answer: A

These permissions are necessary because:

Actions (Read and write): This permission allows you to interact with GitHub Actions, which might be necessary for automated workflows related to your commits.

Contents (Read and write): This permission allows you to read and write repository contents, which is essential for committing code.

upvoted 1 times

You have a Fabric workspace named Workspace1.

You plan to configure Git integration for Workspace1 by using an Azure DevOps Git repository.

An Azure DevOps admin creates the required artifacts to support the integration of Workspace1.

Which details do you require to perform the integration?

- A. the organization, project, Git repository, and branch
- B. the personal access token (PAT) for Git authentication and the Git repository URL
- C. the project, Git repository, branch, and Git folder
- D. the Git repository URL and the Git folder

Correct Answer: A

  **smanzana** 1 week, 6 days ago

Selected Answer: A

Correct is A



upvoted 1 times

  **DarioReymago** 4 weeks ago

Selected Answer: A

folder is important too, but option A is closer

upvoted 3 times



  **Clode** 1 month, 1 week ago

Selected Answer: A

A is correct

<https://learn.microsoft.com/en-us/fabric/cicd/git-integration/git-get-started?tabs=azure-devops%2CAzure%2Ccommit-to-git#connect-to-a-workspace>


upvoted 3 times

  **cayenne06** 2 months, 2 weeks ago

Selected Answer: A

This is correct

upvoted 2 times

  **Azure-JL** 2 months, 1 week ago

Source: <https://learn.microsoft.com/en-us/fabric/cicd/git-integration/git-get-started?tabs=azure-devops%2CAzure%2Ccommit-to-git#connect-to-a-workspace>

upvoted 1 times

You have a Fabric workspace that contains a lakehouse and a semantic model named Model1.

You use a notebook named Notebook1 to ingest and transform data from an external data source.

You need to execute Notebook1 as part of a data pipeline named Pipeline1. The process must meet the following requirements:

- Run daily at 07:00 AM UTC.
- Attempt to retry Notebook1 twice if the notebook fails.
- After Notebook1 executes successfully, refresh Model1.

Which three actions should you perform? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Place the Semantic model refresh activity after the Notebook activity and link the activities by using the On success condition.
- B. From the Schedule settings of Pipeline1, set the time zone to UTC.
- C. Set the Retry setting of the Notebook activity to 2.
- D. From the Schedule settings of Notebook1, set the time zone to UTC.
- E. Set the Retry setting of the Semantic model refresh activity to 2.
- F. Place the Semantic model refresh activity after the Notebook activity and link the activities by using an On completion condition.

Correct Answer: ABC

5e89616 **Highly Voted** 2 months, 3 weeks ago

Selected Answer: ABC

Correct

upvoted 5 times

smanzana **Most Recent** 1 week, 5 days ago

Selected Answer: ABC

ABC is correct

upvoted 1 times

13d2a97 2 months, 1 week ago

Selected Answer: B

@{pipeline().parameters.param1}

This syntax correctly references the parameter param1 and ensures that it is treated as an integer value within the pipeline expression.

upvoted 1 times

13d2a97 2 months, 1 week ago

ignore this as its part of next Q

upvoted 1 times

13d2a97 2 months, 1 week ago

Selected Answer: ABC

Correct

upvoted 2 times

You have a Fabric workspace that contains a lakehouse named Lakehouse1.

You plan to create a data pipeline named Pipeline1 to ingest data into Lakehouse1. You will use a parameter named param1 to pass an external value into Pipeline1. The param1 parameter has a data type of int.

You need to ensure that the pipeline expression returns param1 as an int value.

How should you specify the parameter value?

- A. "@pipeline().parameters.param1"
- B. "@{pipeline().parameters.param1}"
- C. "@{pipeline().parameters.[param1]}"
- D. "@@{pipeline().parameters.param1}"

Correct Answer: B

  **12a2ecc** Highly Voted 2 months, 2 weeks ago

Selected Answer: A

Should be "A"; "B" will be result in string, not int:

<https://learn.microsoft.com/en-us/fabric/data-factory/parameters#parameter-and-expression-concepts>

upvoted 6 times

  **zxc01** 2 months, 2 weeks ago

"@pipeline().parameters.myNumber" Returns 42 as a number. This one comes from that document.

upvoted 1 times

  **smanzana** Most Recent 1 week, 5 days ago

Selected Answer: A

A is correct

upvoted 1 times

  **bettermakeme** 1 month, 1 week ago

Selected Answer: A

<https://learn.microsoft.com/en-us/fabric/data-factory/parameters#parameter-and-expression-concepts>

Parameter value Result

"@pipeline().parameters.myString" Returns foo as a string.

