Microsoft

Analytics for the future

Håkan Forsberg, PSA Data/AI @ Microsoft

SQL & Azure SQL

The family of SQL cloud to edge databases

SQL Server on Azure Virtual Machines

Best for lift and shift and/or workloads requiring OS-level access

Azure SQL Managed Instance

Best for modernizing existing apps

Azure SQL Database

Best for supporting modern cloud apps

Azure SQL Edge

Best for extending apps to IoT edge

Infrastructure-as-a-Service | **Platform-as-a-Service** | **Edge Computing**

Azure SQL enabled by Azure Arc

Run Azure SQL on premises and in multicloud environments

Azure is the cloud that knows SQL Server best

Azure Arc-enabled SQL Server architecture

1. Onboard Arc-enabled server
2. Onboard Arc-enabled SQL Server
3A. Enroll Microsoft Defender and secure SQL Server
3B. Enroll SQL Best Practice assessment

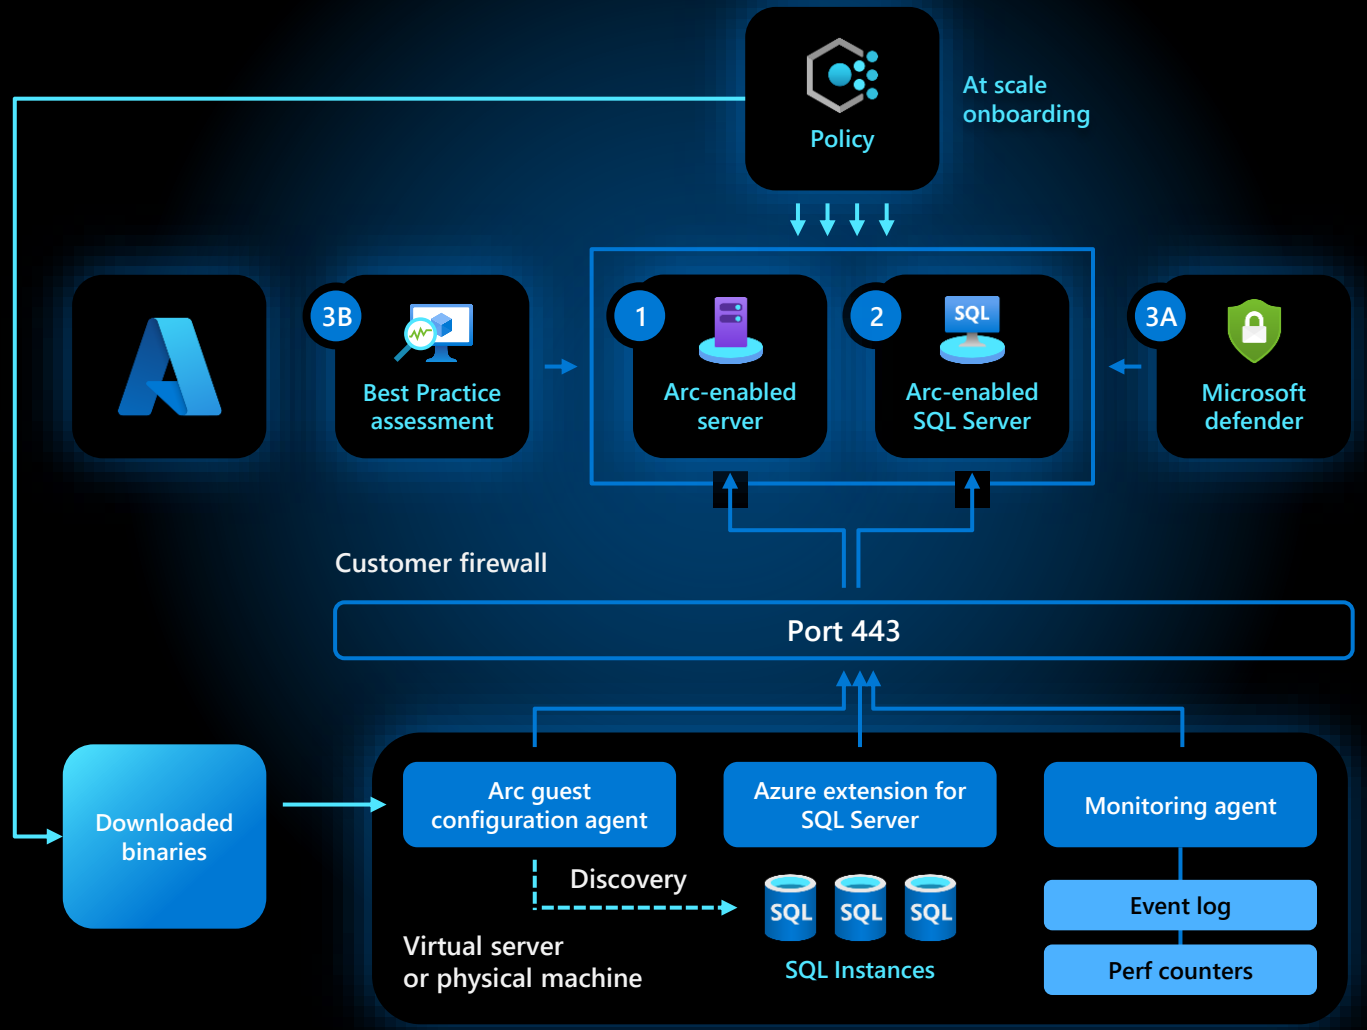

Single view of all SQL Servers from Azure Portal

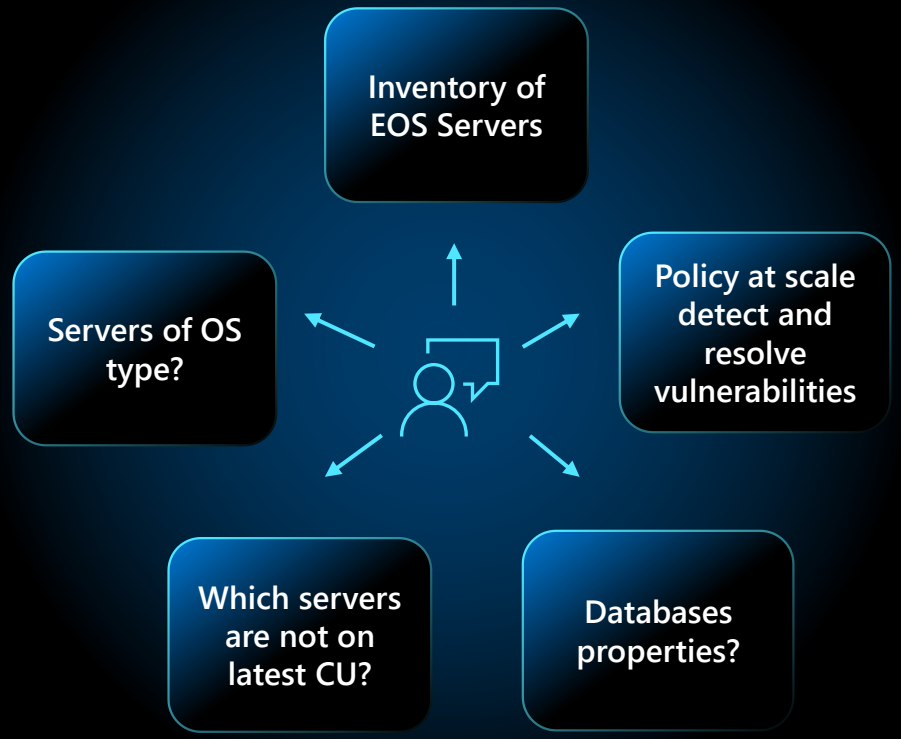

Inventory Management

Single consistent view of all your SQL Servers deployed on-prem, Edge, Multi-cloud

Inventory and tag management using Resource Graph thus increasing the visibility of the entire data estate

License management using Azure portal to review license position and compare with the procurement state

Asset Management

Receive Extended Security Updates (ESU) for reduced price through Arc-enabled SQL Servers.

