Nyheter i Microsoft Fabric

Lasse Arnesen, Technology Strategist @ Microsoft


Good AI, needs
clean data

Since general availability

18 months ago

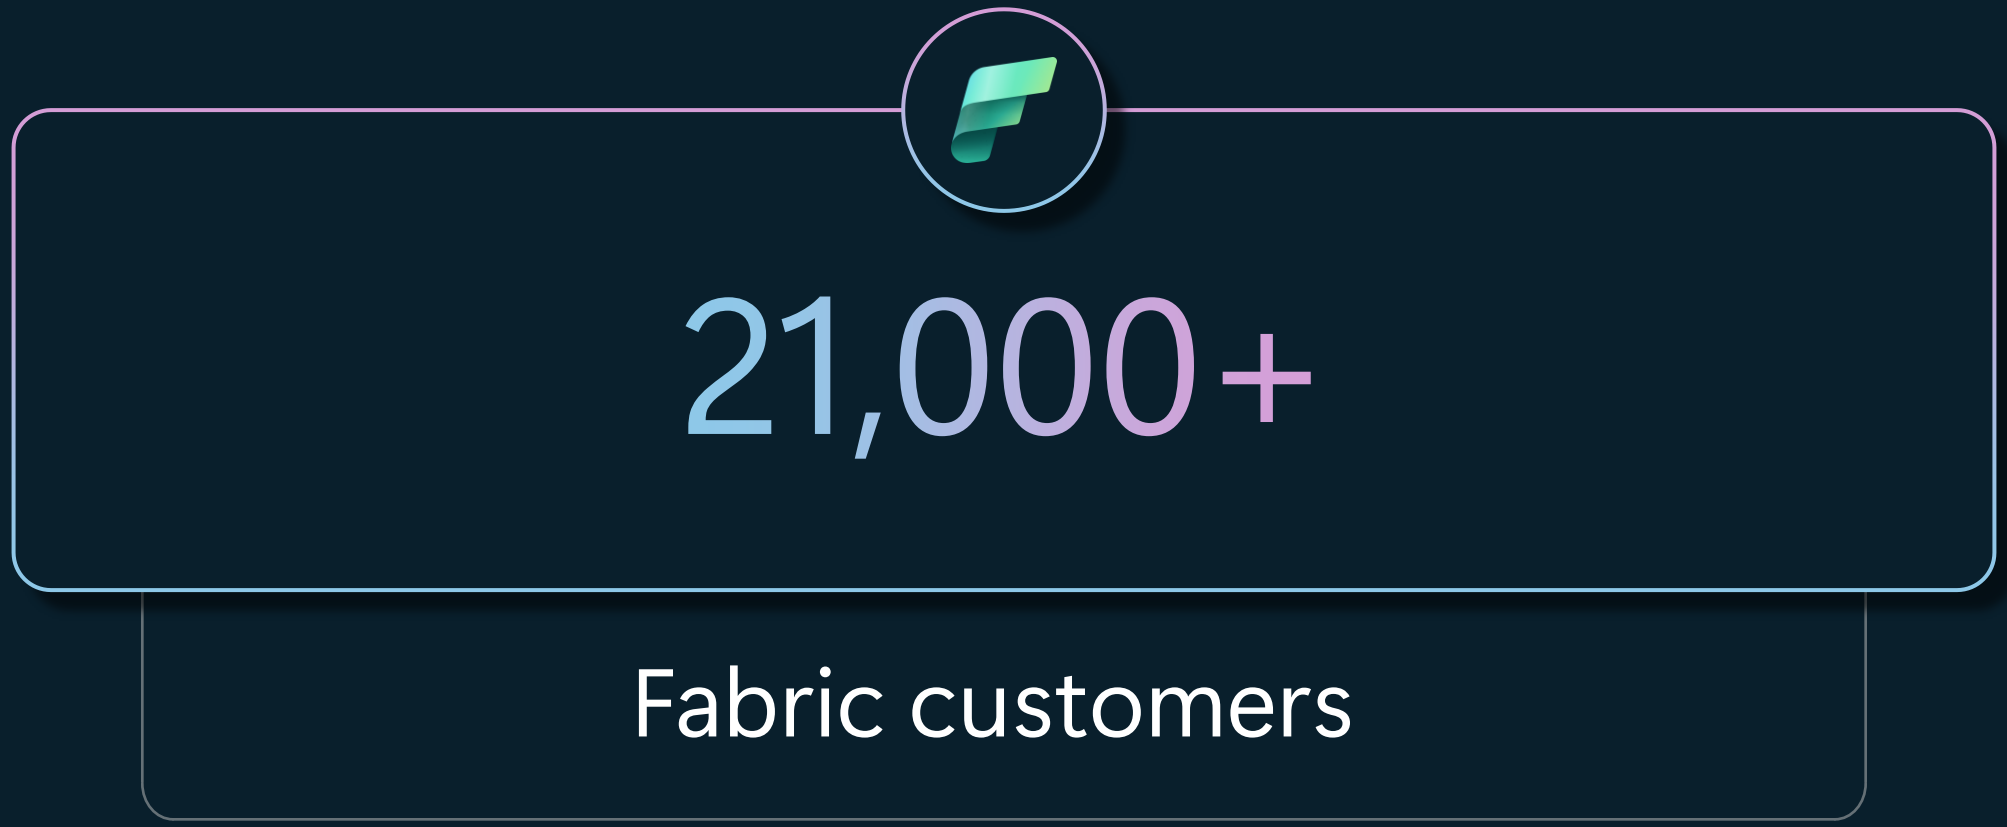
21,000+
Fabric customers

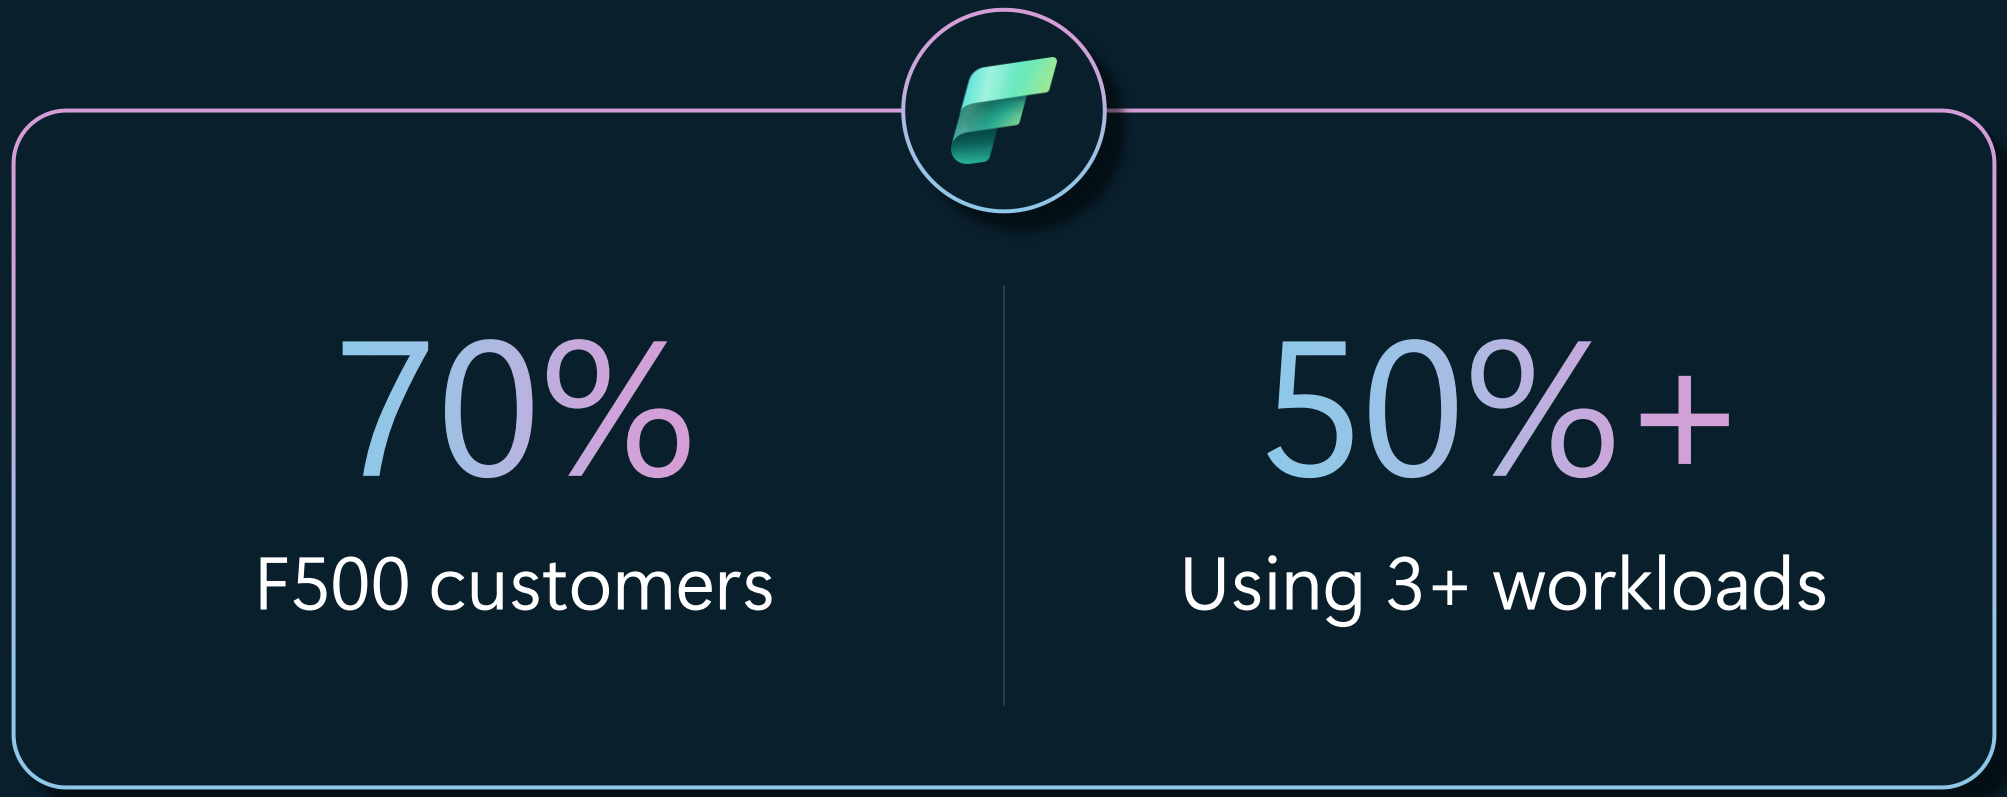
70%
F500 customers
50%+
Using 3+ workloads

Microsoft Fabric

The unified data platform for AI transformation

Data Factory

Analytics

Databases

Real-Time Intelligence

Power BI

Fabric Platform

AI

OneLake

Security & Governance

Fabric Automation & CI/CD Landscape

Git Integration	Variable library	fabric-cicd Library (OSS)
Fabric REST APIs	Command Line Interface	Deployment Pipelines
Sempy Labs Library (OSS)	Terraform Provider	.Net SDK & Workload Dev Kit

Fabric Tooling Layers

High Level Tools

Enhance usability and automation

- Terraform Provider
- fabric-cicd OSS Library
- Sempy Labs OSS Library

Abstraction Layer

Provide direct interaction with the APIs

- Command Line Interface
- .Net SDK & Workload Dev Kit

Foundation Layer

Serve as the base for all interactions

- Fabric REST APIs

Generally available

Fabric CLI is here

Announcing Fabric Command Line Interface is generally available!

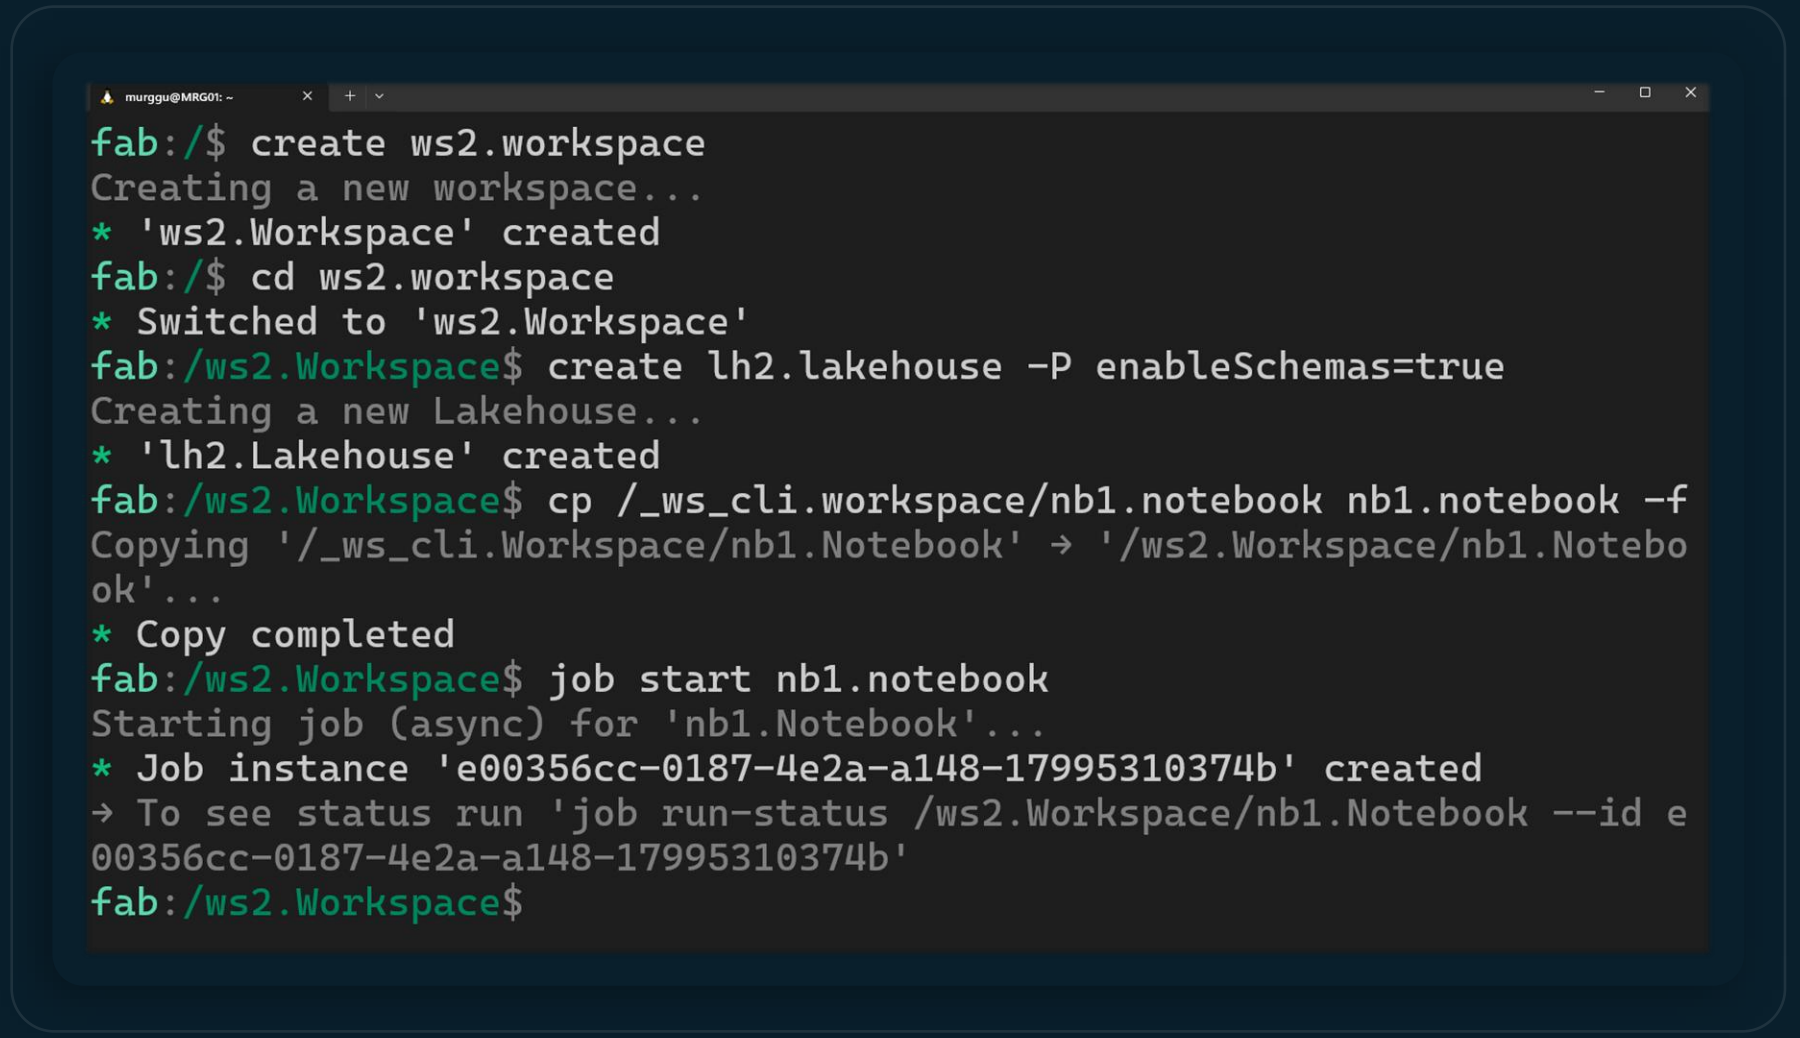

Fabric as a Filesystem

Fabric Hierarchy

Tenant

Workspace

Item

OneLake

Commands

ls/dir | pwd | cd | rm/del | ln/mklink

mkdir/create | cp/copy | mv/move

import | export | get | set | assign

unassign | exists | open | start | stop

acl | label | job | table

api | auth | config | desc

Git Integration

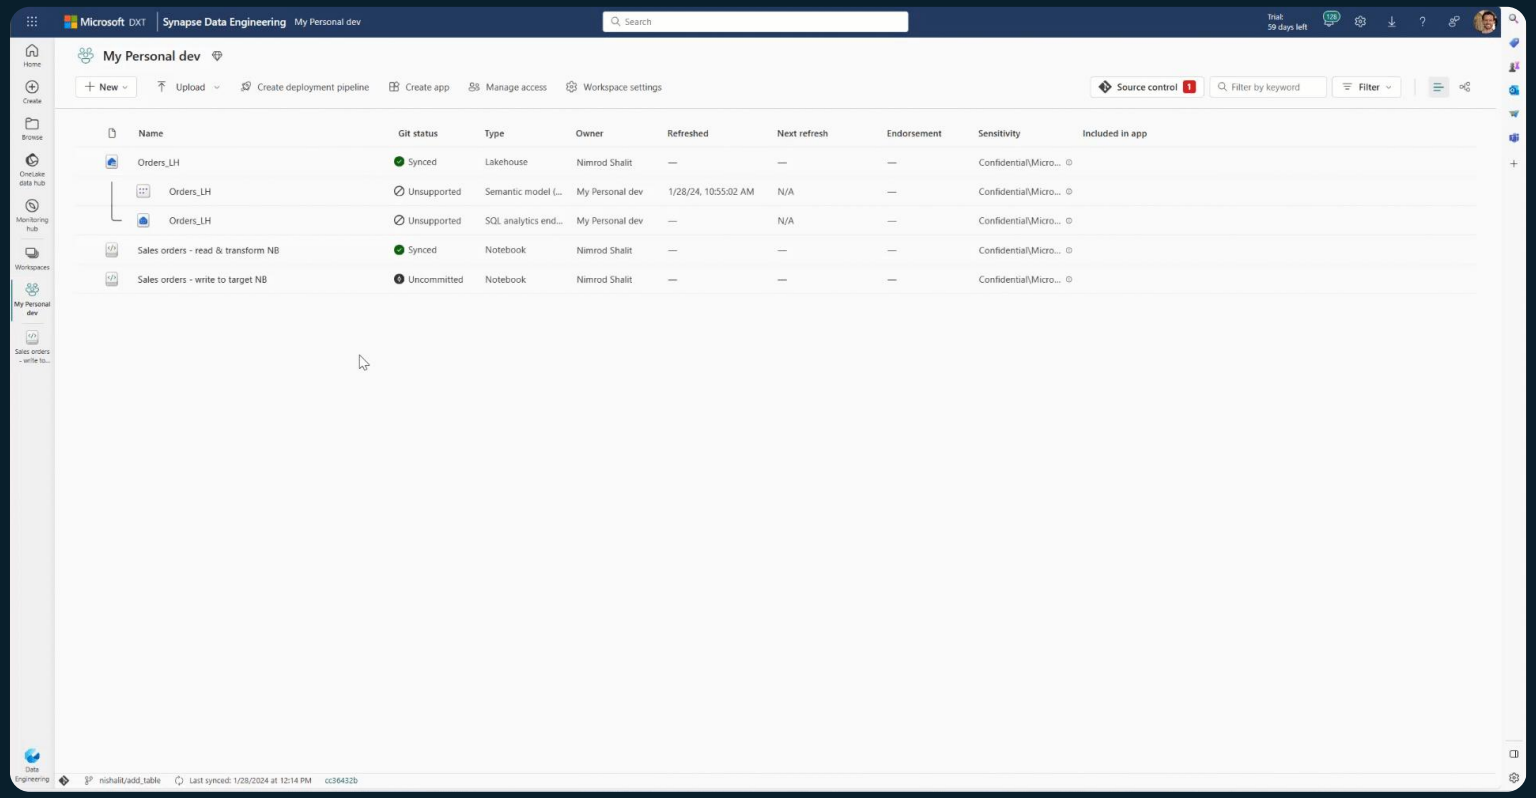

Sync a WS to a Git branch

Supports ~80% of Fabric items

Git providers – Azure DevOps, GitHub, GitHub Enterprise

Fabric APIs – REST APIs, Service principal support

Git Integration Enhancements

NEW!

