



Microsoft Fabric

Unified Analytics for Data-Driven Innovation

A proof of concept framework for data leaders considering Fabric



Contents

Chapter 1 What is Microsoft Fabric?	03
Chapter 2 Your data analytics needs	05
Chapter 3 Fabric success stories	10
Chapter 4 Building a proof of concept for Fabric	12
Take the next step	16

What is Microsoft Fabric?

Your company's data contains untold value—and the tools you use to manage and analyze it are key to unlocking its potential.

Today's data and IT leaders want data solutions that simplify the process of extracting value from data, allowing their teams to work quickly, confidently, and securely. With the right tools, they know they can help teams deliver better insights and maximize the value of their data without unnecessary complexity.

That's where Microsoft Fabric comes in.

Fabric is software as a solution (SaaS) platform that delivers an integrated and simplified experience on an open and governed foundation. It offers intuitive tools for streamlining analytics workflows and serves as a single source of truth for data, making it easier for teams to collaborate and deliver more impactful insights.

Use this e-book to understand Fabric's core capabilities, examine your unique analytics needs, and create a proof of concept for using Fabric to help your team get more valuable insights, faster.

Why Fabric?

Despite collecting vast amounts of data from apps, services, and IoT devices, many businesses struggle to turn it into actionable insights. Data analysis tools often require specialized expertise, creating bottlenecks as users must rely on experts to deliver insights.

Moreover, data silos across lines of business—HR, marketing, sales, finance, etc.—further complicate decision-making, leading to inefficiencies and missed opportunities.

As an organization grows, managing diverse data tools becomes increasingly complex and costly, diverting resources from core priorities. This complexity calls for a streamlined approach that simplifies data management and closes the gap between teams and the insights they need to thrive.

Adopting a unified data platform like Fabric is essential to overcoming the data challenges hindering AI innovation and blocking access to insights. With a data foundation built on Fabric, your organization can better unlock actionable insights, accelerating critical decision-making and driving sustained growth.

What's hindering access to insights?

Data is dispersed across multiple platforms, creating silos and making it difficult to access real-time information.

Teams lack sufficient data literacy and confidence to derive meaningful insights from their data.

Managing a complex web of different data tools and vendors increases operational costs and security risks while complicating governance.

How Fabric helps

Unify your data in Onelake. As Fabric's open and governed data lake foundation, OneLake provides a single source of truth so everyone is working with the same data. With multiple integration options, your teams can bring their data to OneLake without duplication or movement.

Empower your teams with democratized access. Providing secure access to the right data and analytics tools at the right time empowers your teams to gain insights faster, leading to better data-driven decisions.

Integrate your solutions on an open and governed foundation. Fabric unifies all your data, no matter where it's stored, giving your data stewards full control to manage it from start to finish, while addressing the long-standing security and compliance challenges your teams have faced for years.



Your data analytics needs

Every organization has unique data challenges, but they tend to fall into five core categories:

1. Simplifying data handling
2. Preparing data for AI
3. Providing specific experiences for different roles
4. Integrating seamlessly with existing tools
5. Maintaining comprehensive governance and security

Whether you're dealing with data silos, inefficiencies, or complex workflows, addressing these needs is critical to maximizing the value of your data and getting the best AI outcomes. Fabric provides a holistic solution designed to tackle these challenges head-on, equipping teams with the tools they need to work faster and smarter.

Simplified data

[OneLake](#) is a unified data storage solution within Microsoft Fabric that acts as a central hub where data can be stored, managed, and accessed effortlessly.

There are various ways to bring data into Fabric, including traditional ETL pipelines and mirroring, meaning some methods still involve replicating data. However, using shortcuts can help reduce data duplication. OneLake supports managed and unmanaged data storage so users can connect to various data sources without replicating the data. Through symbolic links—or [shortcuts](#)—OneLake enables access to data from different cloud providers within the same environment, removing the need to move or duplicate data.

This setup helps reduce data silos and allows a single copy of data to be used across multiple analytics tools and workflows. For example, data engineers can create real-time analytics dashboards without duplicating data, and data scientists can train machine learning models using the same dataset without creating multiple copies. This approach saves on storage costs while also improving data consistency and integrity across the organization.