Govern, Protect, configure your hybrid and multi-cloud servers with Azure Policy, Defender and Azure Automation, centrally, securely and at scale

Fully automated technical assessment for SQL Server

Evaluate your configuration of SQL Server:

- Security and compliance
- Availability and business continuity
- Performance and scalability
- Operations
- Change and configuration management

Scanned in intervals for most up to date results

Empower DBAs to proactively address any risks

Increases operational stability while reducing routine workloads from DBAs

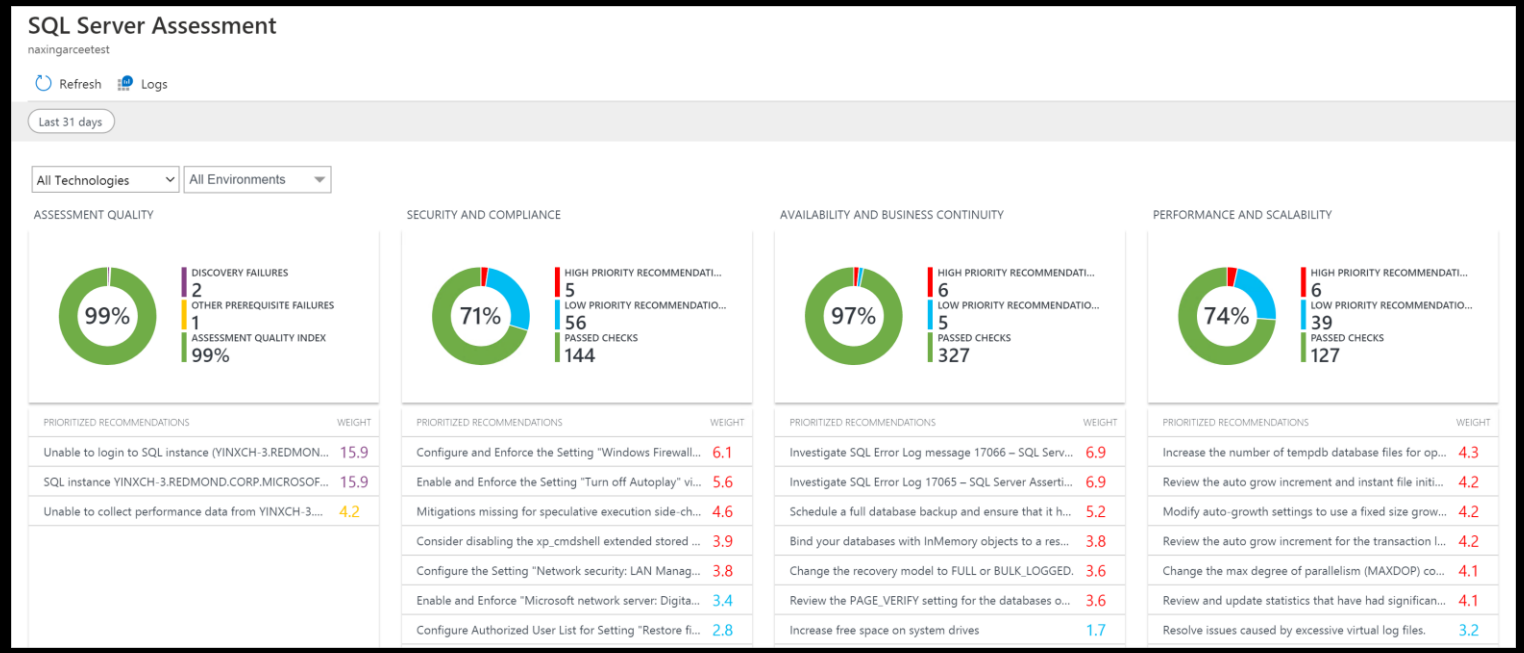

New cloud billing model for SQL Server

Better cost efficiency when paying only for what you use

SQL Server pay-as-you-go licensing enabled by Azure Arc
(per core per month/hour)

Pricing	Monthly rate	Hourly rate
Standard Edition	$73	$0.100
Enterprise Edition	$274	$0.375

Flexible licensing options

Choose from consumption-based licensing and perpetual SQL Server license

Cost efficiency

Pay by the hour for spikes and ad-hoc usage. No need for full upfront investment

Supports hybrid deployment

Consistent purchasing option across on-premise and in 3rd party cloud

Available on SQL Server 2014, 2016, 2017, 2019, and 2022

Azure Data Products and Services

Operational stores

- SQL Server
- Azure SQL DB
- Azure SQL MI
- Azure SQL DB Edge
- Azure Cosmos DB
- Azure Cache
- Azure for PostgreSQL
- Azure for MySQL
- Azure for MariaDB

Analytics

- Microsoft Fabric
- Azure HDInsight
- Azure Synapse
- Azure Data Factory
- Azure Data Explorer
- Azure Stream Analytics
- Azure Databricks

Governance

- Microsoft Purview

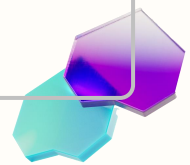

Microsoft Fabric

Data Factory

Data Engineering

Data Warehouse

Data Science

Real-Time Intelligence

Power BI

Partner & Industry workloads

Copilot in Fabric

OneLake

Microsoft Purview

Unifying data in OneLake

Data Factory

Seamlessly connect to more than 170+ data stores

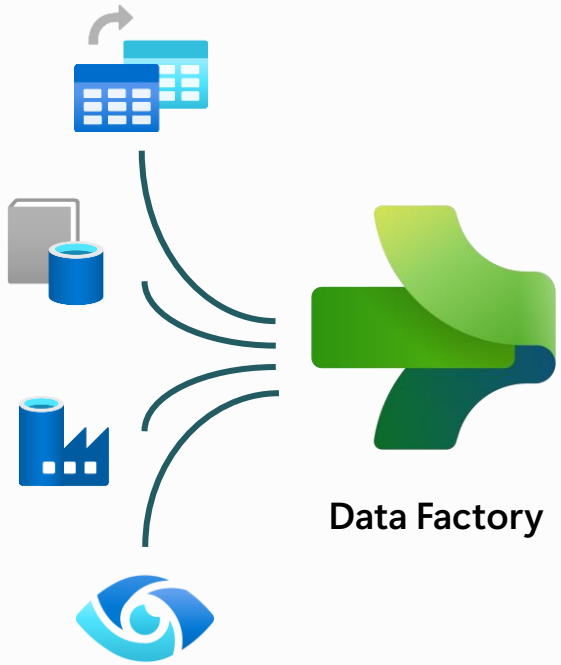

Azure Database for PostgreSQL	Azure Databricks Delta Lake	Amazon RDS for Oracle	Amazon RDS for SQL Server	Amazon Redshift	Phoenix	PostgreSQL	Presto	Magento (Preview)
Azure SQL Database	Azure SQL Database Managed Instance	Apache Impala	Azure SQL Database Managed Instance	DB2	SAP BW Open Hub	SAP BW via MDX	SAP HANA	Oracle Eloqua (Preview)
Azure Table Storage	MongoDB Atlas	Drill	Google AdWords	Google BigQuery	SAP Table	SQL server	Spark	PayPal (Preview)
Azure Cosmos DB (MongoDB API)	Azure Cosmos DB (SQLAPI)	Greenplum	HBase	Hive	Amazon S3	Amazon S3 Compatible	FTP	SAP Cloud For Customer
Azure Data Lake Storage Gen1	Azure Data Lake Storage Gen1 for Cosmos Structured Stream	Informix	MariaDB	Microsoft Access	File system	Google Cloud Storage (S3APD)	HDFS	Salesforce Marketing Cloud
Azure Data Lake Storage Gen1 for Cosmos Structured Stream	Azure Database for MariaDB	MySQL	Netezza	Oracle	HTTP	Oracle Cloud Storage (S3AP)	SFTP	Shopify (Preview)
Teradata	Vertica	ODBC	OData	REST	Amazon Marketplace Web Service	Concur (Preview)	Dataverse (Common Data Service for App)	Web Table
Jira	Kusto	SharePoint Online List	Dynamics 365	Dynamics AX	Dynamics CRM	Cassandra	Couchbase (Preview)	MongoDB

