

Data Engineering újdonságok MS Fabricban



Arató Bence

A BI Consulting igazgatója, két évtizedes szakmai múlttal az adattárházak, BI és analitika területén

- **Adatstratégiai tanácsadás**

Adat- és BI stratégia kialakítása, architektúratervezés, eszközválasztási tanácsadás, pályáztatás, minőségbiztosítás, mentoring.

- **Oktatás**

A BI Akadémia vezető oktatója (DW/BI architektúrák, adatmodellezés, adatvizualizáció). A CEU vendégtanára.

- **Rendezvények**

A Budapest BI Fórum és a Budapest Data+AI Fórum konferenciák főszervezője.

- **Kutatások**

A magyar Data Leaders Survey és a Data Salary Survey szakmai vezetője.



BI Consulting

FABCON Atlanta 2026



https://www.linkedin.com/posts/ruiromano_fabcon-microsoftfabric-powerbi-activity-7444394375939452928-XyKc

Key Announcements

- Engines
 - Multiple Engines
 - Spark: Fabric Runtime 2.0, Materialized Lake Views
- Data Integration
 - Multiple Tools
 - Notebook based approach
 - DevEx: VS Code integration, API calls, Fabric CLI, MCP
- AI Integration
 - AI Functions everywhere
 - Copilots
- Stack Expansion
 - Data Quality Management

FABCON DataEng újdonságok

- **Data Engineering**

- [Materialized Lake Views \(GA\)](#)
- [Data Factory updates](#)
- [Fabric Notebook API](#)
- [Fabric Runtime 2.0](#)

- **Data Warehousing**

- [Custom SQL Pools](#)
- [dbt + Microsoft Fabric](#)
- [AI Functions in Data Warehouse](#)
- [SQL Alerts](#)

Data Engineering

Data Integration

Data Integration Decision Guide

Tartalom

Fejezet

	Data movement				Orchestration		Transformation	
	 Mirroring	 Copy Job	 Copy Activity (Pipeline)	 Eventstream	 Pipeline	 Apache Airflow Job	 Notebooks	 Dataflow Gen 2
Use Case	Data Replication	Data Ingestion & Replication	Data Ingestion	Stream Data Ingestion & Processing	Low Code Orchestration	Code-first Orchestration	Code-first Data Prep / Transform	Code-free Data Prep / Transform
Flagship Scenarios	Near real-time sync with turn-key set up. Replication.	Incremental Copy / Replication (water-mark + Native CDC), Data Lake / Storage Data Migration, Medallion Ingestion, Out-of-the-box multi-table copy.	Data Lake / Storage Data Migration, Medallion Ingestion, Incremental copy via pipeline expressions & control tables (water-mark only)	Incremental processing, event-driven and real-time AI applications.	Logical grouping of several activities together to perform a task.	Python Code-Centric Authoring	Complex Transformations	Transformation & Profiling
Source	6+ connectors  	50+ connectors	50+ connectors	25+ connectors, custom endpoint (AMQP, Kafka, HTTP)	All Fabric compatible sources (depending on selected pipeline activities)	100+ connectors	100+ Spark Libraries	170+ built-in connectors    + Custom SDK
Destination	Mirrored database (stored as read-only Delta table in Fabric OneLake)	40+ connectors 	40+ connectors 	Eventhouse, Activator, Lakehouse, custom endpoint (AMQP, Kafka, HTTP), Spark	All Fabric compatible sources (depending on selected pipeline activities)	100+ connectors	100+ Spark Libraries	7+ connectors   
Type of Incoming Data	Near Real-time	Batch / Incremental Copy (water-mark based & change data capture) / Near Real-time	Batch / Bulk / Manual Watermark-based incremental copy	Real-time streaming data, change data capture feeds	All types	All types	All types	All types
Persona	Business Analyst, Database Administrator	Business Analyst, Data Integrator, Data Engineer	Data Integrator, Business Analyst, Data Engineer	Data analyst, data engineer, & integrator	Data Integrator, Business Analyst, Data Engineer	Apache Airflow Users	Data Scientist, Developer	Data Engineer, Data Integrator, Business Analyst
Skillset	None	ETL, SQL	ETL, SQL	SQL, ETL, KQL	ETL, SQL, Spark (Scala, Py, SQL, R)	Python	Spark (Scala, Py, SQL, R)	ETL, M, SQL
Coding Level	No code	No code / Low code	No code / Low code	No code / Low code	No code / Low code	Code-first	Code-first	No code / Low code
Transformation Support	None	Low	Low	Medium (stream analytics)	None	None	High	High (400+ activities)

Data Factory

- [Migration Assistant](#)
- [MCP Server](#)
- [Expanded connector coverage](#)
- [Copy job](#)
- [Pipeline](#)
- [Dataflow Gen2](#)

Tartalom

Fejezet

Data Factory at FabCon

Fabric Data Factory at FabCon Atlanta: Built for modern data integration

Announcements

Data Factory

Microsoft Fabric



March 18, 2026 by [Faisal Mohamood](#)


👁 10,089 Views

Share

Tweet

Like

If you haven't already, check out Arun Ulag's hero blog ["FabCon and SQLCon 2026: Unifying databases and Fabric on a single, complete platform"](#) for a complete look at all of our FabCon and SQLCon announcements across both Fabric and our database offerings.