"@{pipeline().parameters.myString}" Returns foo as a string.

"@pipeline().parameters.myNumber" Returns 42 as a number.

"@{pipeline().parameters.myNumber}" Returns 42 as a string.

"Answer is: @{pipeline().parameters.myNumber}" Returns the string Answer is: 42.

"@concat('Answer is: ', string(pipeline().parameters.myNumber))" Returns the string Answer is: 42

"Answer is: @@{pipeline().parameters.myNumber}" Returns the string Answer is: @{pipeline().parameters.myNumber}.



upvoted 1 times

  **kim32** 1 month, 2 weeks ago

Selected Answer: A

B is wrong because "@{pipeline().parameters.myNumber}" Returns myNumber as a string.

upvoted 1 times

  **13d2a97** 2 months, 1 week ago

Selected Answer: B

@{pipeline().parameters.param1}

This syntax correctly references the parameter param1 and ensures that it is treated as an integer value within the pipeline expression.



upvoted 2 times

  **9Exam** 2 months, 2 weeks ago

Selected Answer: A

correct answer is a

upvoted 1 times

  **tarmapas** 2 months, 2 weeks ago

Selected Answer: A

B get string output



upvoted 2 times

  **zxc01** 2 months, 3 weeks ago

Selected Answer: A

answer A will get int data output. answer B will get string.

upvoted 4 times

  **12a2ecc** 2 months, 2 weeks ago

Should be "A"; "B" will be result in string, not int:

<https://learn.microsoft.com/en-us/fabric/data-factory/parameters#parameter-and-expression-concepts>

upvoted 2 times

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

- A Dataflow Gen2 dataflow that copies data from an on-premises Microsoft SQL Server database to Lakehouse1
- A notebook that transforms files and loads the data to Lakehouse1
- A data pipeline that loads a CSV file to Lakehouse1

You need to develop an orchestration solution in Fabric that will load each item one after the other. The solution must be scheduled to run every 15 minutes.

Which type of item should you use?

- A. notebook
- B. warehouse
- C. Dataflow Gen2 dataflow
- D. data pipeline

Correct Answer: D

🗲️ 👤 **smanzana** 1 week, 5 days ago

Selected Answer: D

D is correct

upvoted 1 times

🗲️ 👤 **13d2a97** 2 months, 1 week ago

Selected Answer: D

Data Pipeline

upvoted 2 times

🗲️ 👤 **5e89616** 2 months, 3 weeks ago

Selected Answer: D

correct

upvoted 3 times

You are building a Fabric notebook named MasterNotebook1 in a workspace. MasterNotebook1 contains the following code.

```
DAG = {
  "activities": [
    {
      "name": "execute_notebook_1",
      "path": "notebook_01",
      "timeoutPerCellInSeconds": 600,
      "args": {
        "input_value": "999"
      },
      "retry": 1,
      "retryIntervalInSeconds": 30
    },
    {
      "name": "execute_notebook_2",
      "path": "notebook_02",
      "timeoutPerCellInSeconds": 400,
      "args": {
        "input_value": "888"
      },
      "retry": 1,
      "retryIntervalInSeconds": 30
    },
    {
      "name": "execute_notebook_3",
      "path": "notebook_03",
      "timeoutPerCellInSeconds": 600,
      "args": {
        "input_value": "777"
      },
      "retry": 1,
      "retryIntervalInSeconds": 30
    },
    {
      "name": "execute_notebook_3",
      "path": "notebook_03",
      "timeoutPerCellInSeconds": 600,
      "args": {
        "input_value": "777"
      },
      "retry": 1,
      "retryIntervalInSeconds": 30
    }
  ],
  "timeoutInSeconds": 43200,
  "concurrency": 0
}

mssparkutils.notebook.runMultiple(DAG,{"displayDAGViaGraphviz": True})
```

You need to ensure that the notebooks are executed in the following sequence:

1. Notebook_03
2. Notebook_01
3. Notebook_02

Which two actions should you perform? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Move the declaration of Notebook_02 to the bottom of the Directed Acyclic Graph (DAG) definition.

