

Introduction to end-to-end analytics using Microsoft Fabric

1. Introduction

<https://learn.microsoft.com/en-us/training/modules/introduction-end-analytics-use-microsoft-fabric/1-introduction>

Introduction

Completed

Organizations need to ingest, prepare, govern, and analyze data at scale, often across disconnected tools and teams. Increasingly, that same data also needs to be ready for AI workloads like machine learning models, Copilots, and intelligent agents. Managing these tasks across separate systems creates complexity, governance gaps, and duplicated effort.

Microsoft Fabric is an end-to-end analytics platform that provides a single, integrated environment where data professionals and the business collaborate on data projects. Built on a unified data lake called OneLake, Fabric brings together the tools you need across that entire lifecycle.

Because all data is ingested, prepared, and governed within Fabric, the same data that powers your reports and dashboards is also available to AI capabilities like Copilot, data agents, and Fabric IQ. This means the work you do to organize and govern your data directly supports your organization's AI initiatives.

This module introduces the Fabric platform, discusses who Fabric is for, explores Fabric workloads, and examines how Fabric supports both analytics and AI.

2. Explore end-to-end analytics with Microsoft Fabric

Explore end-to-end analytics with Microsoft Fabric

Completed


Scalable analytics can be complex, fragmented, and expensive. Microsoft Fabric simplifies analytics solutions by providing a single, easy-to-use product that integrates various tools and services into one platform.

Fabric is a unified *software-as-a-service* (SaaS) platform where all data is stored in a single open format in OneLake. All analytics engines in the platform can access OneLake, ensuring scalability, cost-effectiveness, and accessibility from anywhere with an internet connection.

OneLake

OneLake is Fabric's centralized data storage architecture that enables collaboration by eliminating the need to move or copy data between systems. OneLake unifies your data across regions and clouds into a single logical lake without moving or duplicating data.

OneLake is built on **Azure Data Lake Storage Gen2** (ADLS Gen2) and supports various formats, including Delta, Parquet, CSV, and JSON. All compute engines in Fabric automatically store their data in OneLake, making it directly accessible without the need for movement or duplication. For tabular data, the analytical engines in Fabric write data in delta-parquet format and all engines interact with the format seamlessly.



Shortcuts are references to files or storage locations within OneLake or external data sources, such as Azure Data Lake Storage, Amazon S3, or Database. Shortcuts allow you to access existing data without copying it, ensuring data consistency and enabling Fabric to stay in sync with the source.

Because all Fabric workloads store data in OneLake using an open format, AI capabilities like Copilot and data agents can access the same governed data as your reports and dashboards without separate data preparation pipelines. The work you do to ingest, prepare, and govern data in Fabric is what makes that data available for AI workloads.

Workspaces

In Microsoft Fabric, workspaces serve as logical containers that help you organize and manage your data, reports, and other assets. They provide a clear separation of resources, making it easier to control access and maintain security.

Each workspace has its own set of permissions, ensuring that only authorized users can view or modify its contents. This structure supports team collaboration while maintaining strict access control for both business and IT users.

Workspaces allow you to manage compute resources and integrate with Git for version control. You can optimize performance and cost by configuring compute settings, while Git integration helps track changes, collaborate on code, and maintain a history of your work.

Administration and governance

Fabric's OneLake is centrally governed and open for collaboration. Data is secured and governed in one place, which allows users to easily find and access the data they need. Fabric administration is centralized in the **Admin portal**.

In the admin portal you can manage groups and permissions, configure data sources and gateways, and monitor usage and performance. You can also access the Fabric admin APIs and SDKs in the admin portal, which can automate common tasks and integrate Fabric with other systems.

The **OneLake catalog** helps you analyze, monitor, and maintain data governance. It provides guidance on sensitivity labels, item metadata, and data refresh status, offering insights into the governance status and actions for improvement.

Note

Review the [Microsoft Fabric administration](#) documentation for more information.

3. Explore data teams and Microsoft Fabric

<https://learn.microsoft.com/en-us/training/modules/introduction-end-analytics-use-microsoft-fabric/3-data-team>

Explore data teams and Microsoft Fabric

Completed

Microsoft Fabric's unified data analytics platform makes it easier for data professionals to collaborate on projects. Fabric increases collaboration between data professionals by removing data silos and the need for multiple systems.