Copy job

Audit columns in Copy job

Tartalom

Fejezet

Lista

Audit columns in Copy job in Fabric Data Factory—every row is traceable for data lineage and compliance

Announcements

Data Factory

Microsoft Fabric



March 19, 2026 by [Ye Xu](#) 3,437 Views

Share

Tweet

Like

If you haven't already, check out Arun Ulag's hero blog "[FabCon and SQLCon 2026: Unifying databases and Fabric on a single, complete platform](#)" for a complete look at all of our FabCon and SQLCon announcements across both Fabric and our database offerings.

Audit columns in Copy job

Select folders or files <<

File path

Search

☐ JSON_Array1.json

☒ column mapping.txt

☐ space 1 delimiter.txt

☐ temp.txt

Audit columns ⓘ

+ New

🗑 Delete

☐ Name	Value
☐ _COPY_JOB_SYNCED	Data extraction time
☐ _COPY_JOB_ID	Copy job ID
☐ _COPY_JOB_SOURCE_FILE_PATH	File path

Column

Workspace ID

Copy job ID

Copy job run ID

Copy job name

Data extraction time

Lower bound

File path

ID	COUNTER	_COPY_JOB_SYNCED	_COPY_JOB_ID	_COPY_JOB_SOURCE_FILE_PATH
a	10	2020-01-01T00:00:00Z	xxxx-xxxx-xxxx-xxxx	Container/Folder/a.csv
a	30	2020-01-02T00:00:00Z	xxxx-xxxx-xxxx-xxxx	Container/Folder/b.csv
b	20	2020-01-01T00:00:00Z	xxxx-xxxx-xxxx-xxxx	Container/Folder/c.csv
c	40	2020-01-02T00:00:00Z	xxxx-xxxx-xxxx-xxxx	Container/Folder/d.csv

New watermark column types in Copy job

Tartalom

Fejezet

Lista

Incremental copy gets more flexible: New watermark column types in Copy job in Fabric Data Factory (Generally Available)

Announcements

Data Factory

Microsoft Fabric



March 24, 2026 by [Ye Xu](#)  3,869 Views

 Share

 Tweet

 Like

If you haven't already, check out Arun Ulag's hero blog "[FabCon and SQLCon 2026: Unifying databases and Fabric on a single, complete platform](#)" for a complete look at all of our FabCon and SQLCon announcements across both Fabric and our database offerings.

Pipeline

Tartalom

Fejezet

Lista

Modernizing pipelines: New activities and innovations in Fabric Data Factory pipelines

Data Factory



March 20, 2026 by [Connie Xu](#)

👁 4,462 Views

Share

Tweet

Like

If you haven't already, check out Arun Ulag's hero blog "[FabCon and SQLCon 2026: Unifying databases and Fabric on a single, complete platform](#)" for a complete look at all of our FabCon and SQLCon announcements across both Fabric and our database offerings.

Lakehouse maintenance activity (Preview)

General

Settings

Connection * ⓘ

 Lakehouse (Preview) ▾

 Refresh  Edit

Lakehouse *

 myLakehouse ▾

 Open

Schema name

Table name *

PUBLIC_PRODUCT_SUBCATEGORIES ▾

 Refresh

Maintenance settings ⓘ

Run optimization ⓘ

☒

Apply v-order ⓘ

☒

Apply z-order ⓘ

☐


Run vacuum ⓘ

☒

Retention period

Wait on completion

☒



Refresh SQL endpoint activity (Preview)

General

Settings

Connection *

FabricSqlEndpointMetadata

Refresh

Edit

Workspace *

DFBugBash

Refresh


SQL endpoint Id *

myLakehouse

Refresh

Open

Timeout



Copilot in the Pipeline expression builder (GA)

Tartalom

Fejezet

Lista

Pipeline expression builder

Add dynamic content below using any combination of **expressions**, **functions** and **system variables**.

 @concat(formatDateTime(convertFromUtc(utcNow(), 'Eastern Standard Time', 'yyyyMMddHHmmss')), 'publicholidays')

 input current time in EST and add string to input folder name "publicholidays"



Accept

Discard



AI-generated content may be incorrect

Dataflow Gen2

New Dataflow Gen2 capabilities

Tartalom

Fejezet

Lista

A wave of new Dataflow Gen2 capabilities at FabCon Atlanta 2026

Data Factory

Microsoft Fabric



March 23, 2026 by [Miguel Llopis](#) 3,195 Views

Share

Tweet

Like

If you haven't already, check out Arun Ulag's hero blog "[FabCon and SQLCon 2026: Unifying databases and Fabric on a single, complete platform](#)" for a complete look at all of our FabCon and SQLCon announcements across both Fabric and our database offerings.

