

Solution Brief

Microsoft Fabric

The Foundation for Data Readiness and Copilot Success



AI Is Only as Good as the Data Behind It

Every Copilot interaction, every agent workflow, every AI-generated insight draws on the data an organization has made available to it. When that data is fragmented across silos, inconsistently governed, or stored in formats the AI platform cannot access, the results reflect those limitations. Outputs are incomplete, confidence is low, and adoption stalls.


This is the data readiness problem, and it is the most common reason enterprise AI programs fail to meet expectations. The technology is AI-ready, but the data infrastructure is not.

Microsoft Fabric solves for that problem by bringing data engineering, data integration, data science, real-time analytics, and business intelligence together in a single environment. For organizations investing in Microsoft Copilot and the broader Microsoft AI stack, Fabric is the foundation that makes those investments perform as intended.

Microsoft Fabric

A unified, AI-powered data platform to simplify data management and analytics

Microsoft Fabric is a unified, cloud-based data platform that delivers services and tools for every data lifecycle stage: ingestion, preparation, storage, analysis, and visualization. It combines components from Microsoft Azure and other services to unify capabilities for data integration, engineering, warehousing, and more.



The top layer includes the core Fabric workloads, each of which provides specialized capabilities and analytics tasks within the same Fabric environment:



Power BI

Connect to data sources, build interactive charts and dashboards, and share insights across the organization for faster, more data-driven decisions.



Databases

A developer-friendly transactional database built on Azure SQL Database that lets teams create and manage operational databases in Fabric and mirror data from external systems into OneLake.



Data factory

A modern data integration service for ingesting, preparing, and transforming data from over 200 native connectors across on-premises and cloud sources.



Industry solutions

Industry-specific data solutions that address unique sector needs across data management, analytics, and decision-making.



Real-time intelligence

Analyzes streaming data as it arrives, from IoT sensors to application logs, handling everything from ingestion and transformation to visualization and AI-driven action.



Data engineering

Apache Spark-powered tooling for processing large datasets, writing and scheduling transformation jobs, and building the infrastructure needed to collect, store, and analyze data at scale.



Fabric data science

Build, deploy, and operationalize machine learning models directly in Fabric, with Azure Machine Learning integration for experiment tracking, model registry, and embedding predictions into BI reports.



Fabric data warehouse

Industry-leading SQL performance at scale, with compute and storage separated for independent scaling and data stored natively in open Delta Lake format.



Fabric IQ

A new workload that unifies business semantics across data, models, and systems to power consistent decisions, reusable metrics, and context-aware automation across the Fabric platform.

The Fabric platform powers these workloads through several shared core services:



OneLake is the centralized, logical data lake for Fabric. All workloads store and access data through OneLake, enabling zero-copy access patterns and allowing data to remain in its original location while being reused across experiences.



Copilot provides AI assistance embedded directly within Fabric workloads to help with authoring, exploration, and routine development tasks, while respecting tenant, data, and permission boundaries.




Governance represents centralized administration and data governance, including permissions, sensitivity labels, and auditing. Purview-backed governance extends to data shared across tenants through OneLake data sharing, so sensitivity labels, access policies, and compliance controls remain enforced regardless of where shared data is consumed.

How Fabric Powers Copilot

Copilot and agentic AI depend on well-structured business data that is contextually relevant, accurate, and trustworthy. When that data is governed, consistent, and accessible through a unified layer, AI performance improves measurably.

Fabric IQ, Foundry IQ, and Work IQ make this connection concrete.

Fabric IQ provides the structural foundation that makes reasoning possible at the data level. Fabric IQ builds an ontology, or shared business language, that defines key business entities, how they relate, the rules governing them, and the actions agents can take on them.



Foundry IQ is the knowledge grounding layer for agents built in Azure AI Foundry. It connects documents, wikis, SharePoint, OneLake, and federated enterprise content into reusable knowledge bases that any custom agent can retrieve.

Work IQ is Microsoft's intelligence framework that powers Copilot and its agents. It operates across three layers: data from emails, chats, and files, organizational memory built from work habits and workflows, and inference that predicts the right action or agent for a given task.

Through these "three IQs," Fabric unifies organizational data and Copilot the foundation needed to perform agentic work consistently, reliably, and at scale. For frontier organizations ready to take the next step with agentic AI, Fabric provides the data infrastructure to help make that aspiration a reality.

Modernize Your Data Estate

with Logicalis and Microsoft Fabric

Organizations that invest in data readiness before scaling Copilot and AI agent deployment realize returns faster and with fewer remediation cycles.

Implementing Microsoft Fabric can help organizations build that foundation, but modernization requires a deliberate architecture, governance design, and adoption program. Logicalis, a Microsoft Global Partner, brings structured methodology to each phase of the Fabric journey, covering:

- **Data estate assessments** that map source systems, existing pipelines, governance gaps, and the data assets Copilot and AI workloads will depend on most
- **Fabric architecture and OneLake design**, including workspace organization, security model, and the semantic layer
- **Migration and pipeline modernization** that moves existing workloads into Fabric
- **Purview integration** that includes sensitivity classification, data lineage, and compliance controls
- **Copilot readiness validation** covering data quality, permission model review, and semantic layer verification
- **Ongoing managed services** that maintain pipeline health, governance policy currency, and platform optimization

[Contact Logicalis to start with a data estate assessment and Fabric readiness review.](#)

