



Recapping Microsoft Ignite Focused on Data & AI

Nicky van Vroenhoven | Powerdobs
Daan Humblé | Microsoft

23 November 2023



Nicky van Vroenhoven

Unit Lead Fabric & Power BI @ Powerdobs

Data Platform MVP





Daan Humblé

Cloud Solution Architect @ Microsoft


Power BI | Fabric

What happened at Ignite?



aka.ms/BookofNews

The 2023 ML, AI, and Data Landscape



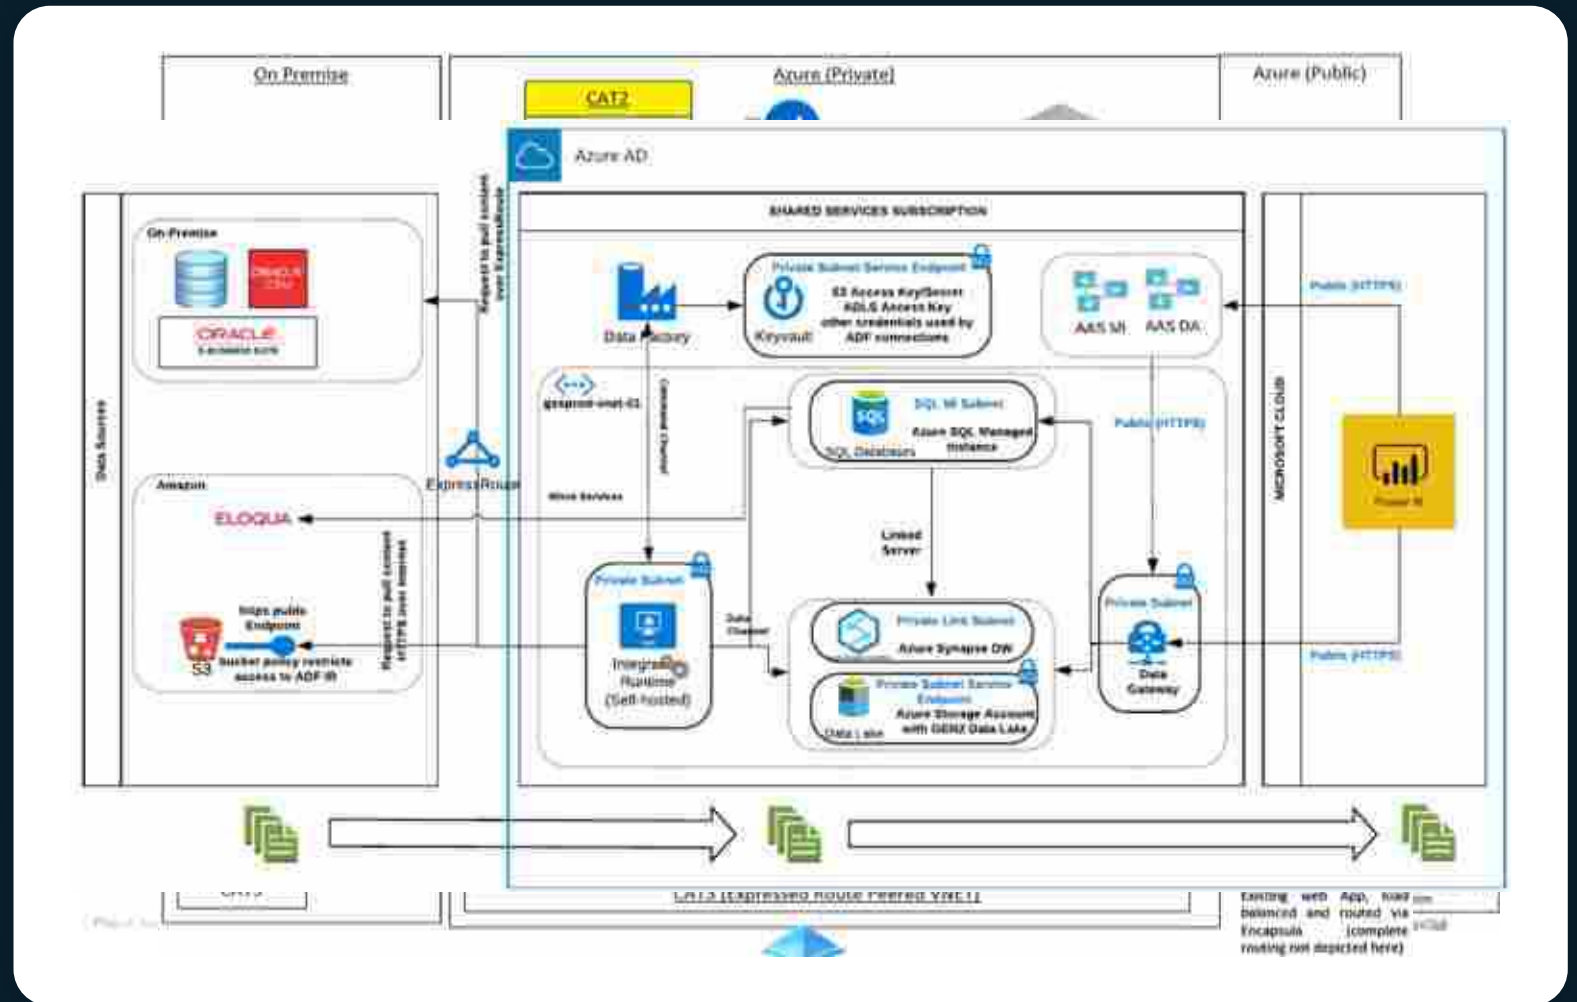
Analytics is complex and fragmented

Every project has many subsystems

Every subsystem need a different class of product

Products often comes from multiple vendors

Integration is complex, fragile and expensive



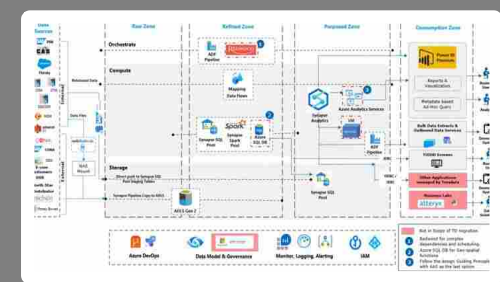

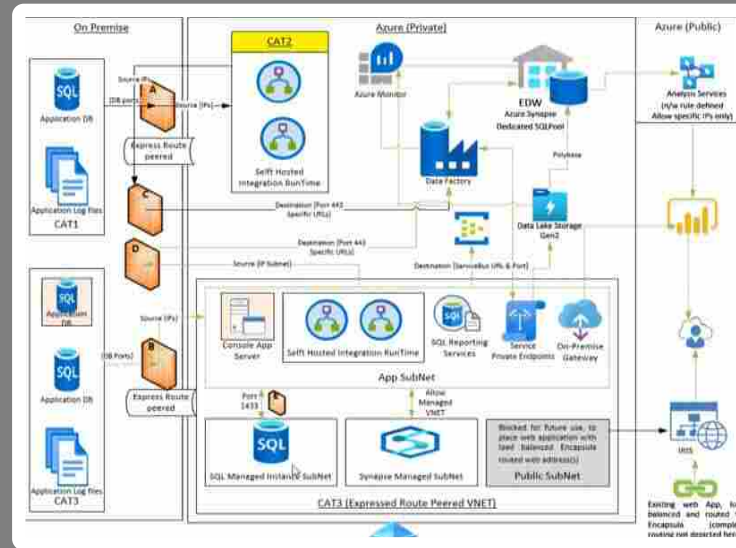

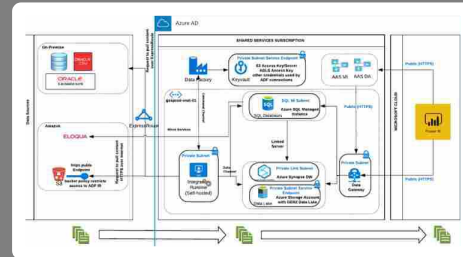
With siloes created everywhere

Applications are built by different teams

With diverse technology stacks and architectures, each designed as a standalone system

Making cross-app integration, and leveraging analytics in operations very hard

This chaotic complexity imperils the transition to the age of AI



//

I am the Chief Data Officer
and **don't want to be the**
Chief Integration Officer.