Learn how to [create a lakehouse in Fabric](#).

A range of analytics experiences

Fabric offers a comprehensive suite of analytics experiences tailored to various data use cases, ensuring users can seamlessly manage, analyze, and visualize their data. Each experience is designed for specific roles and enhances productivity and insights across the entire data lifecycle.

Data engineering: Building efficient data workflows

Data engineers play a critical role in designing, building, and maintaining the infrastructure that enables efficient data collection, storage, and processing. Fabric simplifies these tasks, allowing engineers to focus on getting value from their data.

Data Factory

Use over 150 connectors to design, schedule, and orchestrate data pipelines, enabling seamless Extract, Transform, and Load (ETL) processes for efficient data integration.

Data Engineering

Develop large-scale data solutions using Apache Spark and other big data technologies to handle complex data workloads.

Data Warehouse

Create and manage data warehouses optimized for query performance and data integration.

Data engineering use cases

Speed up Spark applications

Fabric's native execution engine—included at no additional cost—allows data engineers to use the familiar Spark ecosystem while benefiting from faster deployment of Spark clusters. This results in enhanced performance, reducing the time to market for critical applications.

Migrate to Microsoft Fabric

Many organizations are migrating their existing Synapse Spark applications and SQL Server or Synapse Dedicated Pool data warehouses to Fabric, taking advantage of faster deployment and pricing flexibility. Fabric simplifies the migration of T-SQL code, making it a robust option for teams seeking a more streamlined, scalable solution.



Data warehousing: Centralizing and optimizing data

Data warehousing is key for organizations looking to centralize their data for analysis and reporting. Fabric enhances this process by enabling efficient migration and management.

Data Warehouse

Query and analyze large datasets, such as sales data, to identify trends and patterns.

Real-Time Intelligence

Gain access to real-time insights through up-to-the-minute data analysis.

Real-time data analysis: Immediate insights for critical decisions

For businesses requiring real-time data insights, Fabric offers tools that enable instant access to streaming data, empowering faster and more informed decision-making.

Real-Time Intelligence

Build real-time solutions that can ingest and analyze streaming data.

Data Activator

Automate alerts and triggers for in-motion data based on specific criteria.

Real-time analytics use cases

Point-of-sale data streaming

Retail businesses can use Fabric to stream real-time point-of-sale (POS) data from multiple stores into a centralized location. This capability—powered by Event Streams, Kusto DBs, and Data Activator—allows retailers to generate real-time reports, enabling quicker, more strategic decisions.

Time series analysis of manufacturing historian data

When dealing with large volumes of device telemetry, much of it is stored and often aggregated in a manufacturing historian. Fabric allows you to stream this historian data to OneLake, enabling advanced analytics like predictive maintenance, simulations, and cross-site comparisons.

Data governance: Ensuring compliance and security

Fabric's integration with [Microsoft Purview Data Governance](#) provides comprehensive data governance, allowing organizations to classify, protect, and manage their data with ease. You can classify Fabric items, configure data loss prevention, and monitor user actions to identify and mitigate potential security risks.

Governance use cases

Protect sensitive data

Using Microsoft Purview's Information Protection and Data Loss Prevention policies, organizations can automatically detect and protect sensitive data in Fabric, ensuring compliance with regulatory requirements. These tools prevent unauthorized access and restrict the export of sensitive data, safeguarding the organization from potential breaches.

[**Learn more about governance in Fabric.**](#)

Machine Learning: Scaling and automating AI solutions

Fabric empowers data scientists and business users alike to take advantage of machine learning and AI, integrating tools that streamline model development, training, and deployment.

Data Science

Develop and deploy machine learning models within an integrated environment.

Azure Machine Learning

Use advanced machine learning capabilities to experiment with algorithms and optimize models.

Data Science use cases

Transform large-scale unstructured data

Organizations with large volumes of unstructured or semi-structured text, such as customer service transcripts or legal documents, can use Fabric's integration with Azure OpenAI to derive insights. Scalable AI functions enables users to process these data types, extracting complex information such as tables and structured elements from large PDFs or audio recordings.

