



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

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%)

 **mixonfreddy** 1 week ago

Selected Answer: C

Answer is C

upvoted 1 times

 **QAZdbarhate12345678** 2 weeks, 5 days 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 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%)

🗳️ 👤 **QAZdbarhate12345678** 2 weeks, 5 days ago

Selected Answer: A

lakehouse is the most appropriate choice for this scenario

upvoted 1 times

🗳️ 👤 **Tuki93** 3 weeks, 4 days ago

Selected Answer: A

You have semi-structured data.

The data will only be written by using Spark.

upvoted 2 times

🗳️ 👤 **kryper_reg** 3 weeks, 6 days ago

Selected Answer: A

Semi structured data should be stored in a 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%)

🗨️ 👤 **QAZdbarhate12345678** 2 weeks, 5 days 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 2 times

🗨️ 👤 **Tuki93** 3 weeks, 1 day ago

Selected Answer: B

Dataflow Gen 1 is deprecated.

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. 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%)

🗨️ 👤 **QAZdbarhate12345678** 2 weeks, 5 days 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 2 times

🗨️ 👤 **Tuki93** 3 weeks, 4 days 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%)

 **mixonfreddy** 1 week ago

Selected Answer: D

You have to modify the surrogate keys

upvoted 2 times

 **manolet** 1 week, 4 days ago

Selected Answer: C

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

upvoted 1 times

 **QAZdbarhate12345678** 2 weeks, 5 days ago

Selected Answer: D

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

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

upvoted 2 times

 **Tuki93** 3 weeks, 4 days 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 4 times

 **Ghill1982** 1 week, 5 days ago

Would V-Order not already be enabled by default?

upvoted 2 times

 **IshtarSQL** 3 weeks, 6 days ago

Selected Answer: D

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

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

upvoted 2 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

```

(SELECT CAST
SELECT CONVERT
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
(CUBE(YEAR(SO.ModifiedDate), P.Name)
GROUPING SETS ((YEAR(SO.ModifiedDate), P.Name), (YEAR(SO.ModifiedDate)))
ROLLUP(YEAR(SO.ModifiedDate), P.Name)
YEAR(SO.ModifiedDate), P.Name)
ORDER BY OrderDate

```

Correct Answer:

```

(SELECT CAST
SELECT CONVERT
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
(CUBE(YEAR(SO.ModifiedDate), P.Name)
GROUPING SETS ((YEAR(SO.ModifiedDate), P.Name), (YEAR(SO.ModifiedDate)))
ROLLUP(YEAR(SO.ModifiedDate), P.Name)
YEAR(SO.ModifiedDate), P.Name)
ORDER BY OrderDate

```

SELECT YEAR

ROLLUP(YEAR(SO.ModifiedDATE), PName)

upvoted 2 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** 2 weeks, 5 days 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 2 times

🗉 👤 **Tuki93** 3 weeks, 4 days ago

Selected Answer: BCF

I believe the provided answer is correct

upvoted 2 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%)

🗨️ 👤 **QAZdbarhate12345678** 2 weeks, 5 days 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 2 times

🗨️ 👤 **Tuki93** 3 weeks, 4 days 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 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%)

 **Tuki93** 3 weeks, 4 days ago

Selected Answer: DEF

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

upvoted 2 times

 **IshtarSQL** 3 weeks, 6 days 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** 2 weeks, 5 days 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 2 times

🗨️ 👤 **Tuki93** 3 weeks, 4 days ago

Selected Answer: D

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

upvoted 2 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
Data from the semantic models will be deployed to the target stage.	<input type="radio"/>	<input checked="" type="radio"/>
Correct Answer: 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** 1 day, 6 hours ago

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

YES: Gen1 dataflows will be deployed

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

upvoted 1 times

 **mixonfreddy** 2 weeks, 3 days ago

Selected Answer: A

Yes/Yes/No

upvoted 1 times

 **GHill1982** 4 days, 17 hours ago

No/Yes/No

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** 2 weeks, 5 days 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 2 times

🗨️ 👤 **Tuki93** 3 weeks, 4 days 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 1 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** 2 weeks, 5 days 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 2 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



🗨️ 👤 **QAZdbarhate12345678** 2 weeks, 5 days 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 1 times

🗨️ 👤 **IshtarSQL** 3 weeks, 2 days 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 3 times

🗨️ 👤 **Tuki93** 3 weeks, 4 days 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 3 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%)