Every CDO, Every Enterprise

Public Preview May 2023



Microsoft Fabric

Data analytics for the era of AI



Microsoft Fabric

The data platform for the era of AI

From

One component

Single database

Isolated application

Gen AI wired in



To

Unified stack

All the data

Entire estate

Gen AI built in



Customers are using Fabric end to end

3+
Workloads

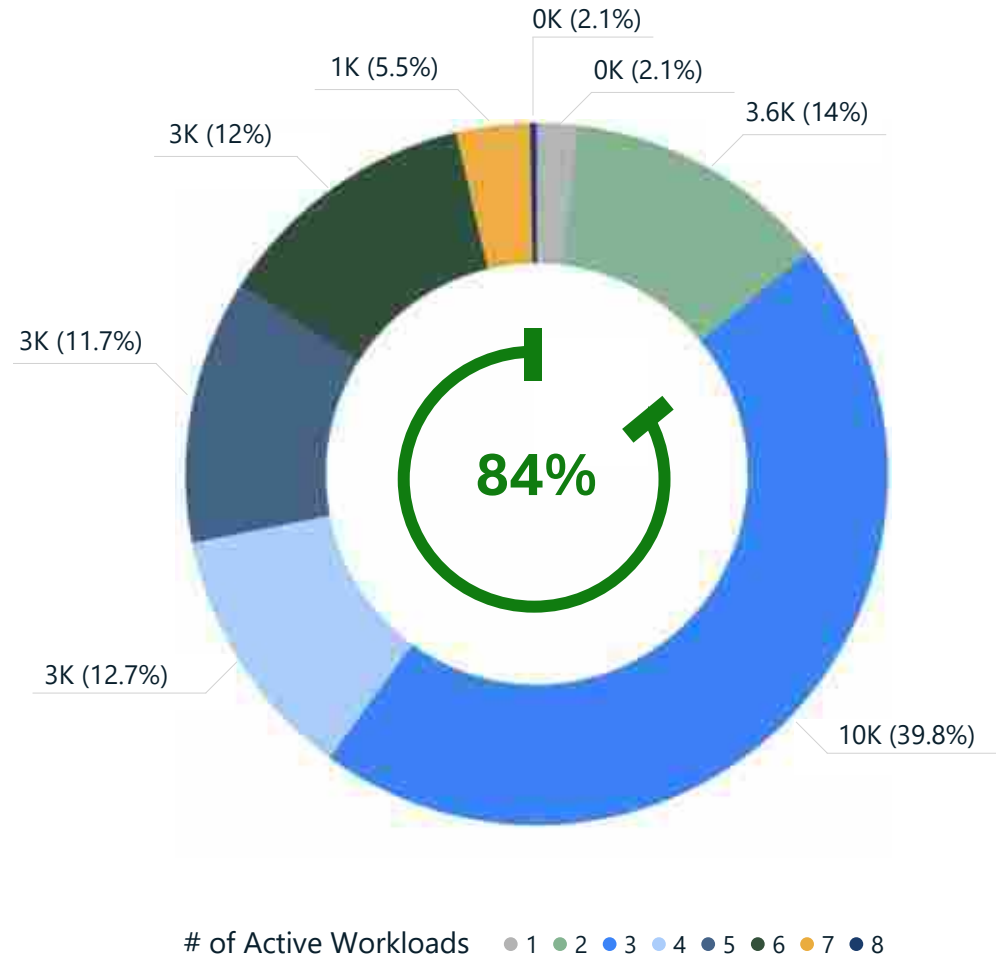


84%
of Customers

4+
Workloads



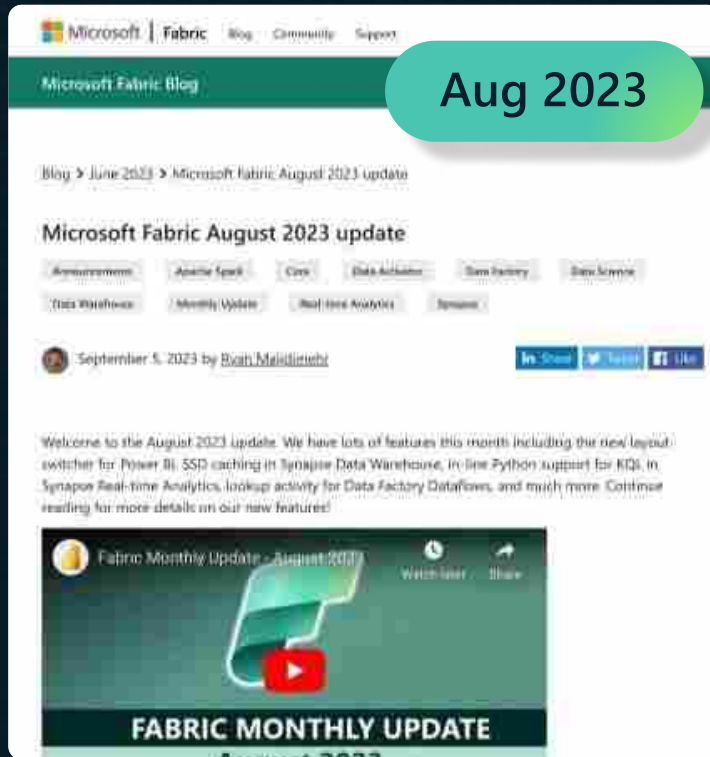
44%
of Customers



Weekly Fabric Releases

Rapid pace of innovation

Aug 2023



Microsoft Fabric Blog

Blog > June 2023 > Microsoft Fabric August 2023 update

Microsoft Fabric August 2023 update

Announcements | Apache Spark | Core | Data Explorer | Data Factory | Data Science | Data Warehouse | Monthly Update | Real-time Analytics | Support

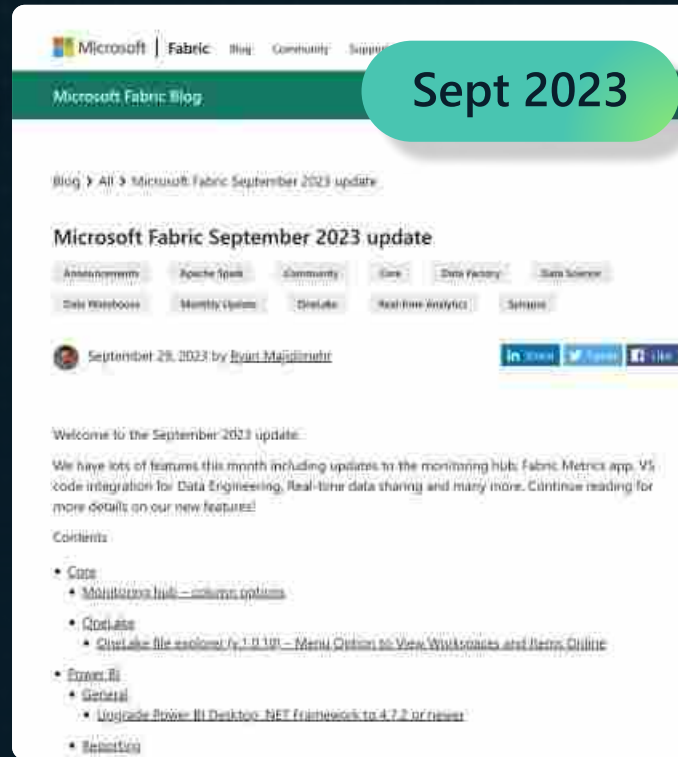
September 5, 2023 by Ryan Majumder

Welcome to the August 2023 Update. We have lots of features this month including the new layout switcher for Power BI, SSD caching in Synapse Data Warehouse, in-line Python support for RQL in Synapse Real-time Analytics, lookup activity for Data Factory Dataflows, and much more. Continue reading for more details on our new features!

Fabric Monthly Update - August 2023

FABRIC MONTHLY UPDATE

Sept 2023



Microsoft Fabric Blog

Blog > All > Microsoft Fabric September 2023 update

Microsoft Fabric September 2023 update

Announcements | Apache Spark | Community | Core | Data Factory | Data Science | Data Warehouse | Monthly Update | Orchestration | Real-time Analytics | Support

September 28, 2023 by Ryan Majumder

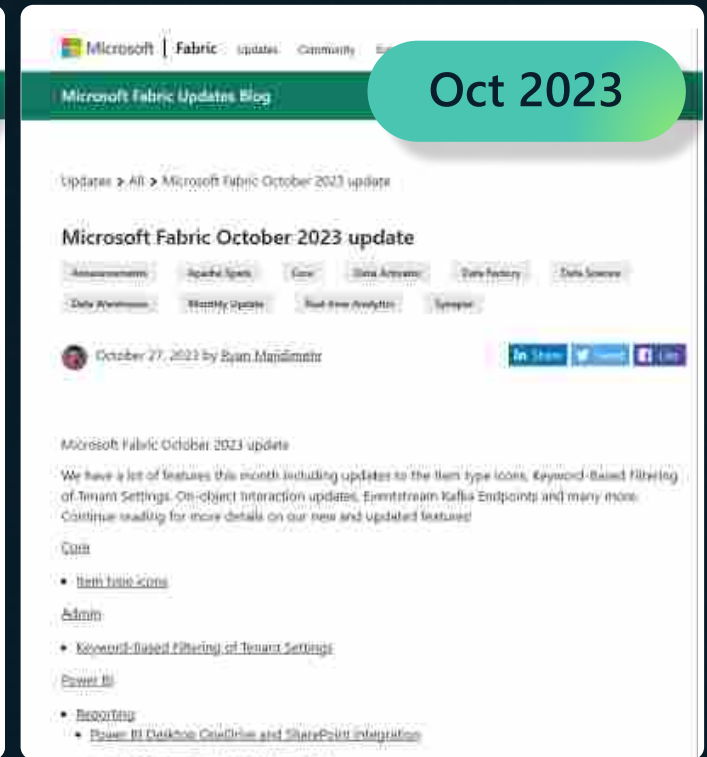
Welcome to the September 2023 update.

We have lots of features this month including updates to the monitoring hub, Fabric Metrics app, VS code integration for Data Engineering, Real-time data sharing and many more. Continue reading for more details on our new features!