Deliver custom generative AI experiences for your data

An AI skill in Fabric allows your users to talk to your data in OneLake from your own applications. This "data agent" is configured as a conversational Q&A chatbot on your data domain. By providing instructions and examples to guide the AI, creators can ensure the "data agent" understands the organization's data context and provides a reliable data-driven response in return.

Reporting and visualization: Driving actionable insights

Non-technical business users often require easy access to data insights. Fabric ensures that data can be visualized and shared seamlessly regardless of technical expertise—or lack of it.



Power BI

Create interactive reports and dashboards, visualize data, and share insights.



Copilot for Power BI

Use natural language prompts to generate reports and insights, making data accessible to non-technical users.

Reporting and visualization use case

Cost savings through tool consolidation

By consolidating multiple data visualization and management tools into Fabric, organizations can save on costs and reduce technical debt. Fabric's Direct Lake mode allows users to mirror data sources, reducing the strain on existing data warehouses while maintaining access to real-time analytics.

Integrations

Your data platform needs to work with the tools your teams already use. Fabric integrates deeply with Microsoft Office, Teams, and other commonly used apps to ensure that insights are easily accessible across your existing workflows.

Whether it's through Excel, Word, or Power BI, you can analyze, share, and collaborate on data directly within the applications your team is already comfortable using, simplifying the data process and enabling quicker decision-making.

Governance and security

Data stewardship and security are critical as data environments become increasingly complex. Fabric simplifies this with built-in security and governance features, minimizing the effort needed to safeguard your entire analytics platform and reducing the risk of data breaches and compliance issues.

OneLake provides governance tools like data lineage, data protection, certification, and catalog integration to help streamline data management. These tools offer full visibility over your entire analytics environment, from data sources to user insights. Meanwhile, your tenant admins can easily maintain control over how your data is accessed and shared.

Fabric success stories

Government

Gamle Oslo enables unified data visibility

The Norwegian government is focused on creating the best possible society for its citizens. To better manage its data, Gamle Oslo sought to centralize and gain insights from it. The district began using Fabric as a unified platform, and after positive results from using it with a few datasets, they now plan to expand it across the entire district.



Fabric makes analytics accessible to people without a technical background. It makes this data accessible and means that we have a steppingstone to adjust, redevelop, and improve our services for residents.”

Aleksander Lorentzen,
Head of Digitization,
District of [Gamle Oslo](#)



The diversity of tools in Microsoft Fabric allows us to choose the right tool for the job. OneLake allows data consolidation from a variety of data sources. Alltech received top-notch technical and best practice support from both Microsoft and Sonata.”

Connie Thompson

Global Director of International
Technology, Data, Analytics,
and Reporting
[Alltech](#)

Agriculture

Alltech centralizes data and streamlines travel spend analysis

Alltech, a global leader in sustainable agriculture, needed to bring together data from its various acquisitions to track travel expenses better and support its sustainability goals. Partnering with Sonata, an early adopter of Fabric, Alltech integrated data from multiple systems into one place, centralizing its data and gaining valuable insights. The pilot project succeeded in managing travel costs and has set the stage for Alltech to continue modernizing its data strategy and empowering users with self-service analytics.

Retail

Chalhoub Group's People Analytics team accelerates reporting

Chalhoub Group, a luxury retailer based in the United Arab Emirates, wanted to streamline its HR approval process and eliminate the need to gather data from multiple platforms. After testing Fabric, the group modernized its HR data analytics by consolidating all data into one platform, improving agility, enhancing analytics, and speeding up processes.



We took full advantage of the six-month trial during the Fabric preview. We had anticipated some issues arising while Microsoft was still upgrading the platform. With the evolution of Fabric and incredible support from the Microsoft technical support teams [and] Fabric specialist[s], all of those challenges disappeared, and we were set up for success."

Mark Hourany

Director, People Analytics,
[Chalhoub Group](#)



The innovative ability of Microsoft Fabric to import data virtually from many sources will establish it as a key platform for supporting diverse applications of generative AI. We want to leverage this foundation to promote generative AI and other advanced technologies over the long term."