At FabCon Atlanta 2026, we're sharing a wave of Dataflow Gen2 announcements designed to help every data team move faster—from self-service analysts to centralized data engineering teams. Dataflow Gen2 brings **low-code data transformation** to Microsoft Fabric with the power of Power Query, enabling you to ingest, shape, and land data into Fabric destinations with built-in scheduling, reuse, and governance through the Fabric workspace experience.

Preview Only Steps

Power Query | Draft saved | Search (Alt + Q)

Home Transform Add column View Help

Table.FirstN("#Navigation 1", 10)

This data preview uses preview-only steps. Results may differ when running the dataflow. [Learn more](#)

	CustomerID	CompanyName	ContactName	ContactTitle	Address
1	ALFKI	Alfreds Futterkiste	Maria Anders	Sales Representative	Obere Str. 57
2	ANATR	Ana Trujillo Emparedados y helados	Ana Trujillo	Owner	Avda. de la Const
3	ANTON	Antonio Moreno Taquería	Antonio Moreno	Owner	Mataderos 2312
4	AROUT	Around the Horn	Thomas Hardy	Sales Representative	120 Hanover Sq.
5	BERGS	Berglunds snabbköp	Christina Berglund	Order Administrator	Berguvsvägen 8
6	BLAUS	Blauer See Delikatessen	Hanna Moos	Sales Repre	
7	BLONP	Blondesddsl père et fils	Frédérique Citeaux	Marketing	
8	BOLID	Bóldo Comidas preparadas	Martin Sommer	Owner	
9	BONAP	Bon app'	Laurence Lebihan	Owner	
10	BOTTM	Bottom-Dollar Markets	Elizabeth Lincoln	Accounting	

Queries [1]

Query settings

Properties

Name: Customers

Applied steps

- Source
- Navigation 1
- Kept top rows
- Filtered rows
- Removed other col...
- Added custom

destination: Lakehouse

Enable only in previews

Properties...

Edit settings

Explain this step

Rename

Delete

Delete until end

Insert step after

Move before

Move after

Extract previous...

View data source query

View query plan

Variable Libraries in Dataflow Gen2

Dataflow Gen2: Variable Library integration in Microsoft Fabric (Generally Available)

Data Factory



March 23, 2026 by [Miguel Escobar](#)  2,462 Views


 Share

 Tweet

 Like

If you haven't already, check out Arun Ulag's hero blog ["FabCon and SQLCon 2026: Unifying databases and Fabric on a single, complete platform"](#) for a complete look at all of our FabCon and SQLCon announcements across both Fabric and our database offerings.

Variable Libraries - Production workloads



AI-Powered Transforms (Prompt)

 **AI Prompt (Preview)** 

Add a column using an AI prompt

 If you don't already have a FabricAI connection, we'll try to create one for you and set its privacy level to Organizational. [Learn more](#)

New column name *

Create a review

Enter your prompt

Using the country and city, write a clear message in their native language inviting this customer to our annual Fabric Conference in Atlanta this year

Select columns for context *

- ☒ City
- ☒ Region
- ☐ PostalCode
- ☒ Country
- ☐ Phone
- ☐ Fax
- ☐ Orders


Data Factory MCP Server

Data Factory MCP Server

Tartalom

Fejezet

Lista



Data Factory MCP Server

Tartalom

Fejezet

Lista

Features

- **Azure AD Authentication:** Interactive and service principal authentication
- **Gateway Management:** List, retrieve, and create Microsoft Fabric gateways (including VNet gateways)
- **Connection Management:** List, retrieve, and create Microsoft Fabric connections (cloud, on-premises, and VNet)
- **Workspace Management:** List and retrieve Microsoft Fabric workspaces
- **Dataflow Management:** List, create, and retrieve Microsoft Fabric dataflows
- **Pipeline Management:** List, create, update, run, monitor, and schedule Microsoft Fabric pipelines
- **Copy Job Management:** List, create, update, run, monitor, and schedule Microsoft Fabric copy jobs
- **Capacity Management:** List and retrieve Microsoft Fabric capacities
- **Microsoft Fabric Integration:** Support for on-premises, personal, and virtual network gateways
-  **NuGet Distribution:** Available as a NuGet package for easy integration
-  **MCP Protocol:** Built using the official MCP C# SDK

Platform

Fabric Runtime 2.0

Fabric Runtime 2.0 (Preview)

Tartalom

Fejezet

Fabric Runtime 2.0 (Preview)

Announcements

Apache Spark

Data Engineering

Data Science



March 19, 2026 by [Arshad Ali](#) 7,352 Views

Share

Tweet

Like

If you haven't already, check out Arun Ulag's hero blog "[FabCon and SQLCon 2026: Unifying databases and Fabric on a single, complete platform](#)" for a complete look at all of our FabCon and SQLCon announcements across both Fabric and our database offerings.