One Copy for all computes

Real separation of compute and storage

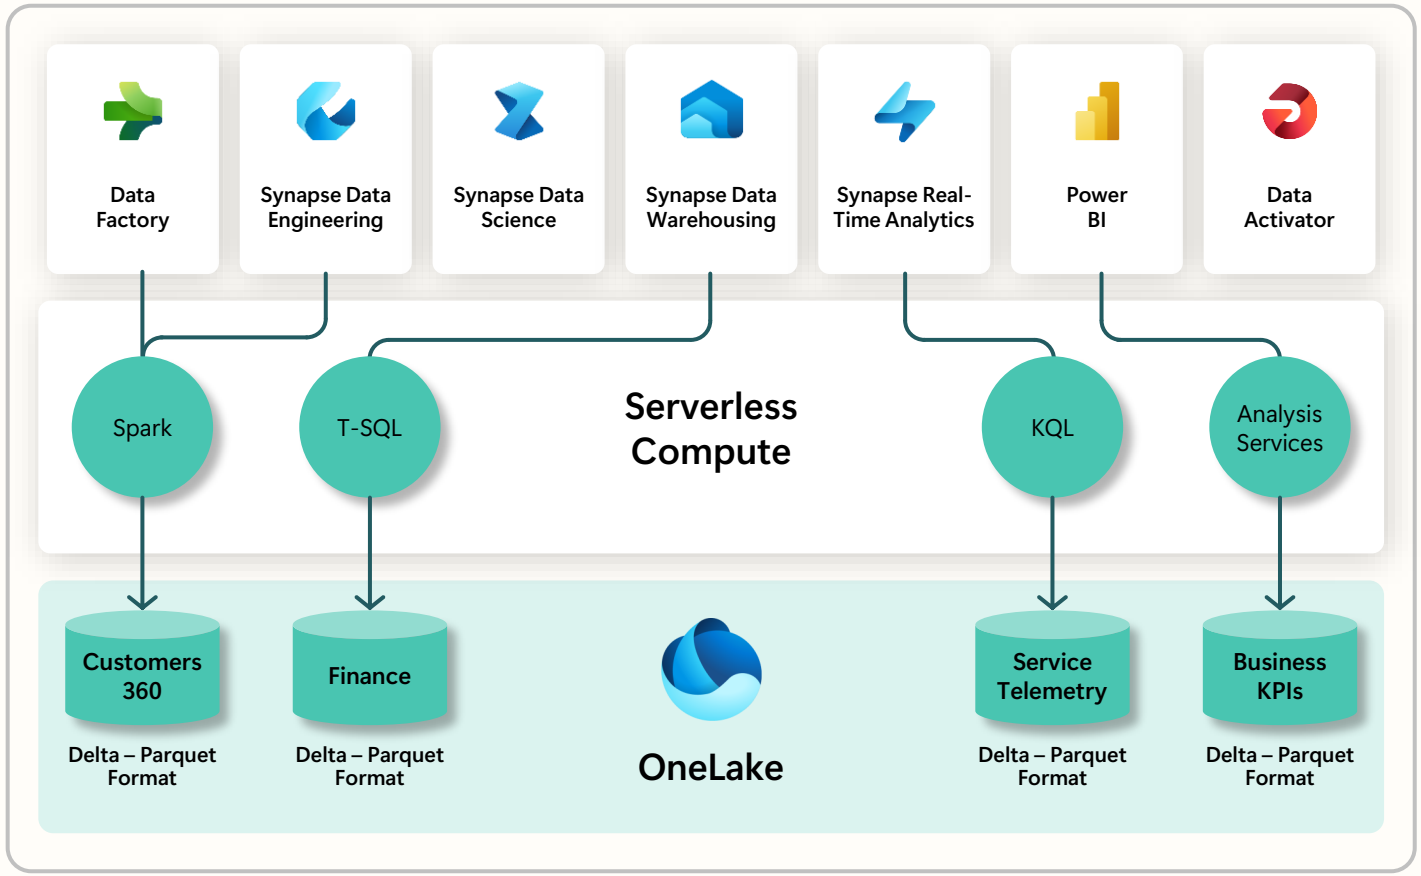

All the compute engines store their data automatically in OneLake

The data is stored in a single common format

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

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