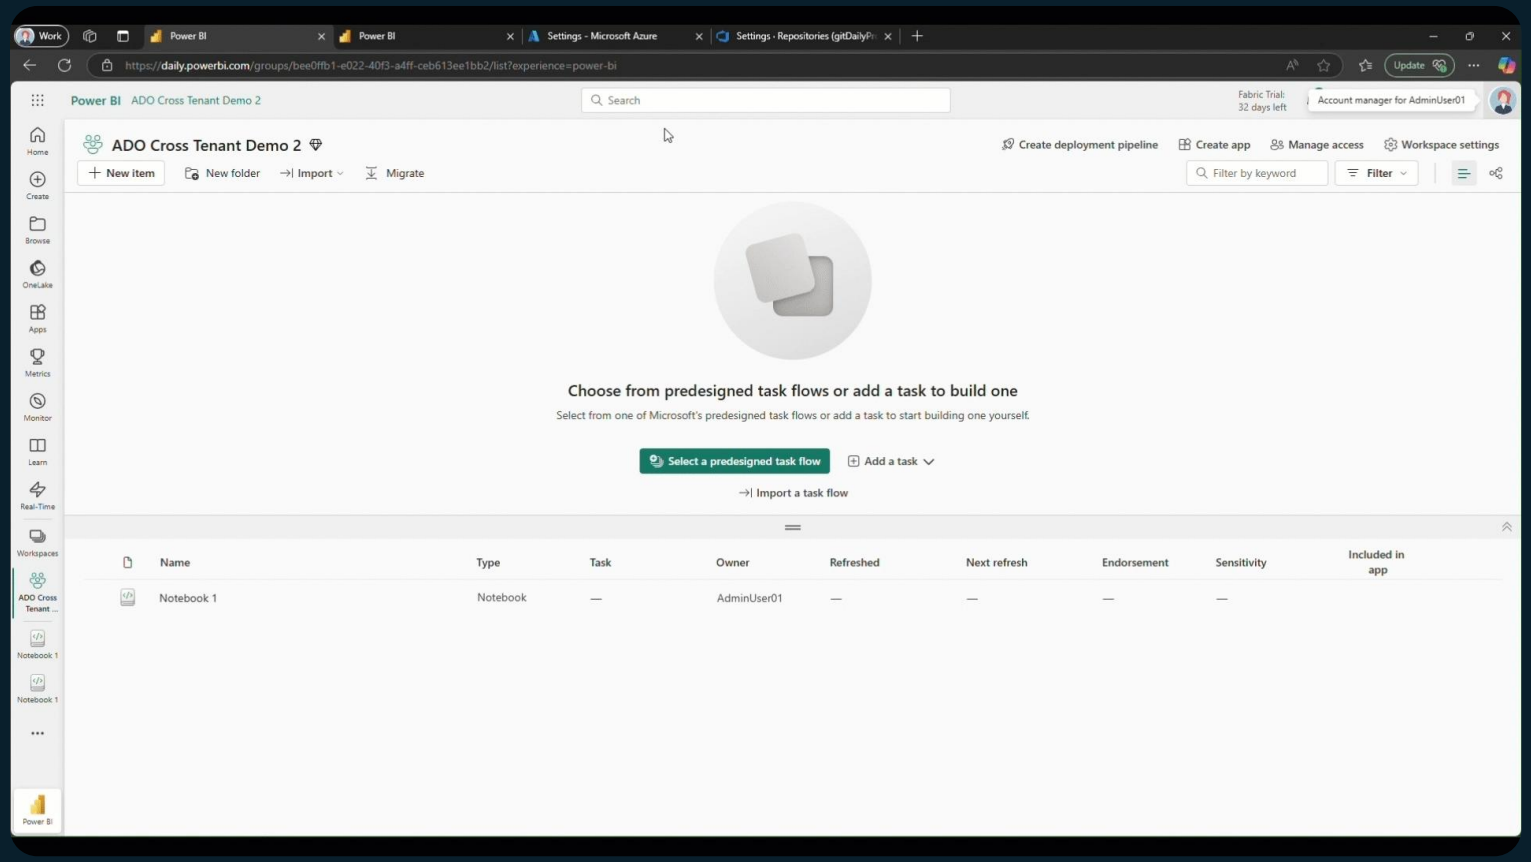

Service Principal support for Azure DevOps *(June)*

Cross-tenant Authentication with Azure DevOps *(June)*

Scheduler CI/CD – Part of an item definition when scheduler is configured *(June)*

Workspace folders retained in repos (*released)*

Variable Libraries

Public preview

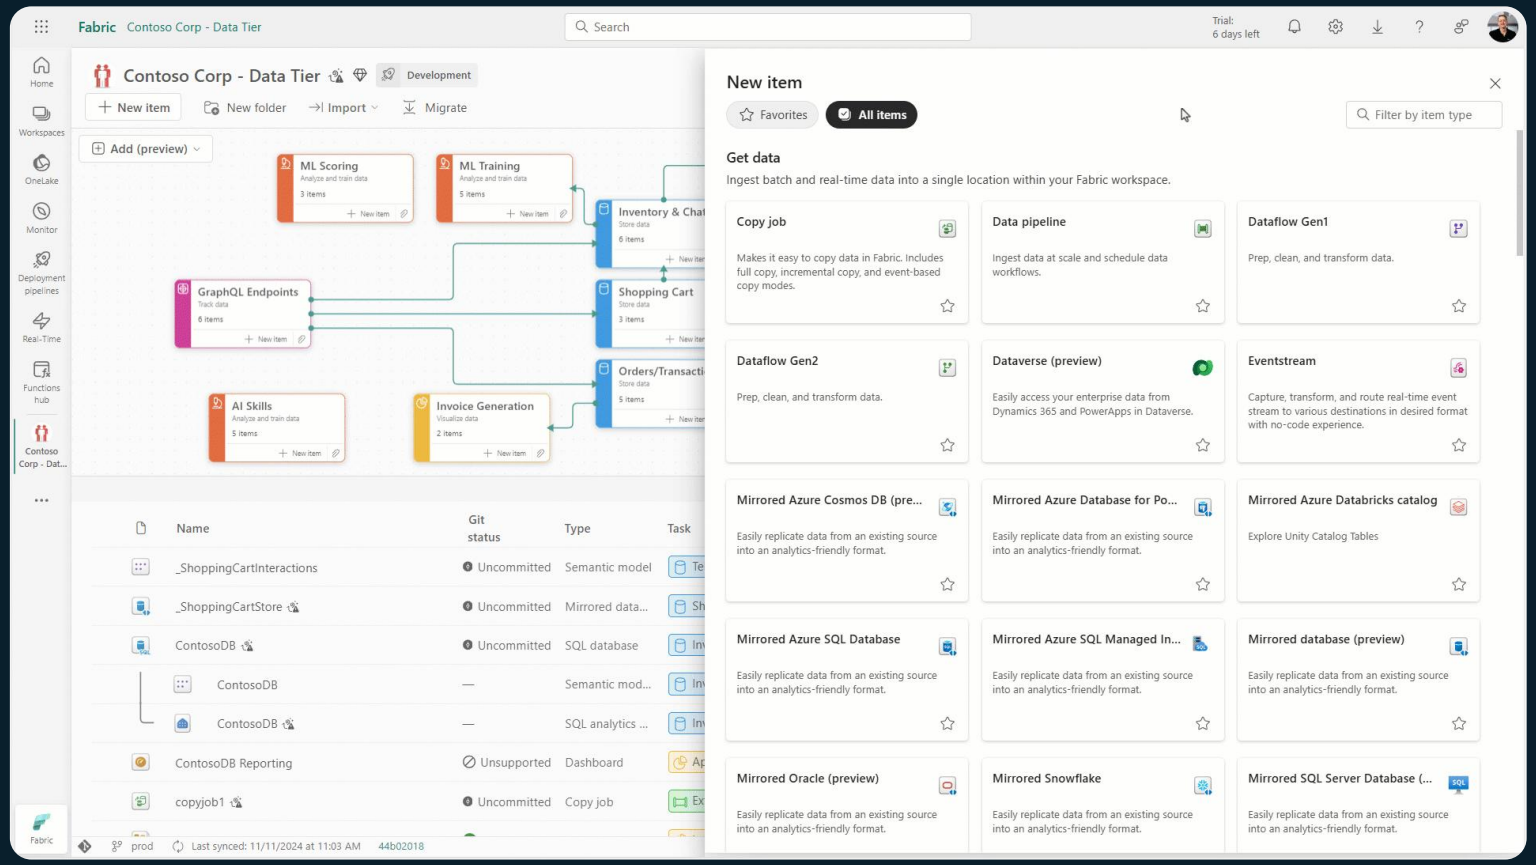

New Fabric item to define and manage variables at the workspace level

Create environment-specific configurations for different deployment stages

Centralized management with support for multiple variable types

Use in Data pipelines

fabric-cicd OSS Library

Public preview

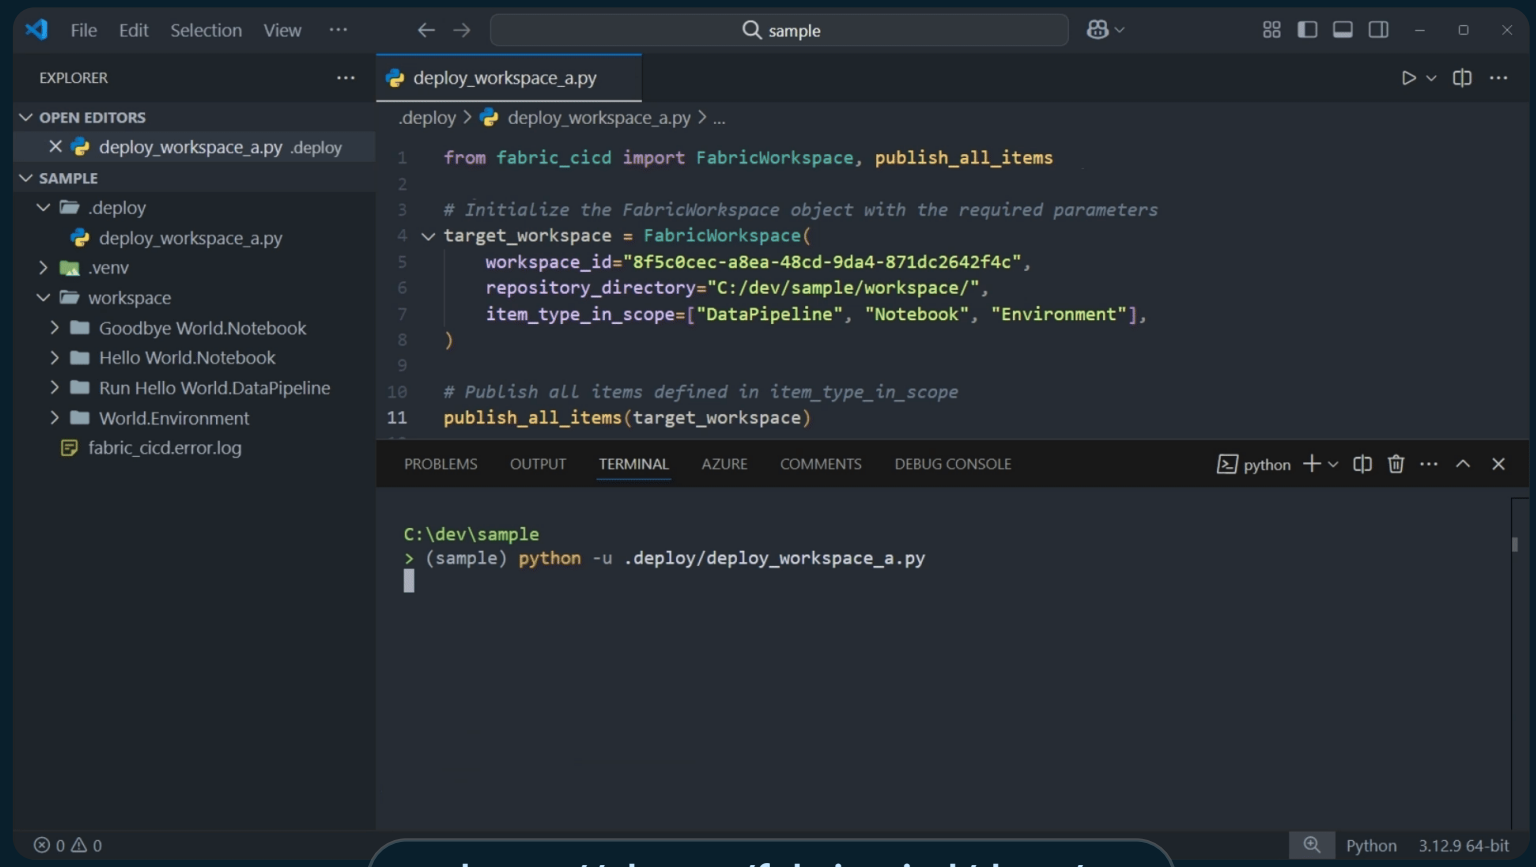

https://aka.ms/fabric-cicd/docs/

- Python Open-Source Library
- Deploy Fabric items with script-based deployments
- Integrates with Azure DevOps and GitHub Actions
- Create environment-specific parameterization for different deployments
- Goal to support all item types with APIs and Source Control

Sample Flow

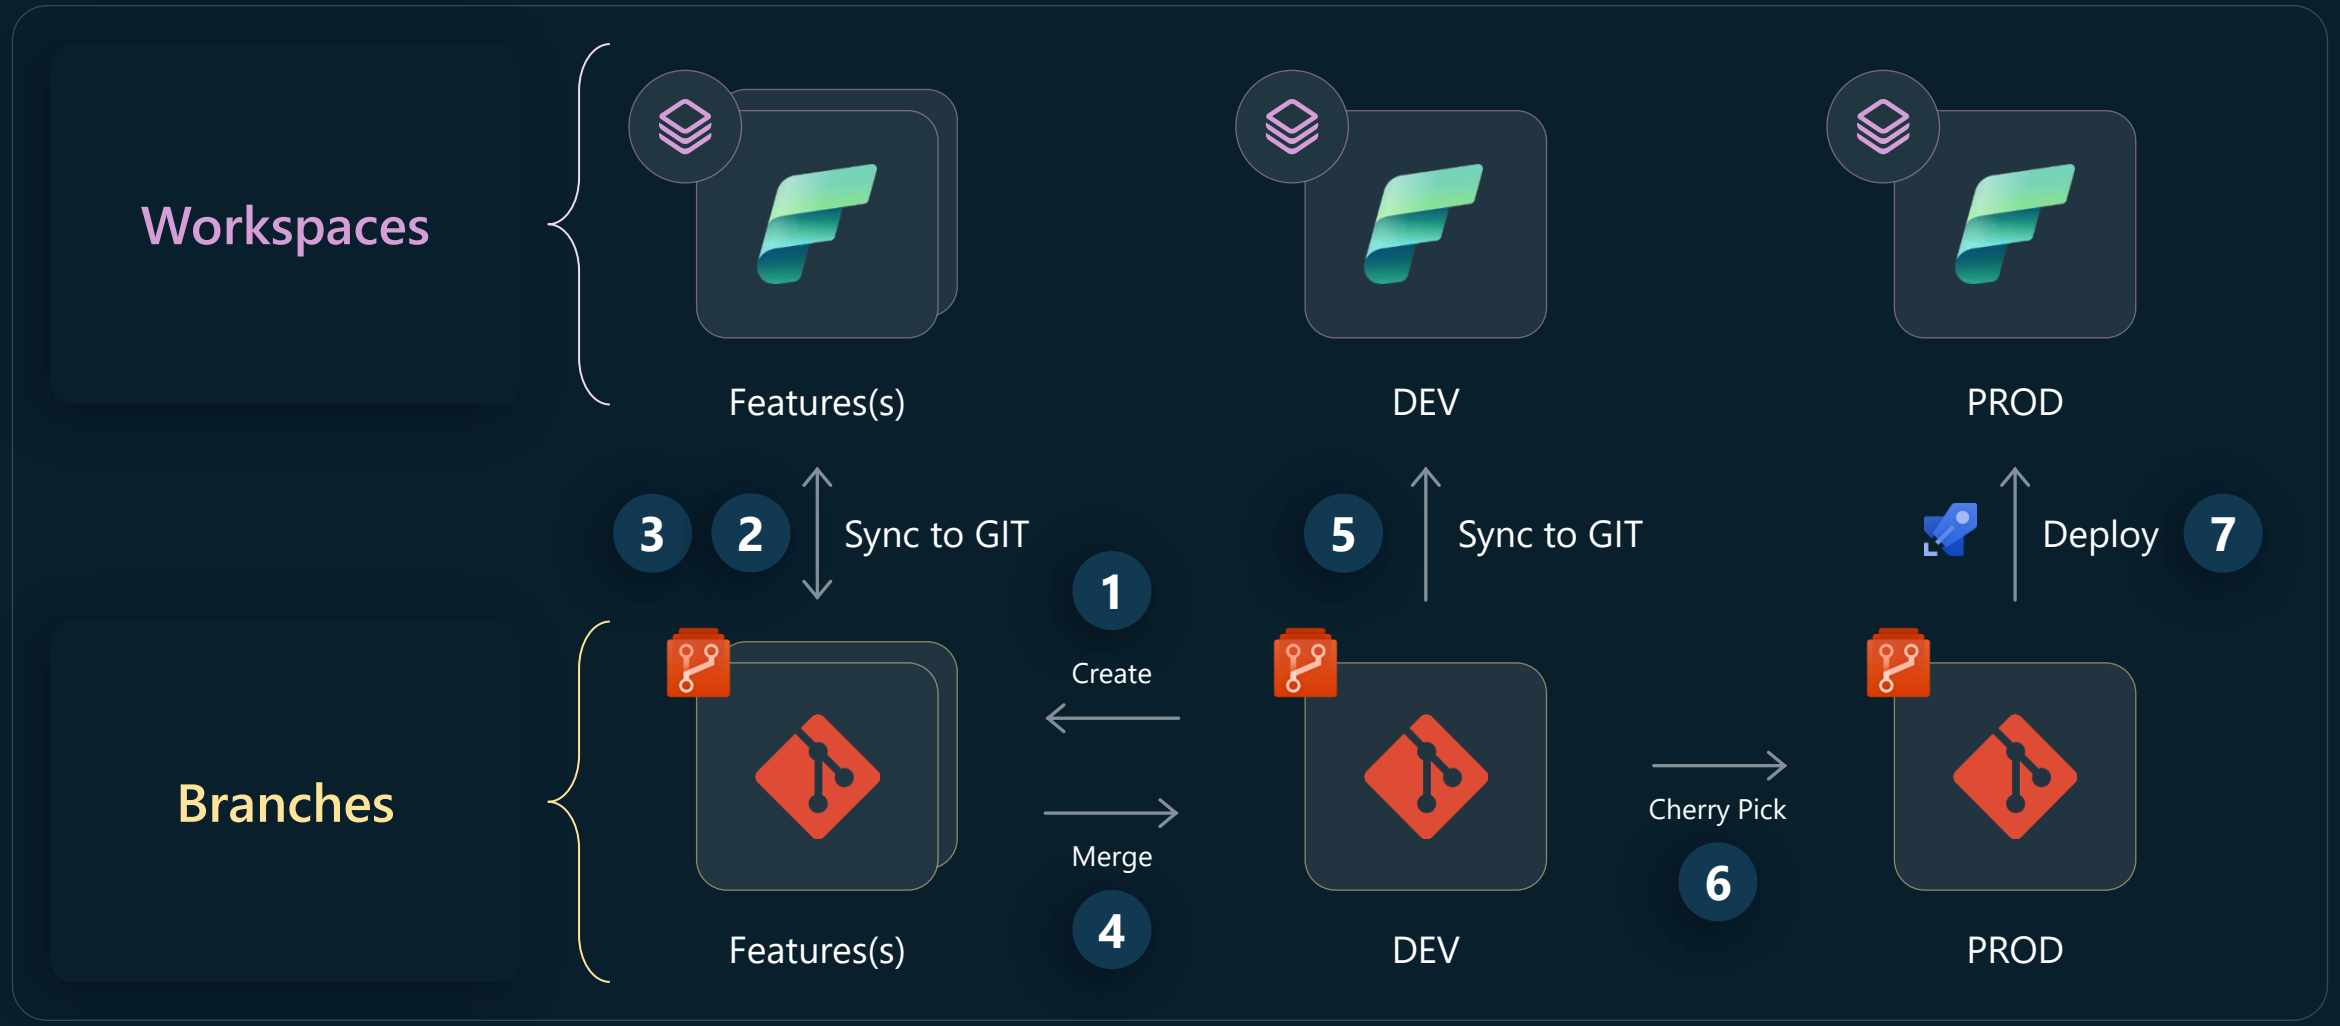
Workspaces
Features(s)
DEV
PROD
3
2
Sync to GIT
5
Sync to GIT
Deploy
7
1
Create
Branches
Merge
4
Cherry Pick
6
Features(s)
DEV
PROD