Yusuke Tsujii

ICT Division task force leader
[ITOCHU Corporation](#)

Professional Services

ITOCHU Corporation launches a data integration platform to accelerate AI

With more than 70 ideas for using generative AI, ITOCHU wanted to develop a data integration platform to bring these ideas to fruition. The company looked to Microsoft Fabric to help aggregate and cleanse data to make it more suitable for AI use.



Building a proof of concept for Fabric

A proof of concept is a trial used to demonstrate the feasibility of a solution in a real-world scenario. It's essential to determine if a platform meets your organization's data management and analytics needs. A successful experimental period will show the potential for production use by validating technical solutions like system integrations or specific configurations.

Follow this basic framework to see how your organization could realistically create and execute a proof of concept for Fabric.

1. Identify sponsors and potential blockers

To ensure your trial proceeds smoothly, it's essential to secure the necessary support and identify any roadblocks early on:

- Identify any organizational restrictions or guidelines for moving data to the cloud.
- Secure executive and business sponsorship for the proof of concept.
- Ensure that Fabric is suitable for your workload.

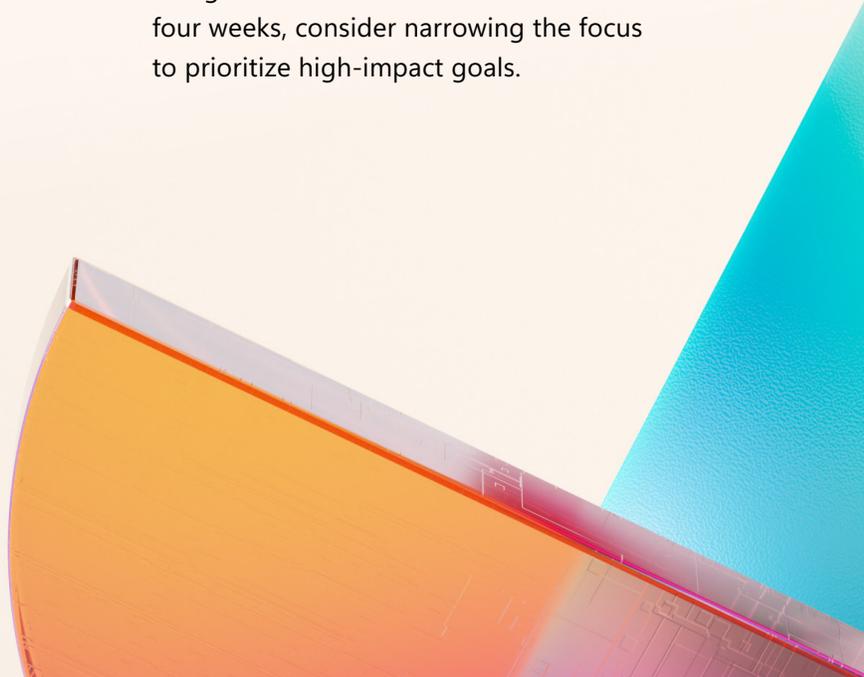
2. Set a timeline

A proof of concept should be a time-bound exercise with specific, measurable goals. Based on best practices, it's recommended that you timebox your trial to about two weeks. This will provide enough time for you to see meaningful results while limiting complexity.

Suggested timeline:

- **Data loading:** 3 days or less
- **Querying:** 5 days or less
- **Additional tests:** 2 days or less

When managing your timeline, making realistic estimates for task completion is important. The scope and timeline can vary based on factors like dataset size, the complexity of database objects, and the number of interfaces being tested. If the estimated time exceeds four weeks, consider narrowing the focus to prioritize high-impact goals.



3. Define the architecture and parameters

Design a high-level architecture for the proof of concept that includes relevant data sources, consumers, and any required AI or machine learning components:

- Decide which components will be part of the trial and exclude those that aren't immediately necessary.
- Identify existing resources (like Azure services) that will support the proof of concept.
- Choose an [appropriate region](#) and subscription for non-production use.
- Ensure your network infrastructure can handle the proof of concept workload without impacting production systems.