Traditional roles and challenges

In a traditional analytics development process, data teams often face several challenges due to the division of data tasks and workflows.

Data engineers process and curate data for analysts, who then use it to create business reports. This process requires extensive coordination, often leading to delays and misinterpretations.

Data analysts often need to perform downstream data transformations before creating Power BI reports. This process is time-consuming and can lack the necessary context, making it harder for analysts to connect directly with the data.

Data scientists face difficulties integrating native data science techniques with existing systems, which are often complex, and makes it challenging to efficiently provide data-driven insights.

Evolution of collaborative workflows

Microsoft Fabric simplifies the analytics development process by unifying tools into a SaaS platform. Fabric allows different roles to collaborate effectively without duplicating efforts.

- **Data engineers** can ingest, transform, and load data directly into OneLake using Pipelines, which automate workflows and support scheduling. They can store data in lakehouses, using the Delta-Parquet format for efficient storage and versioning. Notebooks provide advanced scripting capabilities for complex transformations.
- **Analytics engineers** bridge the gap between data engineering and analysis by curating data assets in lakehouses, ensuring data quality, and enabling self-service analytics. They can create semantic models in Power BI to organize and present data effectively.
- **Data analysts** can transform data upstream using dataflows and connect directly to OneLake with Direct Lake mode, reducing the need for downstream transformations. They can create interactive reports more efficiently using Power BI.
- **Data scientists** can use integrated notebooks with support for Python and Spark to build and test machine learning models. They can store and access data in lakehouses and integrate with Azure Machine Learning to operationalize and deploy models. The predictions they generate can also serve as grounding data for Copilot and AI agents.
- **Low-to-no-code users** and **citizen developers** can discover curated datasets through the OneLake catalog and use Power BI templates to quickly create reports and dashboards. They can also use dataflows to perform simple ETL tasks without relying on data engineers, or ask questions of their data in natural language using Copilot.

Every role in the data team contributes to the organization's ability to use AI effectively. Data engineers who maintain clean, well-governed data in OneLake build the foundation that Copilot and AI agents rely on. Analytics engineers who create consistent semantic models give AI tools the business context needed to generate accurate, meaningful answers.

4. Enable and use Microsoft Fabric

<https://learn.microsoft.com/en-us/training/modules/introduction-end-analytics-use-microsoft-fabric/4-use-fabric>

Enable and use Microsoft Fabric

Completed

Before you can explore the end-to-end capabilities of Microsoft Fabric, it must be enabled for your organization. You might need to work with your IT department to enable Fabric for your organization, including one of the following roles:

- *Fabric administrator*: Manages Fabric settings and configurations.
- *Power Platform administrator*: Oversees Power Platform services, including Fabric.
- *Global administrator*: Has implicit Fabric admin rights through organization-wide permissions.

Enable Microsoft Fabric

Admins can enable Fabric in the **Admin portal > Tenant settings** in the Power BI service. Fabric can be enabled for the entire organization or for specific Microsoft 365 or Microsoft Entra security groups. Admins can also delegate this ability to other users at the capacity level.

Note

If your organization isn't using Fabric or Power BI today, you can sign up for a [free Fabric trial](#) to explore its features.

Create workspaces

Workspaces are collaborative environments where you can create and manage items like lakehouses, warehouses, and reports. All data is stored in OneLake and accessed through workspaces.

Workspaces also support data lineage view, providing a visual view of data flow and dependencies to enhance transparency and decision-making.

In *Workspace settings*, you can configure:

- License type to use Fabric features.
- OneDrive access for the workspace.
- Azure Data Lake Gen2 Storage connection.
- Git integration for version control.
- Spark workload settings for performance optimization.

You can manage workspace access through four roles: *admin*, *contributor*, *member*, and *viewer*. These roles apply to all items in a workspace and should be reserved for collaboration. For more granular access control, use item-level permissions based on business needs.


Note

Learn more about workspaces in the [Fabric documentation](#).

Discover data with OneLake catalog

The OneLake catalog in Microsoft Fabric helps you find and access data sources within your organization. You can explore and connect to data sources, ensuring you have the right data for your needs. You only see items that have been shared with you. Here are some considerations when using OneLake catalog:

- Narrow results by workspaces or domains (if implemented).
- Explore default categories to quickly locate relevant data.
- Filter by keyword or item type.