Contents

- Core
 - [Monitoring hub - columns options](#)
- Orchestration
 - [Orchestration file resolution 1.0.100 - Menu Option to View Workspaces and Items Offline](#)
- Power BI
 - [General](#)
 - [Upgrade Power BI Desktop .NET framework to 4.7.2 or newer](#)
- Reporting

Oct 2023



Microsoft Fabric Updates Blog

Updates > All > Microsoft Fabric October 2023 update

Microsoft Fabric October 2023 update

Announcements | Apache Spark | Core | Data Explorer | Data Factory | Data Science | Data Warehouse | Monthly Update | Real-time Analytics | Support

October 27, 2023 by Ryan Majumder

Microsoft Fabric October 2023 update

We have a lot of features this month including updates to the item type icons, Keyword-based filtering of Tenant Settings, On-object interaction updates, Eventstream Kafka Endpoint and many more. Continue reading for more details on our new and updated features!

Core

- [Batch data sync](#)

Admin

- [Keyword-based filtering of Tenant Settings](#)

Power BI

- [Reporting](#)
 - [Power BI Desktop Credentials and SharePoint integration](#)

Weekly Fabric Releases

Rapid pace of innovation

The image displays three overlapping screenshots of Microsoft Fabric update blogs. The leftmost screenshot is for August 2023, the middle for September 2023, and the rightmost for October 2023. A central white tweet overlay from Andy Cutler (@MrAndyCutler) is positioned over the middle screenshot. The tweet reads: "That's a chonky update! I exported to pdf to read offline when travelling, 50 pages! Great work and thanks to all at MS. Still not a fan of the single blog but I think I'm in the minority #Fabric #Synapse". The background screenshots show various update details, including sections for "Contents" and "New" features.

Aug 2023

Microsoft Fabric Blog

Blog > June 2023 > Microsoft Fabric August 2023 update

Microsoft Fabric August 2023 update

Announcements | Azure Spark | Data | Data Archives | Data Warehouse | Monthly Update | Real-time Analytics | Synapse

September 5, 2023 by Ryan Madsen

Welcome to the August 2023 update. We have lots of features this month including updates to the Item type icons, Keyword-based filtering for interaction updates, Eventstream Kafka Endpoints and many more. For more details on our new and updated features:

FABRIC MONTHLY UPDATE

Sept 2023

Microsoft Fabric Blog

Blog > All > Microsoft Fabric September 2023 update

Contents

- Core
 - Monitorize task - column updates
- Querying
 - OnLake file explorer v1.0.130 - Menu Option to View Workspaces and Items Online
- Power BI
 - General
 - Upgrade Power BI Desktop .NET framework to 4.7.2 or newer
- Reporting

Oct 2023

Microsoft Fabric Updates Blog

Updates > All > Microsoft Fabric October 2023 update

October 2023 update

Announcements | Azure Spark | Data | Data Archives | Data Warehouse | Monthly Update | Real-time Analytics | Synapse

Team Management

October 2023 update

This month including updates to the Item type icons, Keyword-based filtering for interaction updates, Eventstream Kafka Endpoints and many more. For more details on our new and updated features:

Contents

- Core
 - Item type icons
- Admin
 - Keyword-based filtering of Tenant Settings
- Power BI
 - Reporting
 - Power BI Desktop Condition and SharePoint integration



Microsoft Fabric

Generally Available, Nov 15, 2023

aka.ms/IgniteFabric2023



Microsoft Fabric

The data platform for the era of AI



Data
Factory



Synapse Data
Engineering



Synapse Data
Science



Synapse Data
Warehouse



Synapse Real
Time Analytics



Power BI



Data
Activator



OneLake



Microsoft Fabric

Data analytics for the era of AI

Complete Analytics Platform

Everything, unified

SaaS-ified

Secured and governed

Lake Centric and Open

OneLake

One copy

Open at every tier

Empower Every Business User

Familiar and intuitive

Built into Microsoft 365

Insight to action

AI Powered

Copilot accelerated

Gen AI on your data

AI-driven insights



Microsoft Fabric

Data analytics for the era of AI

Complete Analytics Platform

Everything, unified

SaaS-ified

Secured and governed

Lake Centric and Open

OneLake

One copy

Open at every tier

Empower Every Business User

Familiar and intuitive

Built into Microsoft 365

Insight to action

AI Powered

Copilot accelerated

Gen AI on your data

AI-driven insights

SaaS

"It just works"