- B. Add dependencies to the execution of Notebook_03.
- C. Split the Directed Acyclic Graph (DAG) definition into three separate definitions.
- D. Add dependencies to the execution of Notebook_02.
- E. Change the concurrency to 3.
- F. Move the declaration of Notebook_03 to the top of the Directed Acyclic Graph (DAG) definition.

Correct Answer: *DF*

🗳️ 👤 **e35e079** 4 days, 11 hours ago

Selected Answer: BD

Explanation:

To enforce the required sequence:

Set Notebook_03 to execute first (no dependencies needed).

Then set Notebook_01 to depend on Notebook_03.

Then set Notebook_02 to depend on Notebook_01.

This is achieved by adding the "dependsOn" key to each not
upvoted 1 times

🗳️ 👤 **Stef6108** 1 week, 4 days ago

Selected Answer: BD

The execution order in a DAG is only ensured by dependencies, not the position in the defining data structure (JSON).
upvoted 1 times

🗳️ 👤 **Stef6108** 1 week, 4 days ago

Wrong, ignore.

upvoted 1 times

🗳️ 👤 **smanzana** 1 week, 5 days ago

Selected Answer: BD

BD is correct

upvoted 1 times

🗳️ 👤 **moons_123** 2 weeks, 6 days ago

Selected Answer: DF

D,F - Concurrency set to 0 means they will start all at once. So add dependencies and change the order. <https://learn.microsoft.com/en-us/azure/synapse-analytics/spark/microsoft-spark-utilities?pivots=programming-language-python#reference-run-multiple-notebooks-in-parallel>
upvoted 2 times

🗳️ 👤 **Sr18** 1 month ago

Selected Answer: BD

Why BD is best answer:

Sequence of activity declaration wont matter till the time you have dependencies written for Notebook1 in Notebook 03, and Notebook 02 dependency on Notebook 01. So no need to move notebook 03 on top
upvoted 2 times

🗳️ 👤 **Sr18** 1 month ago

Another thing they will execute in parallel if there are no dependencies mentioned.

upvoted 1 times

🗳️ 👤 **Sr18** 1 month ago

DF: Finally after long thinking and being lost in verbiage.

Concurrency is set to 0: They will start executing all together

So best option is to move notebook 3 on top

Add dependency to notebook 2.

But this will still not solve the problem, as we need to add dependency in Notebook1 for notebook 2, or change concurrency to 1. And moreover

delete the duplication of notebook 3 execution.

So overall this question is either stupid or wrong options are mentioned. Does not make sense at all. Whatever you choose from here it will never work.

upvoted 1 times

🗨️ 👤 **8d1e8fe** 1 month, 2 weeks ago

Selected Answer: F

What am i missing here? Why is notebook_03 mentioned twice? Also moving Notebook_03 to the top already solves the entire problem since the concurrency is set to 1. There is no need to add a dependency for execution order since 1 and 2 are already ordered correctly.

upvoted 1 times

🗨️ 👤 **kim32** 1 month, 2 weeks ago

Selected Answer: BD

Why Not Other Options?

A/F (Moving declarations): Order in the DAG array doesn't control execution sequence; dependencies do.

C (Splitting DAG): Unnecessary; a single DAG with dependencies works.

E (Concurrency=3): Irrelevant for sequential execution (use concurrency=1 if strict order is needed).

upvoted 2 times

🗨️ 👤 **AnetaK** 2 months ago

Selected Answer: DF

Why is Notebook_03 listed two times in the code?

upvoted 1 times

🗨️ 👤 **5e89616** 2 months, 3 weeks ago

Selected Answer: DF

I would also add dependencies to Notebook_01 but it was not listed

upvoted 3 times

You have a Fabric workspace that contains a data pipeline named Pipeline1 as shown in the exhibit. (Click the Exhibit tab.)

The screenshot shows the Microsoft Fabric Pipeline1 interface. The pipeline consists of two activities: 'Execute procedure1' (a Stored procedure activity) and 'Copy_kdi' (a Copy data activity). Both activities are marked with green checkmarks, indicating they are active. The 'Output' tab is selected, showing a table with 2 items. The first item is 'Copy_kdi' with a status of 'Succeeded' and a duration of 33s. The second item is 'Execute procedure1' with a status of 'Inactive' and a duration of 'Less than 1s'.

What will occur the next time Pipeline1 runs?