Medallion reference architecture in Microsoft Fabric

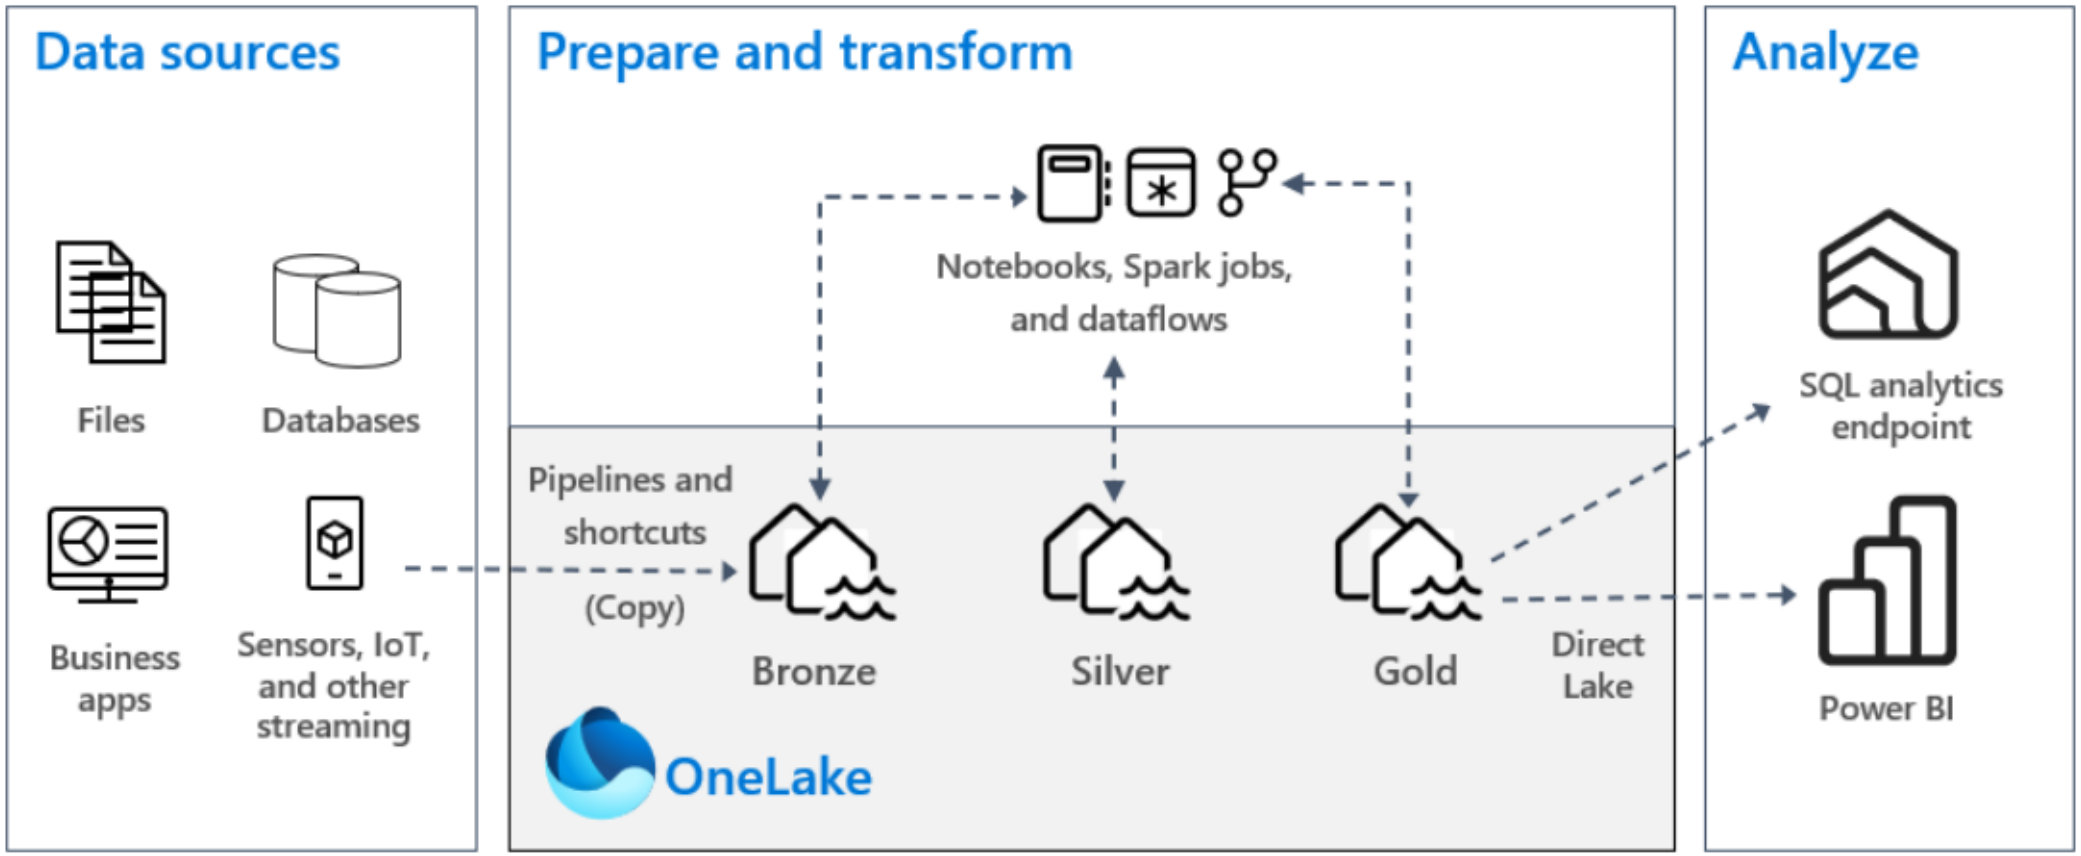

All roads lead to OneLake

Creating Data Gravity in OneLake

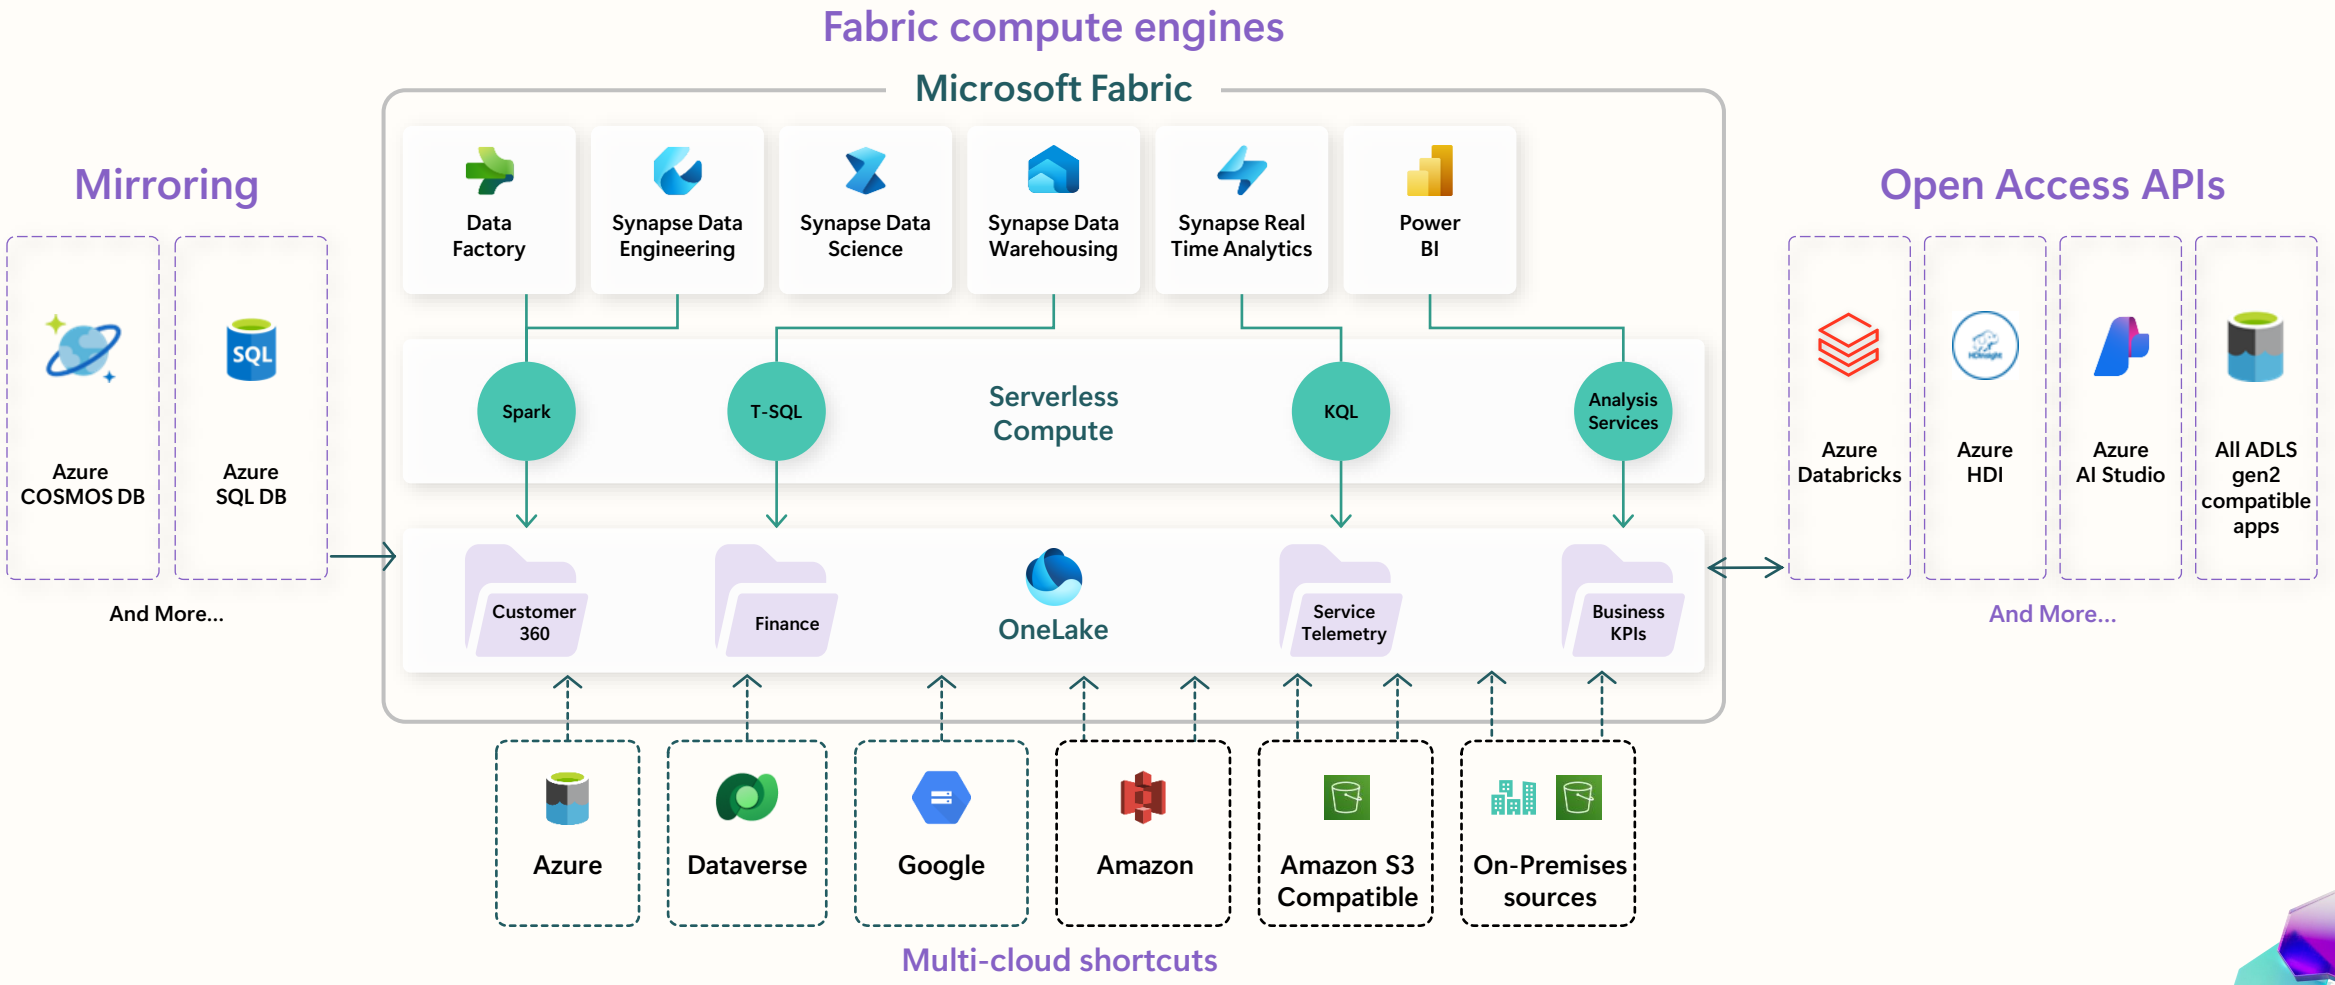

Power BI: The bridge between data and decisions

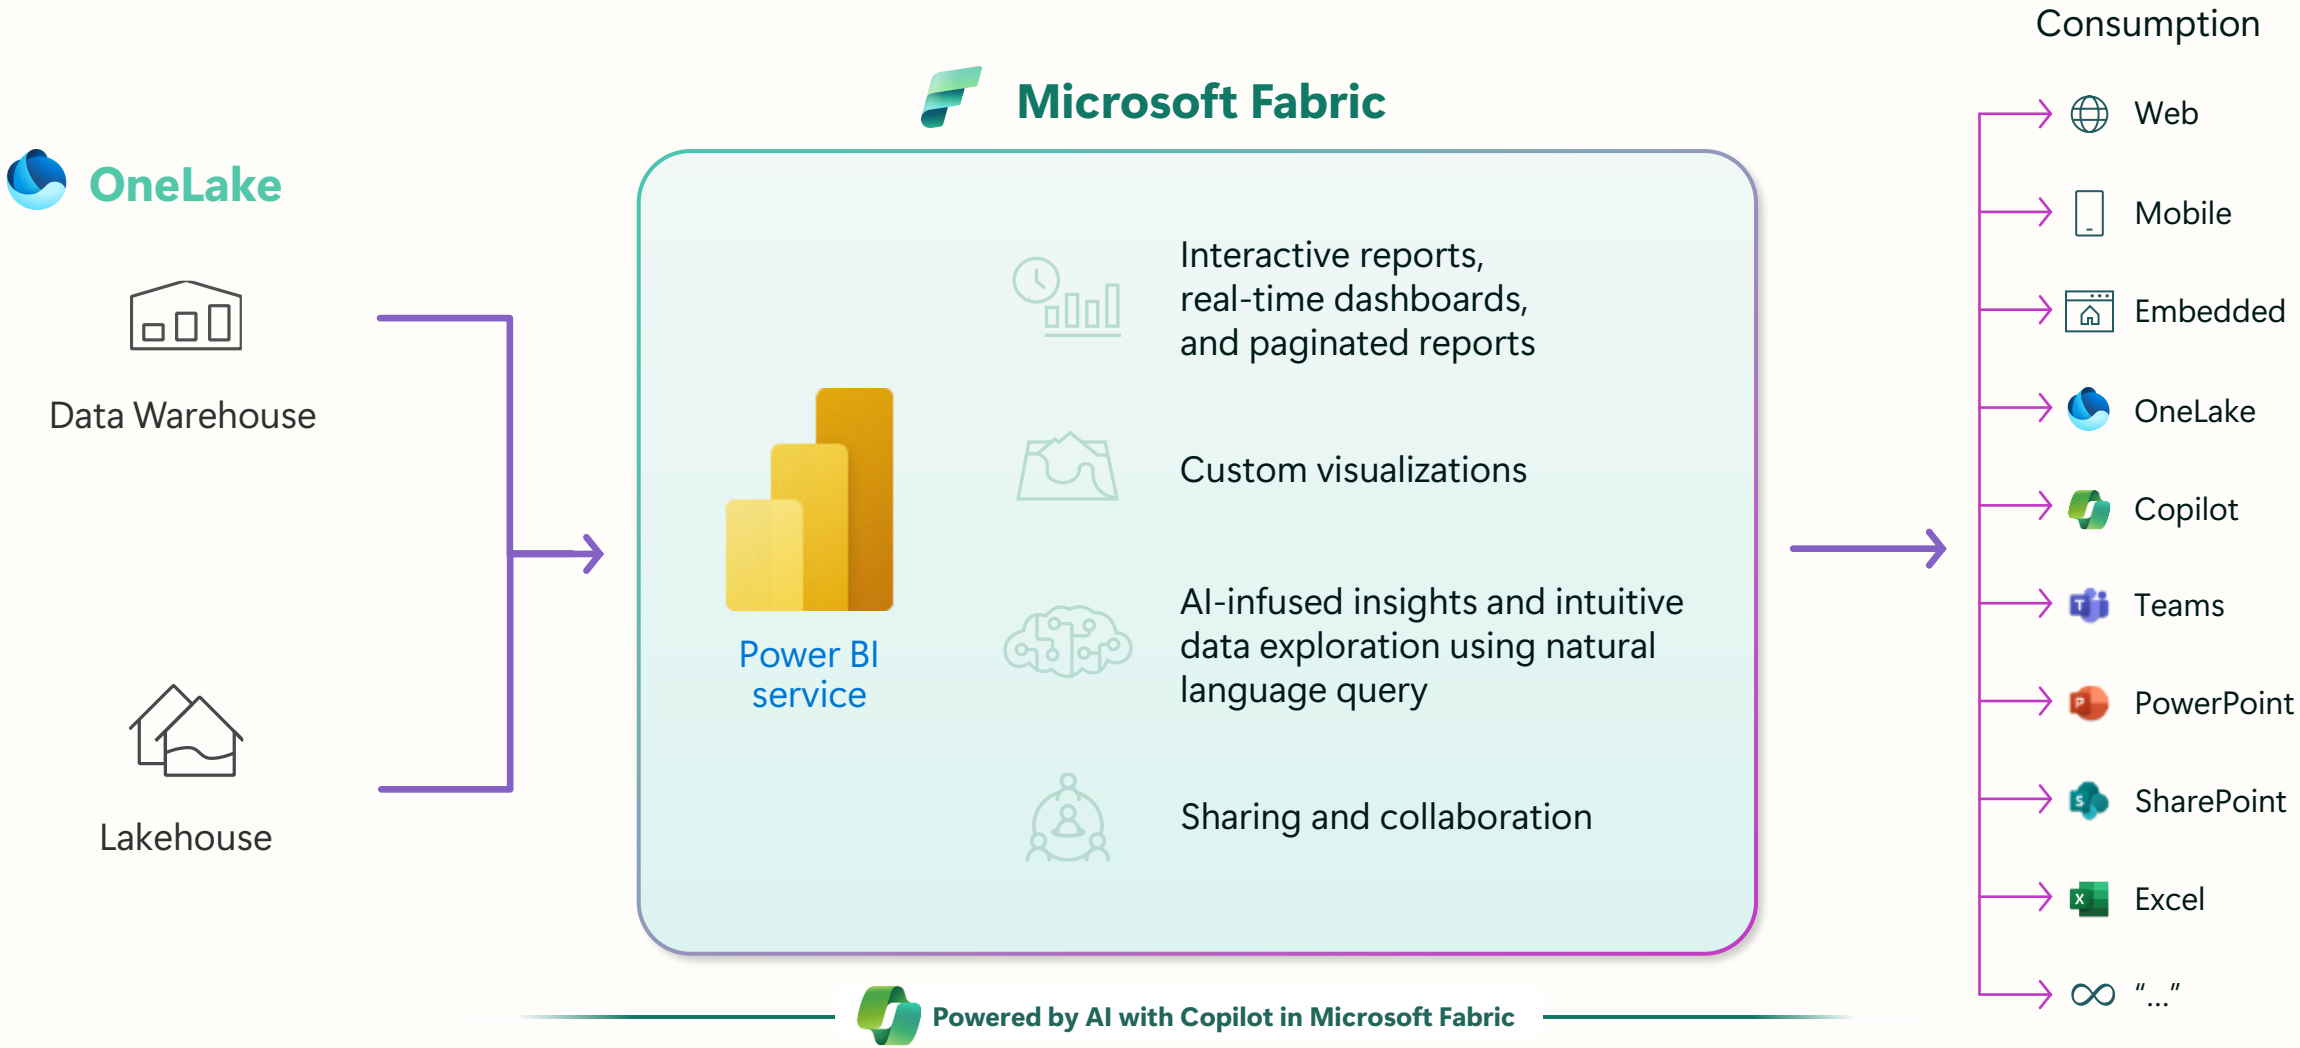

Power BI | Direct Lake Mode

Direct Lake is a fast-path to load the data from the lake straight into the Power BI engine, ready for analysis

Direct Lake is based on loading parquet-formatted files directly from a data lake without having to query a Lakehouse endpoint, and without having to import or duplicate data into a Power BI dataset