5x5

Frictionless onboarding

Instant provisioning

Quick results w/ Intuitive UX

Success by default

Minimal knobs

Auto optimized

Auto integrated

Centralized administration

Tenant-wide governance

Centralized
security management

Compliance built-in



Microsoft Fabric



Data
Factory



Synapse Data
Engineering



Synapse Data
Science



Synapse Data
Warehouse



Synapse Real
Time Analytics



Power BI



Data
Activator

AI Assisted

Shared Workspaces

Universal Compute Capacities

One Security

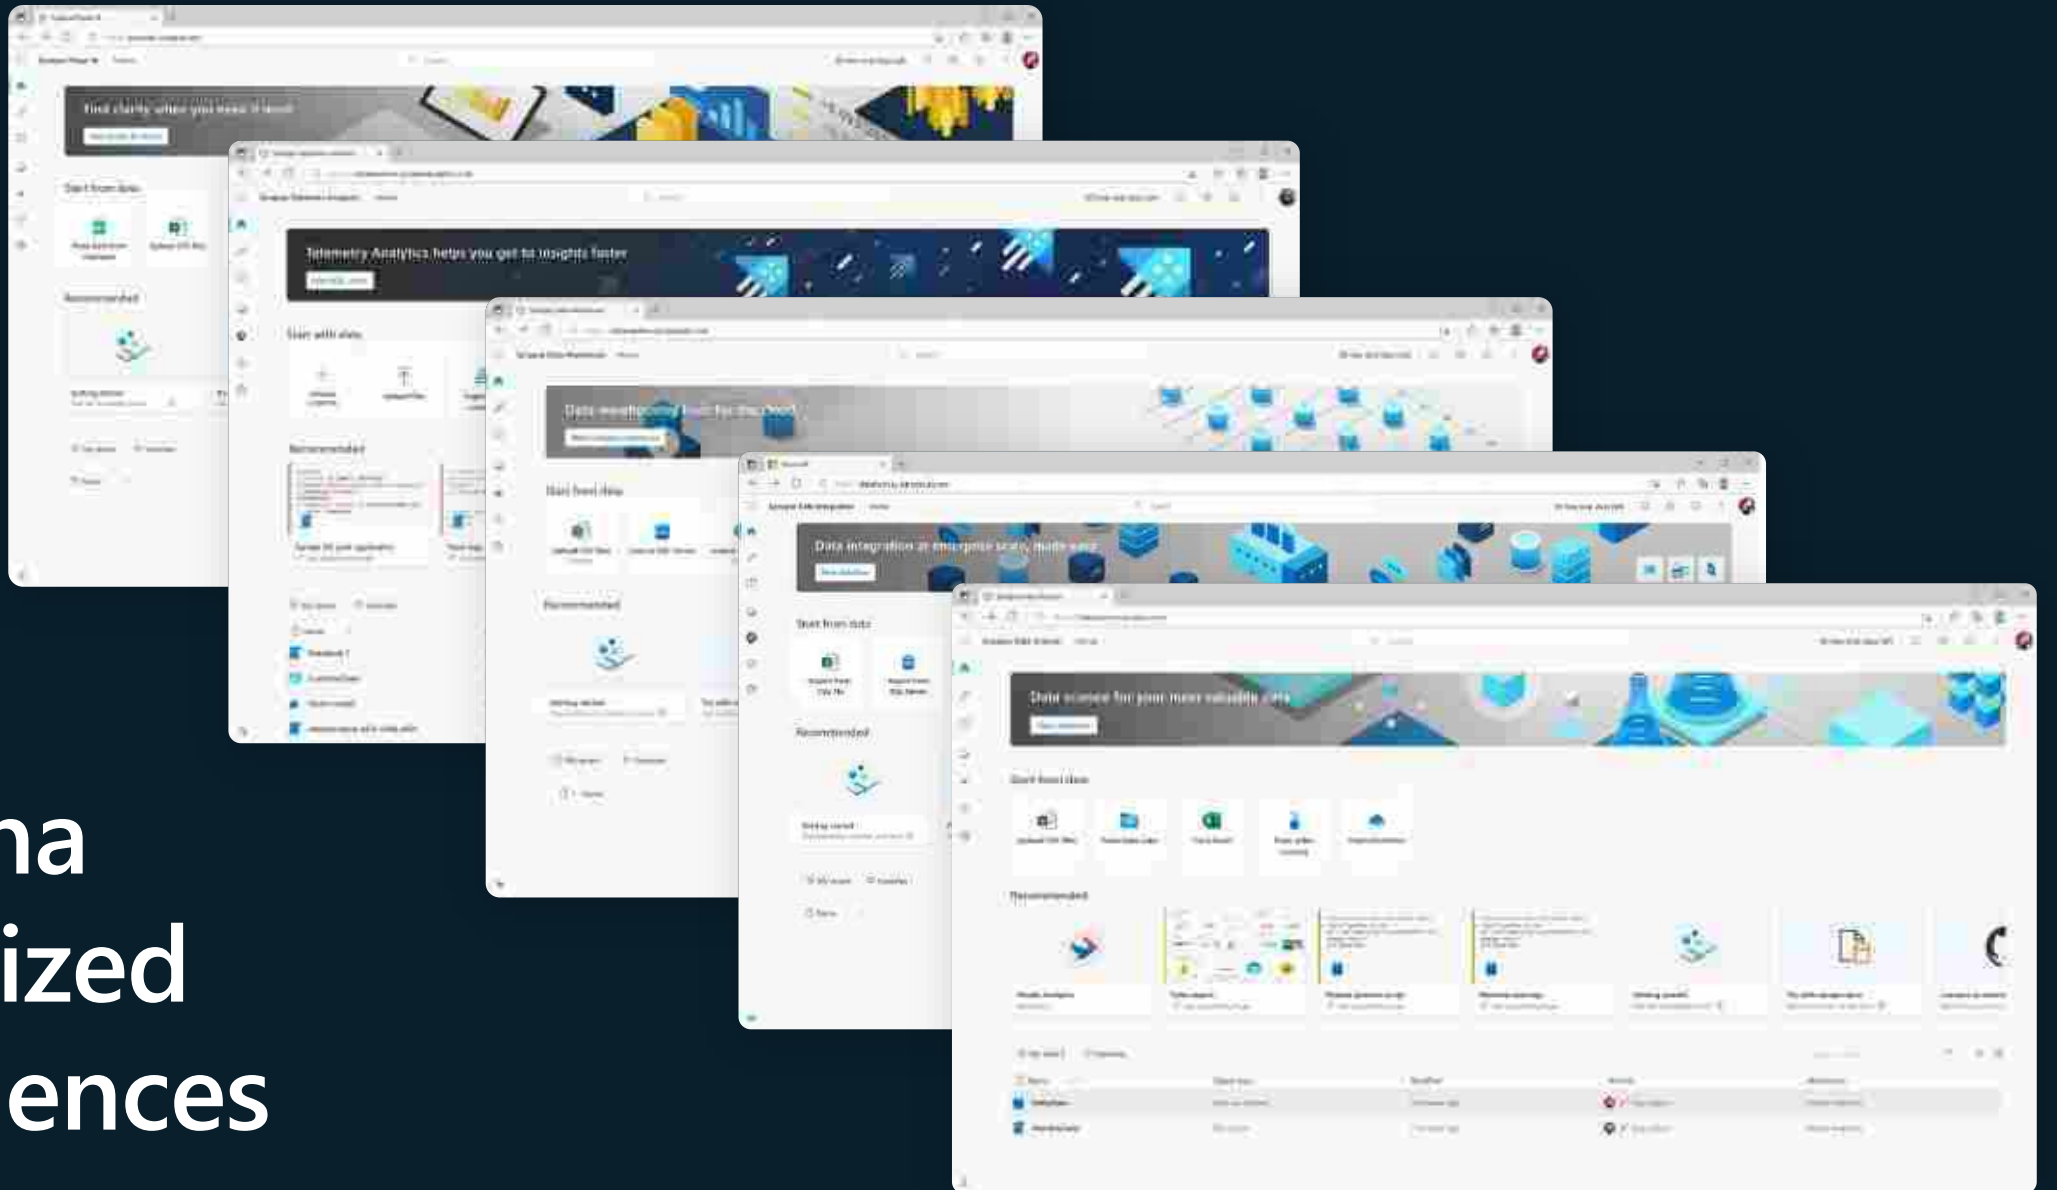
OneLake

Intelligent data foundation

Single...

- Onboarding and trials
- Sign-on
- Navigation model
- UX model
- Workspace organization
- Collaboration experience
- Data Lake
- Storage format
- Data copy for all engines
- Security model
- CI/CD
- Monitoring hub
- Data Hub
- Governance & compliance

Persona
optimized
experiences



**Demo: Every data workload in one SaaS
experience**



Microsoft Fabric

Data analytics for the era of AI

Complete Analytics Platform

Everything, unified

SaaS-ified

Secured and governed

Lake Centric and Open

OneLake

One copy

Open at every tier

Empower Every Business User

Familiar and intuitive

Built into Microsoft 365

Insight to action

AI Powered

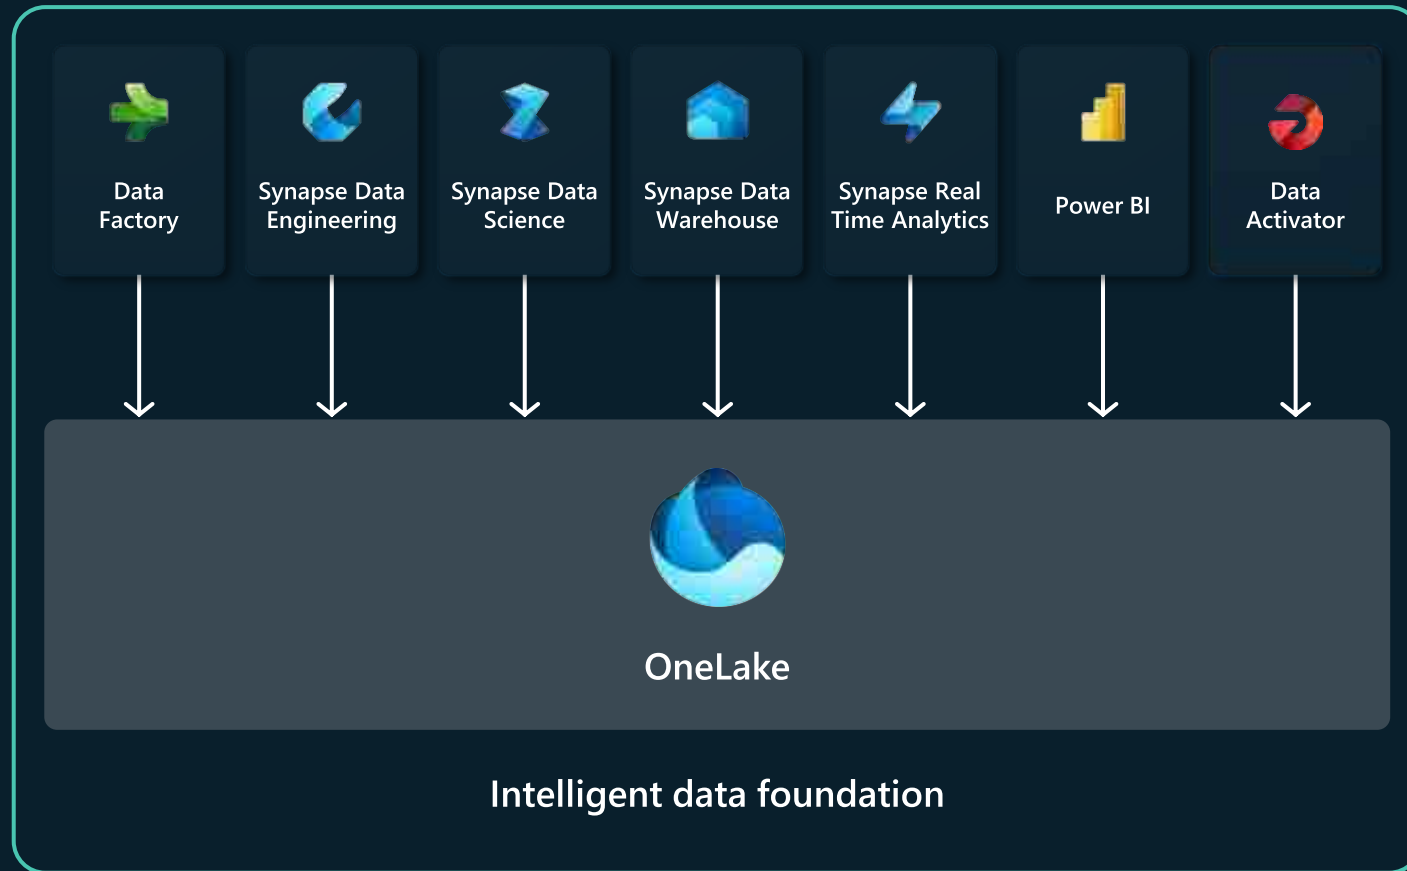
Copilot accelerated

Gen AI on your data

AI-driven insights

OneLake for all Data

“The OneDrive for Data”