- A. Copy_kdi will run first, and then Execute procedure1 will run.
- B. Execute procedure1 will run first, and then Copy_kdi will run.
- C. Execute procedure1 will run and Copy_kdi will be skipped.
- D. Copy_kdi will run and Execute procedure1 will be skipped.
- E. Both activities will run simultaneously.
- F. Both activities will be skipped.

Correct Answer: F

tczubat 1 month, 2 weeks ago

Selected Answer: D

I just tested this. If you set the activity state of any activity to deactivated, it will show as inactive when run since it will be skipped over.
upvoted 3 times

kim32 1 month, 2 weeks ago

Selected Answer: D

No need to complicate things, Stored procedure activity is inactive
upvoted 3 times

f6f3e50 2 months ago

Selected Answer: D

Here's what happens:

If one activity (e.g., a Copy activity) is active and connected, and another (like a Stored Procedure activity) is inactive and unconnected, only the active one runs.

Inactive activities do not execute, just like in Azure Data Factory.

Unconnected activities are not part of the execution graph, so Fabric does not trigger them.

upvoted 1 times

Tamele001 2 months ago

Selected Answer: E

Both activities will run simultaneously.

The Execute procedure1 activity was inactive in the previous run, it does not mean that on the next one it will be inactive again. There's no way of disabling an activity

upvoted 2 times

12a2ecc 2 months, 2 weeks ago

Selected Answer: E

Is that just the history showing inactive for copy_kdi? Isn't it active on the diagram? I was thinking they would both run(E)

upvoted 1 times

🗨️ 👤 **12a2ecc** 2 months, 2 weeks ago

I incorrectly viewed this originally; D does look correct since the Stored Proc step appears disabled in the pipeline

upvoted 1 times

🗨️ 👤 **5e89616** 2 months, 3 weeks ago

Selected Answer: D

(D) Copy_kdi will run and Execute procedure1 will be skipped (because it's inactive).

I do not see why both would be skipped?...

upvoted 1 times

🗨️ 👤 **5e89616** 2 months, 3 weeks ago

or (E) Both would run simultaneously, even thou Stored Procedure is inactive

upvoted 1 times

🗨️ 👤 **12a2ecc** 2 months, 2 weeks ago

Is that just the history showing inactive for copy_kdi? Isn't it active on the diagram? I was thinking they would both run(E)

upvoted 1 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. 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 WorkspaceA can be configured for source control.

Which two actions should you perform? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. From Tenant setting, set Users can synchronize workspace items with their Git repositories to Enabled.
- B. From Tenant setting, set Users can sync workspace items with GitHub repositories to Enabled.
- C. Configure WorkspaceA to use a Premium Per User (PPU) license.
- D. Assign WorkspaceA to Cap1.

Correct Answer: AD

  **5e89616** 2 months, 3 weeks ago

Selected Answer: AD

(A)

https://learn.microsoft.com/en-us/fabric/admin/git-integration-admin-settings?utm_source=chatgpt.com#users-can-synchronize-workspace-items-with-their-git-repositories

(D)

https://learn.microsoft.com/en-us/fabric/cicd/git-integration/git-get-started?utm_source=chatgpt.com&tabs=azure-devops%2CAzure%2Ccommit-to-git#prerequisites

upvoted 2 times

HOTSPOT

-

You have a Fabric workspace that contains a warehouse named Warehouse1. Warehouse1 contains a table named Customer. Customer contains the following data.

CustomerID	FirstName	LastName	Phone	CreditCard
1	John	Doe	555-123-4567	1234567812345670
2	Jane	Smith	555-987-6543	8765432187654320
3	Michael	Johnson	555-555-5555	1234987654321230
4	Emily	Davis	555-222-3333	4321123456789870
5	David	Brown	555-444-5555	5678123498761230

You have an internal Microsoft Entra user named User1 that has an email address of user1@contoso.com.

You need to provide User1 with access to the Customer table. The solution must prevent User1 from accessing the CreditCard column.

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

NOTE: Each correct selection is worth one point.

