





Accelerating Business Insights with **Unified Analytics**



Contents

Introduction	3
Chapter 1	
The limits of innovating with a disconnected data estate	4
Chapter 2	
Unify your analytics tools for the era of AI	7
Chapter 3	
Real-world outcomes with Fabric	10
Conclusion	
The impact of embracing a “one-stop shop” for analytics	12



Introduction

Your data holds the potential to revolutionize your business at every level. Today, AI and machine learning are giving rise to new opportunities in data analytics that provide deeper insights into how to drive growth, improve operational efficiencies, and deliver exceptional customer experiences. These experiences are delivered through applications and services, where unified data powers real-time decision making and customer engagement.

To fully embrace the power of AI for data analytics, you need the right tools working in unison within a cohesive cloud environment.

This e-book explores how using a unified analytics ecosystem helps remove common barriers to using AI and machine learning to power faster insights and get more value from your data.



Chapter 1

The limits of innovating with a disconnected data estate

Data analytics helps fuel strategies for growth and optimization by providing deep insights that reveal how your processes operate and perform. For example, marketing professionals can use data analytics to examine the correlation between social media impressions and spending to help them maximize impressions per spend. Meanwhile, analytics performed on sales, competitor, and promotional data helps implement more profitable pricing strategies. In finance, it's used to optimize expense management by tracking spending and budgets and flagging outliers that might indicate fraud.

However, the reality is that many organizations face several barriers trying to extract these deep insights from their data. These challenges arise primarily from not having a unified platform that integrates all of their data and analytics tools. This fragmentation not only limits analytics outcomes but also makes it significantly difficult to build, deploy, and scale data driven AI applications.

Disconnected and duplicated data

Data silos take up more resources while producing lower-quality results

Organizations often gather data from various sources, including multiple clouds, on-premises, and third-party systems. A common issue arises when these sources aren't effectively integrated, resulting in unmanaged data silos. When this happens, organizations must spend more time and effort to store and manage multiple versions of the same data sets. Data silos also make it harder to ensure data is properly governed, leading to potential compliance issues or security breaches.

Meanwhile, data scientists working with silos often spend hours chasing requests throughout the organization to extract, load and transform data in preparation for analysis.

Concerning analytics, data silos can drastically limit the quality of an organization's reporting. Without a single source of truth for trustworthy data, organizations tend to rely more on descriptive analytics—that is, examining what's already happened—as opposed to exploring predictive analytics, which can help determine what's going to happen. These extra steps increase labor costs, slow decision-making and often require organizations to invest in additional tools, storage and processing capacity to support duplicate datasets and redundant workflows.

Ultimately, data silos and duplicate data end up using more organizational resources while simultaneously making it harder to deliver reliable and impactful insights.

Increased tech complexity

Managing different analytics solutions creates unnecessary complexity

Analysts and data scientists require various tools and services to uncover business insights. However, acquiring and maintaining these resources can become extremely costly, especially as more are added. Often, these tools are from different vendors, resulting in a challenging and lengthy integration process to guarantee seamless compatibility and functionality. The result is that tech teams are often focused on getting their tools to work together, rather than using their data for goal-oriented purposes.

This complexity of managing analytics solutions makes it harder for teams to be agile and responsive in a constantly changing market. Without a lean and integrated solutions stack, organizations are more likely to miss opportunities for improvement when they would have the greatest impact.

Increased security risk

Fragmented environments make sensitive data harder to protect and govern

Security and governance become exponentially harder to manage when data and analytics tools are scattered across multiple clouds, on-premises systems, and third-party platforms. Each additional tool or storage location introduces its own access controls, encryption standards, and compliance requirements, leaving security teams to stitch together a patchwork of policies rather than enforce one consistent standard. As more copies of the same data set get created and shared across these environments, the number of potential entry points for a breach grows right along with them.

This patchwork approach also limits visibility. Without a single source of truth, it's difficult for organizations to know exactly where sensitive data lives, who has access to it, and how it's being used—a problem that compounds quickly as AI applications and agents enter the picture. These systems often need broad access to data to power predictions and automation, and without strong governance in place, it becomes difficult to track what data a model has touched or whether it's exposing regulated information.