Announcing

Cosmos DB in Fabric

Public Preview

Cosmos DB in Fabric

Public preview

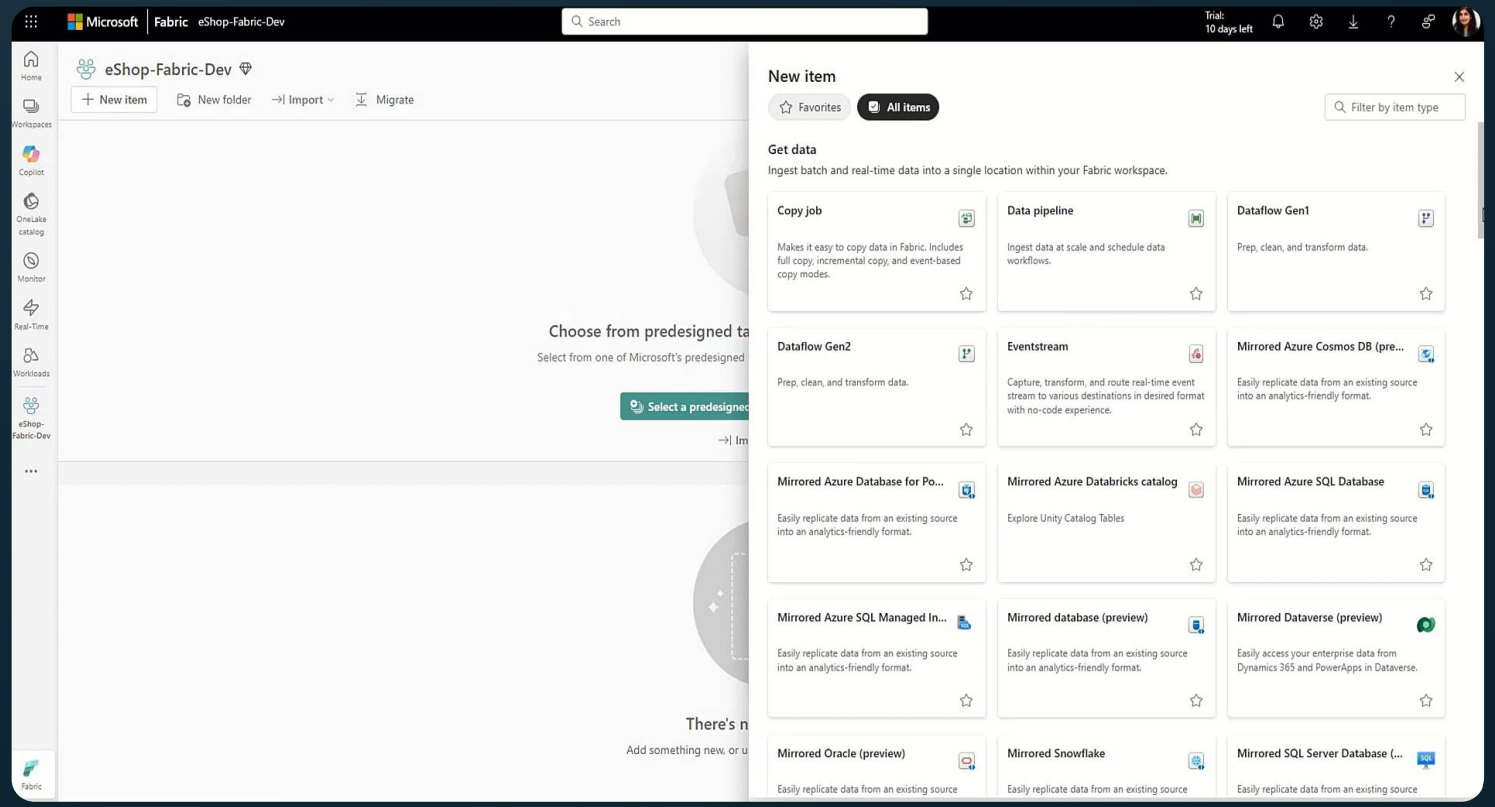

Enterprise grade Azure Cosmos DB built into the Fabric unified SaaS platform

The most scalable NoSQL database with extreme low latency servicing of semi-structured data

All the data is available in OneLake unifying operational and analytical data

Available for all Fabric capacities

Achieve the best price performance with the Native Execution Engine

Generally available

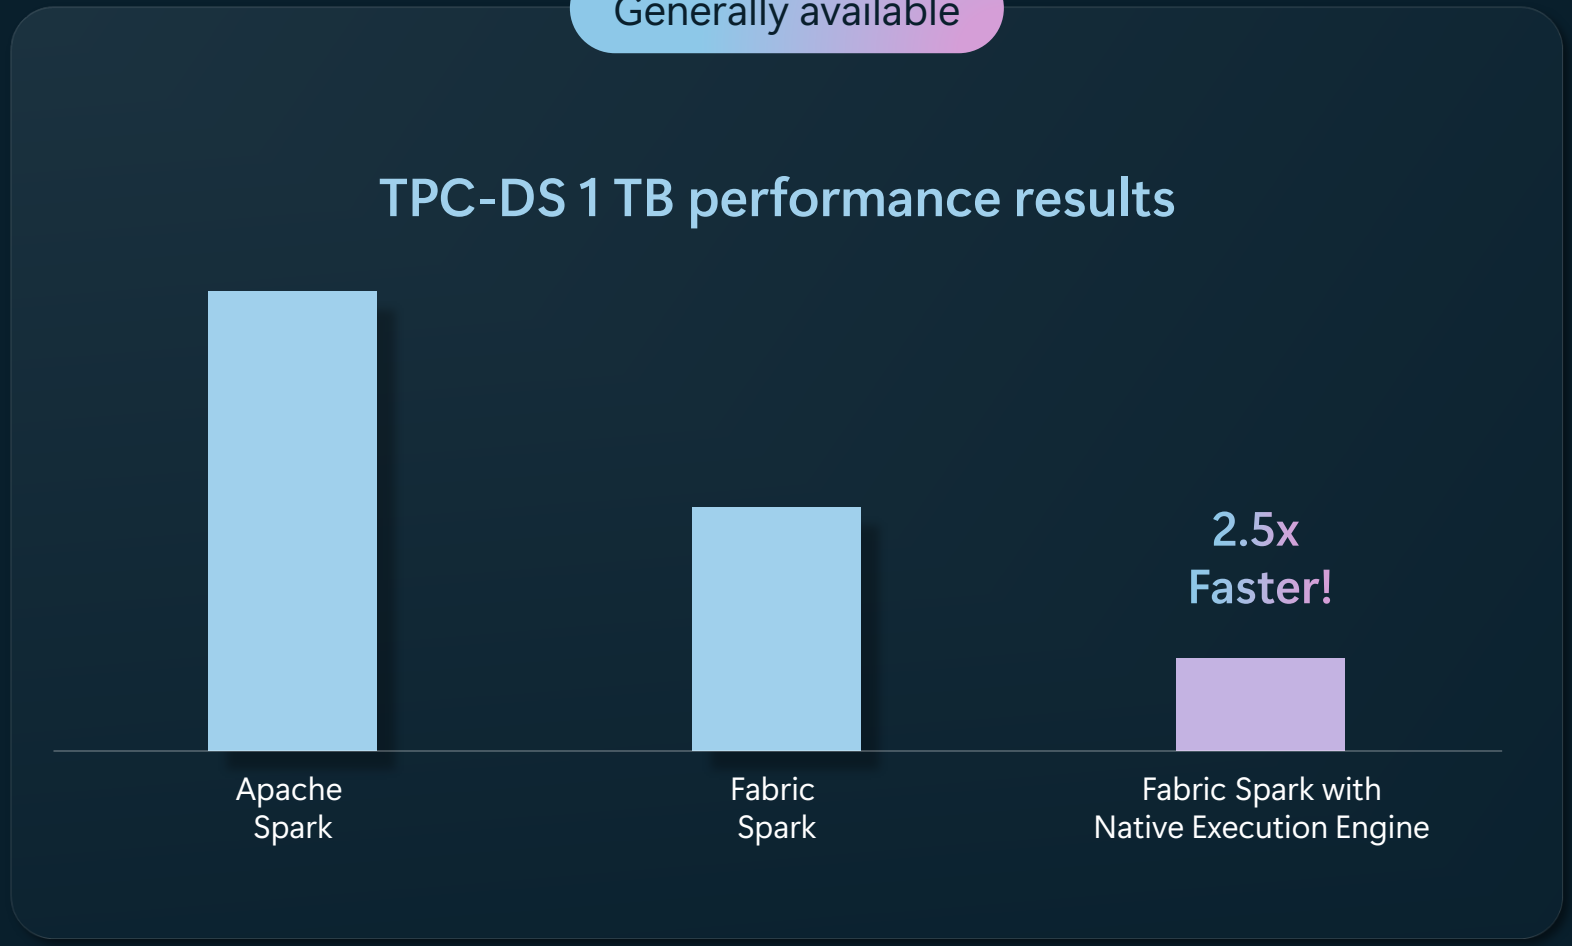

~2.5X faster execution over Fabric Spark

Entirely in C++ based on OSS projects Velox & Gluten

Fully compatible with Apache Spark and Delta Operators – no vendor lock-in

Available at no additional cost

ML Model endpoints

Public preview

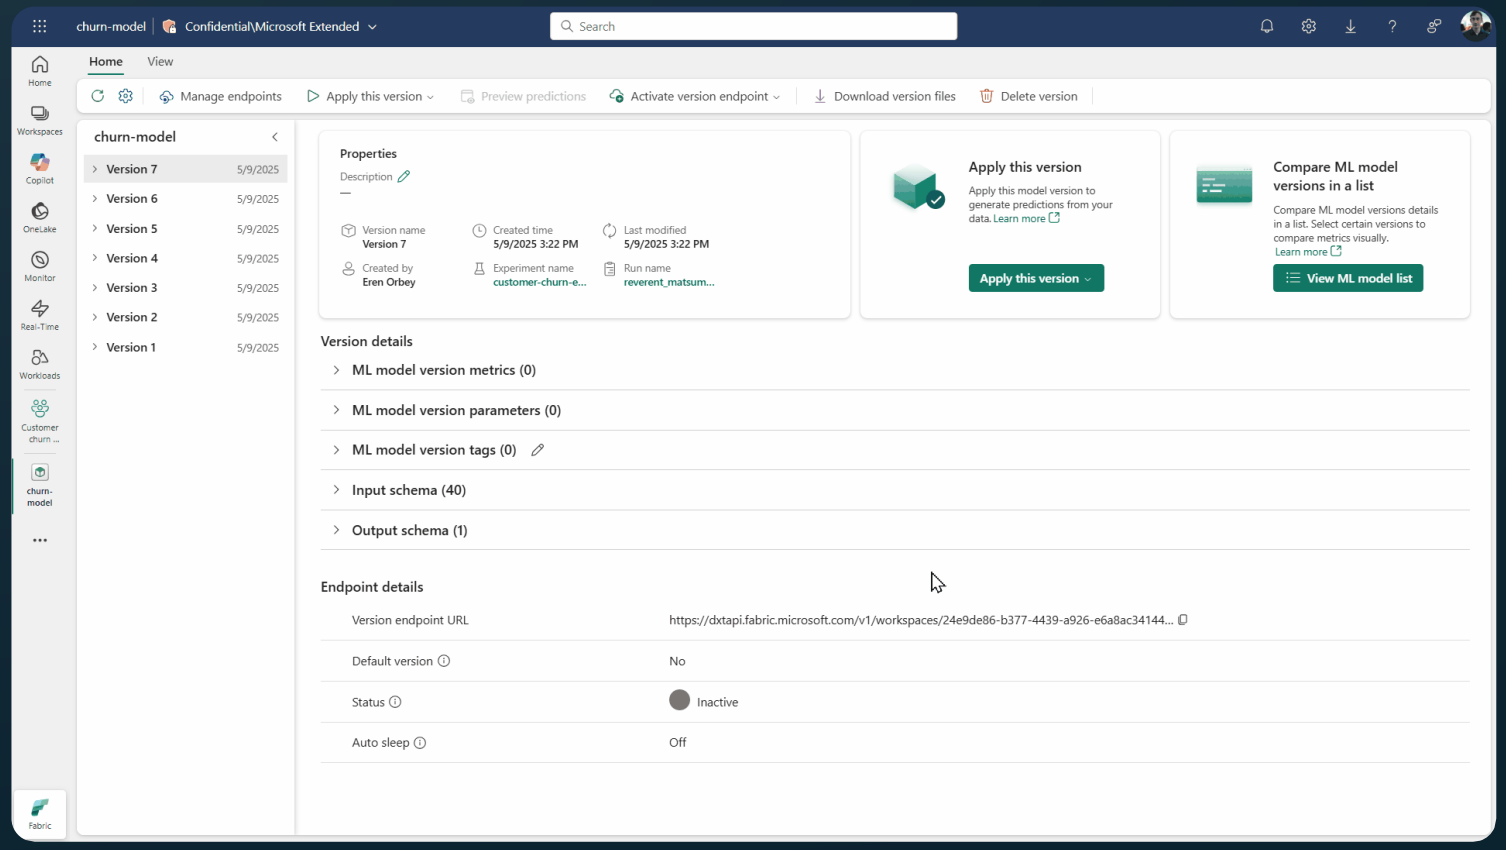

Serve real-time predictions with secure, scalable, and fully managed endpoints

Activate, customize, and query endpoints with a REST API or a low-code interface

Save costs with auto-scaling enabled out of the box

Consume predictive insights from custom apps or other Fabric experiences

Announcing

Digital twin builder

Public Preview

Digital twin builder

Public preview

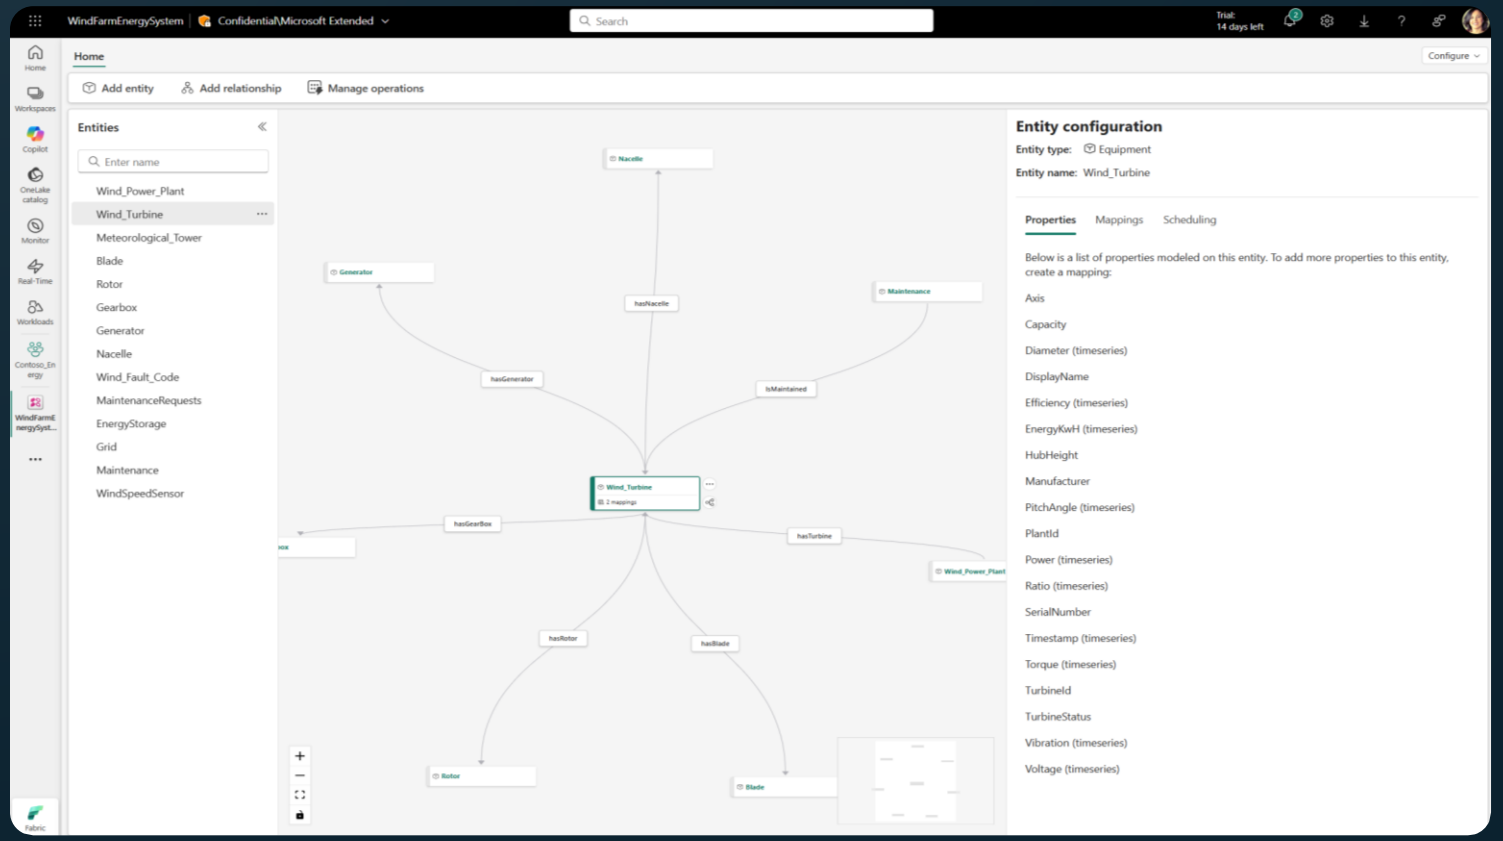

Simpler, faster way to build and manage digital twins with a low-code/no-code approach

Easily contextualize and map data to model assets, processes and systems and create a digital twin

Drill down to explore relationships and derive contextualized high value insights

Unlock insights with digital twin data in Power BI Copilot and Real-Time Dashboards