Direct Query Mode. Slow, but real time

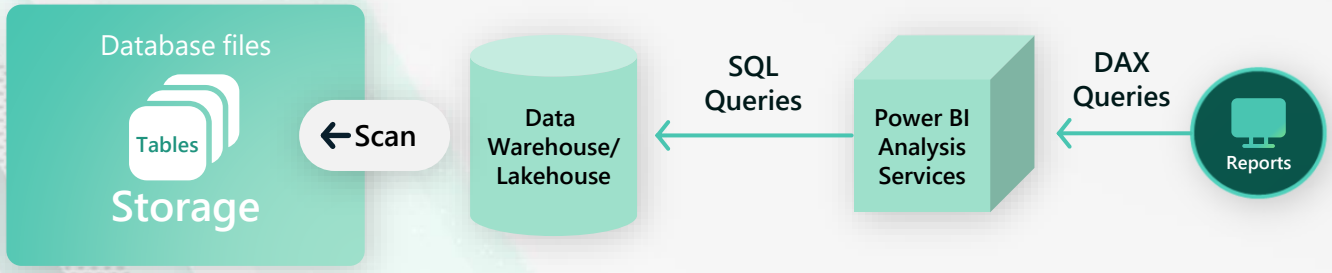

Import Mode. Fast, but not latest updates

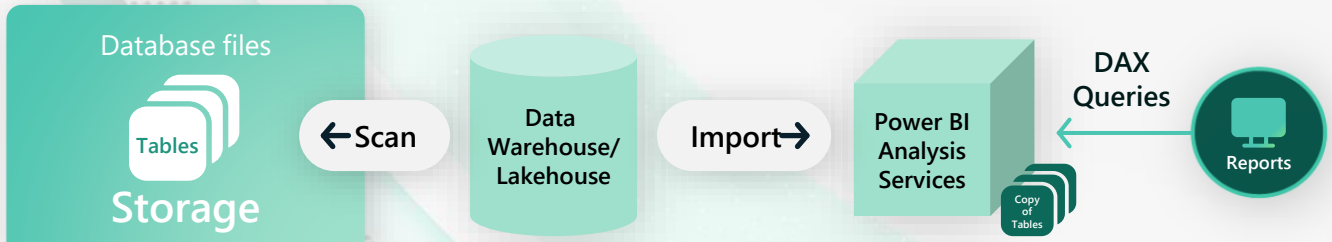

Direct Lake Mode. Fast and real time

Real-Time Intelligence scenario

End-to-end analytics scenario

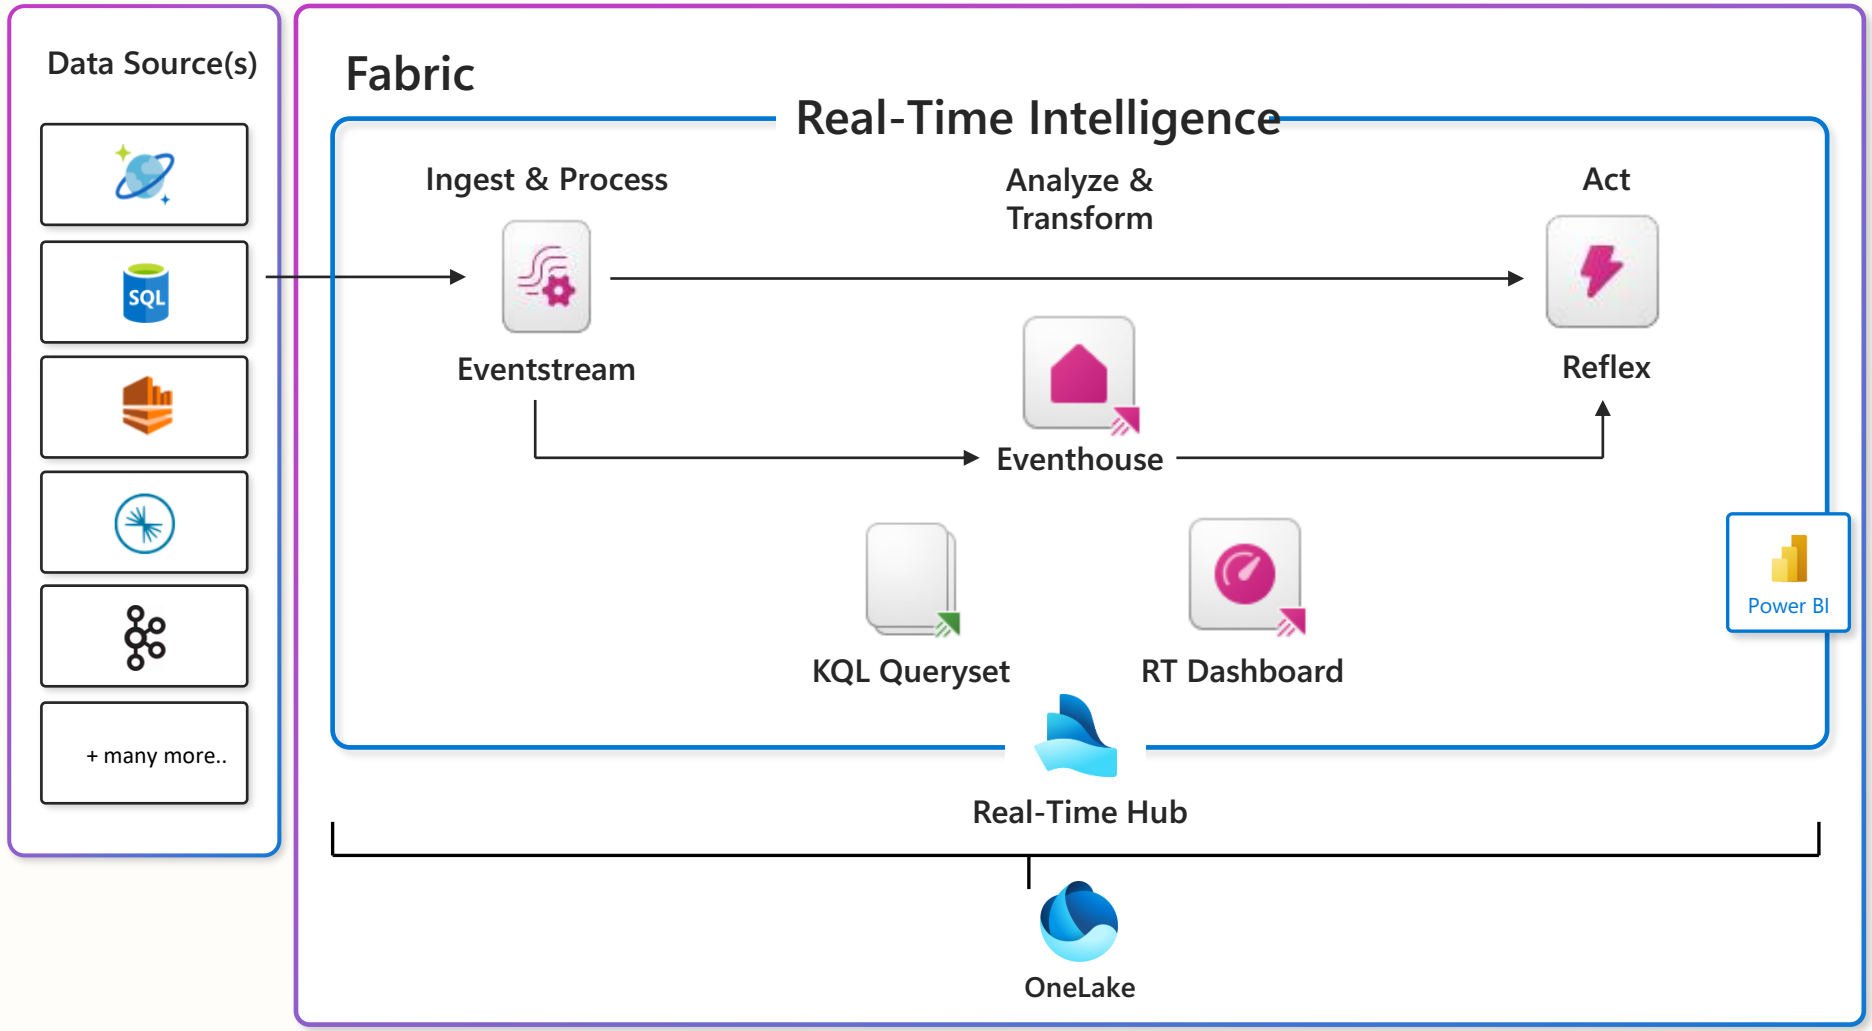

Build your own powerful Gen AI experiences

Microsoft Fabric

Your unified solution for analytics

Enhance data insights

Build enterprise chat applications to uncover insights using natural language from structured, unstructured, and real-time data

Analyze customer interactions

Build speech analytics applications to enhance customer service, tailor support responses, and make data-driven decisions

Customize machine learning models

Train, deploy, and orchestrate machine learning models tailored to specific business needs—from predictive maintenance to customer sentiment analysis

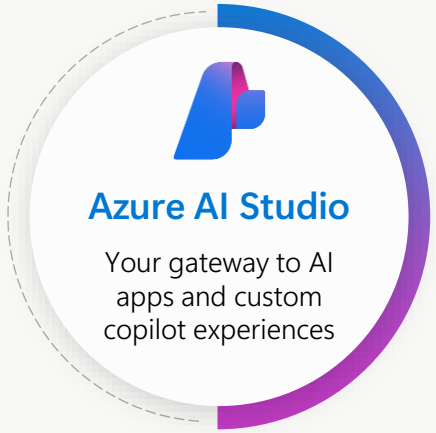

Azure AI Studio

Your gateway to AI apps and custom copilot experiences

OneLake serves as the connective tissue to build generative AI apps powered by your data