Meir 1 week, 2 days ago

Selected Answer: A

See <https://learn.microsoft.com/en-us/answers/questions/2074603/run-pipeline-via-azure-function-by-rest-api>
upvoted 2 times

GHill1982 4 days, 17 hours 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 1 times

mixonfreddy 2 weeks, 3 days ago

Selected Answer: D

Use of service principal is discouraged
upvoted 1 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

C (80%)

D (20%)

🗨️ **Bharat** 2 weeks, 5 days 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 1 times

🗨️ **Bharat** 2 weeks, 4 days 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 1 times

🗨️ **Tuki93** 3 weeks, 4 days 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 4 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** 3 days, 21 hours 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 1 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%)

 **01ceacf** 3 weeks ago

Selected Answer: B

File greater than 1GB in size are not cached.

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

upvoted 2 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.

 GHill1982 3 days, 21 hours ago

All four answers are correct.

upvoted 1 times

  **boonfidelio** 4 days, 22 hours ago

Selected Answer: A

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

upvoted 1 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%)

🗨️ 👤 **Zekes** 1 week, 4 days 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

🗨️ 👤 **Bharat** 2 weeks, 5 days ago

Selected Answer: B

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

upvoted 2 times

🗨️ 👤 **IshtarSQL** 3 weeks, 1 day 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** 3 days, 21 hours 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** 17 hours, 59 minutes 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 1 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%)

 **gfors** 2 weeks ago

Selected Answer: D

This is wrong in the docs. Read this <https://community.fabric.microsoft.com/t5/Service/Create-or-modify-warehouse-items-can-do-with-Contributor-role/m-p/4318708>

upvoted 2 times

 **GHill1982** 3 days, 21 hours ago

Thanks for the link, very useful

upvoted 1 times

 **01ceacf** 3 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 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** 3 weeks ago

Selected Answer: A

Read Access to Table Data.

B/D - Viewer/Member can access other items

C - cannot access SQL DB

upvoted 2 times

 **IshtarSQL** 3 weeks, 1 day 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 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 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"/>

Currently there are no comments in this discussion, be the first to comment!

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.	<input type="radio"/>	<input type="radio"/>
User3 has Domain contributor access to Domain1.	<input type="radio"/>	<input type="radio"/>
User2 has Contributor role access to Workspace3.	<input type="radio"/>	<input type="radio"/>

Statements	Yes	No
User3 has Viewer role access to Workspace3.	<input checked="" type="radio"/>	<input type="radio"/>
User3 has Domain contributor access to Domain1.	<input type="radio"/>	<input checked="" type="radio"/>
User2 has Contributor role access to Workspace3.	<input type="radio"/>	<input checked="" type="radio"/>

Correct Answer:

Currently there are no comments in this discussion, be the first to comment!

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

 **GHill1982** 3 days, 21 hours ago

Selected Answer: D

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

upvoted 1 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

Currently there are no comments in this discussion, be the first to comment!

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.	

Currently there are no comments in this discussion, be the first to comment!


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

  **GHill1982** 3 days, 20 hours ago

Selected Answer: C

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

upvoted 1 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** 3 days, 20 hours 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 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** 2 days, 20 hours 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 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



The Answer Area contains three empty rectangular boxes. Red lines indicate the correct sequence of code segments to be placed in them: the first box should contain 'ALTER TABLE dbo.DimCustomer', the second box should contain 'ADD CONSTRAINT PK_DimCustomer PRIMARY KEY CLUSTERED (CustomerKey)', and the third box should contain 'ENFORCED'.

Currently there are no comments in this discussion, be the first to comment!

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** 3 days, 7 hours 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 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** 2 days, 21 hours 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 2 times

 **GHill1982** 3 days, 7 hours ago

Selected Answer: B

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

upvoted 1 times

 **GHill1982** 3 days, 7 hours 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** 2 days, 20 hours 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** 2 days, 20 hours ago

The duplicated project, even though not required, does however work and does not throw any errors.
upvoted 1 times

  **QAZdbarhate12345678** 5 days, 1 hour 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%)

 **QAZdbarhate12345678** 5 days, 1 hour 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** 2 weeks, 5 days 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** 2 days, 20 hours 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 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 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** 3 days, 7 hours 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 2 times

 **Bharat** 2 weeks, 4 days 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

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** 1 week, 6 days ago

Selected Answer: A

SQL By default shows ASC order
upvoted 2 times

 **Bharat** 2 weeks, 5 days 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 2 times

 **GHill1982** 2 days, 19 hours 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 1 times