OneLake is the OneDrive for data

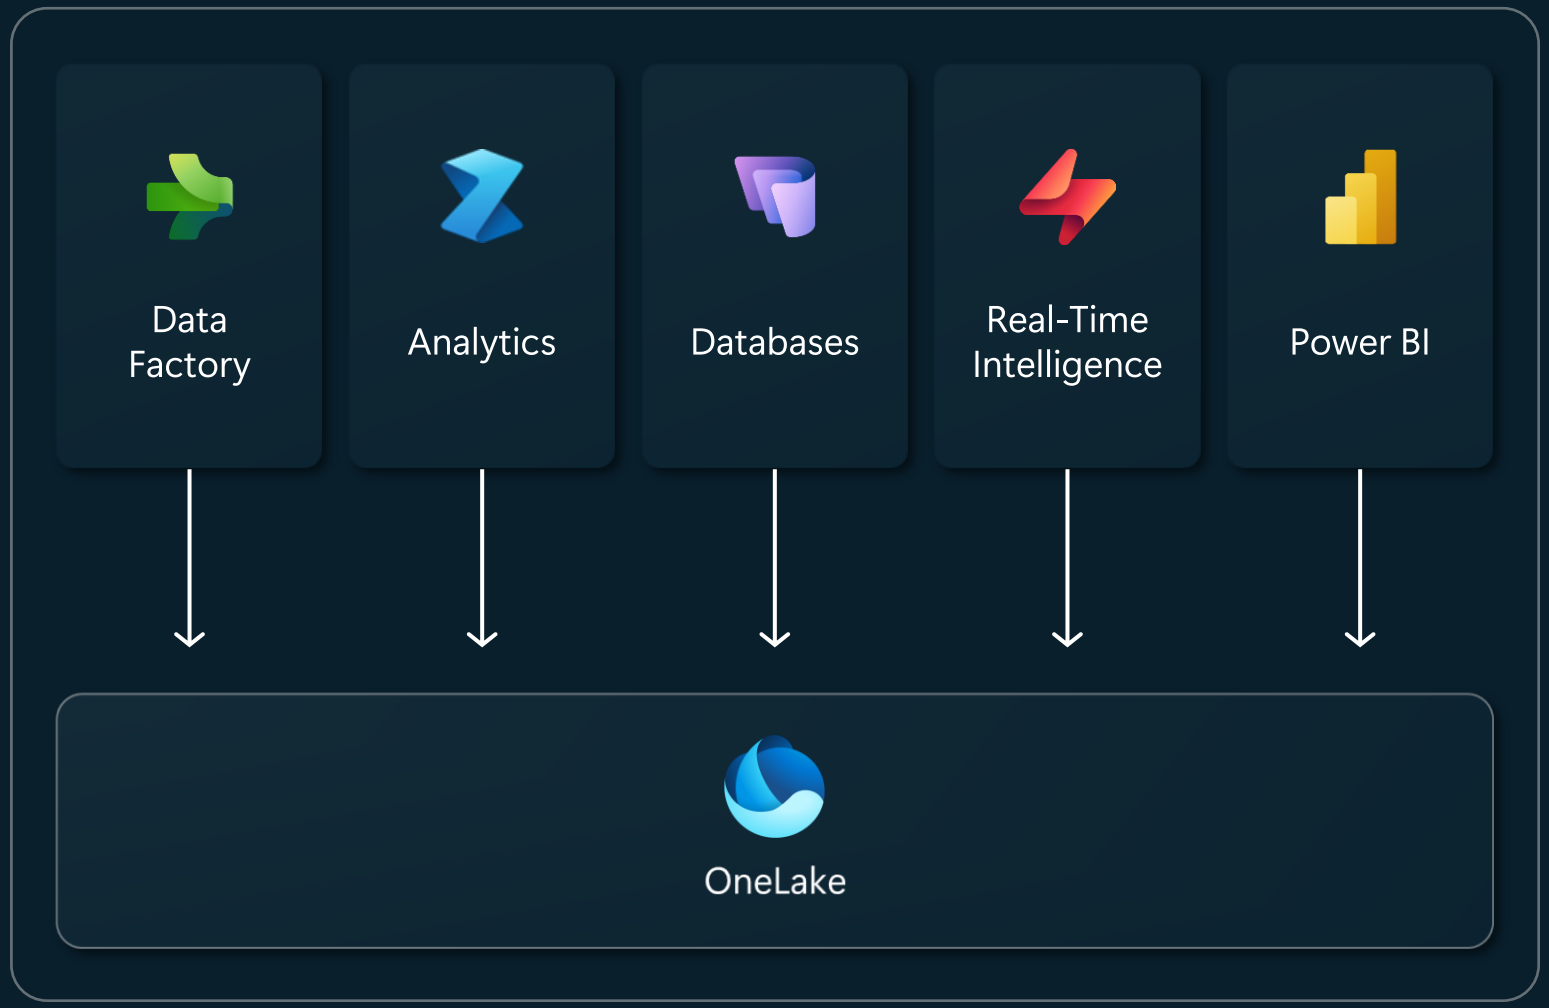

Single SaaS lake for the whole organization

Single open format used by all workloads to store and query data

Automatically indexed for discovery, MIP labels, lineage, PII scans, sharing, governance and compliance

Unifying data in OneLake

Cross-cloud shortcuts & mirroring

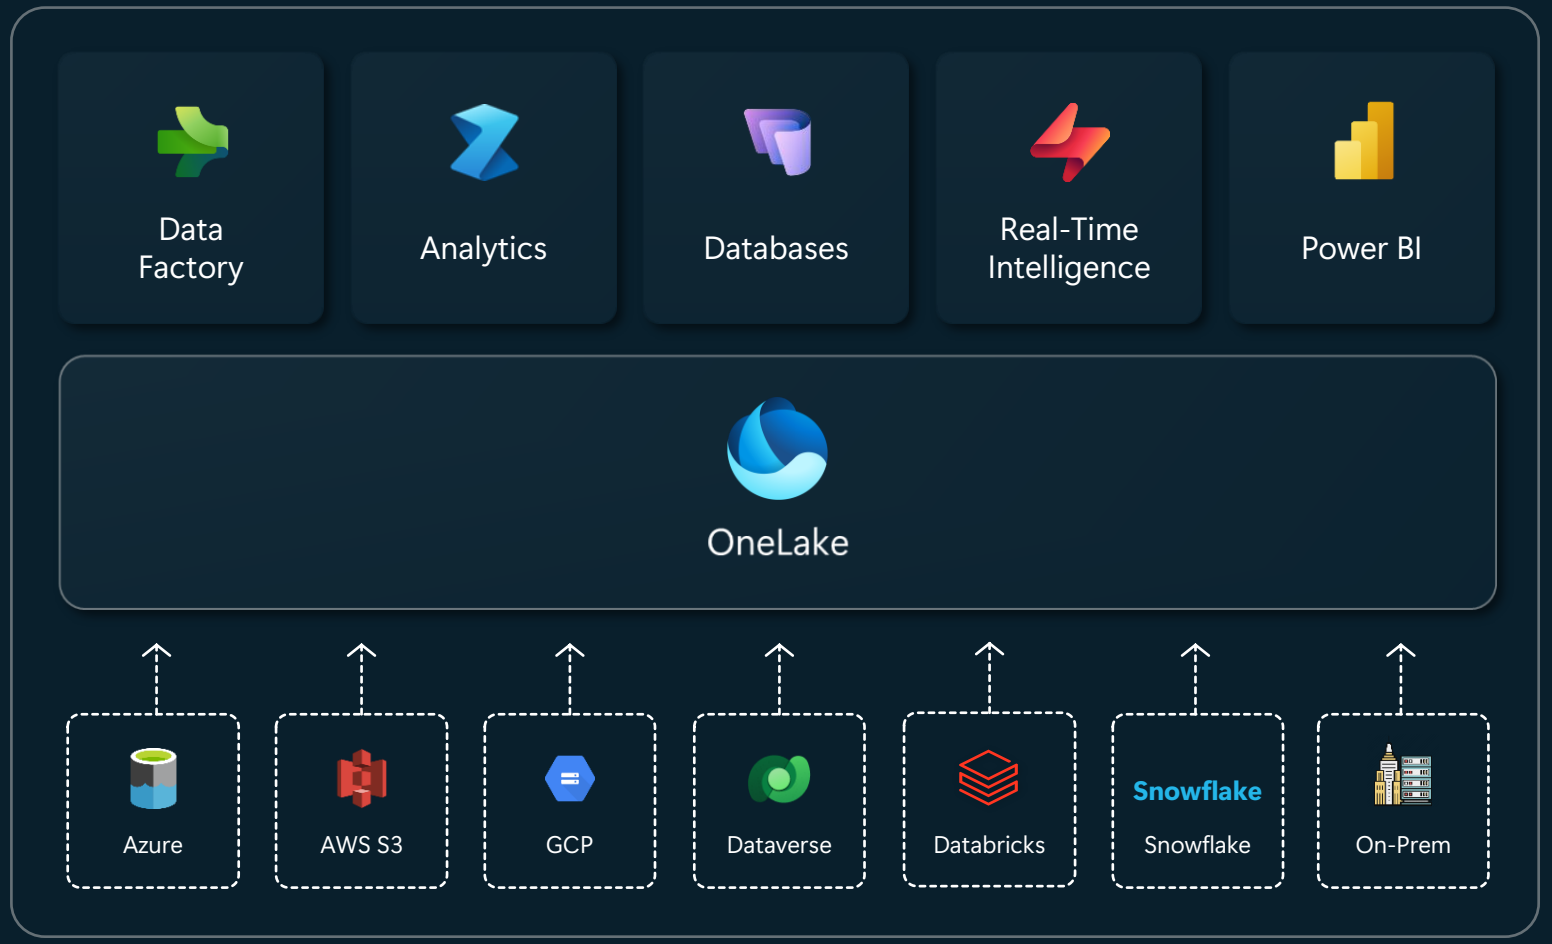

First multi-cloud SaaS data lake

Shortcuts for existing data cross clouds and on-premises

Shortcuts are instantaneous, with no data duplication or movement

Mirroring creates a reflection of an entire database in OneLake

Unify data in OneLake with zero ETL

Shortcut & mirroring sources

Generally available

- Azure SQL DB
- Azure Data Lake store
- Microsoft OneLake
- Google Cloud Storage
- Snowflake
- Amazon S3
- Microsoft Dataverse
- S3 Compatible (cloud/On-prem)

Public preview

- Azure Cosmos DB
- Azure SQL MI
- Azure PostgreSQL
- Databricks Catalog
- SQL Server 2025 — NEW!
- Azure Blob storage — NEW!

Coming soon

- Oracle DB
- Google BigQuery

Open Mirroring

Generally available

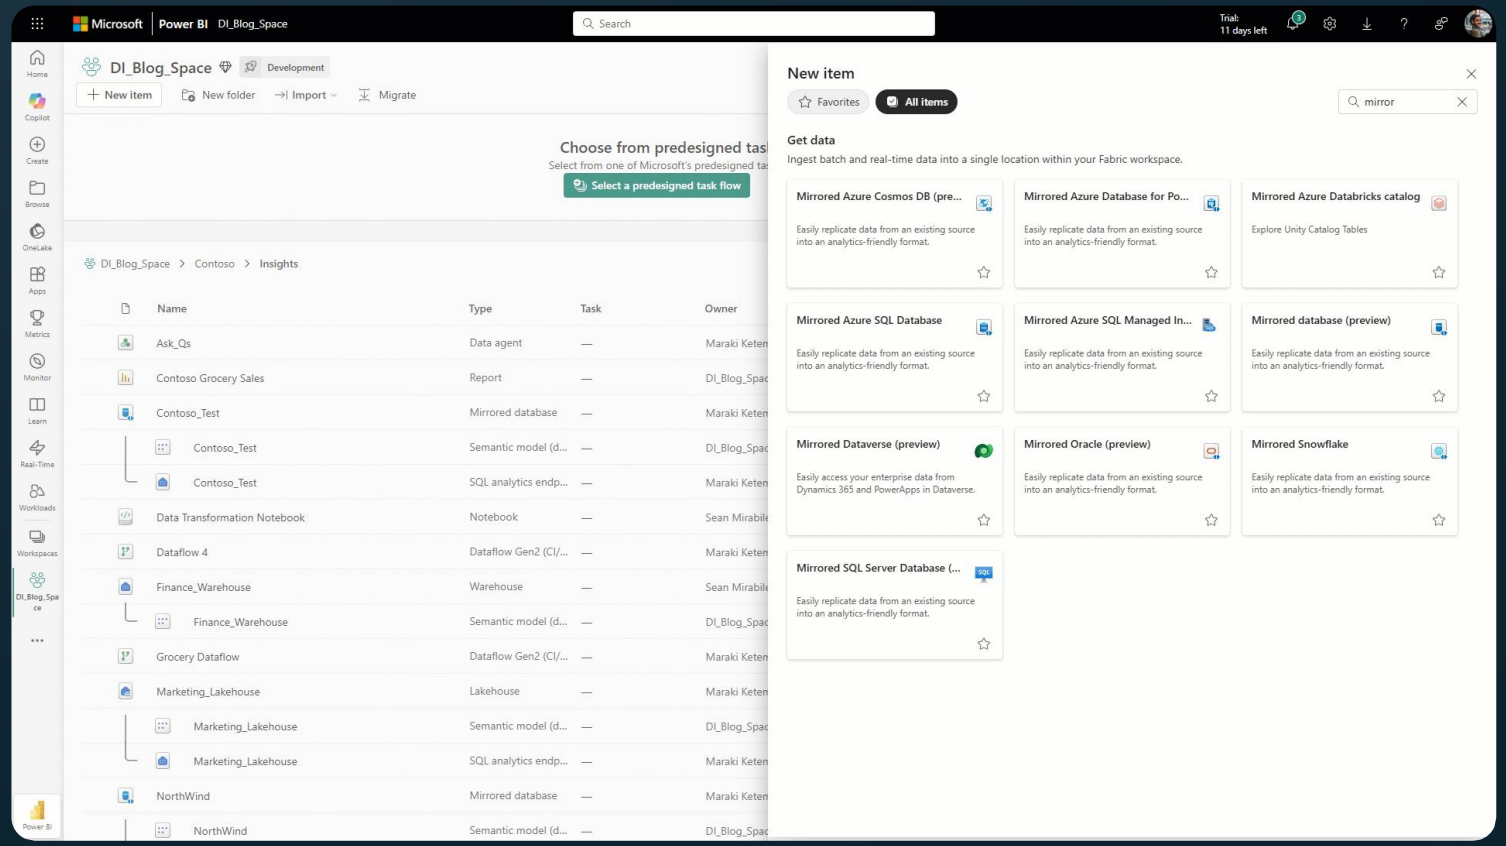

Simplifies building custom mirroring solutions and handling of complex data changes for easy data integration into OneLake via Mirroring

Open Mirroring enables landing parquet or CSV data via API or manually via drag-and-drop

Automatically converts parquet and CSVs into Delta before replication to optimize your BI and AI workloads in Fabric

Shortcut Transformations

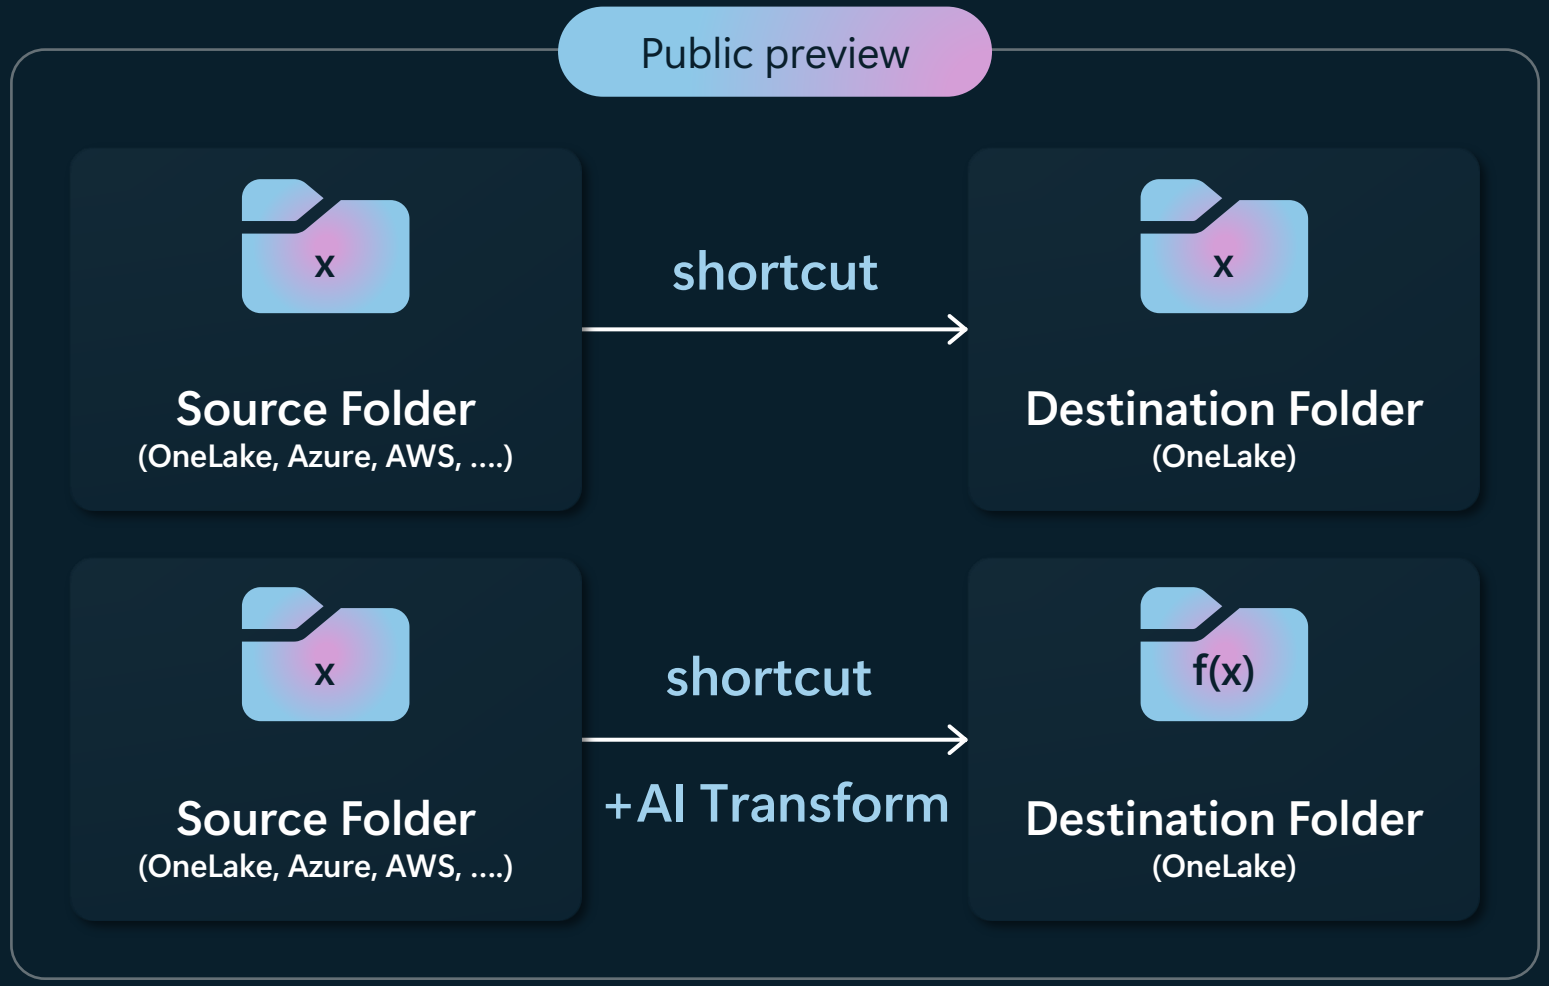

AI Transformations built directly into OneLake

Available as an extension of Shortcuts; simply shortcut to a source folder and add a transformation definition