A single SaaS lake for the whole organization

Provisioned automatically with the tenant

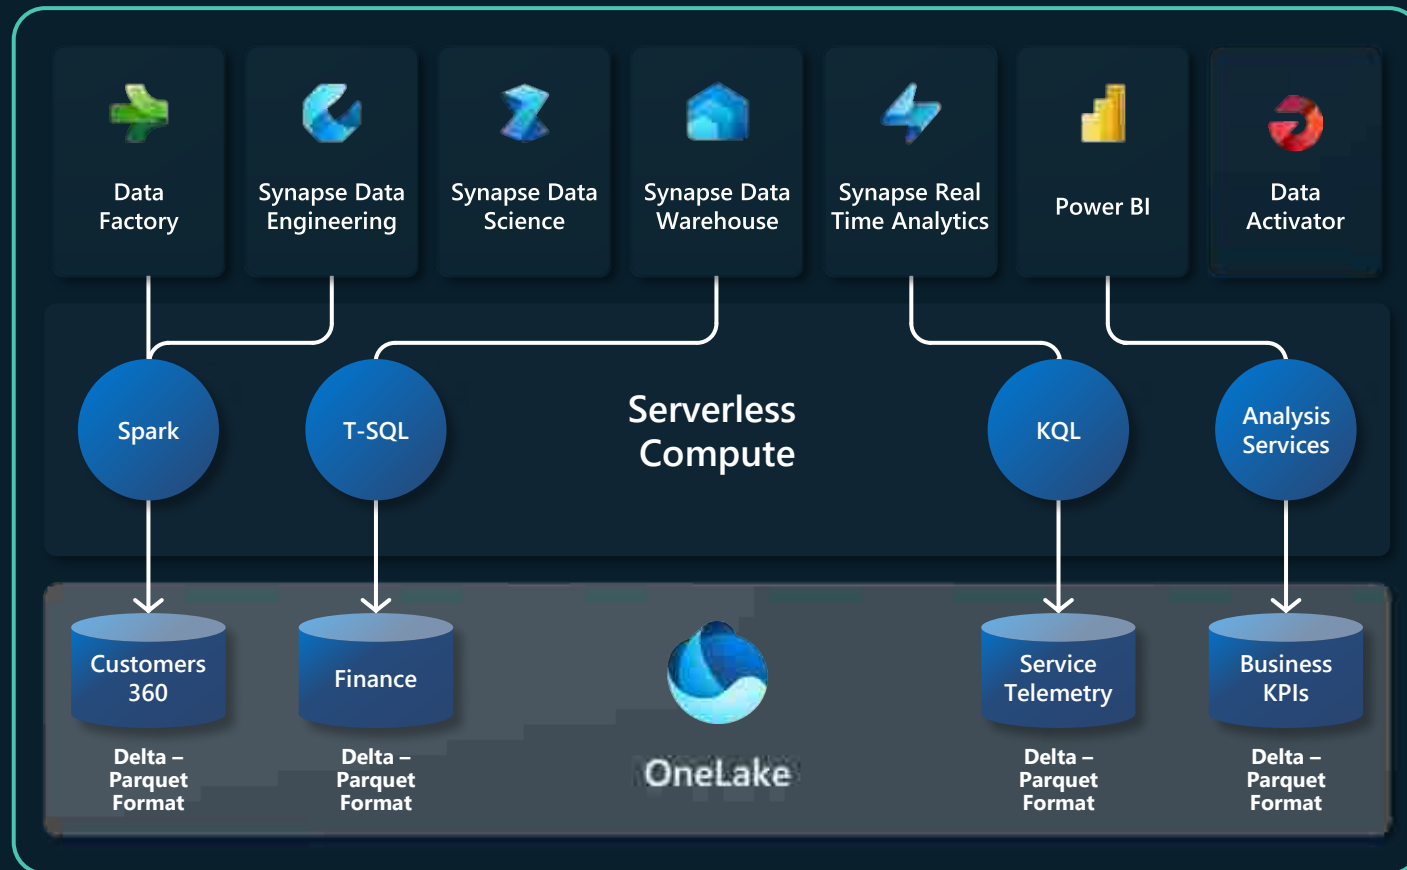
All workloads automatically store their data in the OneLake workspace folders

All the data is organized in an intuitive hierarchical namespace

The data in OneLake is automatically indexed for discovery, MIP labels, lineage, PII scans, sharing, governance and compliance

One Copy for all computes

Real separation of compute and storage



All the compute engines store their data automatically in OneLake

The data is stored in a single common format

Delta – Parquet, an open standards format, is the storage format for all tabular data in Analytics vNext

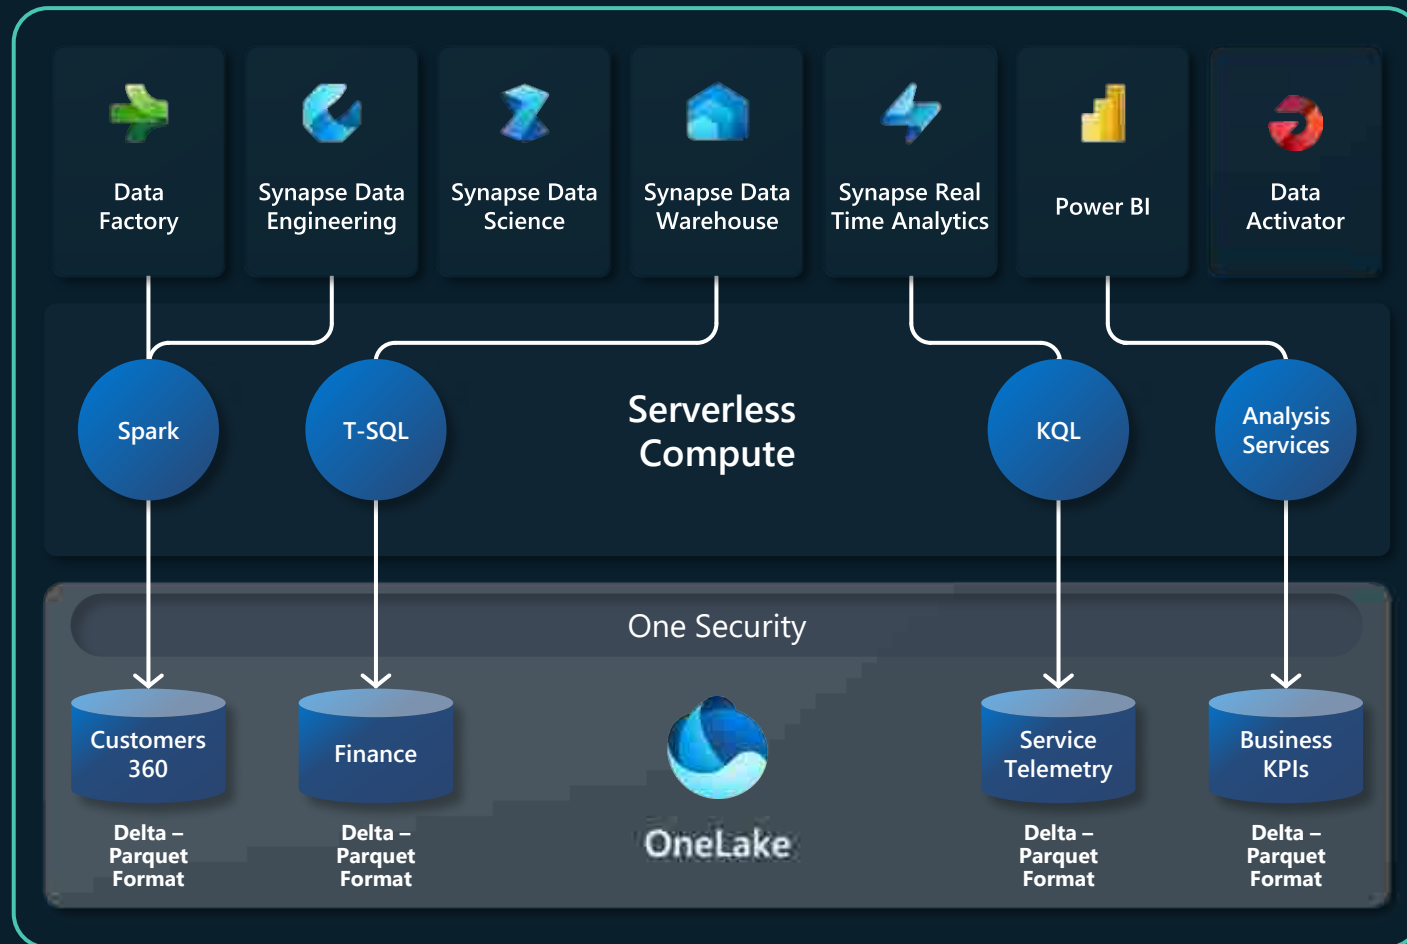
Once data is stored in the lake, it is directly accessible by all the engines without needing any import/export

All the compute engines have been fully optimized to work with Delta Parquet as their native format

Shared universal security model is enforced across all the engines

One Copy for all computes

Universal security makes it real



All the compute engines store their data automatically in OneLake

The data is stored in a single common format

Delta - Parquet, an open standards format, is the storage format for all tabular data in Analytics vNext