Fabric Runtime provides seamless integration within the Microsoft Fabric ecosystem, delivering a powerful environment for data engineering and data science projects built on Apache Spark.

The next-generation runtime is purpose-built for large-scale data computations in Microsoft Fabric and introduces key features and components that mark a significant leap forward in enabling scalable analytics and advanced workloads.

Fabric Runtime 2.0 (Preview)

Tartalom

Fejezet

Workspace settings

Search

- General
- Workspace type
- Azure connections
- System storage
- Git integration
- OneLake
- Workspace identity
- Outbound networking

Spark settings

Configure and manage settings for Spark workloads and the default environment for the workspace.

Optimize for your use case

Answer a few quick questions and get a recommended compute configuration to optimize your resource usage.

Get started

Pool **Environment** Jobs High concurrency

Set default environment

The default environment will provide Spark properties, libraries, and settings for notebooks and Spark job definitions in this workspace select a different environment. [Learn more about Set default environment](#)

Inbound networking

Encryption

Monitoring

Power BI

Delegated Settings

Extension API

playground

Data
Engineering/Science

Spark settings

ML tracking

Data Factory

Data Warehouse

Runtime

Runtime version

Runtime version defines which version of Spark your Spark pool will use. [Learn more about Runtime version](#)

2.0 Public Preview (Spark 4.0, Delta 4.0)

1.2 (Spark 3.4, Delta 2.4)

1.3 (Spark 3.5, Delta 3.2)

2.0 Public Preview (Spark 4.0, Delta 4.0)

Materialized Lake Views

Materialized Lake Views in Microsoft Fabric (Generally Available)

[Announcements](#)[Data Engineering](#)[Fabric platform](#)[Lakehouse](#)[Microsoft Fabric](#)[Roadmap](#)

March 18, 2026 by [Balaji Sankaran](#)

👁 9,510 Views

[in Share](#)[🐦 Tweet](#)[👍 Like](#)

If you haven't already, check out Arun Ulag's hero blog "[FabCon and SQLCon 2026: Unifying databases and Fabric on a single, complete platform](#)" for a complete look at all of our FabCon and SQLCon announcements across both Fabric and our database offerings.

Tartalom

Fejezet



Fabric Notebook API-k

Fabric Notebook Public APIs (Generally Available)

[Announcements](#)[Data Engineering](#)[Data Science](#)

March 20, 2026 by [Jene Zhang](#)

👁 15,896 Views

[in Share](#)[🐦 Tweet](#)[👍 Like](#)

If you haven't already, check out Arun Ulag's hero blog "[FabCon and SQLCon 2026: Unifying databases and Fabric on a single, complete platform](#)" for a complete look at all of our FabCon and SQLCon announcements across both Fabric and our database offerings.

Fabric Notebook Public APIs (Generally Available), enabling data engineers and data scientists to programmatically manage and execute notebooks as part of modern, automated data workflows.

Fabric Notebook Public APIs

- Fabric notebooks addressable through public APIs
- Full CRUD support
- On-demand execution
- Exit values for better orchestration
- Service principal authentication supported

Tartalom

Fejezet

Fabric CLI

Fabric CLI v1.5 is here (Generally Available)

Microsoft Fabric



March 24, 2026 by [Hasan Abo Shally](#) 1,133 Views

Share

Tweet

Like

Co-author: Jeremy Hoover

If you haven't already, check out Arun Ulag's hero blog "[FabCon and SQLCon 2026: Unifying databases and Fabric on a single, complete platform](#)" for a complete look at all of our FabCon and SQLCon announcements across both Fabric and our database offerings.

Platform

Since the [open-source launch](#) in October, the Fabric CLI has gained one-command deployments, first-class Power BI support, an AI agent execution layer, interactive REPL mode, Python 3.13 support, and contributions from community members whose code ships in the CLI you install today.

Fabric CLI in Azure DevOps (Preview)

Tartalom

Fejezet

Fabric CLI in Azure DevOps: automation without friction (Preview)

Fabric platform

Microsoft Fabric



March 23, 2026 by [Evelina Alroy-Brin](#)

👁 2,357 Views

Share

Tweet

Like

If you haven't already, check out Arun Ulag's hero blog "[FabCon and SQLCon 2026: Unifying databases and Fabric on a single, complete platform](#)" for a complete look at all of our FabCon and SQLCon announcements across both Fabric and our database offerings.

Modern analytics platforms demand modern DevOps practices. As organizations standardize Azure DevOps for CI/CD, they expect seamless integration with the services that power their data and analytics workloads. With Microsoft Fabric, that expectation is fully met.