Automatically track all changes in source and sync the destination

Demo

Shortcut transformations

Public Preview

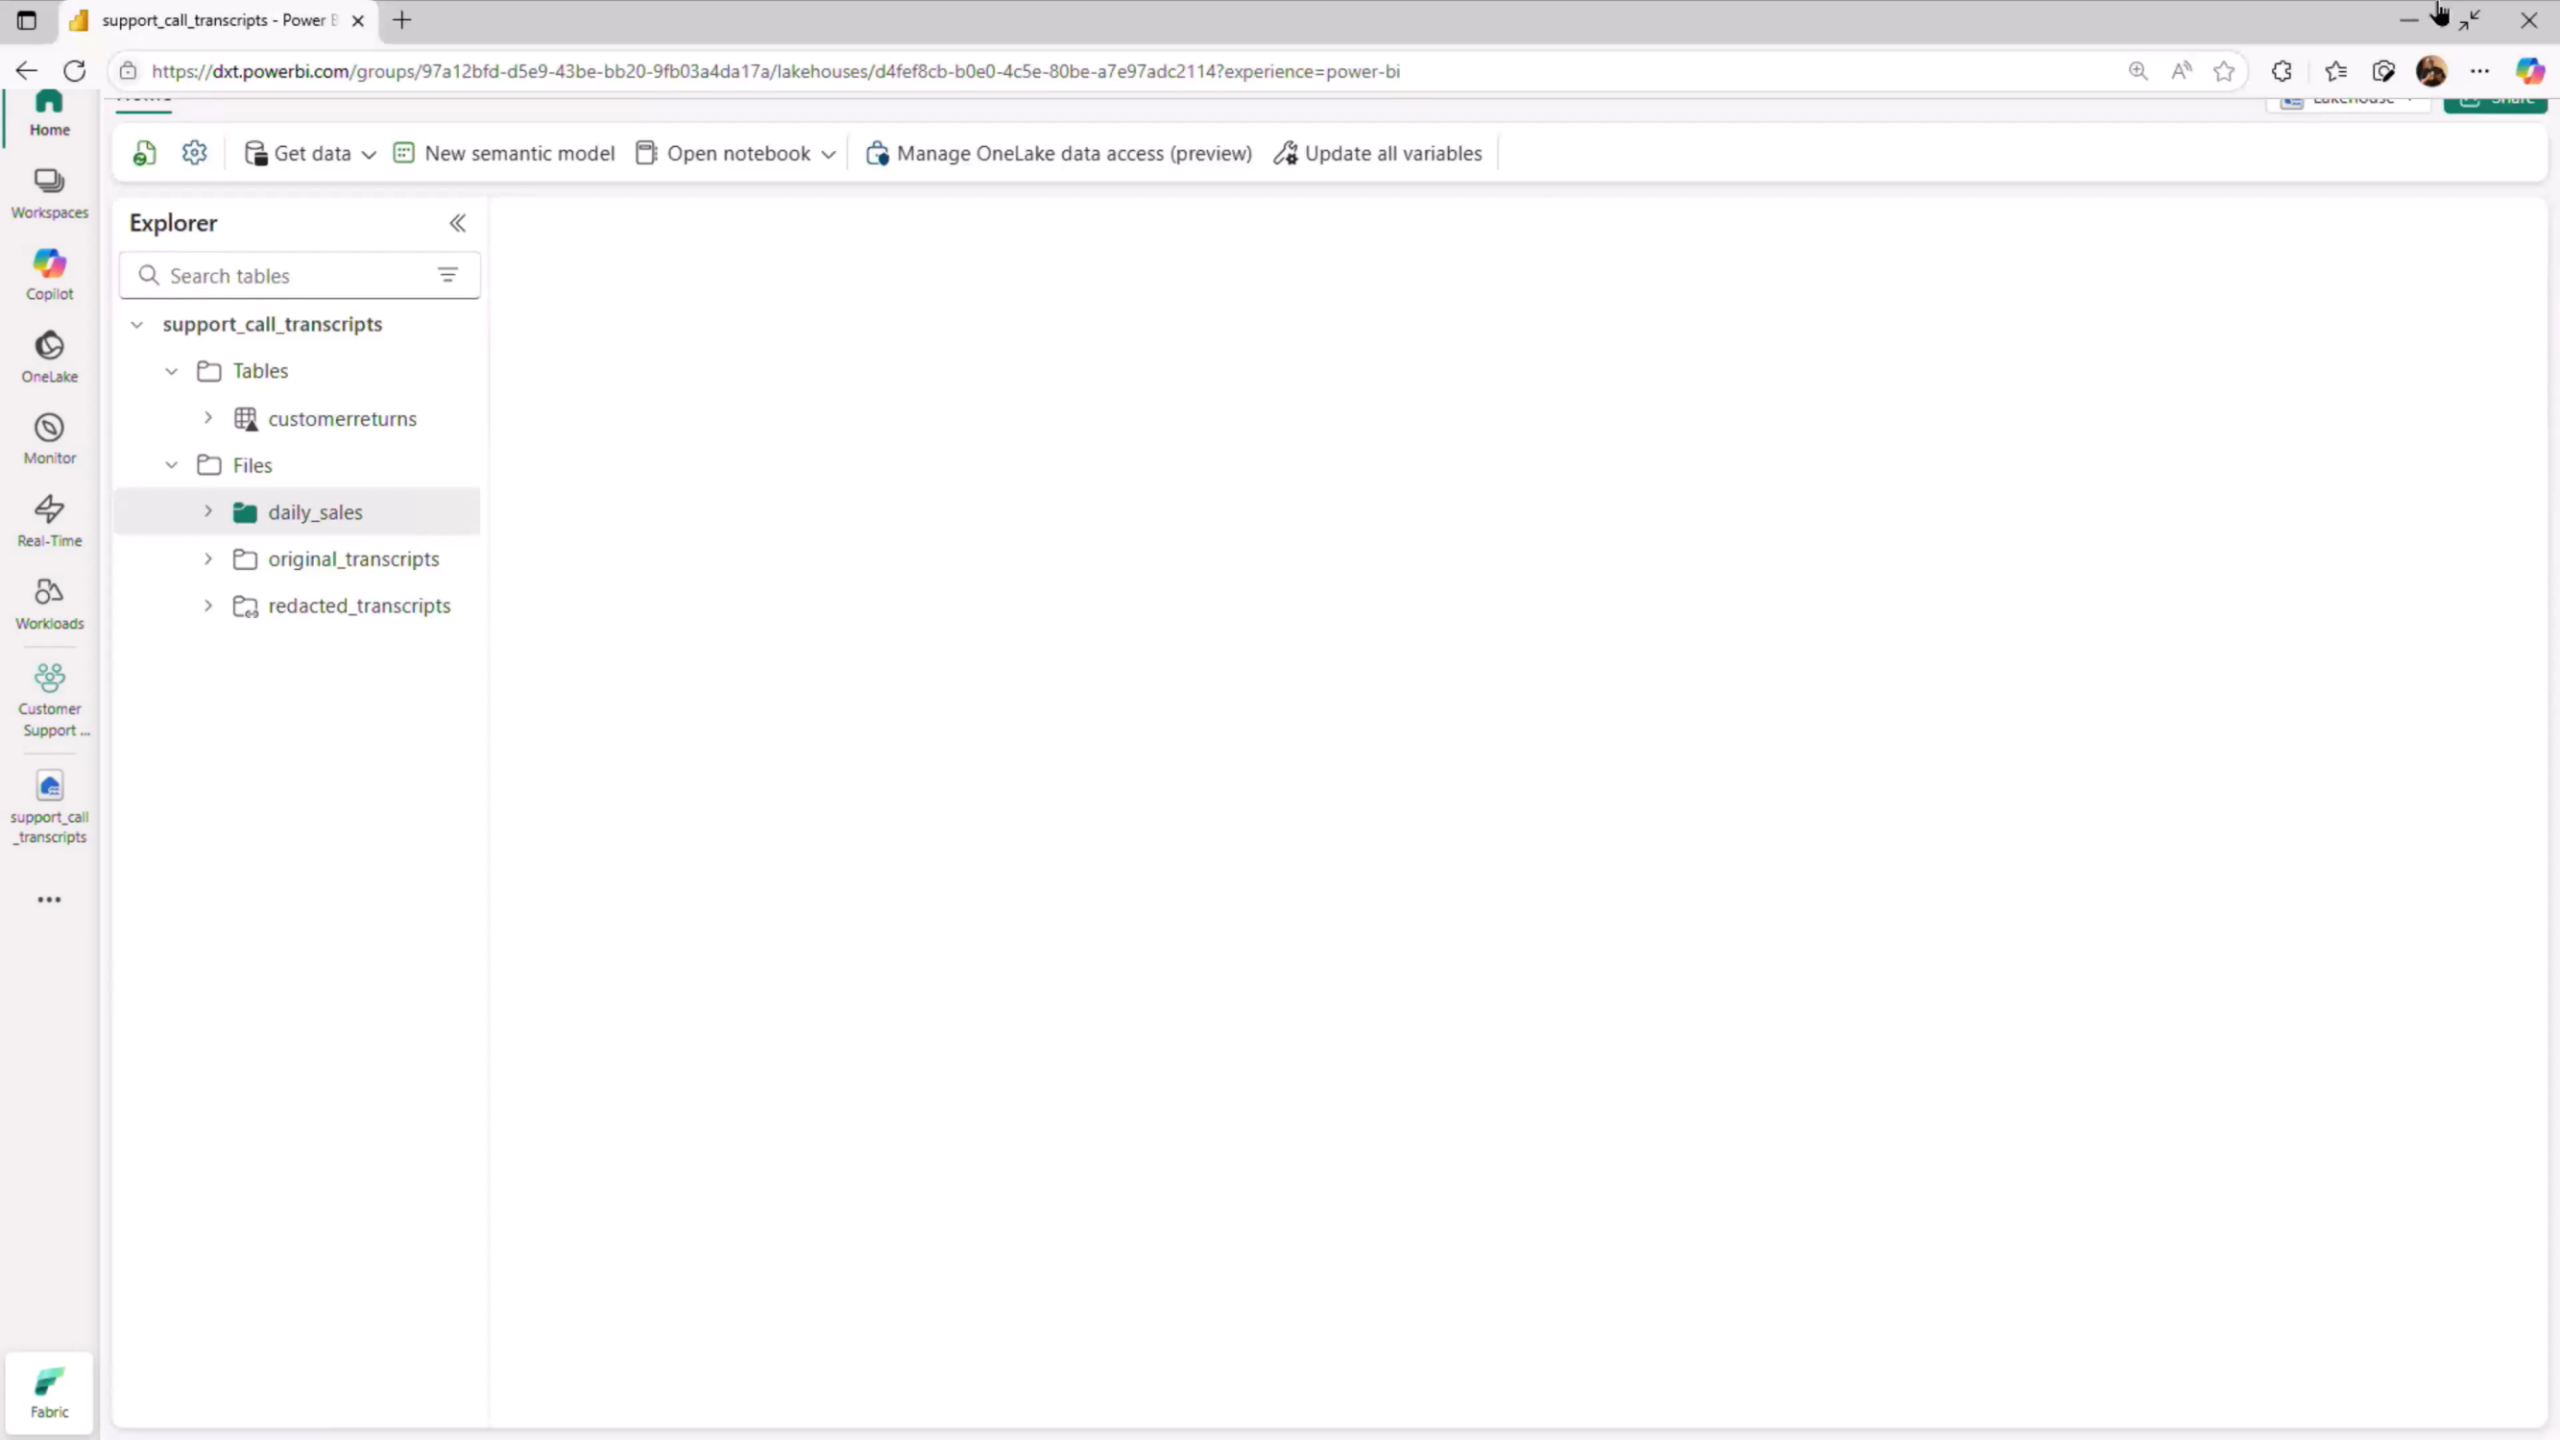
support_call_transcripts - Power B
https://dxt.powerbi.com/groups/97a12bfd-d5e9-43be-bb20-9fb03a4da17a/lakehouses/d4fef8cb-b0e0-4c5e-80be-a7e97adc2114?experience=power-bi
Home
Workspaces
Copilot
OneLake
Monitor
Real-Time
Workloads
Customer Support ...
support_call_transcripts
Fabric
Get data
New semantic model
Open notebook
Manage OneLake data access (preview)
Update all variables
Explorer
Search tables
support_call_transcripts
Tables
customerreturns
Files
daily_sales
original_transcripts
redacted_transcripts

Explorer

Search tables

- support_call_transcripts
  - Tables
    - customerreturns
  - Files
    - original_transcripts

Files > original_transcripts

Showing 6 items

Search files

Name	Date modified	Type	Size
support_call_1 1.txt	5/14/2025, 6:30:56 PM	txt	331 B
support_call_2 1.txt	5/14/2025, 6:47:29 PM	txt	327 B
support_call_3 1.txt	5/14/2025, 6:47:29 PM	txt	330 B
support_call_4 1.txt	5/14/2025, 6:30:56 PM	txt	329 B
support_call_5 1.txt	5/14/2025, 6:47:29 PM	txt	307 B
support_call_6 1.txt	5/14/2025, 6:47:29 PM	txt	318 B

Translytical task flows

Bridging the gap between data and decisions

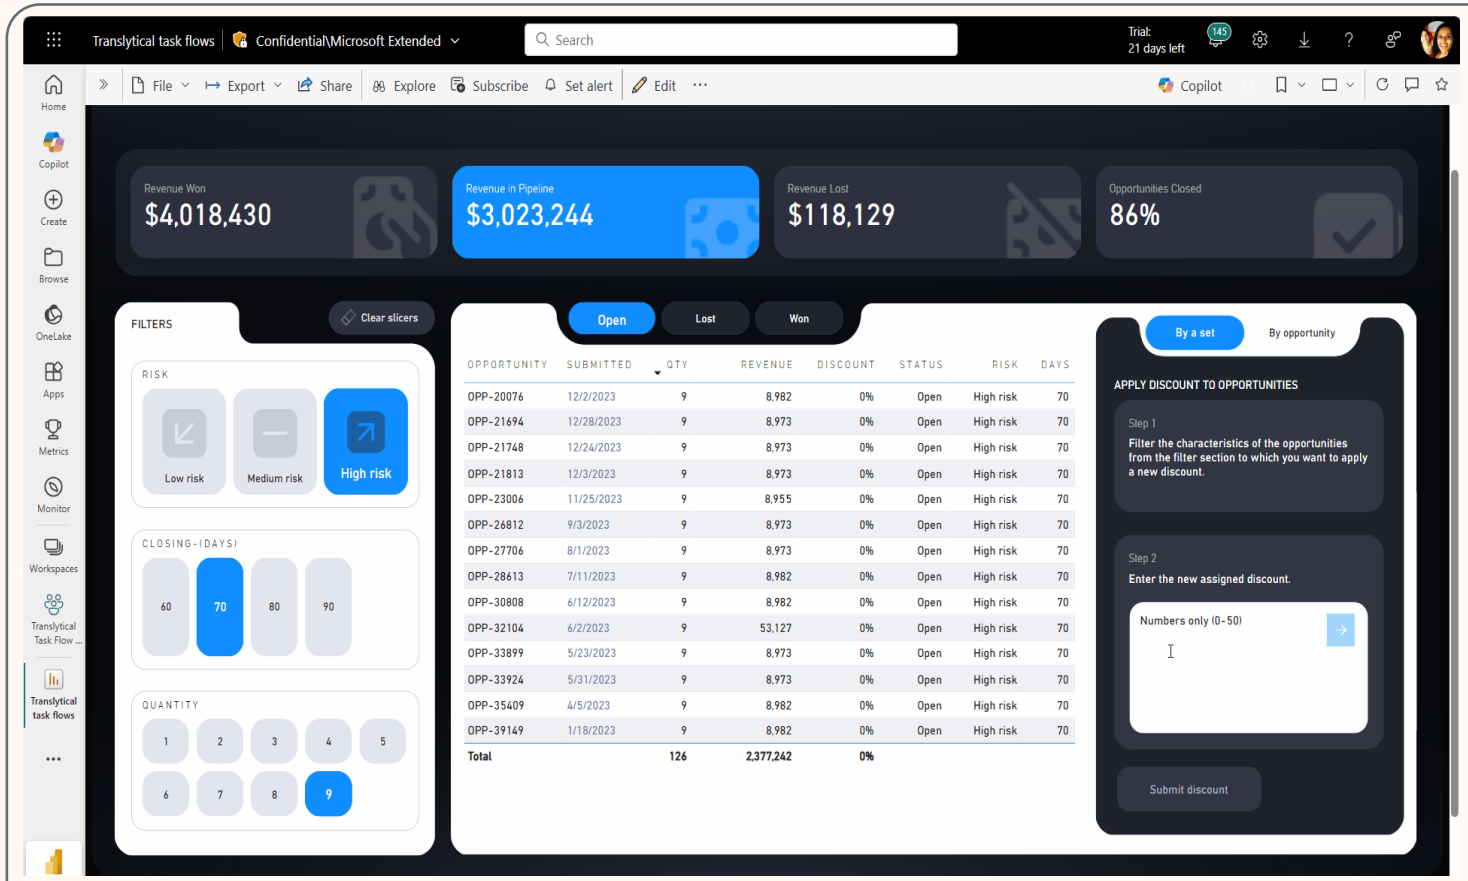

Announcing Public Preview!

Translytical task flows marks a major evolution in Power BI reports, enabling automated action and data writeback directly within the report—streamlining decision-making and operational follow-through.

Powered by Fabric User data functions, this new capability allows users to automate tasks such as updating records, dynamic notifications, and even triggering workflows across other systems.

Shipped

aka.ms/FabricSecurityWhitepaper

Network Security

GA

Entra ID conditional access

Service tags

User audit logs

Enterprise data gateway

VNET Data Gateway Support with Private Links for Dataflows Gen2 & Semantic models