When considering migration scenarios from a legacy data warehouse, evaluating your approach is important. Are you aiming to make minimal changes to your existing ETL processes, or are you open to making significant improvements during the migration?

Alternatively, you might be building an entirely new environment, known as a greenfield project, from scratch. Each scenario requires careful planning to ensure a smooth transition.

4. Identify current pain points and set goals

Your experimental period should address specific pain points or gaps in your current system.

Consider the following questions:

- What gaps do you expect Fabric to fill?
- What new business needs are emerging?
- What service-level agreements (SLAs) must be met?
- What types of workloads (e.g., ETL, analytics, reporting) will the proof of concept focus on?

After you've identified your pain points, the next step is to define the goals of your trial clearly:

- Ensure that query performance meets new Service Level Agreements (SLAs).
- Assess system responsiveness for end users.
- Evaluate if current ETL processes work well with Fabric.
- Ensure robust governance and security compliance.



5. Create a test plan

Based on your established goals, design specific tests to validate the performance and capabilities of Fabric. Each test should align with a particular goal, with clear expectations for the outputs. Plan for specific queries, reports, ETL jobs, and other processes to be executed on a defined dataset. It's also essential to involve all stakeholders in agreeing on the test plan and success criteria to ensure alignment and clarity throughout the process.

6. Identify and validate the proof of concept dataset

From the tests you've planned, identify the data required in Fabric to execute them effectively.

Take some time to review this dataset with the following considerations:

- Ensure that the dataset accurately represents the future processing on Fabric in terms of content, complexity, and scale.
- Avoid using a dataset smaller than 1 TB, as it may not provide representative performance, and avoid overly large datasets, unless the proof of concept is specifically testing processing capacity.
- Determine the appropriate distribution pattern and indexing option for each table. If there are any uncertainties regarding distribution, indexing, or partitioning, include tests in your proof of concept to address those questions.

Here are some examples of Fabric trial goals and how they can be tested:

> Goal: Ensure that query performance meets new SLAs.

Test complex queries under varying loads to verify they can handle the expected data volume and complexity within the required timeframes.

> Goal: Assess system responsiveness for end users.

Measure latency and throughput of queries executed by users to ensure a smooth, efficient experience.

> Goal: Evaluate if current ETL processes work well with Microsoft Fabric.

Analyze ETL workflows for compatibility and opportunities for improvement, focusing on transformation logic and data loading processes.

> Goal: Ensure robust governance and security compliance.

Test data access controls, encryption, and compliance with regulatory standards to ensure that the platform meets organizational security policies and industry regulations.

7. Assemble your team

Ensuring that the right people are involved in the experimental period is critical to its success, as each role uniquely guides the project from planning to execution.

A diverse team brings together the necessary skills and expertise to tackle both technical and business challenges:

- A project manager to lead the effort.
- A business representative to oversee requirements and results.
- A data expert familiar with sourcing the data for the proof of concept.
- A Fabric specialist and potentially an expert advisor to optimize tests are needed.
- Other specialists, such as network or Azure administrators, for specific components of the trial.

Since you are evaluating a new platform, consider engaging a [Microsoft partner](#) or consultant with expertise in Fabric to assist with the proof of concept.

8. Run key workflows

Test the essential workflows that are part of your trial:

- Data engineers should test data integration methods and transformations—for example, as in medallion architectures.
- Data scientists can train machine learning models using the same dataset to ensure team efficiency.
- Business analysts can use generated reports and insights, testing ease of use, customization options, and the ability to drill down into their data for deeper analysis.
- IT and security teams should test role-based access controls and data masking features to ensure data access permissions are correctly configured.
- End users should interact with dashboards and reports to test their usability and accessibility, taking note of the ease of navigation and clarity of information.



9. Measure results

To measure the results of your trial, track key metrics such as data processing speed, capacity levels, and workflow efficiency. These metrics will provide valuable insights into how well Fabric performs in your environment. Focus on identifying improvements in operational efficiency, cost savings, and data consistency—especially the advantages of unifying all your disparate data to one central data lake.