Data Warehousing

Custom SQL Pools

Custom SQL Pools for Fabric Data Warehouse (Preview)

[Data Warehouse](#)[Microsoft Fabric](#)

March 19, 2026 by [Bret Myers](#)

👁 1,456 Views



Share



Tweet




Like

If you haven't already, check out Arun Ulag's hero blog "[FabCon and SQLCon 2026: Unifying databases and Fabric on a single, complete platform](#)" for a complete look at all of our FabCon and SQLCon announcements across both Fabric and our database offerings.

Custom SQL Pools for Microsoft Fabric Data Warehouse (Preview) is a new capability that gives workspace administrators fine-grained control over how SQL compute resources are allocated for SQL workloads in a Fabric Data Warehouse and SQL analytics endpoint.

Fabric Data Warehouse Pool Isolation



Custom SQL Pool Configuration

Tartalom

Fejezet

Workspace settings

Search

General

Workspace type

Azure connections

System storage

Git integration

OneLake

Workspace identity

Outbound networking

Inbound networking

Encryption

Monitoring

SQL pools (preview)

SQL pools divide your compute for query processing. Query runs are assigned to pools using a classifier. Choose the setup you want: built-in (automatic) or custom (you control the compute divisions and allocations) for all warehouses and SQL analytics endpoints in your workspace.

[Learn more about SQL pools](#)

SQL pool setup

Using custom setup

Revert to built-in setup

The built-in setup divides your compute evenly into two pools. Query runs are classified (assigned) to pools based on the query statement type (SELECT vs non-SELECT). To change the classifier (from "query statement type" to "query submitting application") and divide your compute differently, set up custom pools.

Edit

Workspace SKU is F32. 100% of compute resources are allocated.

Compute for Pipelines: 30%

Compute for PowerBI: 60%

Compute for Adhoc: 10%

Unused: 0%

<https://blog.fabric.microsoft.com/en-gb/blog/custom-sql-pools-for-fabric-data-warehouse-preview?ft=Data-warehouse:category>

dbt + Microsoft Fabric

dbt+ Microsoft Fabric: A strategic investment in the modern analytics stack

[Data Engineering](#)[Data Factory](#)[Data Warehouse](#)

March 24, 2026 by [Pradeep Srikakolapu](#) 3,747 Views

[in](#) Share[Twitter](#) Tweet[Facebook](#) Like

If you haven't already, check out Arun Ulag's hero blog "[FabCon and SQLCon 2026: Unifying databases and Fabric on a single, complete platform](#)" for a complete look at all of our FabCon and SQLCon announcements across both Fabric and our database offerings.


Co-Author: Abhishek Narain

Why dbt + Microsoft Fabric matters

Modern analytics teams are adopting open, SQL-first data transformation, robust CI/CD and governance, and seamless integration across lakehouse and warehouse platforms. **dbt** is now the standard for analytics engineering, while **Microsoft Fabric** unifies data engineering, science, warehousing, and BI in OneLake.

dbt: The Standard for SQL Data Transformation

Trusted worldwide to model, test, and document data with SQL



Industry momentum

100,000+

Active community
members

60,000+

Teams
transforming data

Core capabilities

Model reliable pipelines with modular,
version-controlled SQL

Test data quality through built-in,
automated checks

Document and govern data products
at enterprise scale

Tartalom

Fejezet

Microsoft + dbt Labs

Coming soon

Microsoft
Fabric



dbt Fusion
in Fabric



dbt Labs

SQL Alerts

Alerts and actions

Microsoft Fabric Data Warehouse provides operational intelligence closer to the data by integrating SQL queries with [Fabric Activator rules](#). Traditionally, identifying an issue in query results is only the first step. Teams then need separate systems or manually follow up to notify the right people and act. With this integration, Data Warehouse makes it possible to define rules directly from SQL query outputs, so changes in data can trigger alerts and downstream actions automatically.

This unlocks a simpler way to monitor business-critical conditions using familiar SQL workflows. Teams can create queries that detect scenarios such as SLA risks, failed processes, unusual trends, or threshold breaches, then attach rules that evaluate results continuously and respond in real time. The result is a more proactive analytics experience, where insights move beyond the warehouse and are acted on immediately.

SQL Alerts

Home

Copilot

Create

Browse

OneLake catalog

Apps

Metrics

Monitor

Learn

Real-Time

Workspaces

SalesContos

WorldWide Sales

...