The result is that security teams spend more time on manual audits, access reviews, and incident response than on enabling the business to move faster. And as data privacy regulations continue to evolve, organizations with fragmented environments will find it increasingly difficult to prove compliance—putting customer trust and revenue at risk.


Limited real-time intelligence

Insights that arrive too late are insights that can't be acted on

Many analytics environments are still built around batch processing—pulling data into a warehouse on a schedule, then running reports or models against it after the fact. This approach works well for understanding what happened last week or last quarter, but it leaves a gap when the moment to act is now. By the time data has been extracted, loaded, transformed, and finally surfaced in a dashboard, the business conditions it describes may have already changed. A fraudulent transaction clears before the pattern is flagged, a pricing window closes before the adjustment is approved, a supply chain disruption is identified only after the shipment has already missed its delivery date.

This gap becomes even more costly as organizations shift from analytics that simply inform people to AI and agents meant to act on their behalf. Agentic AI depends on a constant stream of accurate, current context to make good decisions autonomously—if the data feeding it is hours or days old, or scattered across the same disconnected systems and tools described above, the agent ends up acting on a stale or incomplete picture of the business. The risk here isn't just a missed opportunity; it's an automated decision made confidently on the wrong information.

Ultimately, real-time intelligence requires more than faster dashboards—it requires unified, governed data that's available the instant it's created. Without that foundation, organizations are left automating yesterday's decisions, no matter how sophisticated their analytics and AI investments become.



Five key benefits of integrated analytics

Taken together, these challenges make it difficult for organizations to trust their data, simplify their technology, secure their environment, and act in the moments that matter most. But that's beginning to change.

By integrating your analytics tools and establishing a single source of truth for your data, Fabric makes it easier to find out what's not working optimally within your organization. From there, AI-powered analytics help you make timely, informed decisions to address those issues and track the progress of your improvement strategy. Meanwhile, additional benefits help ensure that your analytics processes remain lean, efficient, and adaptable to emerging technologies.

Benefit #1

Enhanced visibility and governance capabilities

[Microsoft Purview](#) works with Fabric to provide governance, risk, and compliance solutions that help ensure data is protected and consistent. Purview creates a unified map of your entire data estate that gives you a bird's-eye view of sensitive data and allows you to provision access only to the right users. Knowing that your analytical insights are based on high-quality data that adhere to even the most stringent compliance regulations gives you greater peace of mind.

Benefit #2

Higher analytics performance

Using multiple tools together allows you to process complex analytics workloads faster and more efficiently. You can use one tool to collect and store your data, another to process and transform it, and another to create charts and reports—all without the downtime normally resulting from disconnected solutions.

Benefit #3

Lower costs

Whereas stitching together different analytics tools creates additional costs through buying and managing multiple licenses, integrating your solutions helps lower costs. It eliminates the need for additional licenses and cuts costs by reducing time-intensive integration efforts that come with manually connecting those tools. Plus, thanks to OneLake, you don't have to pay more to store multiple versions of the same dataset.

Benefit #4

Ease of AI adoption

The compatibility between Fabric, Microsoft Foundry and Azure Application Platforms simplifies AI adoption by tightly integrating data, AI and application layers into a unified workflow. Microsoft Fabric provides a governed, unified data foundation where data is ingested, prepared, and made AI-ready. Microsoft Foundry directly builds on this foundation, enabling teams to develop, customize, and operationalize models and agents using Fabric data. Azure's application platform natively connects to both Fabric and Foundry, making it easy to deploy, run, and scale AI applications that embed real-time insights into business workflows, APIs, and user experiences.

Benefit #5

Real-time analytics

Without real-time analytics capabilities, businesses may have difficulty identifying the root of their biggest challenges. For example, if a company suddenly starts losing sales, it needs fast and complete insight into its processes to determine the cause and correct it quickly. With the integration of Synapse Real-Time Analytics, Fabric makes it easier to find and act on opportunities for improvement in minutes as opposed to hours or days.