Private endpoint, tenant level

Shortcuts for ADLS Gen2 in VNET

Spark connectivity to Azure data services in a VNET

Trusted workspace access

Data Security

GA

Workspace roles

SQL Object-level security

SQL Column-level security

SQL Row-level security

SQL Dynamic data maillling

SQL granular permissions

Lockbox

Limited Public preview

CMK for OneLake

OneLake security

Governance

GA

Purview Information protection

Purview Sensitivity labels

Purview Audit

Endorsement

Lineage

Impact analysis

Domains & Sub Domains

Metadata scanning

Resiliency

Purview data catalog

Public preview

Admin Monitoring and Insights

Fabric Monitoring

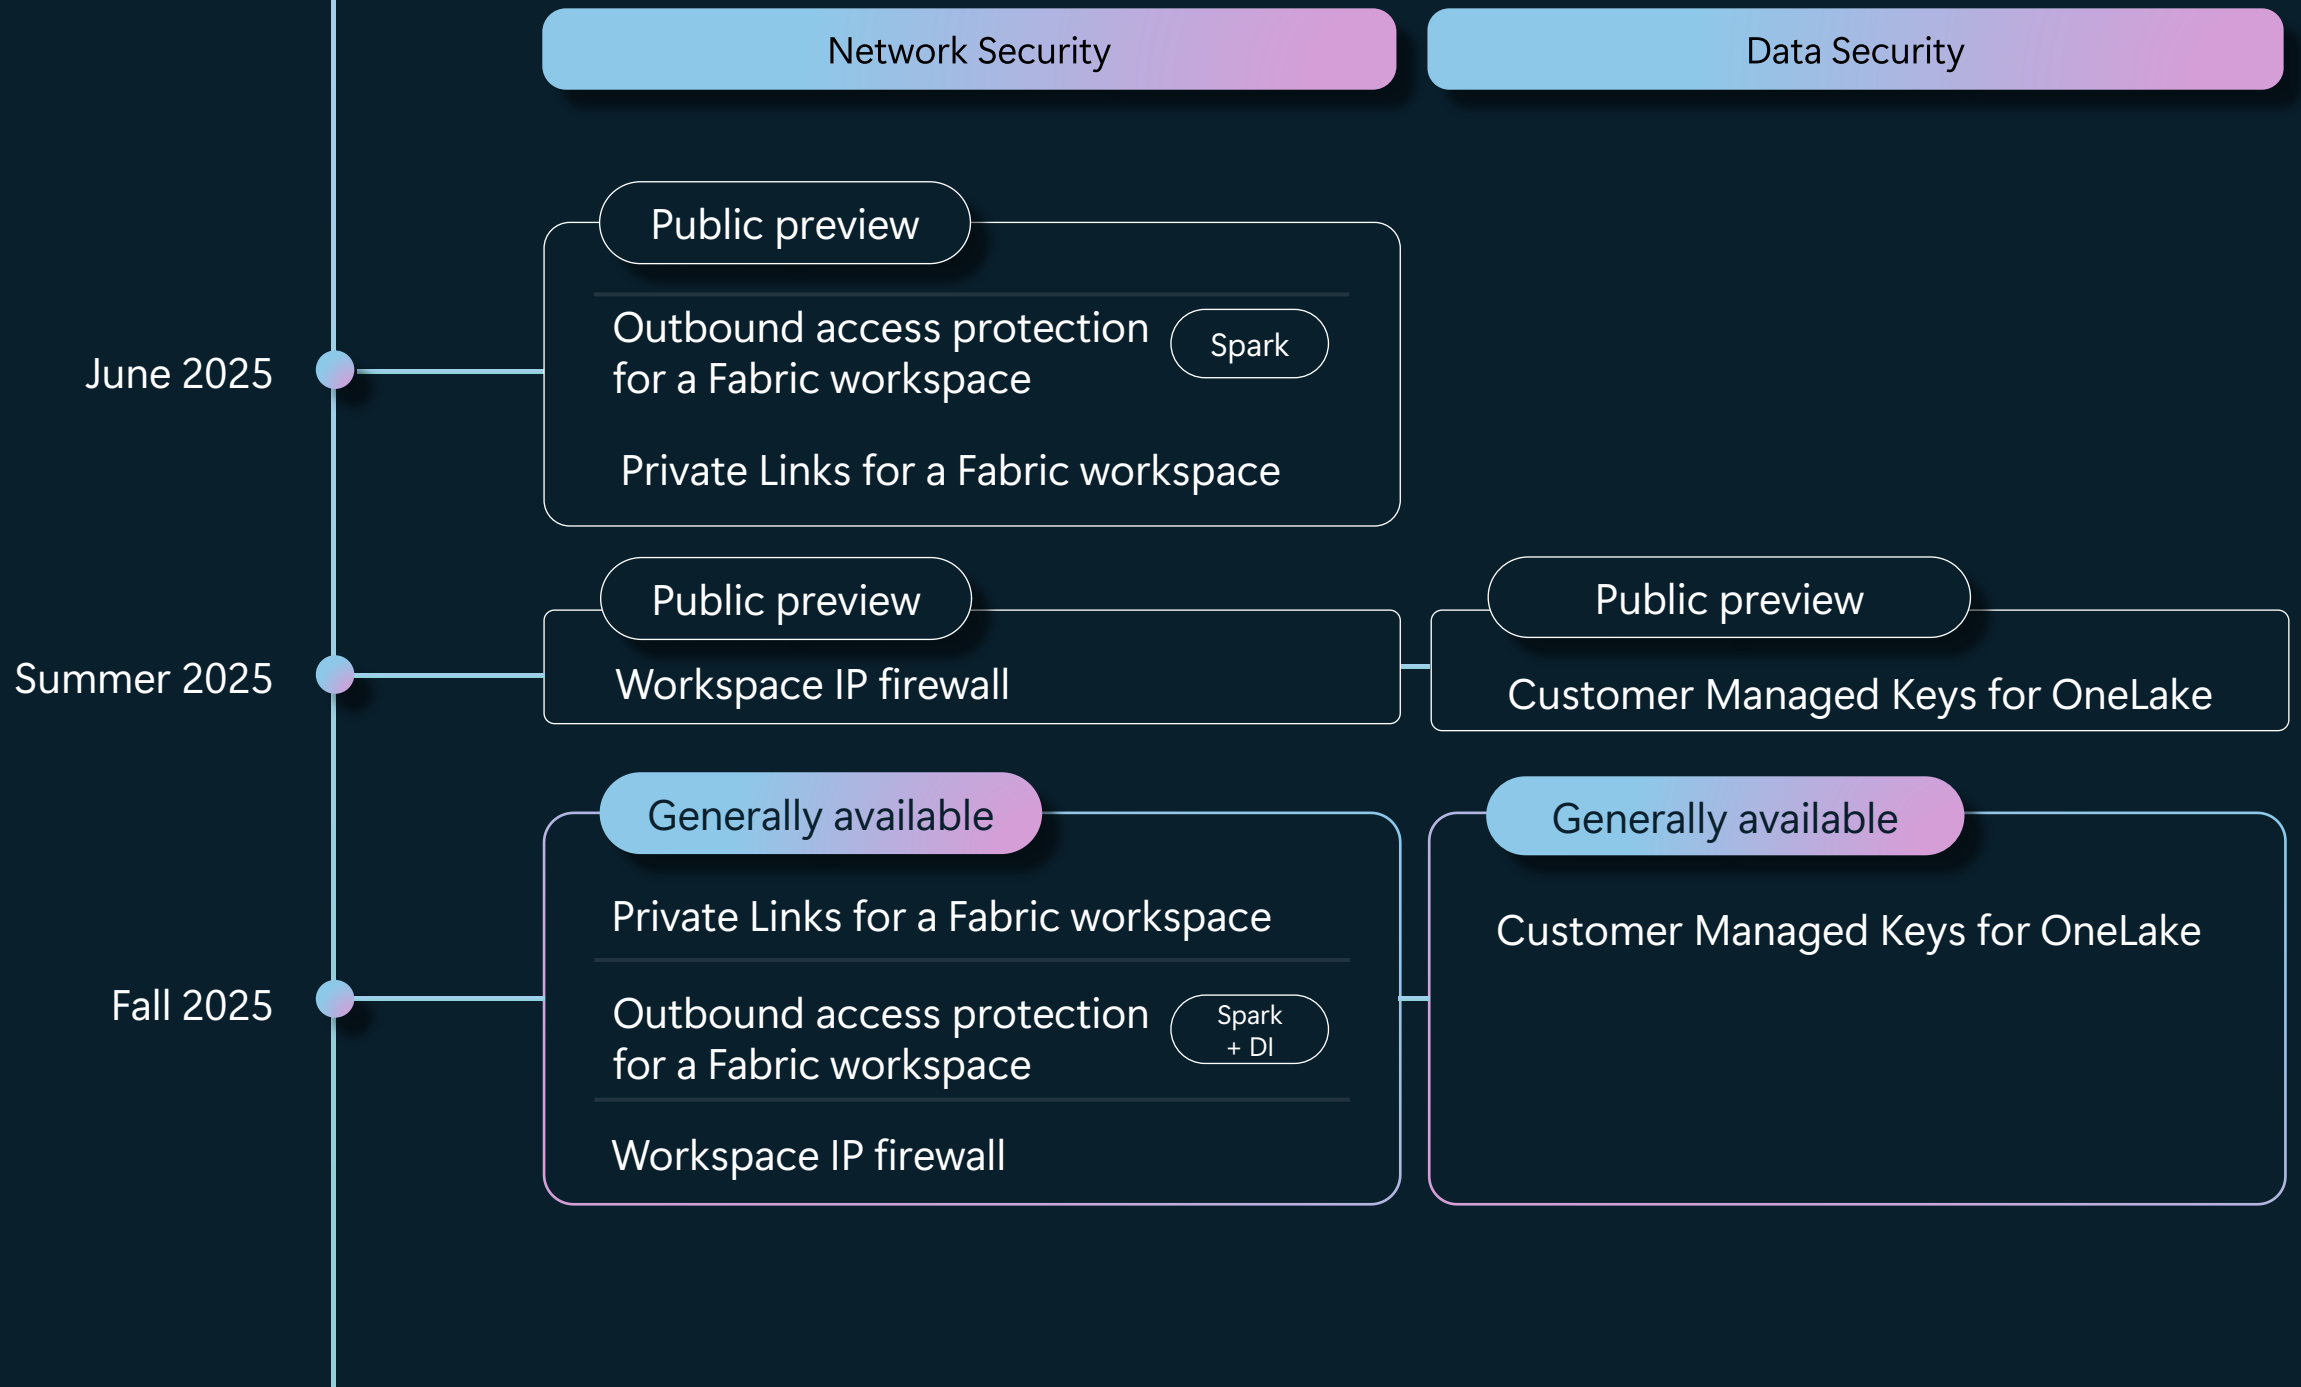
Network Security
Data Security
June 2025
Public preview
Outbound access protection for a Fabric workspace
Spark
Private Links for a Fabric workspace
Summer 2025
Public preview
Workspace IP firewall
Public preview
Customer Managed Keys for OneLake
Fall 2025
Generally available
Private Links for a Fabric workspace
Outbound access protection for a Fabric workspace
Spark + DI
Workspace IP firewall
Generally available
Customer Managed Keys for OneLake