Power BI

Home

Reporting

Security

Management

Help

Get data

New SQL query

SQL templates

Query activity

Model layouts

Download SQL database project

Open in

New API for GraphQL

Add to data agent

Copilot

SQL query 8

SQL query 13

Run

Save as view

Explain query

Fix query errors

Preview

Copilot uses AI. Mistakes can happen. Verify code suggestions before running them. [Review terms](#)

-- Step 3: Use SUM([LatestRecordedPopulation]) to get total population per city/state

-- Step 4: Filter groups where total population exceeds 2,500,000

-- Step 5: Sort results by total population descending

SELECT

[City],

[StateProvince],

SUM([LatestRecordedPopulation]) AS [TotalPopulation]

FROM

[dbo].[dimension_city]

GROUP BY

[City],

[StateProvince]

HAVING

SUM([LatestRecordedPopulation]) > 2500000

ORDER BY

[TotalPopulation] DESC;

Messages

Results

Grid

Table

Chart

Download

Search

ABC	City	ABC	StateProvince	123	TotalPopulation
1	New York		New York		24525399
2	Los Angeles		California		15170484
3	Chicago		Illinois		8086794
4	Brooklyn		New York		7696905
5	Queens		New York		6818313
6	Houston		Texas		6298353
7	San Diego		California		5229608
8	Manhattan		New York		4857270
9	Philadelphia		Pennsylvania		4578018
10	Phoenix		Arizona		4336896
11	The Bronx		New York		4225419
12	San Antonio		Texas		3982221

Succeeded (1 sec 822 ms)

Copilot completions: On

Copilot completions: Ready

Copilot

Preview

result set. For example, you could use WHERE to only include rows where City = 'Seattle'.

- HAVING filters groups after the GROUP BY operation and after any aggregate functions (like SUM or AVG) have been calculated. It applies to the grouped results, not the raw rows. Use HAVING when you want to restrict results based on the outcome of aggregate functions, for example, only showing groups where SUM(Population) > 2,500,000.

You can use both WHERE and HAVING in the same query—WHERE limits rows before aggregation, HAVING limits the groups after aggregation.

Would you like to see an example using both WHERE and HAVING in the same query?

Is this response helpful?

View prompts

Ask me anything about this warehouse, or tell me how I can help.

0 / 500

AI-generated content can have mistakes. Make sure it's

Columns: 3 Rows: 18

<https://blog.fabric.microsoft.com/en-us/blog/fabric-march-2026-feature-summary>

Alert rules

HomeSecurityManagementHelp

SQL query 44

SQL query 43

RunSave asSave as viewCopy queryCreate ruleExplain queryFix query errors

Preview

Copilot uses AI. Mistakes can happen. Verify code suggestions before running them. [Review terms](#)

1

-- Data-driven rule: cities with >30% revenue jump vs previous day

2

WITH DailyRevenue AS (

3

SELECT

4

c.Date,

5

city.City,

6

SUM(fs.TotalIncludingTax) AS Revenue

7

FROM fact_sale fs

8

JOIN dimension_city city

9

ON fs.CityKey = city.CityKey

10

JOIN dimension_date c

11

ON fs.InvoiceDateKey = c.Date

12

GROUP BY c.Date, city.City

13

MessagesResults

Date

ABC City

e^x Revenue

e^x PrevRevenue

1

2000-11-30 00:00:00.0000000

Sherwin

747349.12

313388.80

2

2000-11-30 00:00:00.0000000

Severna Park

174956.40

1076.40

3

2000-11-30 00:00:00.0000000

Sea Island

120777.60

38345.60

4

2000-11-30 00:00:00.0000000

Frankewing

61272.00

14490.00

5

2000-11-30 00:00:00.0000000

Bourbonnais

51152.00

1150.00

6

2000-11-30 00:00:00.0000000

Moquino

44012.80

12217.60

7

2000-11-30 00:00:00.0000000

Alden Bridge

31960.80

7948.80

8

2000-11-30 00:00:00.0000000

Toughkenamon

20240.00

883.20

9

2000-11-30 00:00:00.0000000

Oval

6844.80

478.40

AutoSave: On (Last saved: 42m ago)

Succeeded (7 sec 927 ms)

Copilot completions: On

Copilot completions: Ready

Columns: 4 Rows: 1961

Rules pane

Rules

 My workspace /  WorldWideSalesActiva... [Edit](#)

More than 30 percent jump in a day New  Running 

 Edit
 Delete
 Open in Activator

Sample alert

Tartalom

Fejezet

John Doe via Fabric Activator 5:03 PM



Revenue jump for Sherwin more than 30%

Revenue for @City jumped more than 30% from yesterday.

Manage rule

View details

Activation Time:

Mar 12, 2026 21:03:04 UTC

Last modified by John Doe on Mar 12, 2026 20:08:57 UTC, UID: 75c66f8f

If you have any questions or feedback on Activator, please share it on <https://aka.ms/ActivatorCommunity>.