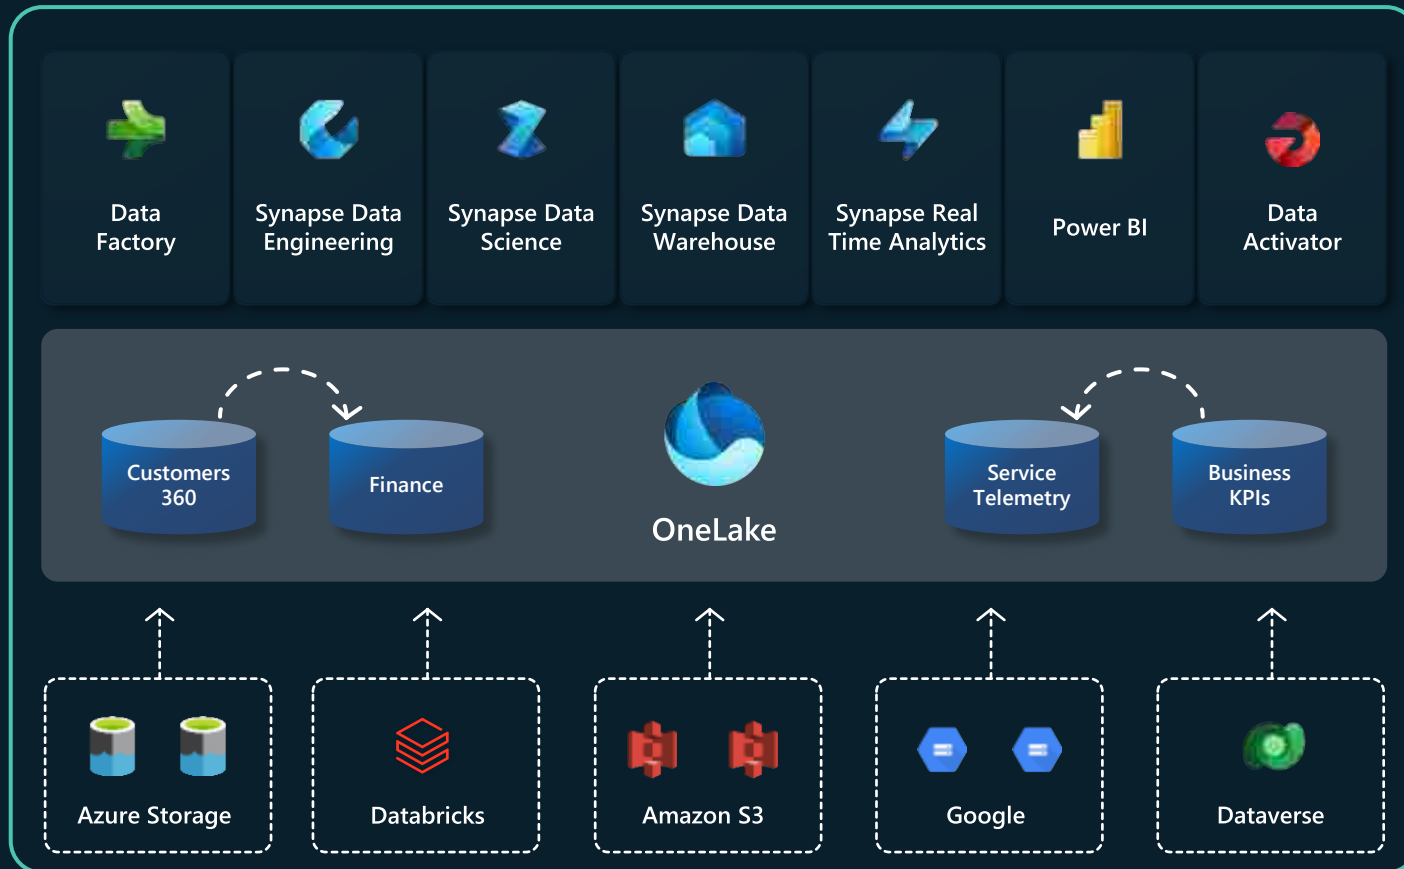
Once data is stored in the lake, it is directly accessible by all the engines without needing any import/export

All the compute engines have been fully optimized to work with Delta Parquet as their native format

Shared universal security model is enforced across all the engines

Unifying data in OneLake

Shortcuts



Sharing data in OneLake is as easy as sharing files in OneDrive, removing the needs for data duplication

With **shortcuts**, data throughout OneLake can be composed together without any data movement

Shortcuts also allow instant linking of data already existing in Azure and in other clouds, without any data duplication and movement, making **OneLake the first multi-cloud data lake**

With support for industry standard APIs, OneLake data can be directly accessed by any application or service

Announcing



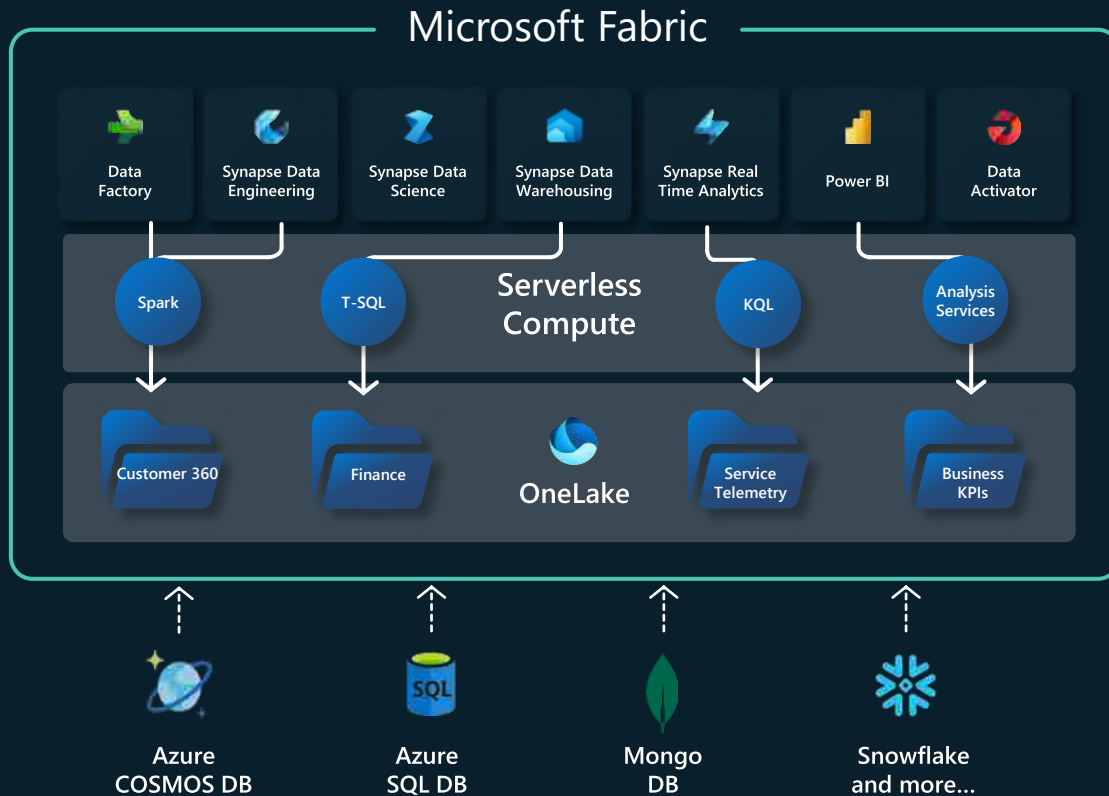
Mirroring

aka.ms/IgniteFabric2023

Mirroring in Microsoft Fabric

Simplify near real-time analytics

Fabric compute engines



Mirroring

Fabric Mirroring enables adding existing databases and data warehouses to Fabric without any ETL.

A full editing experience of the source database is available for the Mirrored database.

Data is replicated into OneLake in Delta format and kept up-to-date in near-real-time.

All the Fabric experiences instantly work with the OneLake replica.

Analysts and Data Scientists can work with real-time data.

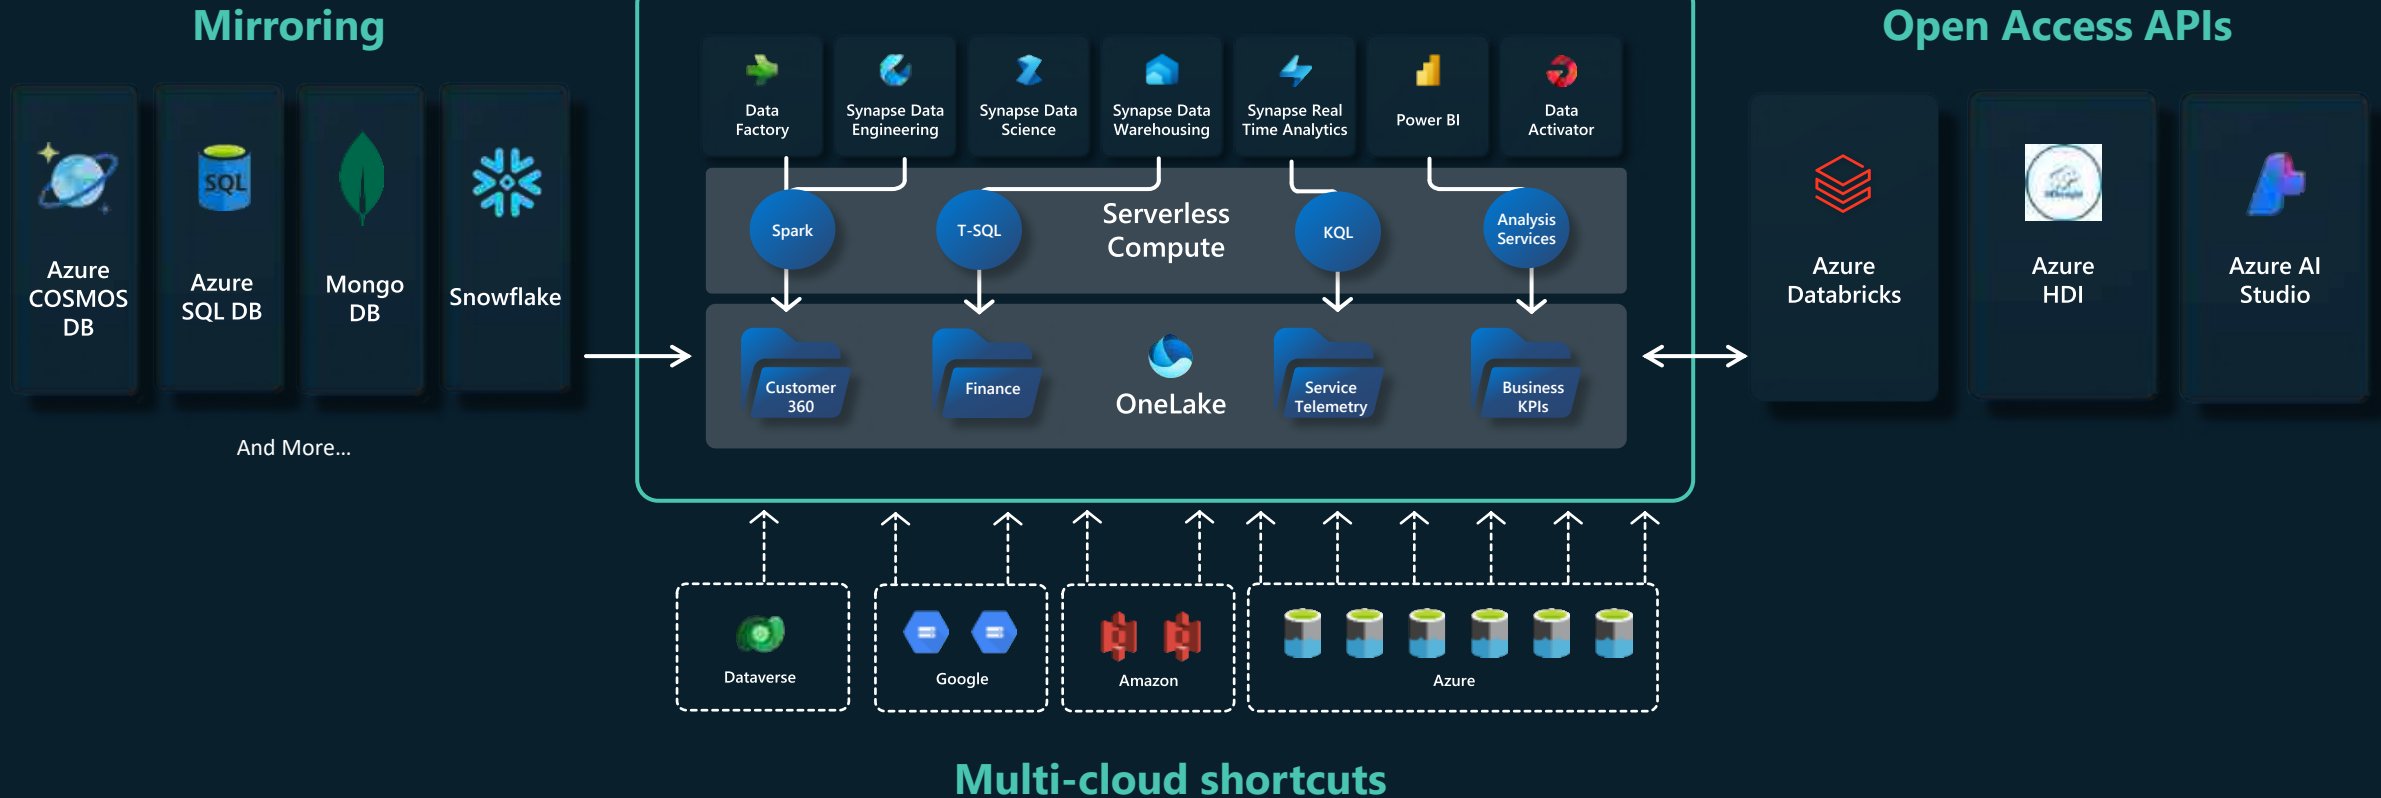
The replica protects operational databases from analytical queries.