Chapter 2

Unify your analytics tools for the era of AI

For leaders building the data and AI foundation of their organization who need to bring clarity to data sprawl and accelerate agent growth without slowing frontier ambitions, Microsoft's unified data and AI platform offers the only solution that is unified across your estate, provides context at every layer, and is secure and trusted by design end to end. Unlike fragmented stacks of disconnected tools, Microsoft gives you one foundation where data, AI, context, and governance work together as a single open, multi-cloud experience for the future of AI.

What is Microsoft Fabric?

Microsoft Fabric is an end-to-end analytics platform for ingesting, storing, processing, and analyzing data within a single integrated environment. Built on a software foundation as a service, Fabric eliminates the need to stitch solutions together in preparation for analytics workloads. Instead, Fabric users access their tools within an integrated interface optimized for their unique experience, so they can jump right into their analytics tasks without waiting for data to be copied, transformed, or moved. Essentially, Fabric makes it easier for data professionals and citizens to derive fast and meaningful insights from their data without running up against the limits of non-integrated tools mentioned earlier.

Components of Fabric

Integrated experiences for data professionals

Fabric encompasses a suite of analytics compute experiences. Bringing together these different experiences helps remove common roadblocks that traditionally slow down time to insights, providing data professionals with a clearer path to the data they need based on their roles and objectives.

Organizations can drastically increase the value of their insights by:

- Having a single source of truth for their data.
- Using an open and governed environment that keeps their data secure and compliant.
- Access to analytics tools for non-technical users.
- Integrating AI and machine learning capabilities in analytics workloads.

Capability	Experience	Objective
Data integration	 Data Factory	Automate the movement and transformation of data to a centralized location for analysis.
Data engineering	 Synapse data engineering	Design, build, and maintain infrastructure that makes data usable for analysis.
Data warehousing	 Synapse data warehouse	Store and organize data for fast querying, analysis, and reporting.
Data science	 Synapse data science	Interpret data to help businesses understand their customers, markets, products, processes, and performance more accurately and clearly.
Real-time analytics	 Synapse real-time analytics	Quickly collect, query, and process real-time data from various sources like websites, sensors, or devices.
Business intelligence	 Power BI	Create and share reports and dashboards that simplify and visualize data into actionable insights.

Unified storage built for modern analytics

All of these Fabric experiences are built on a cloud-based data lake foundation called OneLake. Sometimes referred to as “the OneDrive for data,” OneLake is a unified storage system that lets multiple analytical engines use a single copy of data, rather than requiring multiple sets across different groups. When one compute engine stores data in OneLake, every other compute engine can instantly access and use the same data without copying, moving, importing, or exporting it.

Just as Word, Excel, PowerPoint and other Office software automatically store documents in OneDrive, all the workloads of Fabric store their data in OneLake. It automatically indexes that data and scales to provision as much storage as needed, making it the only storage platform you’ll ever need for analytics.

OneLake shortcuts eliminate the need to move and duplicate data

Shortcuts are a key feature of OneLake. If your organization stores or uses data across different clouds, domains, subscriptions, and accounts, shortcuts allow you to bring that data into OneLake without any data movement or duplication.

Without OneLake shortcuts, you have to schedule periodic jobs to copy data from other locations into your platform. Depending on the data size, this could add significant time and overhead costs just to make your data accessible. OneLake shortcuts help you avoid these issues by letting you access data from different storage locations without copying it, giving your data professionals a head start on their analytics workflows.

OneLake shortcuts act like shortcut links on your desktop, allowing you to access data from other sources outside of Fabric, including:



Azure Data Lake Storage Gen2



Amazon S3



Dataverse



Fabric + Databricks: Making your data ready for AI

The Magic behind Fabric is in its ability to connect solutions that make it easier to apply AI and machine learning to your data, driven with class-leading Spark engines that can process data sets quickly and efficiently.

Because OneLake supports different data formats, users can easily access them from Fabric and Azure Databricks. This allows users to move data from one platform to another without losing quality or compromising security. Fabric and Azure Databricks help leverage AI and derive deeper business insights at an accelerated pace.

Bidirectional read/write interoperability between Microsoft OneLake and Azure Databricks is achieved natively through Delta Lake protocol sharing. This enables both environments to read and store the same data on OneLake simultaneously, leaving OneLake as a unified, single source of data discovery. If your end users are already on Databricks, they can stay there.