AI és Copilot

AI Functions

AI functions for Data Engineering

Tartalom

Fejezet

Announcing AI functions for seamless data engineering with GenAI

AI

Announcements

Apache Spark

Community

Data Engineering

Data Science

Microsoft Fabric



March 10, 2025 by [Eren Orbey](#)



98,008 Views

Share

Tweet


Like

If you saw the keynote at FabCon last fall, you may remember [an early demo of AI functions in Fabric](#), a new feature that makes it easy to apply LLM-powered transformations to your OneLake data. We're thrilled to announce that [AI functions are now in preview](#).

With AI functions, you can harness the power of GenAI for summarization, classification, text generation, and so much more—**all with a single line of code**. It's seamless to incorporate AI functions as part of data-science and data-engineering workflows, whether you're transforming Spark or pandas DataFrames. There is no complex setup, no tricky syntax, and, hopefully, no hassle.

<https://blog.fabric.microsoft.com/en-US/blog/announcing-ai-functions-for-easy-llm-powered-data-enrichment/>

AI functions for Data Engineering



The screenshot shows a JupyterLab environment with a notebook titled "Getting started with AI functions". The interface includes a top bar with tabs for Home, Edit, Run, and View, and buttons for Comments, History, Develop, and Share. Below the top bar is a toolbar with icons for file operations, a Run all button, a Standard session indicator, a PySpark (Python) environment selector, and a Workspace default dropdown. The left sidebar shows the Explorer view with a file tree. The main area displays the notebook content, which includes a code cell with the following Python commands:

```
1 # Install fixed version of packages
2 %pip install -q openai==1.30
3 %pip install -q --force-reinstall httpx==0.27.0
4
5 # Install latest version of SynapseML-core
6 %pip install -q --force-reinstall https://mmlspark.blob.core.windows.net/pip/1.0.9/synapseml_core-1.0.9-py2.py3-none-
7
8 # Install SynapseML-Internal .whl with AI functions library from blob storage:
9 %pip install -q --force-reinstall https://mmlspark.blob.core.windows.net/pip/1.0.10.0-spark3.4-5-a5d50c90-SNAPSHOT/sy
```

Below the code cell, a status bar indicates that the command was executed successfully in 30 seconds and 194 milliseconds by Eren Orbey on 9:18:45 PM, 3/03/25. The output is hidden.

AI functions for Data Engineering

```
1 df["Translated_Text"] = df["Foreign_Text"].ai.translate("English")
2 display(df)
```

✓ Executed

PySpark (Python) ▾

Apply LLM-powered operations in a single line of code ...



Table		Chart	Download ▾	Showing rows 1 - 5	Inspect	Search
ABC Foreign_Text	ABC Translated_Text					
1 Agent: Välkommen till Contoso Outdoors. Hur kan jag ...	Agent: Welcome to Contoso Outdoors. How can I assist you today? Customer: Hello, I am really di...					
2 Agent: Tack för att du ringer Contoso Outdoors. Det hä...	Agent: Thank you for calling Contoso Outdoors. This is Elsa speaking. How can I assist you today?...					
3 Agent: Tack för att du ringer Contoso Outdoors. Hur ka...	Agent: Thank you for calling Contoso Outdoors. How can I assist you today? Customer: Hello, this...					
4 Agent: Det här är Bjorn som talar. Tack för att du ringer ...	Agent: This is Bjorn speaking. Thank you for calling Contoso Outdoors. How can I assist you today...					
5 Agent: Tack för att du ringer Contoso Outdoors. Det hä...	Agent: Thank you for calling Contoso Outdoors. This is Elsa speaking. How can I assist you today?...					

Choose from eight AI functions out of the box, including custom text generation ...

```
1 df["Summarized_Text"] = df["Translated_Text"].ai.
```

- analyze_sentiment
- classify
- extract
- fix_grammar
- gen
- similarity
- stats
- summarize**
- translate

PySpark (Python) ▾

AI functions for Data Engineering

AI function	Description
<u>ai.similarity</u>	Compare the meaning of input text with a single common text value, or with corresponding text values in another column.
<u>ai.classify</u>	Classify input text values according to labels you choose.
<u>ai.analyze_sentiment</u>	Identify the emotional state expressed by input text.
<u>ai.extract</u>	Find and extract specific types of information from input text, for example locations or names.
<u>ai.fix_grammar</u>	Correct the spelling, grammar, and punctuation of input text.
<u>ai.summarize</u>	Get summaries of input text.
<u>ai.translate</u>	Translate input text into another language.
<u>ai.generate_response</u>	Generate responses based on your own instructions.

Analyze unstructured text using T-SQL AI functions (Preview)

Microsoft Fabric Data Warehouse extends modern analytics beyond structured and semi-structured data by introducing built-in AI functions for working directly with the unstructured text. Traditionally, processing free-form content such as notes, logs, or comments requires external services or complex pipelines. With these new capabilities, Fabric Data Warehouse enables text extraction, classification, sentiment analysis, and transformation directly in T-SQL language, allowing data engineers and analysts to keep AI-driven text processing inside the warehouse.