Microsoft Fabric

Connect your entire data estate

Fabric compute engines

Microsoft Fabric




Microsoft Fabric

Introducing Direct Lake Mode


“Direct Query Mode”

Slow, but real time



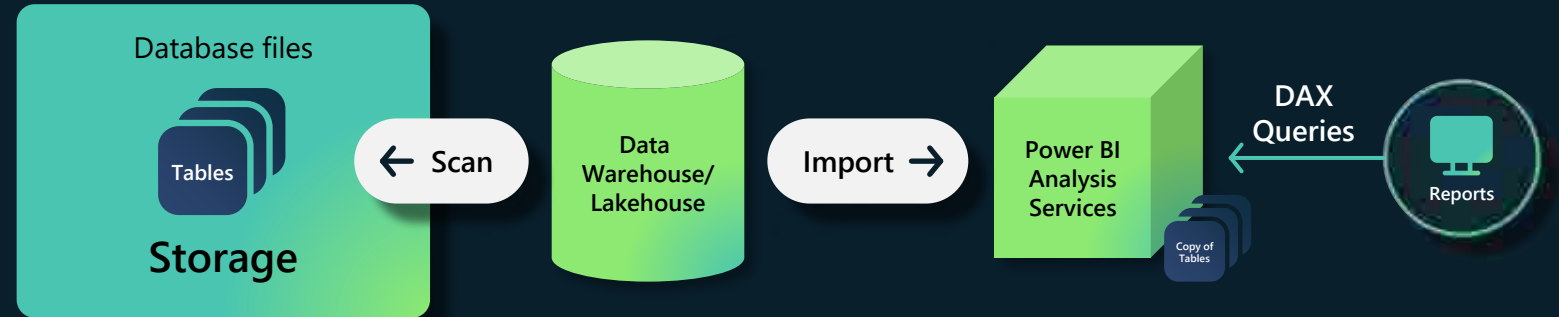
"Direct Query Mode"

Slow, but real time



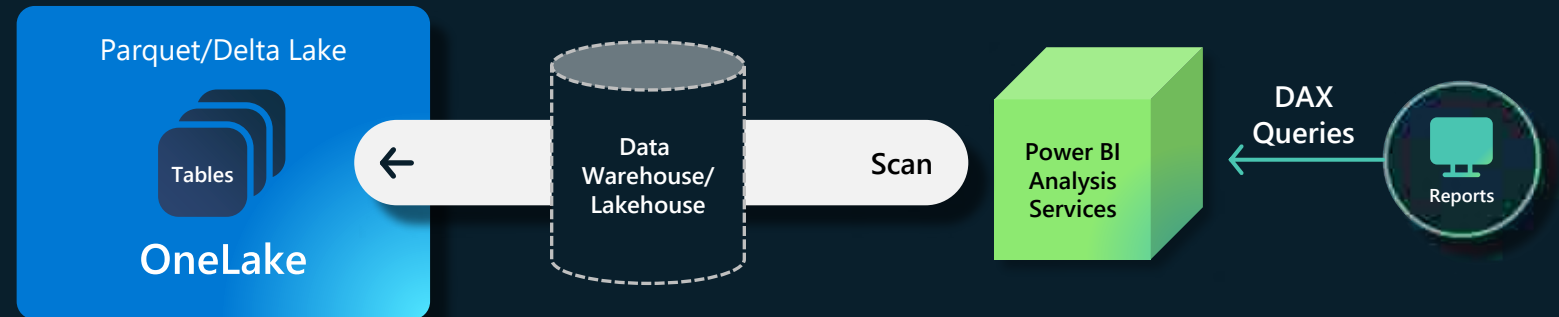
"Import Mode"

Latent & duplicative but fast



"Direct Lake Mode"

Perfect!



Direct Lake mode

On start, no data is loaded in-memory

Column data is transcoded from Parquet files when queried

Tables can have mix of transcoded (resident) and non-resident

Column data can get evicted over time

Direct Lake fallback to Direct Query for suitable sub-queries

"Framing" of dataset determines what gets loaded from DeltaLake

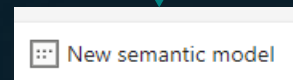
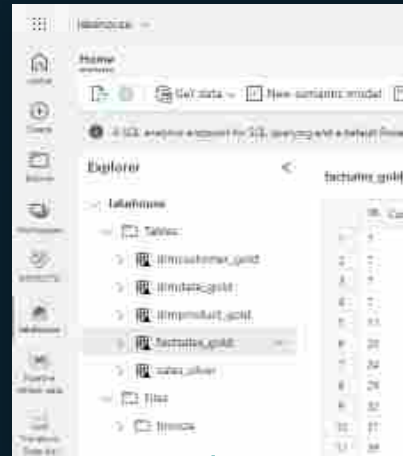
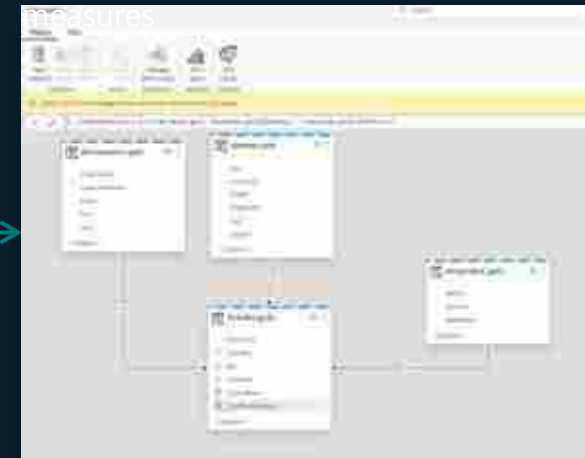



Table selection



Edit model : create relationships and measures



Create the report



Why Delta (Parquet)?