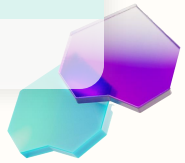

Build your copilot in Azure AI Studio

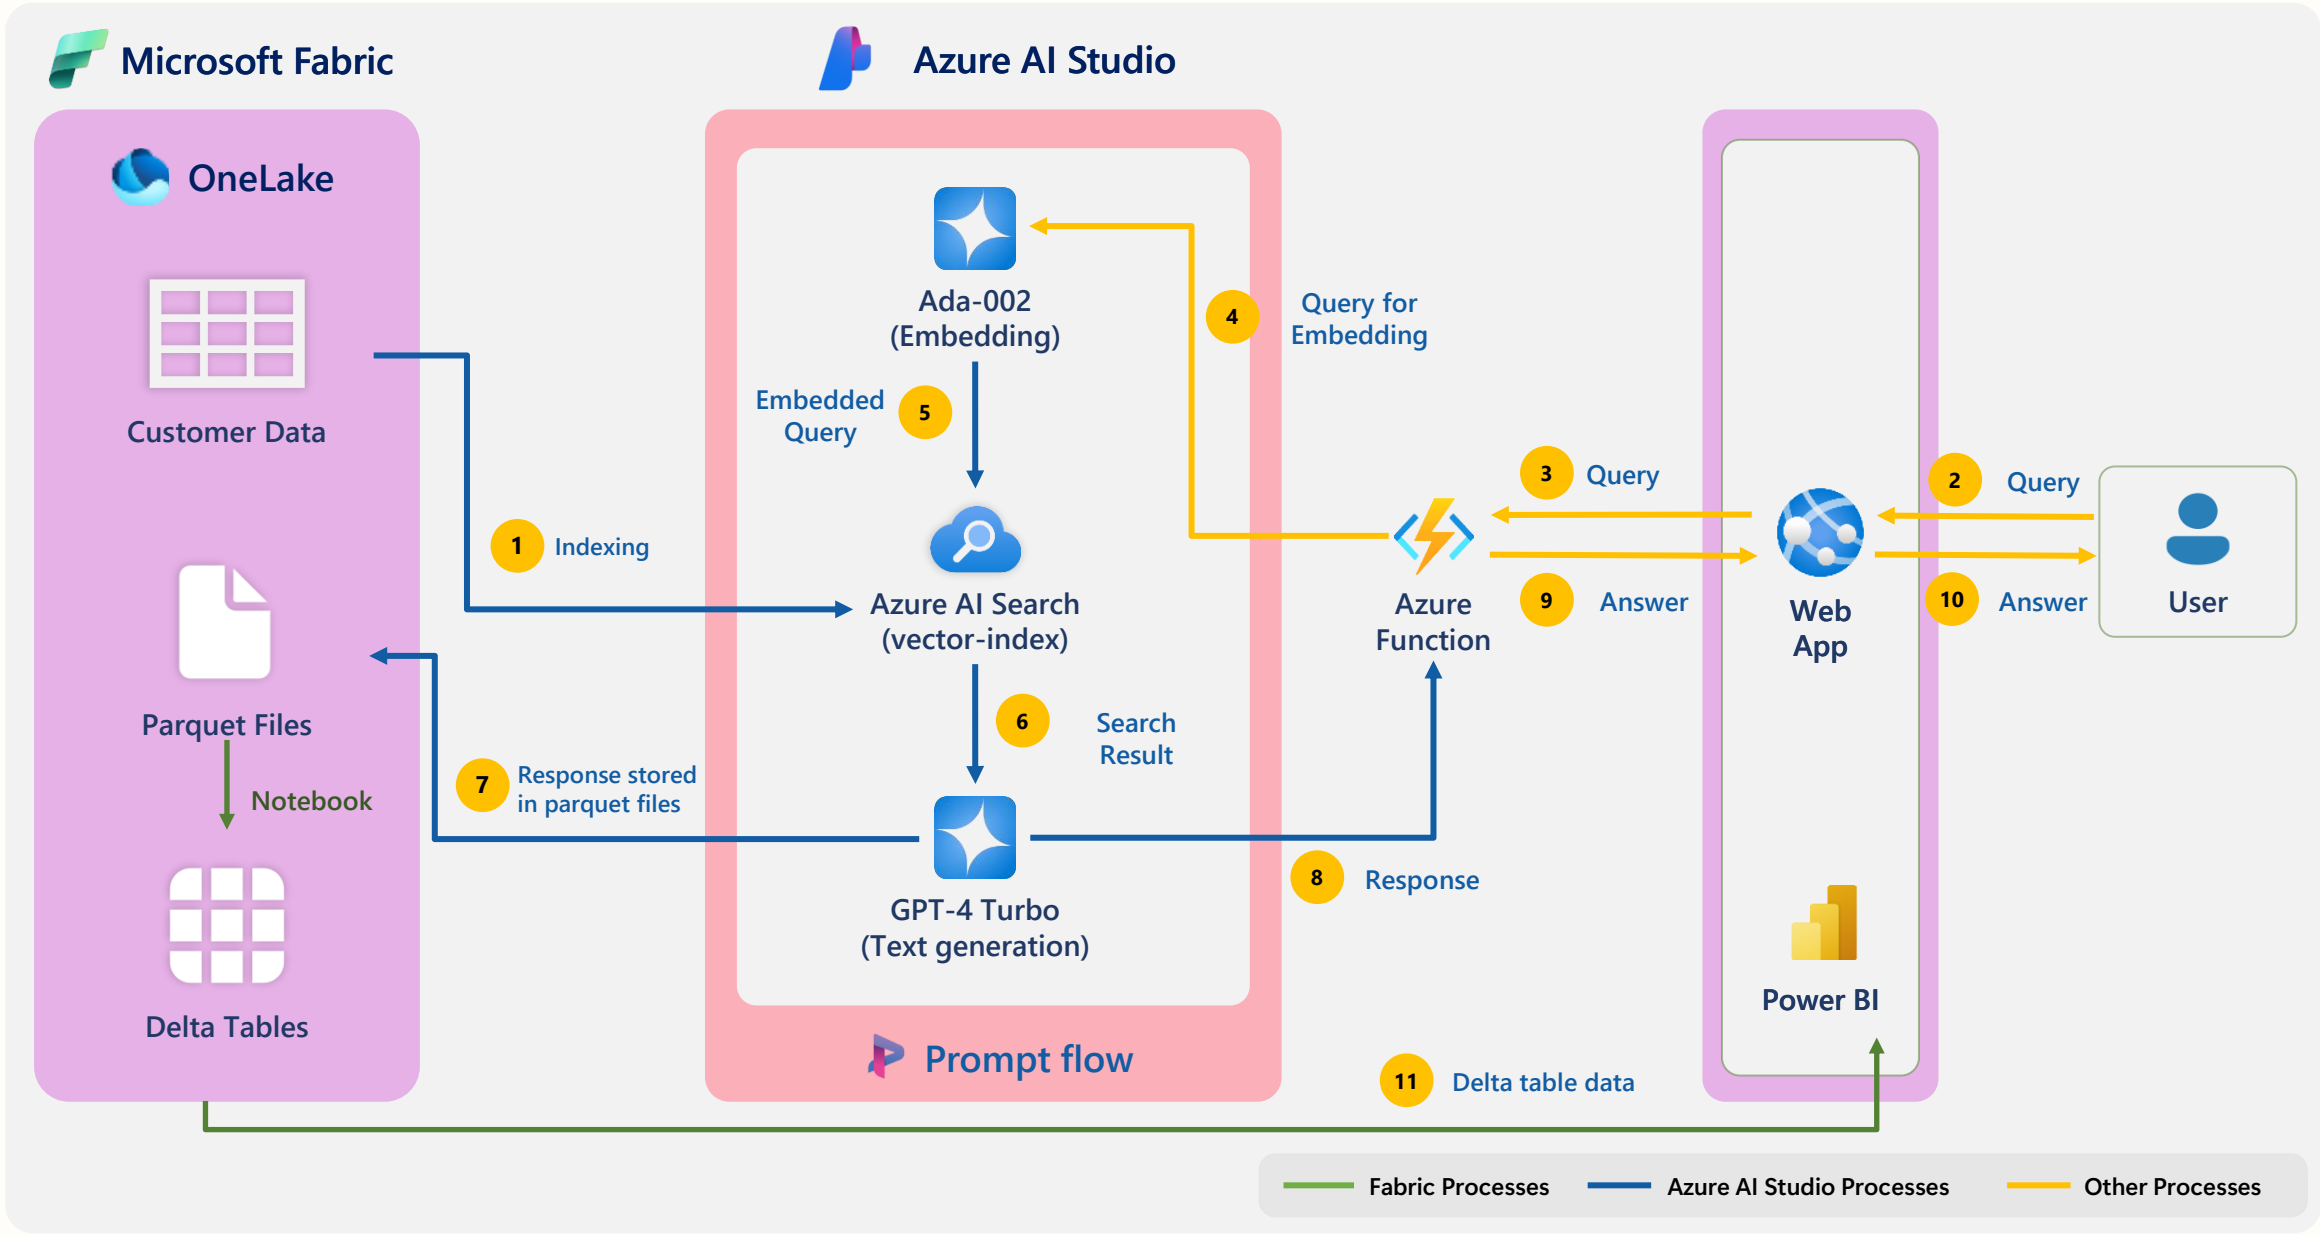

Microsoft Fabric business model overview

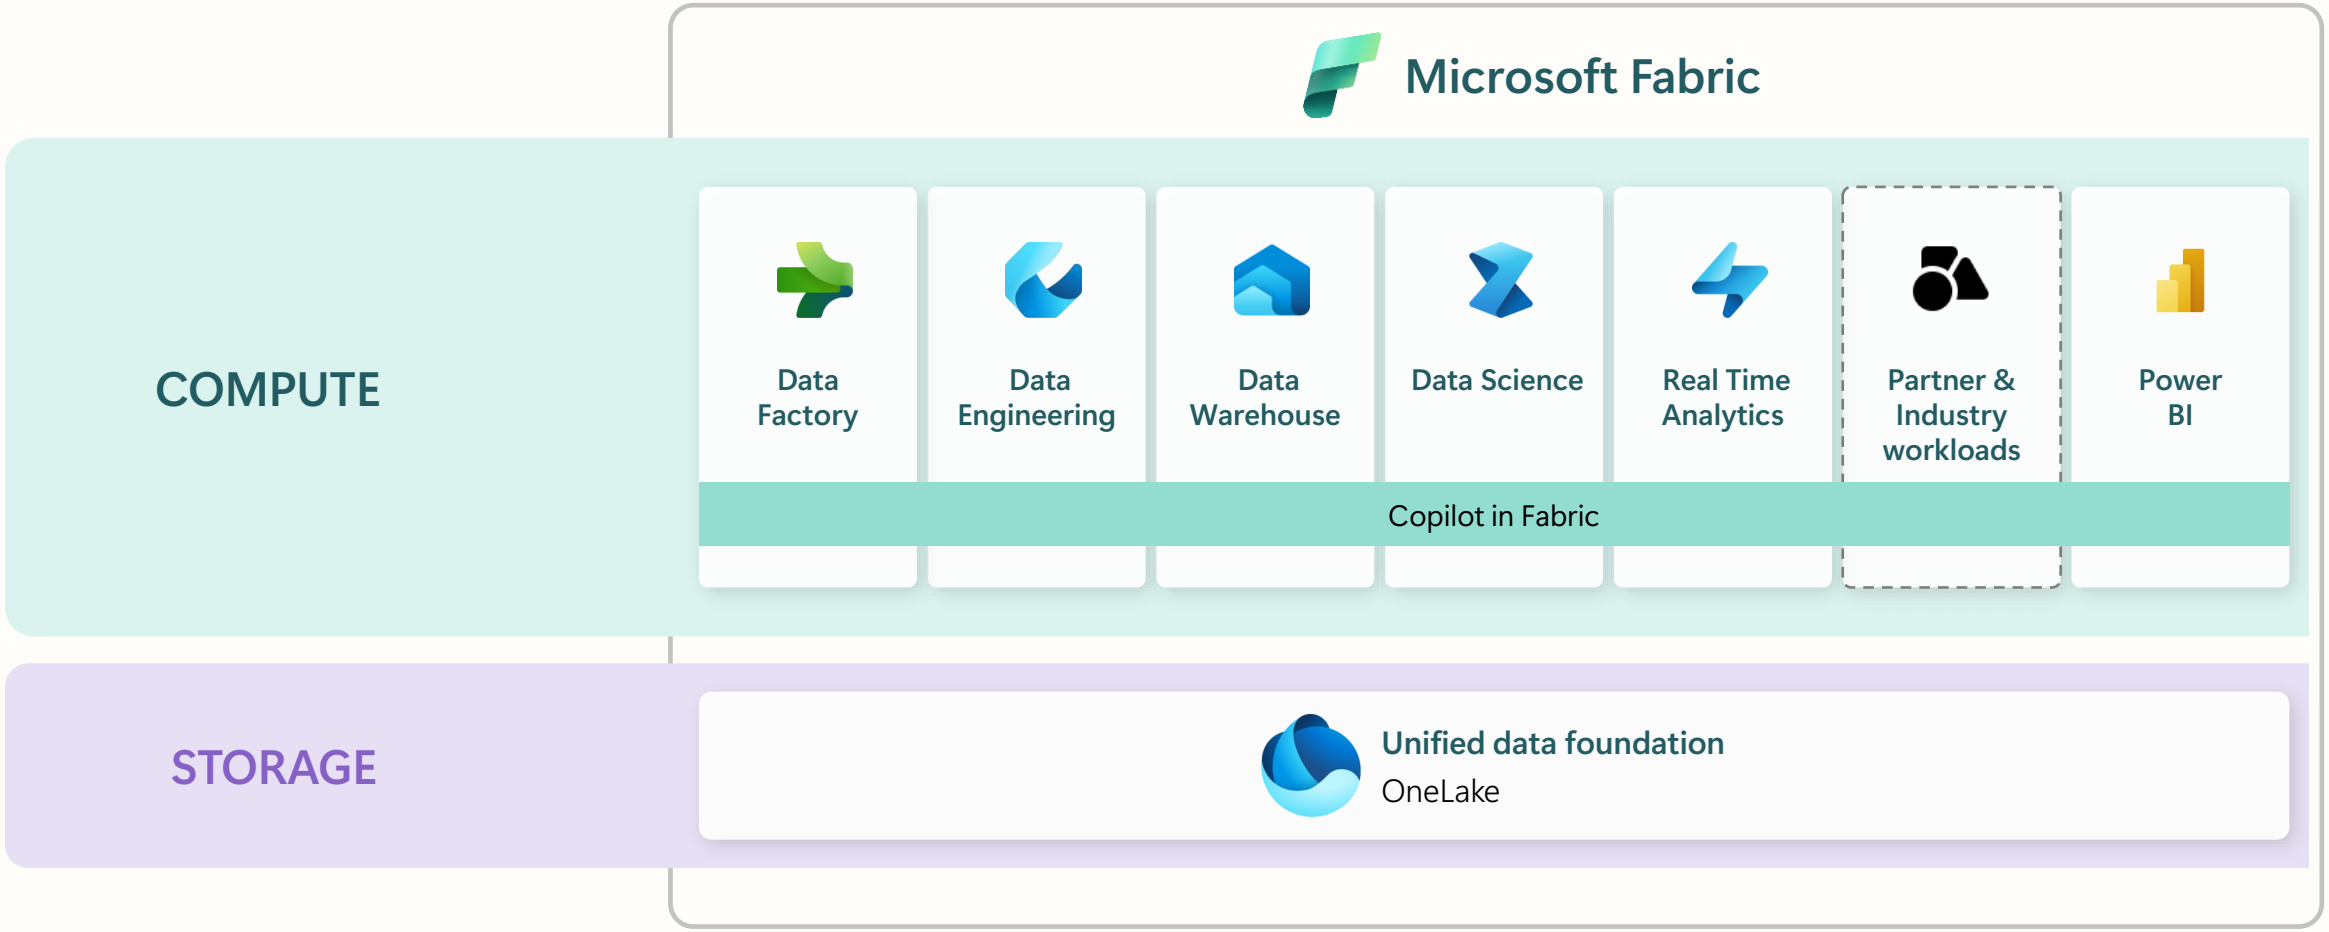

Microsoft Fabric simplicity

Microsoft Fabric is a unified product for all your data and analytics workloads. Rather than provisioning and managing separate compute for each workload, with Microsoft Fabric, your bill is determined by two variables: the amount of compute you provision and the amount of storage you use.

COMPUTE

A shared pool of capacity that powers all capabilities in Microsoft Fabric.

Pay-as-you-go and 1-year Reservation.

STORAGE

A single place to store all data.

Pay-as-you-go ($ per GB/month).

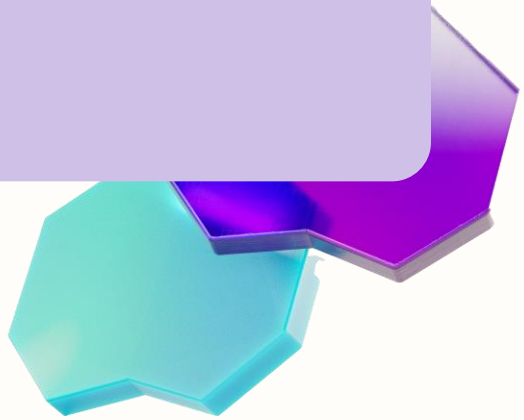

Fabric capacity pricing

Fabric capacity is priced uniquely across regions. The following table shows the pricing at US West 2 for reference. Fabric capacity can be purchased at Azure portal. Visit Fabric pricing page for more details.

1 CU pay-as-you-go price at US West 2 *$0.18/hour*

SKU	Capacity unit (CU)	Pay-as-you-go (monthly)	Reservation (monthly) *~40.5% saving over Pay-as-you-go*
F 2	2	$262.8	$156.334
F 4	4	$525.6	$312.667
F 8	8	$1,051.2	$625.334
F 16	16	$2,102.4	$1,250.667
F 32	32	$4,204.8	$2,501.334
F 64	64	$8,409.6	$5,002.667
F 128	128	$16,819.2	$10,005.334
F 256	256	$33,638.4	$20,010.667
F 512	512	$67,276.8	$40,021.334
F 1024	1024	$134,553.6	$80,042.667
F 2048	2048	$269,107.2	$160,085.334

Note:

1. 1 CU PAYGO monthly rate calculation: $0.18*730 =$131.4. F2 =$131.4*2=$262.8
2. 1 CU RI monthly rate calculation: Round ($0.18* (1-0.405)*730*12,0)/12=~$78.166…F2 RI = ~$78.166…*2=~$156.334
3. Power BI Pro license is required for all Power BI Premium (“P”) and Fabric Capacity (“F”) SKUs to publish Power BI content to Microsoft Fabric. Enabling content consumers to review and interact with Power BI reports without additional paid per-user licenses is available at P1 and above (and F64 and above).