Answer Area

GRANT ON

ALTER
EXECUTE
READ
SELECT
VIEW

Customers(CustomerID, FirstName, LastName, Phone)

TO

User1
[User1]
[user1@contoso.com]

Answer Area

Correct Answer:

GRANT ON

ALTER
EXECUTE
READ
SELECT
VIEW

Customers(CustomerID, FirstName, LastName, Phone)

TO

User1
[User1]
[user1@contoso.com]

  **12a2ecc** Highly Voted 2 months, 2 weeks ago

answer should be GRANT SELECT ON Customers(CustomerID, FirstName, LastName, Phone) TO [User1@contoso.com]

<https://learn.microsoft.com/en-us/fabric/data-warehouse/column-level-security>

upvoted 20 times

  **Biju1** Highly Voted 2 months, 2 weeks ago



GRANT SELECT ON Customers(CustomerID, FirstName, LastName, Phone) TO [User1@contoso.com] is corrected. I tested

upvoted 7 times

  **zxc01** Most Recent 2 months, 2 weeks ago



answer should be GRANT SELECT ON Customers(CustomerID, FirstName, LastName, Phone) TO [User1@contoso.com]

upvoted 5 times

  **12a2ecc** 2 months, 2 weeks ago

Correct, since Entra user: <https://learn.microsoft.com/en-us/fabric/data-warehouse/column-level-security>

upvoted 1 times

  **5e89616** 2 months, 3 weeks ago

Correct.

GRANT SELECT ON Customers(CustomerID, FirstName, LastName, Phone) TO [User1]

Alternatively:

DENY SELECT ON Customers(CreditCard) TO [User1]

upvoted 4 times

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. an Azure DevOps pipeline
- B. an eventstream
- C. GitHub Actions

Correct Answer: A

Community vote distribution



🗉 👤 **PBridge** 2 weeks, 4 days ago

Selected Answer: A

To automate deployment of such components across multiple Fabric workspaces (Dev → Test → Prod), you'd use deployment automation tools like Azure DevOps pipelines .

upvoted 1 times

🗉 👤 **moons_123** 2 weeks, 6 days ago

Selected Answer: B

Azure DevOps pipeline and GitHub Actions are both CI/CD tools. An eventstream is the only unique choice and is used with an Eventhouse

upvoted 1 times

🗉 👤 **kaushikcbs** 3 weeks ago

Selected Answer: A

Devops Pipeline

upvoted 2 times

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. TRIGGER
- B. SCHEMA
- C. FUNCTION
- D. DATABASE ROLE

Correct Answer: C

Community vote distribution

C (83%)

D (17%)

🗳️ 👤 **rehoboth2165** 2 weeks, 3 days ago

Selected Answer: C

FUNCTION ✓ Used to define the logic to filter rows based on user identity (predicate function) ✓ Yes

D. DATABASE ROLE Can assign permissions, but not used for filtering rows dynamically ✗ No
upvoted 1 times

🗳️ 👤 **Jay63002** 2 weeks, 3 days ago

Selected Answer: C

Correct answer is function not database role
upvoted 1 times

🗳️ 👤 **kaushikcbs** 3 weeks ago

Selected Answer: C

Function and Security Policy is used. So as per the options available it is Function.
upvoted 3 times

🗳️ 👤 **gretito_g** 3 weeks ago

Selected Answer: D

D. DATABASE ROLE
upvoted 1 times

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. SECURITY POLICY
- B. TABLE
- C. TRIGGER
- D. STORED PROCEDURE

Correct Answer: A

Community vote distribution

A (100%)

  **kaushikcbs** 3 weeks ago

Selected Answer: A

Function and Security Policy is used. So as per the options available it is Security Policy.
upvoted 2 times

  **gretito_g** 3 weeks ago

Selected Answer: A

Without a SECURITY POLICY, the filtering logic won't be applied, even if the function exists.
upvoted 1 times

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. Create views.
- B. Modify the surrogate keys to use a different data type.
- C. Change the MD5 hash to SHA256.
- D. Increase the capacity.
- E. Disable V-Order on the warehouse.

Correct Answer: B

Community vote distribution

B (100%)

 **PBridge** 2 weeks, 4 days ago

Selected Answer: B

Changing surrogate keys from MD5 (string) to integers improves query speed, storage efficiency, and reduces costs.

upvoted 1 times

 **kaushikcbs** 3 weeks ago

Selected Answer: B

Modify the surrogate keys to use a different data type

Why?

- MD5 hash surrogate keys are 128-bit values, which can be inefficient for large datasets.
- Switching to a more efficient data type (such as BIGINT or INT) reduces storage and processing overhead, improving query performance.
- Smaller data types are faster to process, leading to better indexing and query execution.

Answer validated with ChatGPT

upvoted 2 times

 **gretito_g** 3 weeks ago

Selected Answer: B

Modify the surrogate keys to use a different data type.

upvoted 2 times