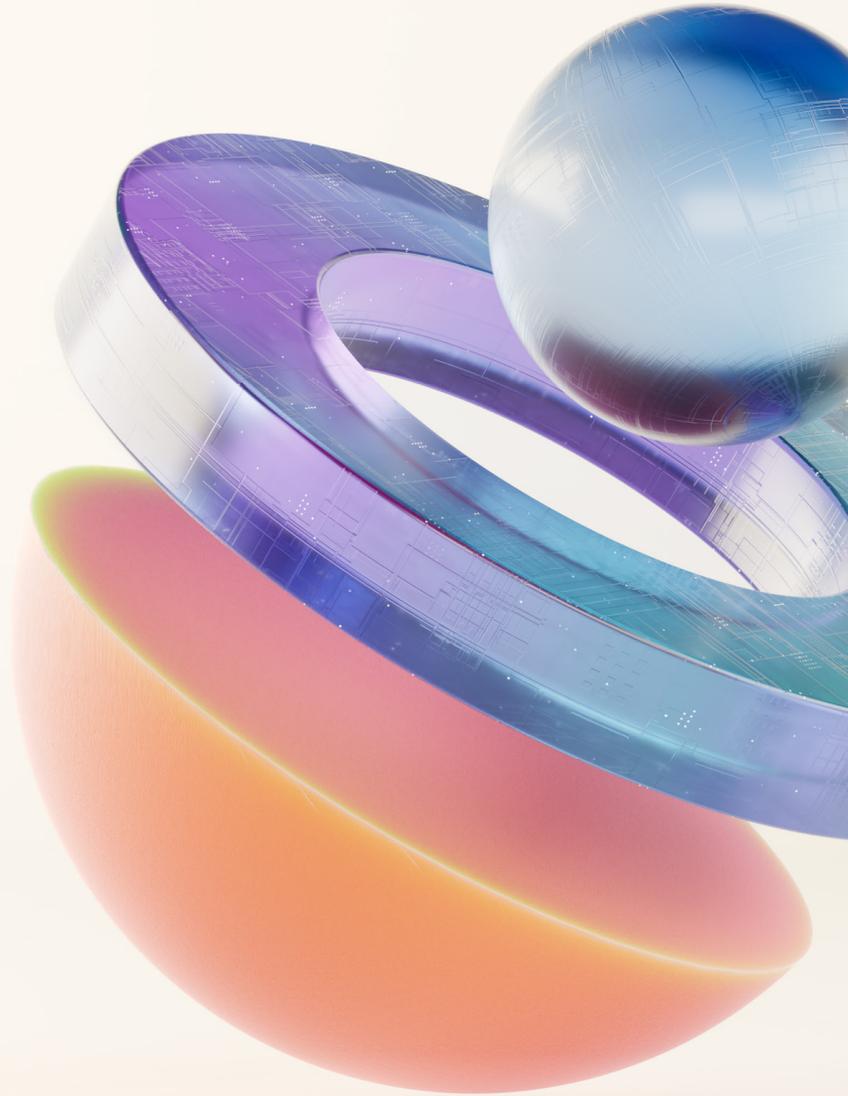
By assessing these factors, you'll be able to determine how effectively the platform streamlines processes and reduces the complexity and costs associated with data management.

10. Evaluate results

Once your trial is complete, review the outcomes against the initial goals:

- Have the goals been met?
- Were the expected outputs achieved?
- Identify areas where further testing or evaluation may be required and document any new questions raised during the experimental period.

This structured approach to a Fabric proof of concept will ensure a clear validation of the platform's capabilities in addressing your organization's specific challenges and goals.





Take the next step

Now that you understand what Microsoft Fabric is and how it helps organizations get more value from their data, it's time to take the next step—testing your Fabric proof of concept.

Whether you're facing issues like data silos, bottlenecks in decision-making, or high operational costs due to complex data systems, Fabric has the tools to overcome these hurdles.

- **Establish a single source of truth** for your data to break down silos.
- **Democratize data access** and empower teams to make data-driven decisions.
- **Connect and govern your data foundation** to reduce complexity and ensure security and compliance.

Ready to unlock your data's potential?

Start a trial run of Microsoft Fabric today and see how it can transform your organization's decision-making process.

[Sign up for a free trial](#) >