Create items with Fabric workloads

After you create your Fabric enabled workspace, you can start creating items in Fabric. Each workload in Fabric offers different item types for storing, processing, and analyzing data. Fabric workloads include:

- **Data Engineering:** Create lakehouses and operationalize workflows to build, transform, and share your data estate.
- **Data Factory:** Ingest, transform, and orchestrate data.
- **Data Warehouse:** Combine multiple sources in a traditional warehouse for analytics.
- **Real-Time Intelligence:** Process, monitor, and analyze streaming data.
- **Industry Solutions:** Use out-of-the-box industry data solutions.
- **Data Science:** Detect trends, identify outliers, and predict values using machine learning.
- **Databases:** Create and manage databases with tools to insert, query, and extract data.
- **IQ (preview):** Unify data across OneLake and organize it according to the language of your business using ontologies, graphs, and semantic models.

- **Power BI:** Create reports and dashboards to make data-driven decisions.

Fabric integrates capabilities from existing Microsoft tools like Power BI, Azure Synapse Analytics, and Azure Data Factory into a unified platform. Fabric also supports a data mesh architecture, allowing decentralized data ownership while maintaining centralized governance. This design eliminates the need for direct Azure resource access, simplifying data workflows.

AI capabilities in Microsoft Fabric

Fabric includes features that support AI development as well as AI-powered productivity across workloads.

Fabric IQ (preview) is a Fabric workload for unifying data across OneLake and organizing it according to the language of your business. Its core item is the **ontology**, which defines your business concepts, relationships, and rules so that AI agents can reason across domains using consistent business language rather than raw table schemas.

Fabric IQ is one of three IQ workloads that Microsoft provides to give agents access to different aspects of your organization:

- **Fabric IQ** models business data (ontologies, semantic models, and graphs) so agents can reason over analytics in OneLake and Power BI.
- **Foundry IQ** connects structured and unstructured data across Azure, SharePoint, OneLake, and the web so agents can access permission-aware enterprise knowledge.
- **Work IQ** captures collaboration signals from documents, meetings, chats, and workflows, providing agents with insight into how your organization operates.

Each IQ workload is standalone, but you can use them together to provide comprehensive organizational context for agents.

Fabric data agents let you build conversational interfaces where users ask questions about organizational data in natural language. Agents translate those questions into structured queries across your lakehouses, warehouses, and semantic models.

In the Fabric IQ workload, data agents can connect to your ontology as a source, enabling them to understand and use your business concepts when answering questions.

Copilot across workloads

Microsoft Copilot in Fabric is a generative AI assistant available across all Fabric workloads. Copilot helps data professionals and business users complete common tasks more efficiently. Key capabilities include:

- **Code completion and generation:** Copilot provides intelligent code suggestions in notebooks, generates SQL queries from natural language descriptions, and translates questions into Kusto Query Language (KQL) for real-time analysis.

- **Data transformation guidance:** In Data Factory, Copilot supports both citizen and professional data wranglers with code generation for data transformation and plain-language explanations of complex logic.
- **Report and insight generation:** In Power BI, Copilot generates reports automatically, creates page summaries, and lets business users ask questions about their data in natural language.

Note

Copilot in Microsoft Fabric is enabled by default. Administrators can disable Copilot from the **Admin portal > Tenant settings** or control access for specific security groups or at the capacity level.

5. Module assessment

<https://learn.microsoft.com/en-us/training/modules/introduction-end-analytics-use-microsoft-fabric/5-knowledge-check>

Module assessment

Completed

6. Summary

<https://learn.microsoft.com/en-us/training/modules/introduction-end-analytics-use-microsoft-fabric/6-summary>

Summary

Completed

Organizations need to ingest, prepare, govern, and analyze data at scale. They also need that data to be ready for AI workloads like copilots, agents, and machine learning models.

Microsoft Fabric provides a unified foundation for both. In this module, you explored Fabric's OneLake storage architecture, the workloads that serve different analytics needs, and how data teams collaborate within the platform.

You also learned how AI capabilities build on top of well-governed data. Copilot provides intelligent assistance across every workload, data agents let users query organizational data in natural language, and Fabric IQ is a workload that helps agents reason across domains using consistent business language. Together with Foundry IQ and Work IQ, these capabilities position Fabric as both a data platform and an intelligence platform.

Learn more

- [What's new in Microsoft Fabric?](#)
- [Migrate to Microsoft Fabric](#)