The new AI functions simplify common text analytics scenarios using familiar SQL patterns. You can extract structured insights from unstructured text, analyze sentiment in feedback or messages, and classify content such as application logs or incident reports using contextual understanding rather than fragile rules or expressions. Fabric Data Warehouse also supports text transformation scenarios, including summarization, grammar correction, and translation, making it easier to standardize and enrich text data as part of existing data preparation workflows.

T-SQL AI functions

Tartalom

Fejezet

SQL query 15

Run

Save as

Save as view

Copy query

Create rule

Explain query

Fix query errors

```
2 SELECT
3     c.CommentId,
4     ai_analyze_sentiment(c.CommentText) AS Sentiment,
5     ai_classify(c.CommentText, 'praise', 'complaint', 'question', 'suggestion' ) AS Label,
6     ai_extract( c.CommentText, 'topic', 'user_intent', 'requested_action' ) AS ExtractedInfo
7 FROM dbo.Comments AS c;
```

Messages

Results

	123 CommentId	ABC Sentiment	ABC Label	ABC ExtractedInfo
1	1	positive	praise	{"topic":"dashboard perform...
2	2	negative	complaint	{"topic":"billing","user_intent"...
3	3	neutral	question	{"topic":"permission model","..."

Copilot

Copilot for Data Engineering

Introducing the updated Copilot for data engineering and data science (Preview)

Data Engineering

Microsoft Fabric



March 19, 2026 by [Deepa Rajagopal](#)  5,185 Views

 Share

 Tweet

 Like

If you haven't already, check out Arun Ulag's hero blog "[FabCon and SQLCon 2026: Unifying databases and Fabric on a single, complete platform](#)" for a complete look at all of our FabCon and SQLCon announcements across both Fabric and our database offerings.

Copilot for Data Engineering

Tartalom

Fejezet

The screenshot displays the Copilot for Data Engineering interface. The main editor shows a code cell with a SparkSQL query and comments. A red box highlights the "Fix with Copilot" button. The output area shows an `AnalysisException` with a traceback. The Copilot sidebar on the right provides an "Error Summary" and a "Suggested Code Fix".

```
69     "product_id", "product_name", "product_status",
70     "ending_quantity", "sales_velocity_daily", "days_of_supply", "revenue_at_risk"
71   ).orderBy(F.desc("revenue_at_risk"))
72 )
73
74 # Key Fixes:
75 # - Switched from `spark.read.table('dbo...')` to loading with abfss absolute paths and Delta format, matching the actual location and format of a
76 # - This avoids AnalysisException since the tables exist in the RetailLH lakehouse, not the current default lakehouse.
```

[1] 20 sec - Session ready in 11 sec 145 ms. Command executed in 9 sec 820 ms by Deepa Rajagopal on 3/6/2025, 8:28:29 AM

Fix with Copilot

Log

```
-----
AnalysisException                                Traceback (most recent call last)
Cell In[8], line 6
      3 from pyspark.sql.window import Window
      5 # 1. Load tables
----> 6 inventory = spark.read.table("dbo.inventory_history")
      7 sales = spark.read.table("dbo.sales_transactions")
      8 products = spark.read.table("dbo.products")

File /opt/spark/python/lib/pyspark.zip/pyspark/sql/readwriter.py:484, in DataFrameReader.table(self, tableName)
    450 def table(self, tableName: str) -> "DataFrame":
    451     """Returns the specified table as a :class:`DataFrame`.
    452
    453     .. versionadded:: 1.4.0
    454     (...)
    482     >>> _ = spark.sql("DROP TABLE tblA")
    483     """
--> 484     return self._df(self._jreader.table(tableName))

File ~/cluster-env/trident_env/lib/python3.11/site-packages/py4j/java_gateway.py:1322, in JavaMember.__call__(self, *args)
    1316 command = proto.CALL_COMMAND_NAME + \
    1317     self.command_header + \
```

Copilot Preview

Code cell updated successfully.

Error Summary

Error Type: SparkSQL Table or View Not Found

Root Cause: The code attempted to load tables using `spark.read.table("dbo.inventory_history")`, which looks for the table in the current default lakehouse. However, these tables are stored in the attached "RetailLH" lakehouse, so specifying their full abfss absolute paths is required.

Suggested Code Fix

```
Python

# Identify Product-Store Stockout Risk and Revenue
from pyspark.sql import functions as F
from pyspark.sql.window import Window

# Reference RetailLH tables using abfss paths from
abfss_prefix = "abfss://faa77a06-6d21-4945-a4ae-bac

inventory_path = abfss_prefix + "inventory_history"
sales_path = abfss_prefix + "sales_transactions"
```

1 cell changed

Keep Undo

Ask questions and work with this notebook

Agent

AI-generated content may be incorrect

Selected Cell 4 of 4 cells

Thank You