Real-Time Intelligence is the difference between leading and lagging

Microsoft Fabric

The unified data platform for AI transformation

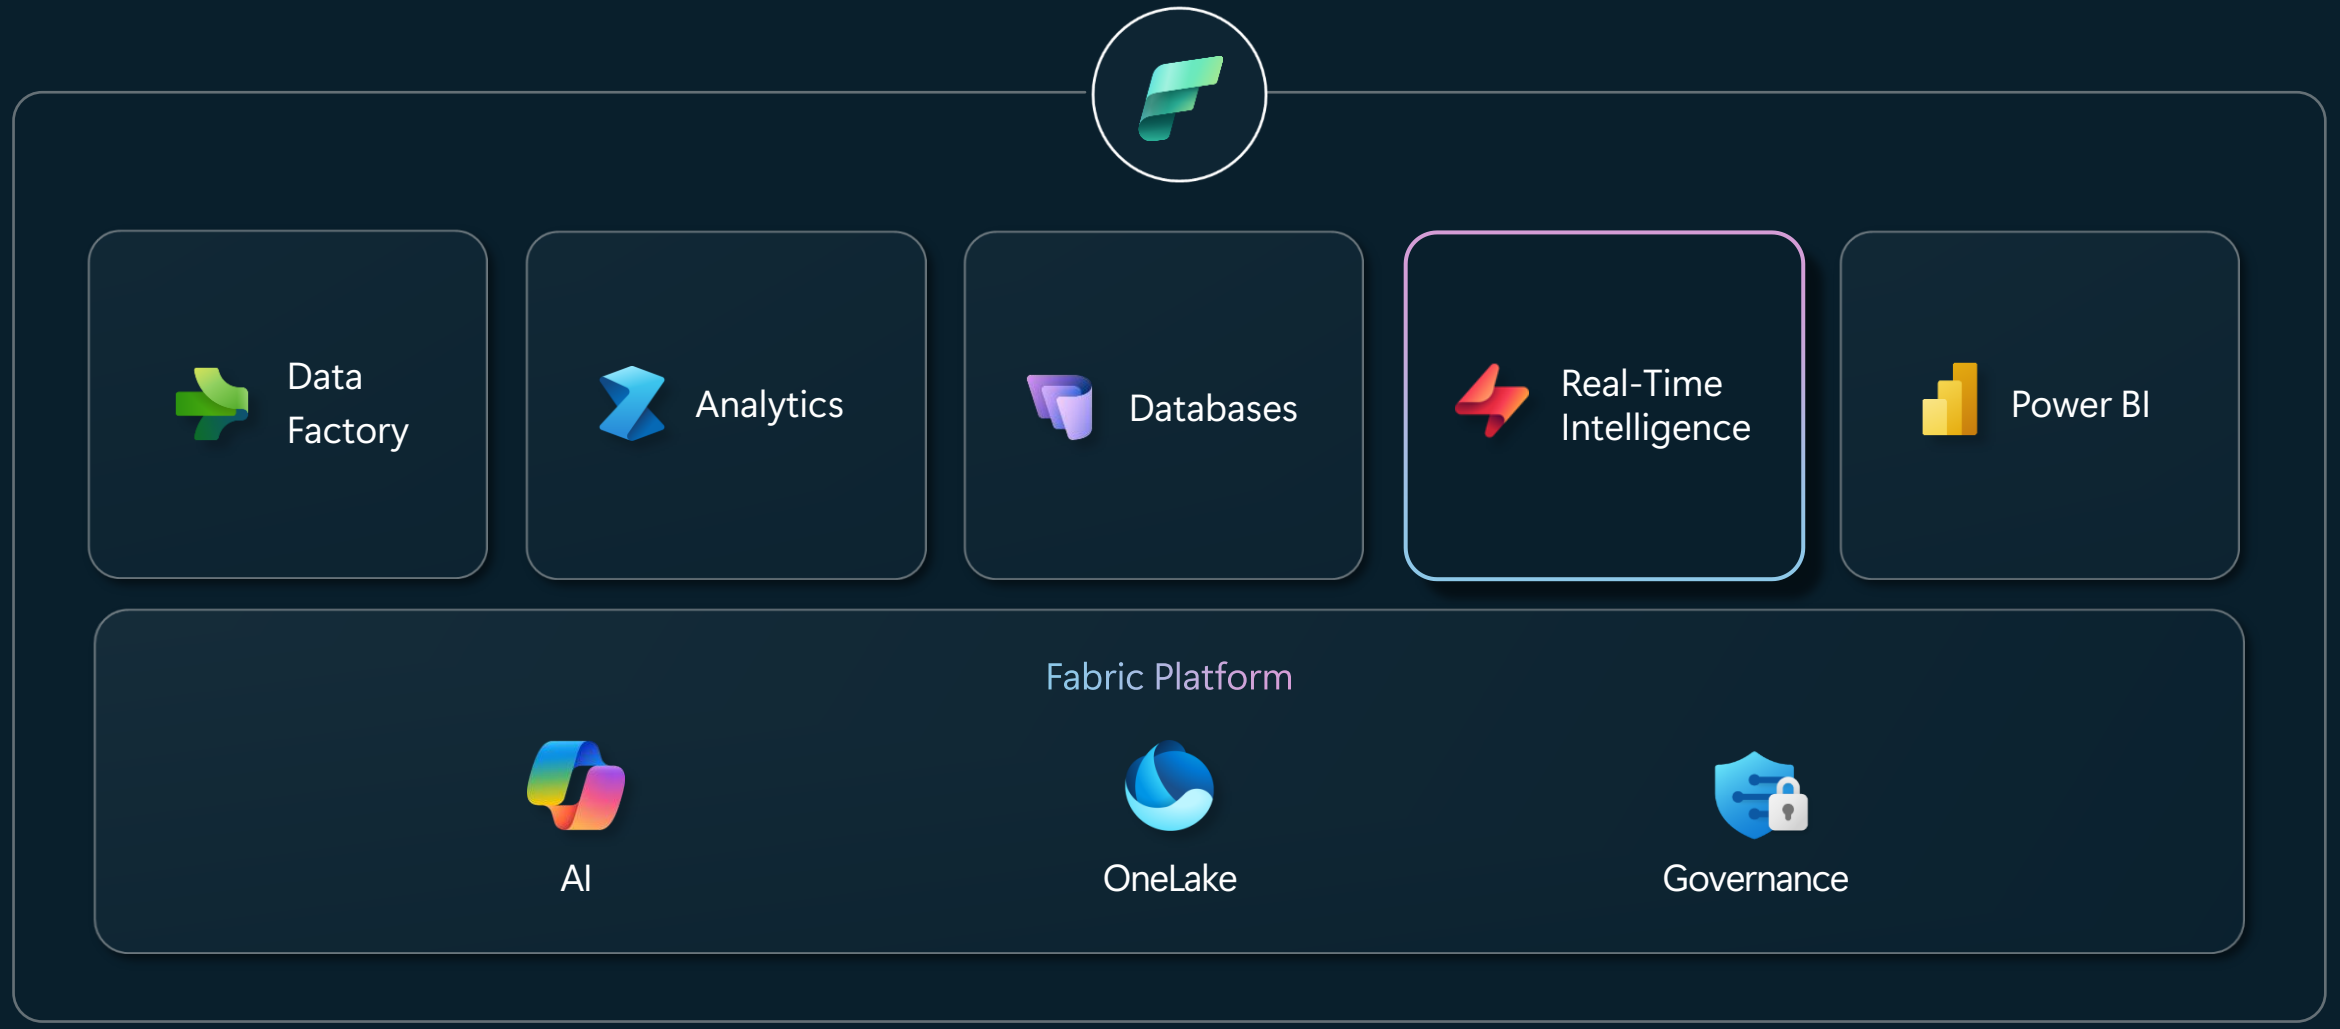

Real-Time Intelligence in Microsoft Fabric

Planet-scale infrastructure for real-time data

10 EB
Events and logs per month

350 PB
Ingested daily

5.1 B
Real-time queries per day

19.2 T
Streaming events processed monthly

99.9992%
Success rate across messaging

Real-Time Intelligence in Microsoft Fabric

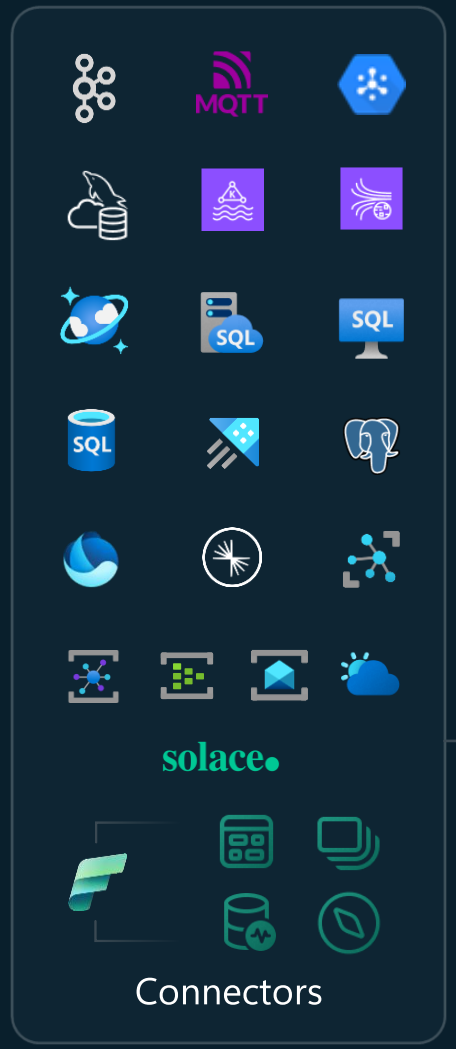

Ingestion

Analytics

Dashboards

Rules

Actions

AI

Real-Time Hub

OneLake

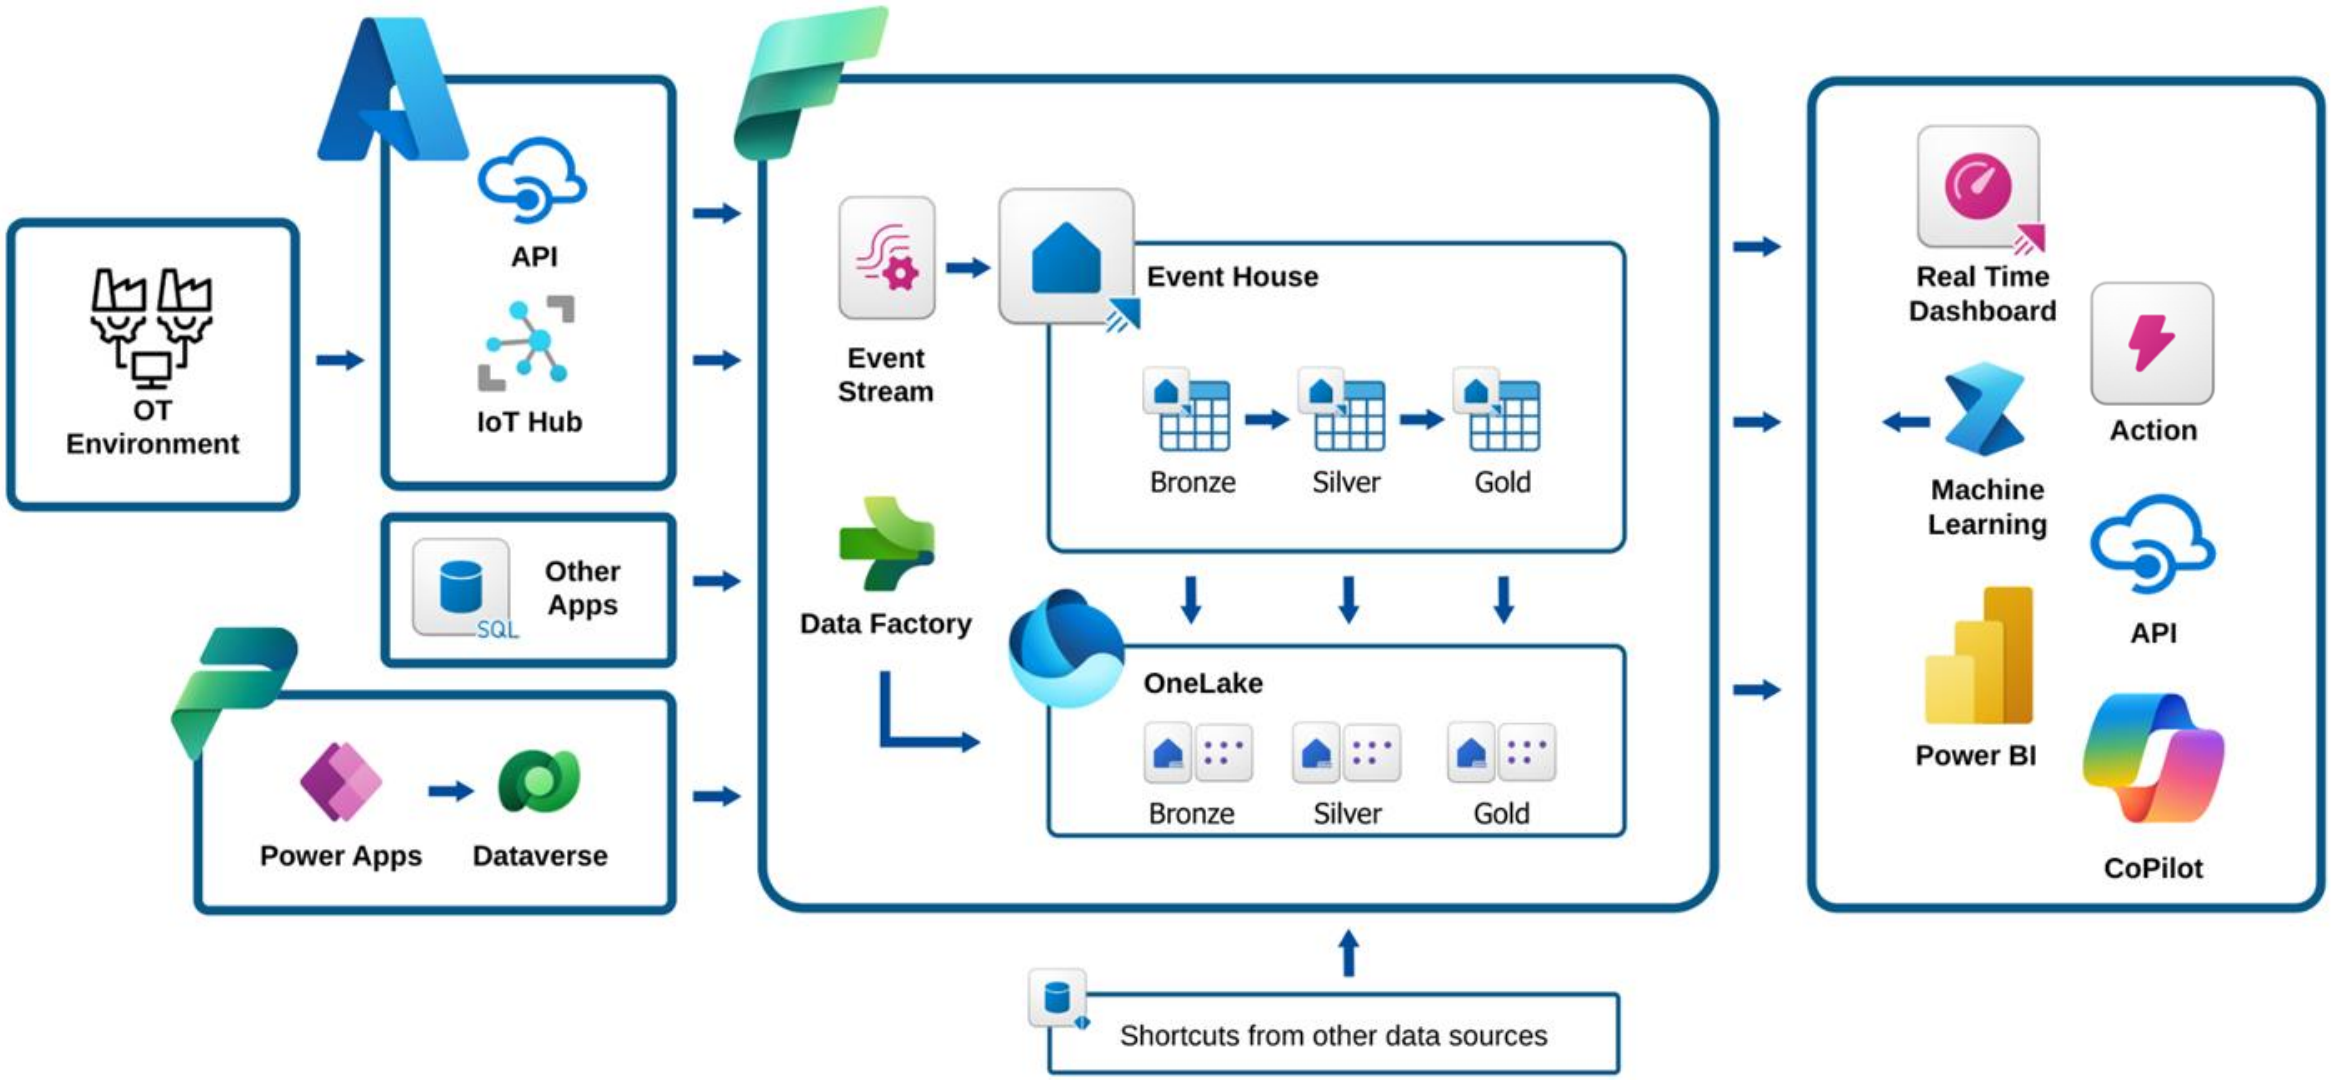
OT Environment
API
IoT Hub
Other Apps
SQL
Power Apps
Dataverse
Event Stream
Event House
Bronze
Silver
Gold
Data Factory
OneLake
Bronze
Silver
Gold
Shortcuts from other data sources
Real Time Dashboard
Action
Machine Learning
API
Power BI
CoPilot

Event-driven architecture

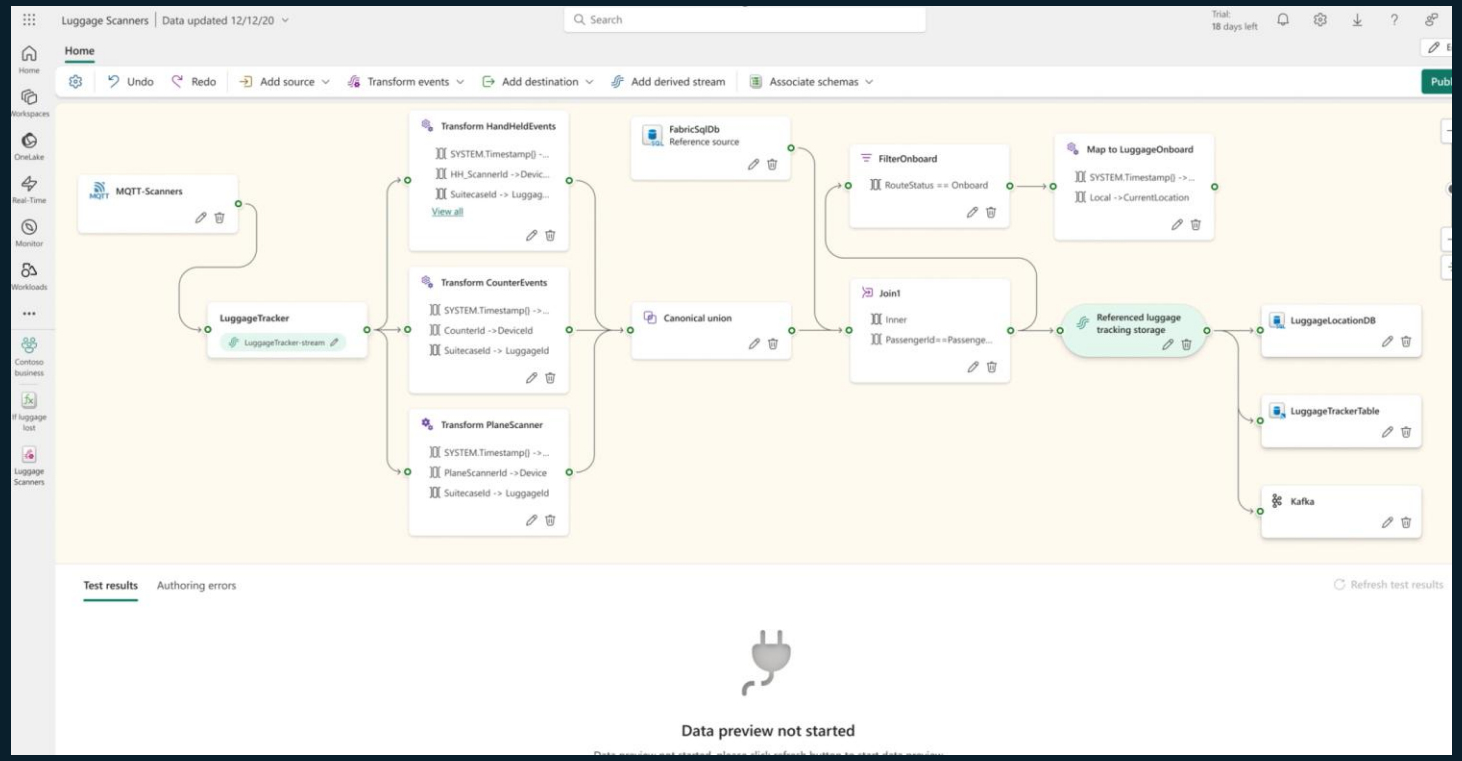

Use events to trigger applications instead of relying on schedules or manual actions

Separate event consumers and event producers, making the system more flexible and scalable

Process events as they occur, reducing the time to value for the rest of the system

Real-Time Intelligence in Microsoft Fabric

PrPr – Private Preview
PuPr – Public Preview
GA – Generally available

Current Releases

- Digital Twin Builder PuPr
- SQL custom code support in Eventstreams PuPr
- Connectors PuPr: Weather, Solace PubSub+, ADX Table Streamify, MQTT v5, Event Grid Namespaces
- Activator Fabric jobs w. parameters PuPr
- Event Schema Set PuPr
- Multiple Schema Inference PuPr
- Eventstreams CI/CD & API GA
- Managed Private Endpoint PuPr
- Data Preview – Kafka, Confluent, Kinesis PuPr
- Entra for custom endpoint PuPr
- Fabric Events GA: OneLake events, Job events, Workspace events, Azure Storage events
- Synapse Data Explorer -> Fabric migration tools PuPr
- Trigger pane management in Power BI GA
- Eventhouse DB entity map view PuPr

Roadmap

- Workspace level RTI private link PrPr
- Task-oriented real-time hub
- Event Schema Set PuPr
- Connectors PuPr : Confluent w. Schema Registry, etc.
- Azure SQL CES to Eventstreams w. custom endpoint PuPr
- MSFT sources: Dataverse events PrPr , Azure Monitor in RT hub PrPr
- Connector on-prem/Vnet PrPr
- Custom business events PrPr
- Storage events to Eventhouse GA
- Anomaly detection PuPr
- Eventhouse Workspace level network isolation PuPr
- Activator Fabric Functions PuPr
- Activator Trigger Pane for DW
- Activator support for Power BI tables
- MSFT sources: M365 Graph data in RT hub PrPr
- Improved data exploration PuPr

Public Preview

Digital Twin Builder

Real-Time Intelligence

Real-Time Intelligence in Microsoft Fabric

Public Preview

Digital Twin Builder

Accelerate and simplify building digital representation of the physical world.

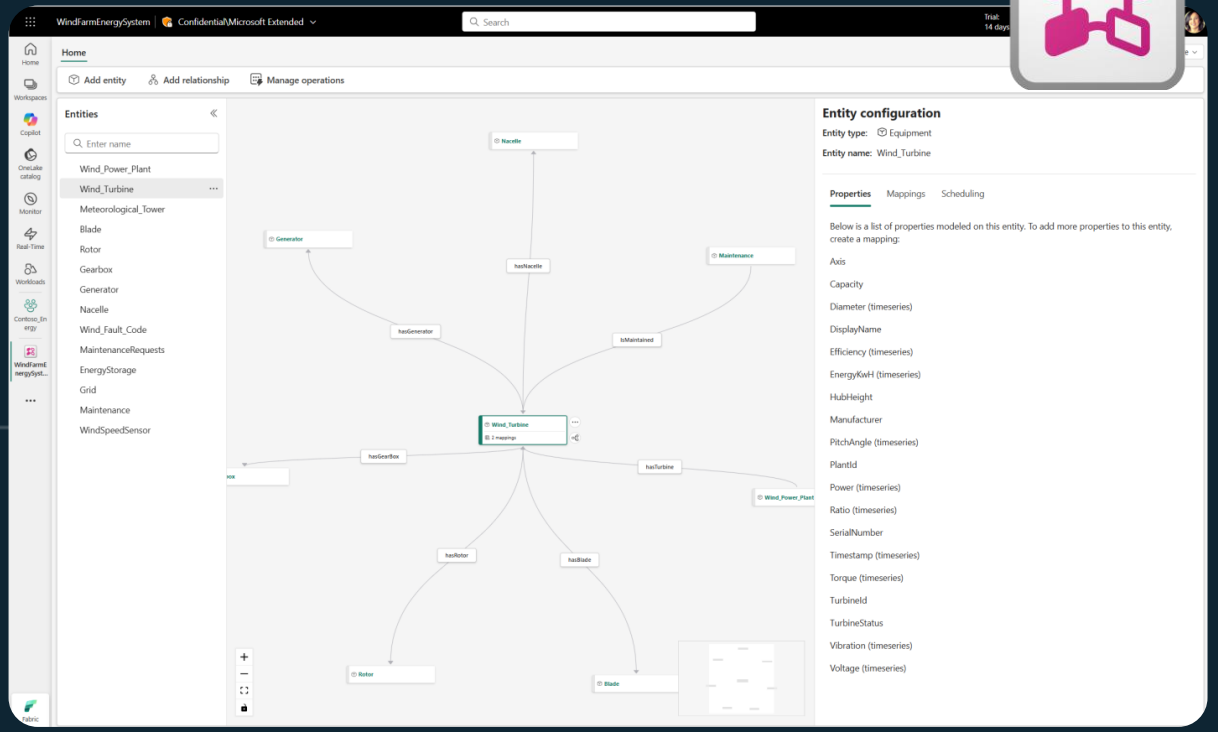

Digital Twin

A digital twin is an integrated data-driven virtual representation of real-world entities and processes, with synchronized interaction at a specified frequency and fidelity.