With Fabric, we've gained the speed to data we were hoping for, we've expanded data access, and we've even reduced the cost of generating that data by 50%."

Ramakrishna Burugupalli, Engineering Lead for Commercial Finance Data & Experiences, Microsoft

Chapter 3

Real-world outcomes with Fabric

Avanade cuts data duplication and frees teams to focus on business outcomes

Avanade, the global consultancy formed as a joint venture between Accenture and Microsoft, relies heavily on data and analytics to support both internal decisions and client engagements. With different teams across the company having different needs, Avanade had built a fragmented data stack over time, relying separately on Azure Synapse Analytics and Azure SQL Database for the backend, with Power BI for reporting. Whenever someone needed information that lived in a different tool, IT had to manually duplicate that data into the requester's preferred system—creating delays, adding complexity, and making the environment harder to secure and govern.

Choosing Microsoft Fabric let Avanade consolidate the tools its teams already used into a single SaaS platform built around OneLake, a unified data store. By moving to this centralized setup, Avanade gave employees the ability to work directly with data through whichever Fabric tool they preferred, without IT needing to copy or move it first.

Rather than attempting a full migration at once, Avanade rolled Fabric out in stages, starting with its analytics layer and adding further layers over time. This let individual teams, such as finance reporting,

start benefiting right away instead of waiting for a multi-year rollout to finish. The company also began using Copilot in Power BI to give executives plain-language summaries of their data instead of requiring them to navigate dashboards manually.

Through this shift, Avanade cut down on data duplication, sped up access to information, and strengthened its compliance with internal security policies. It also freed up its technical staff from spending most of their time maintaining

“

With Fabric, I can spend less time focused on project-oriented work and move the team to focus on other priority items. That's really powerful.”

Ron White, Global CIO, Avanade

[Read the full story →](#)




FoodPharma reduces reporting from two days to 90 minutes

FoodPharma is a contract manufacturer of functional food products, including soft chews, bars, and its own UltraChews format, for nutrition brands. As an FDA-registered facility with an “Excellent” SQF rating, accurate and traceable data is core to its operations.

The company ran six disconnected systems covering finance, production, manufacturing execution, maintenance, payroll, and communications. Cross-functional questions, like a plant manager investigating a yield dip, required the BI team to manually pull and reconcile data from each system, a process that took about two business days and often closed the decision window. Finance teams faced the same lag answering leadership questions about cost per unit and margin.

After validating the approach with a small proof of value, FoodPharma worked on a seven-week project to build a foundation layer on Microsoft Fabric, consolidating more than 50 tables and a terabyte of historical data into OneLake and automating daily refreshes. Power BI gave teams a consistent place to report, while Fabric’s built-in governance let every number be traced back to its source, an important feature for a regulated manufacturer.

The result: reporting time dropped from two days to about 90 minutes, and the BI team got back roughly 15 hours a week previously spent on manual work. Plant managers now get same-day answers instead of multi-day investigations, quality teams have full audit traceability, and the unified foundation positions FoodPharma to pursue future AI use cases like predictive maintenance and direct querying through Copilot.



Conclusion

The impact of embracing a “one-stop shop” for analytics

Adopting Fabric helped Microsoft’s Finance Data and Experiences (FD&E) team reduce process times by two-thirds, land reports five hours sooner, and cut data generation costs by 50%.

[Read the full story →](#)

With Fabric and other Microsoft tools, Heritage Grocers Group sped up data acquisition, scaled to 115 stores, and handled 1.3 terabytes of POS data through automated sorting, saving at least \$500,000 along the way.

[Read the full story →](#)

Launch your journey to unified analytics
Take the next step to enabling business-driven insights.

Speed up your AI transformation.
[Frontier Accelerate →](#)

Discover the AI app and agent factory.
[Build with Foundry →](#)

Get started with integrated analytics.
[Try Fabric for free →](#)

©2026 Microsoft Corporation. All rights reserved. This document is provided “as is.” Information and views expressed in this document, including URL and other internet website references, may change without notice. You bear the risk of using it. This document does not provide you with any legal rights to any intellectual property in any Microsoft product. You may copy and use this document for your internal, reference purposes.