Open Standard for file format

Column oriented, efficient data storage and retrieval

Efficient Data Compression and Encoding

Becoming the Industry Standard

Thrives on bulk operations

Well suited for pruning (Column, rowgroup)

Introducing V-Ordering

Write time optimization to parquet files


Sorting, row group distribution, dictionary encoding, and compression (Shuffling)

Complies to the open standard

Z-Order, compaction, vacuum, time travel, etc. are compatible with V-Order

V-ordering in action

Microsoft Internal DB (162 tables)



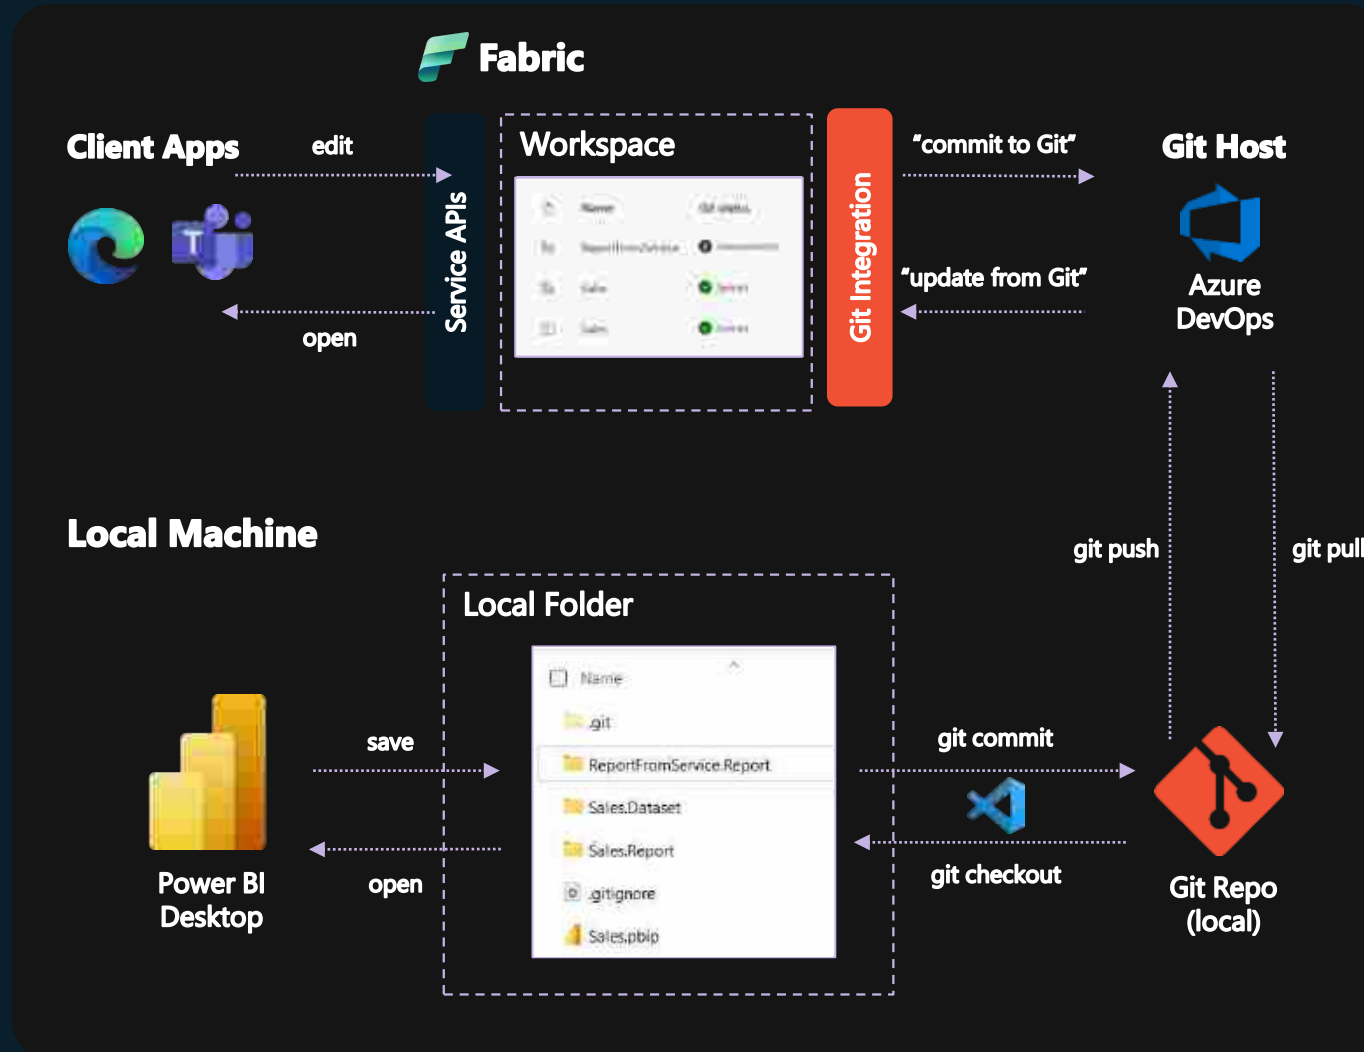
x3.2

Reduced IO for workloads

Demo: Direct Lake



Microsoft Fabric Developer Experience



Demo: Developer experiences



Microsoft Fabric

Data analytics for the era of AI

Complete Analytics Platform

Everything, unified

SaaS-ified

Secured and governed

Lake Centric and Open

OneLake

One copy

Open at every tier

Empower Every Business User

Familiar and intuitive

Built into Microsoft 365

Insight to action

AI Powered

Copilot accelerated

Gen AI on your data

AI-driven insights



Microsoft Fabric

Data analytics for the era of AI

Complete Analytics Platform

Everything, unified

SaaS-ified

Secured and governed

Lake Centric and Open

OneLake

One copy

Open at every tier

Empower Every Business User

Familiar and intuitive

Built into Microsoft 365

Insight to action

AI Powered

Copilot accelerated

Gen AI on your data

AI-driven insights

Announcing



Copilot in Microsoft Fabric

aka.ms/IgniteFabric2023

Copilot in Microsoft Fabric

Generative AI for every step of your data journey



Unlock Your Data

Intelligent data prep and
integration for every data
source



Enhance Productivity

Harness the power of AI
to turn your ideas into
solutions



Accelerate Insight Discovery

Analyze and explore
your data faster



Microsoft Fabric

A recap

Generally Available

OneLake

Data Factory

Synapse Data Engineering

Synapse Data Science

Synapse Data Warehousing

Synapse Real Time analytics

OneLake Shortcut: ADLS2; AWS-S3; DV

Fabric SKU 1Y RI (F-SKU)

Power BI

Public Preview

Copilot in Microsoft Fabric**

Data Activator

Private Preview

Snowflake Mirroring

Azure SQL Mirroring

Azure Cosmos DB Mirroring

Next steps



Attend Fabric sessions

aka.ms/IgniteData2023



Read all the announcements

aka.ms/bookofnews



Try Microsoft Fabric

aka.ms/try-Fabric




Explore the Fabric Roadmap

aka.ms/FabricRoadmap

Bonus!

Unified business model

- "Pay for the average, not the peak"



How it works

Capacity units (CUs) = Compute power

Capacity units (CUs) are units of measure that represent a pool of compute power needed. Compute power is required to run queries, jobs, or tasks.

CU Consumption

The CU consumption is highly correlated to the underlying compute effort needed for the tasks performed by the capability during the processing time.

Each capability, such as Power BI, Spark, Data Warehouse, with the associated queries, jobs, or tasks has a unique consumption rate.

Note: For additional documentation on purchasing experience and capacity models, please see here: aka.ms/fabriclicensing

Translation for existing Power BI Premium per capacity Customers


Today	
SKU	PBI vCore
P1	8
P2	16
P3	32
P4	64
P5	128



Microsoft Fabric Generally Available	
SKU	CU
F 64	64
F 128	128
F 256	256
F 512	512
F 1024	1024

- Power BI Premium P SKUs get **access to the full Fabric capabilities**
- 8 CUs are providing the same compute power of 1 Power BI Premium v-core for Power BI workloads. That is, **F64 provides the equivalent compute power of Power BI Premium P1**

For additional documentation on Fabric licensing types, please see here: [Microsoft Fabric concepts – Microsoft Fabric | Microsoft Learn](#)



Translation for existing Power BI Premium per capacity Customers

COMPUTE

Today	
SKU	PBI vCore
<i>not available</i>	
A1	1
P1	8
P2	16
P3	32



Microsoft Fabric Generally Available	
SKU	CU
F 2	2
F 8	8
F 64	64
F 128	128
F 512	512




Microsoft Fabric Pricing*	
SKU	CU
€ 303	€ 180
€ 1.214	€ 722
€ 9.715	€ 5.777
€ 19.431	€ 11.554
€ 77.726	€ 46.219

1 CU pay-as-you-go price at West Europe
~€ 0.21/hour

Visit [Fabric pricing page](#) for more details.

For additional documentation on Fabric licensing types, please see here: [Microsoft Fabric concepts – Microsoft Fabric | Microsoft Learn](#)



Demo: Capacity Metrics

OneLake pricing

OneLake is a data lake built into Microsoft Fabric and provides a single place to store all organizational data. Data storage is charged at a rate of \$ per GB per month and priced uniquely across regions.



Data Storage

Type

Pay-as-you-go price at West US 2

OneLake storage

\$0.023 per GB / month

OneLake BCDR storage

\$0.0414 per GB / month

OneLake cache

\$0.2 per GB / month



Data Transfer & Internet Egress

Cross-region data transfer network charges may apply based on source/destination of each storage access. Learn more at the [Bandwidth Pricing](#).

Note: OneLake cache is billed for KQL cache storage and Data Activator data retained.

New Microsoft Credentials



Get ready!

Try the
Fabric Cloud
Skills Challenge

aka.ms/igniteCSC



Microsoft Certifications

Microsoft Certified: Fabric
Analytics Engineer Associate

Exam DP-600: Implementing
Analytics Solutions
Using Microsoft Fabric

BETA in mid-Jan



Microsoft Applied Skills

Fabric
credentials

Covering scenarios like real-time
analytics, data lakehouses, and data
warehouses using Microsoft Fabric

NEW in early 2024