Build the data foundation for digital twins and AI solutions

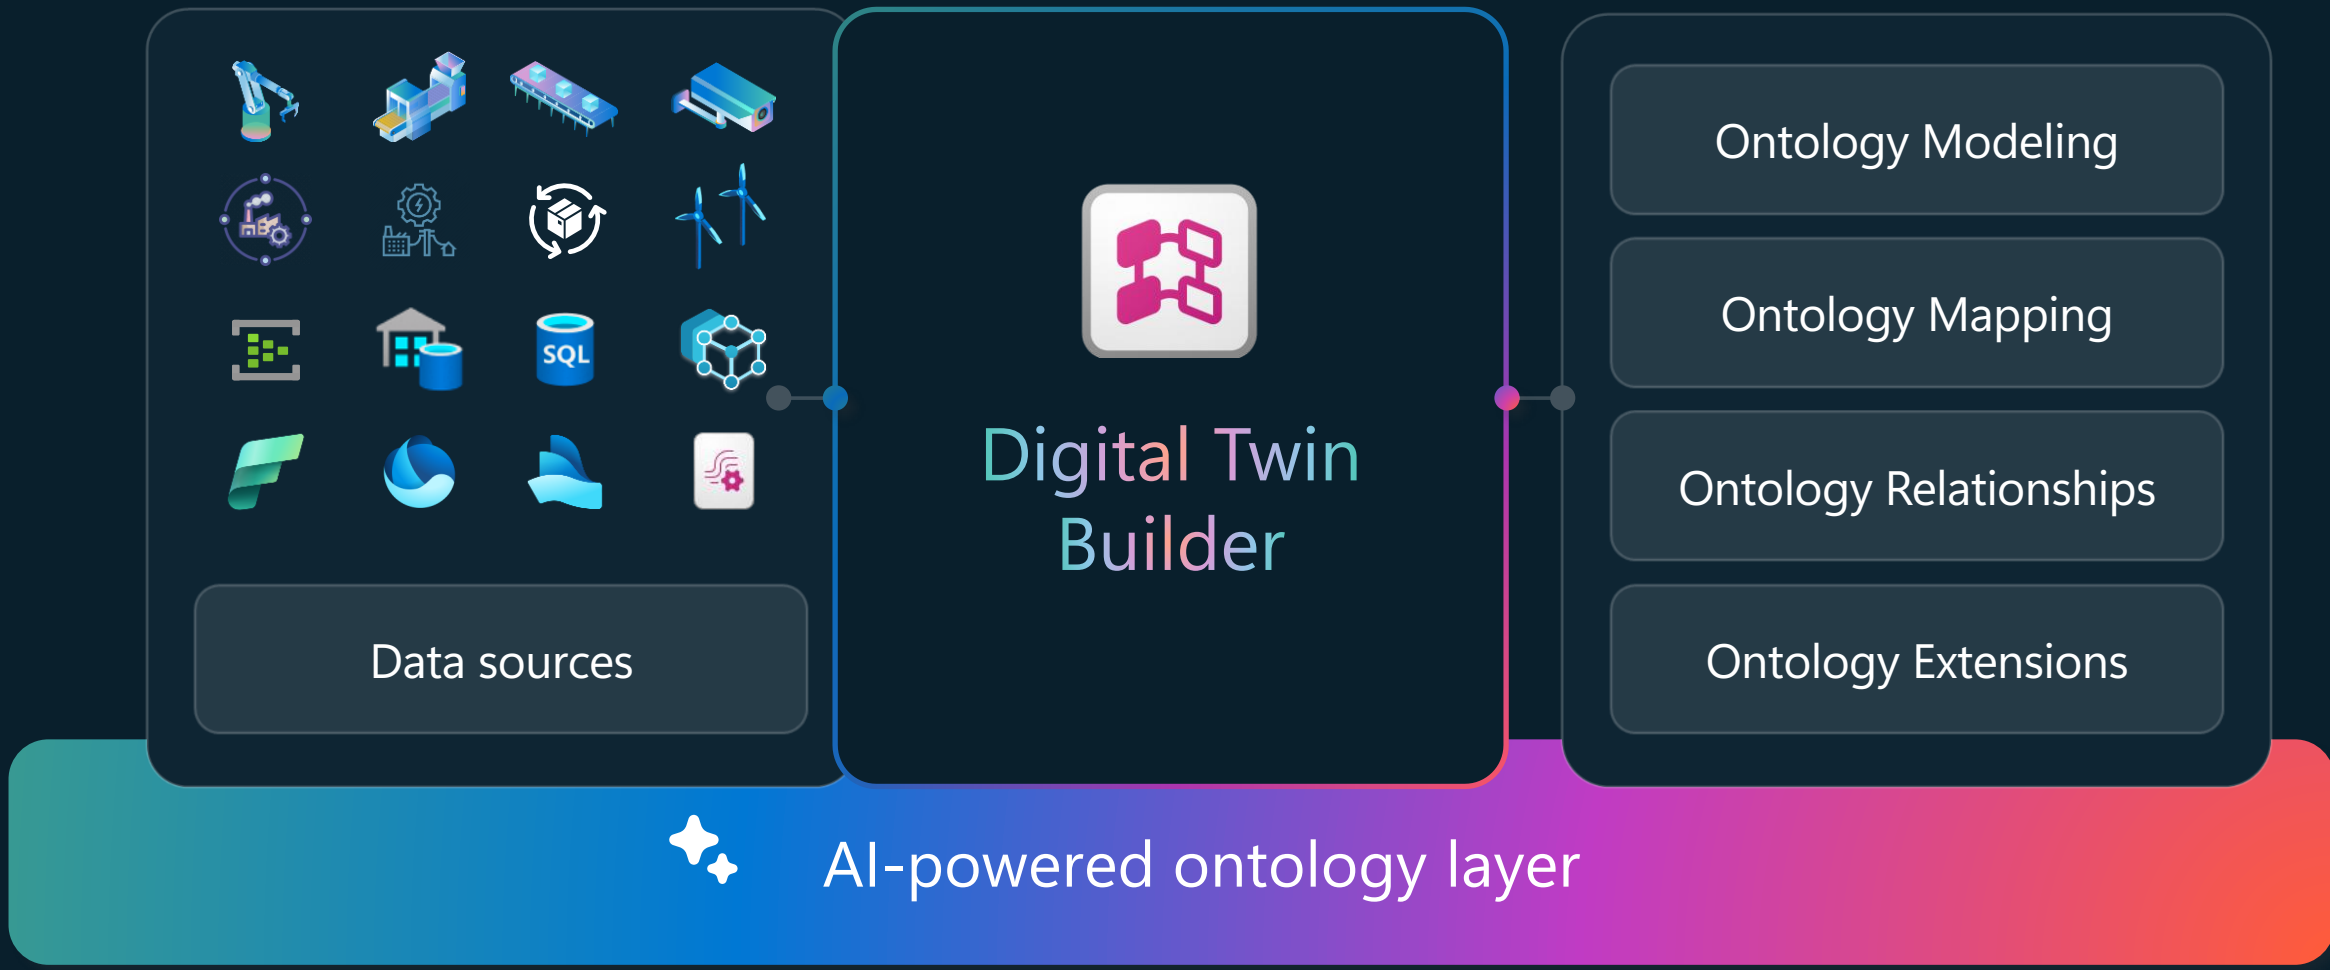

Digital twin builder

Public Preview

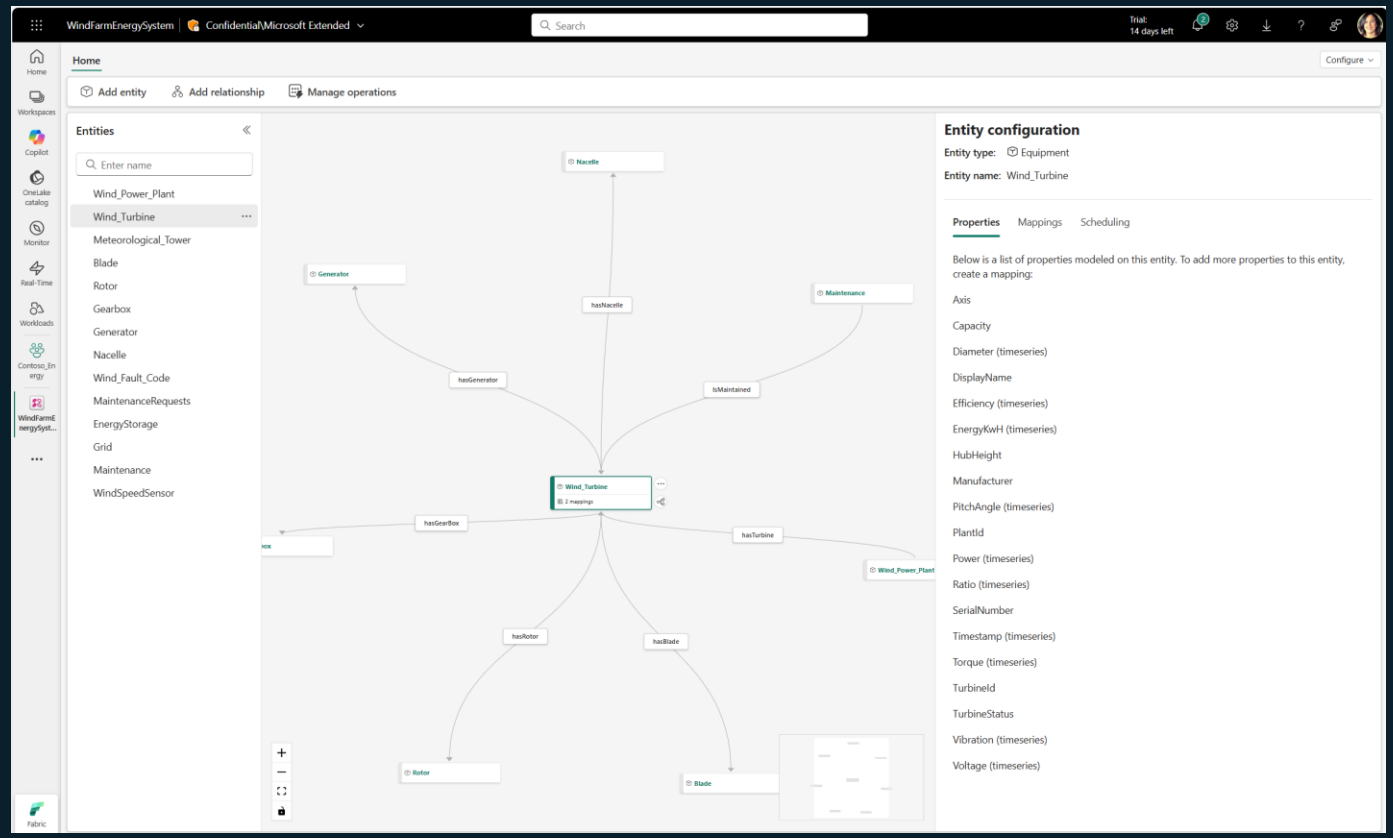

Simpler, faster way to build and manage digital twins with a low-code/no-code approach

Easily contextualize and map data to model assets, processes and systems and create a digital twin

Drill down to explore relationships and derive contextualized high value insights

Unlock insights with digital twin data in Power BI Copilot and Real-Time Dashboards

Energy Management

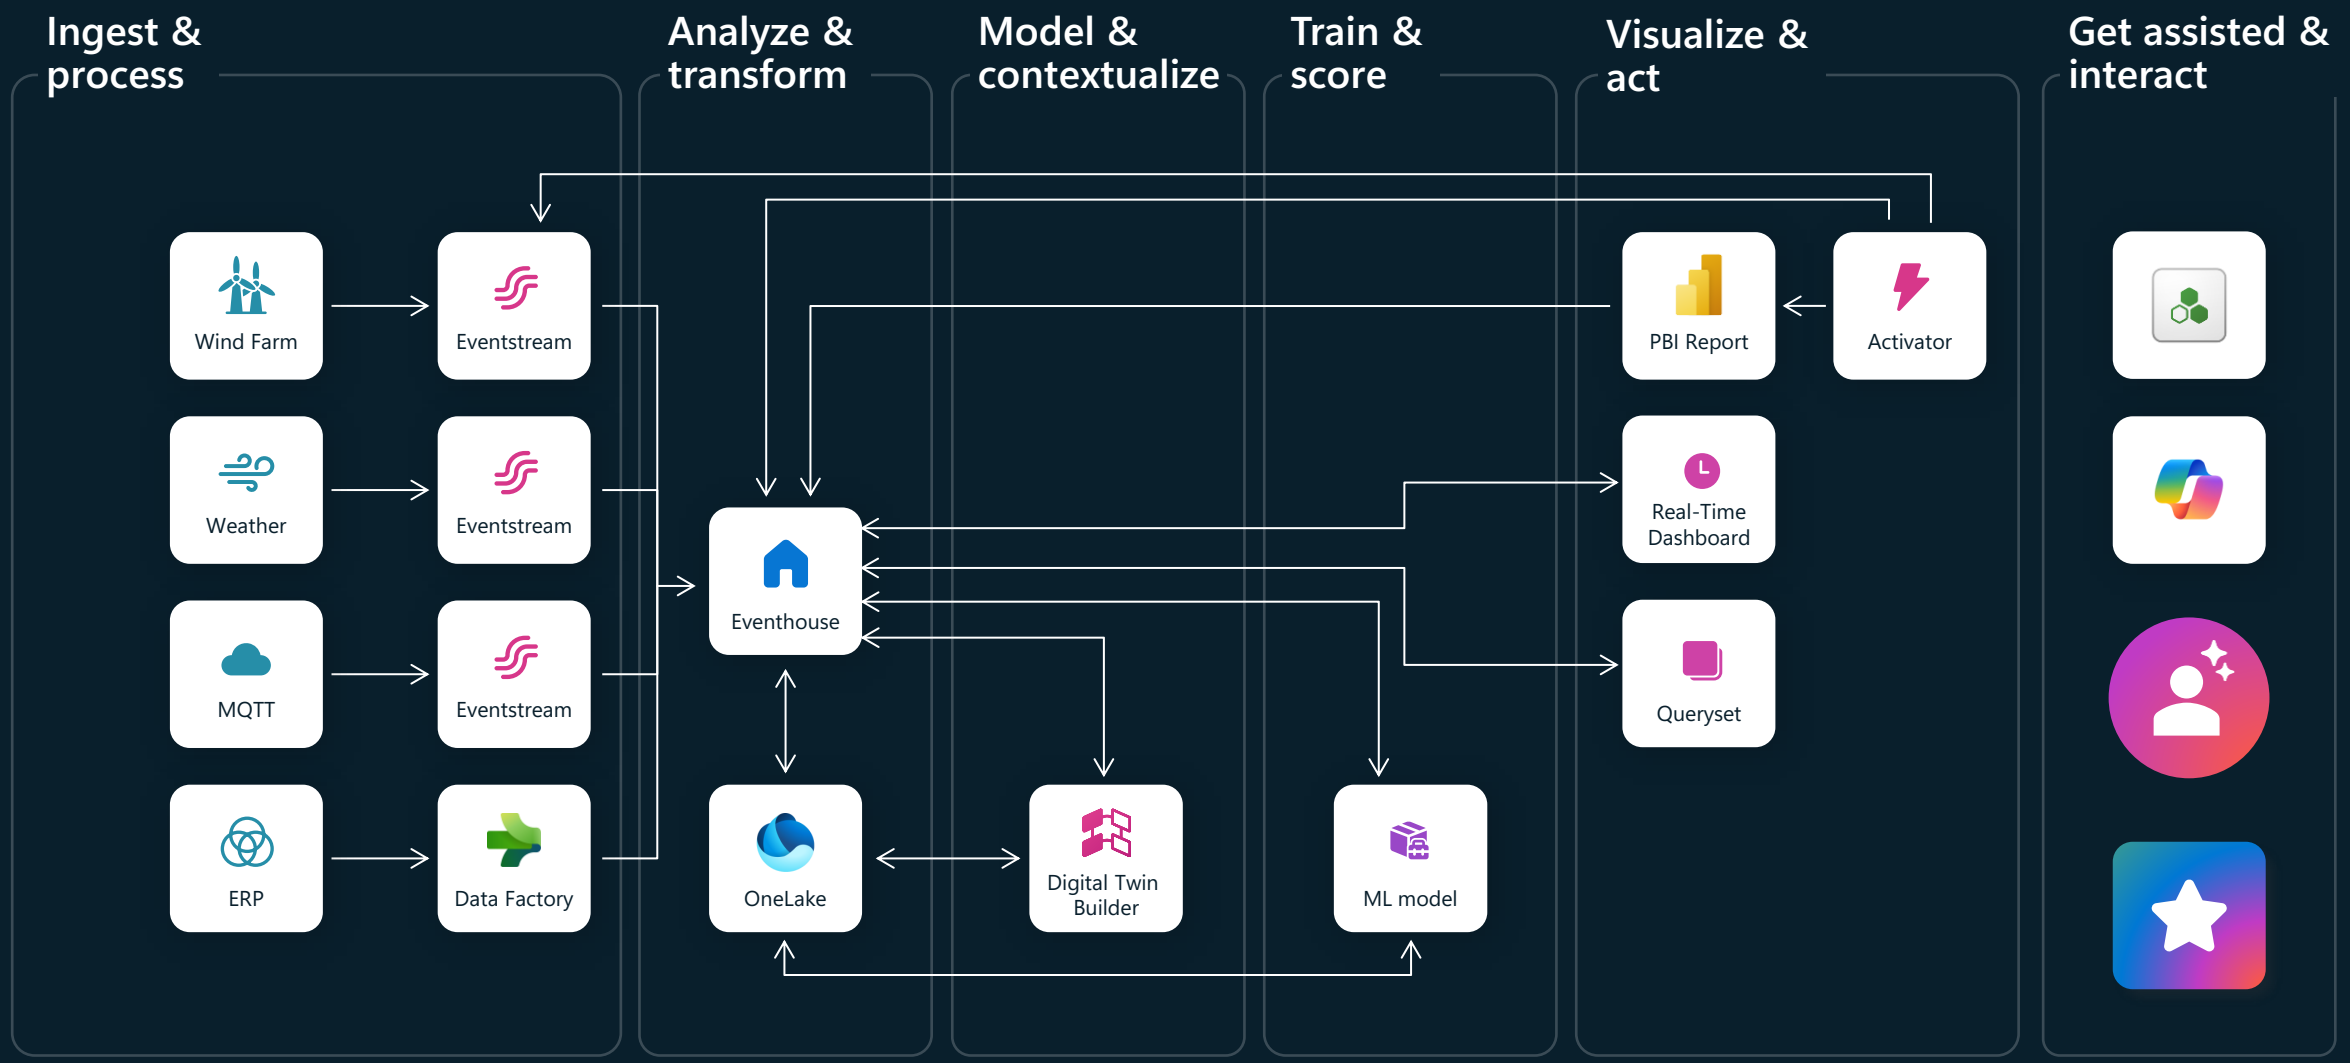

Private Preview

Anomaly Detection

Real-Time Intelligence

Finding high-value needles in a haystack

Anomaly detection

Monitor thousands, or even millions, of data points 24/7, well beyond human capacity

Tailor monitoring for each instance and uncover hidden trends

Detect and act on high-value anomalies in time to make a difference

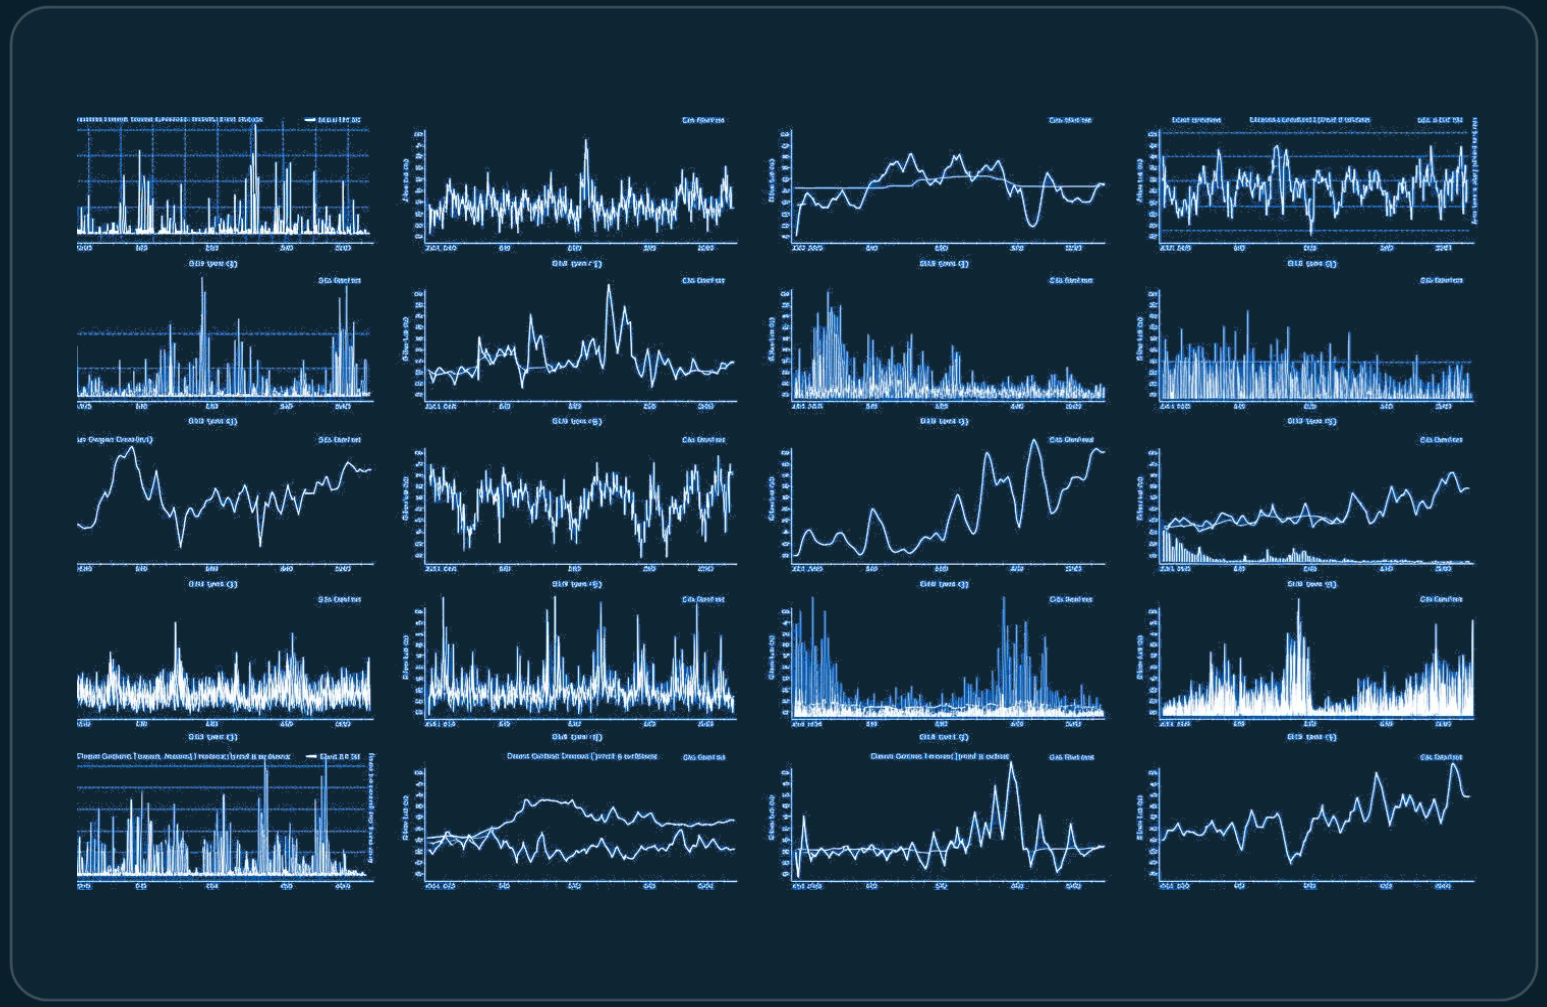

Elevating your operations to a whole new level

Sneak Peak

Operational agents

Teach AI-powered agents to observe, decide, notify, and act on your operations 24/7

Augment your operations team with AI agents, freeing them up from repetitive, reactive tasks to drive strategy and innovation

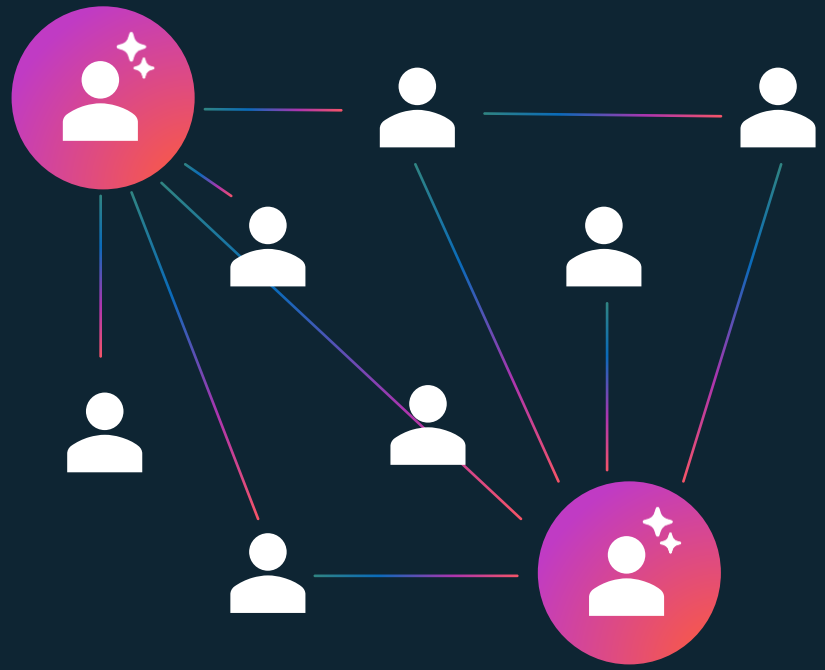

Building agents using retrieval-augmented generation

Real-Time RAG

Develop agents that use fresh, relevant knowledge on the business

Connect real-time data directly into RAG loop using Real-Time Intelligence

Combine the Eventhouse and its vector search with other capabilities in Fabric