Microsoft Fabric

Auto Elasticity: Bursting

Job acceleration

- Bursting provides extra compute resources to jobs and queries to accelerate their completion

Go beyond

- The extra resources of bursting allow jobs to utilize far more resources than “face value”
- Instead of running a job on 64 CU and completing in 60 seconds, bursting could use 256 CUs to complete the job in 15 seconds
- Burstable capacity is finite, depending on your Fabric Capacity SKU size (SKU guardrails)

No hassle, no overload

- Bursting is automatic when the systems reasons it can accelerate the job by applying extra resources
- Bursting never causes an overload as the smoothing mechanism will always flatten the resource burst

Microsoft Fabric

Auto Elasticity: Smoothing

Load stabilization

- Load smoothing helps capacities self-stabilize by flattening large spikey loads into a smooth load profile, eliminating temporal spikes

Scheduling contention elimination

- Large/schedule jobs are smoothed over 24 hours, eliminating the need to decide the timing and order of job execution
- Operations are always executed for maximum performance

Streamline Capacity Management

- Size your capacity based on average, rather than peak usage

Copilot in Microsoft Fabric

Data Factory

Get intelligent code generation to transform data with ease and code explanations to help you better understand complex tasks


Data Engineering and Data Science

Quickly generate code in Notebooks to help work with Lakehouse data and get insights.

Data Warehouse

Write and explain T-SQL queries, or even make intelligent suggestions and fixes while you are coding

Real-Time Intelligence

Translate questions into KQL queries that you can execute.

Power BI

Quickly create report pages, natural language summaries, and generate synonyms.

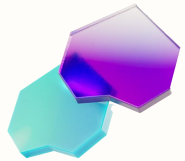

Thank you!