




The Unified Data Platform for Trusted Agents

How Microsoft OneLake, Microsoft Fabric and Microsoft
Foundry activate multicloud, real-time data for AI

An abstract graphic on the left side of the page. It features a large, blue, translucent cube with a grid-like pattern on its top surface. Surrounding the cube are various blue, translucent, pill-like or capsule-like shapes of different sizes and orientations. The background is a light blue gradient with a soft orange glow in the upper right corner.

Contents

3

Introducing a unified data strategy for the AI era

4

From trusted data to trusted action

6

Unify: OneLake for your entire organisation

9

Govern: Consistent controls for analytics and agents

11

Activate: AI agents and real-time execution with Foundry

13

Roadmap to execution

14

Next steps

Introducing a unified data strategy for the AI era

In a world of AI-powered apps and agents, data is driving decisions and workflows at scale. Most organisations are still early in operationalising AI, while surveys from Deloitte and others show that a smaller group is accelerating rapidly.¹ The difference is having a foundation of trusted, connected, current data.

That's the challenge for business leaders. Data sprawl creates multiple versions of truth and multiple copies of the same datasets. Extract-transform-load (ETL) pipelines become a patchwork of point-to-point jobs that break whenever a source system changes. Governance and lineage gaps slow decisions, because teams can't confidently explain how metrics were built. Real-time signals stay trapped in operational systems while analytics run on delayed extracts. Without the right foundation, teams ship slower, adoption stalls and costs rise as more tools are added to compensate.



Unified, governed data is the shortest path to AI that people trust and use.

Today more than 30,000 customers rely on Microsoft Fabric to reduce silos, simplify their data architecture and support AI-ready experiences and agent workflows. Fabric is the unified data platform that brings operational, analytical and AI workloads together on a single, governed foundation. The platform runs on Microsoft OneLake, the logical data lake that gives teams access to the fresh, trusted data they need across data engineering, data warehousing, data science, real-time analytics, databases and business intelligence.

And with the integration of OneLake and Microsoft Foundry, the AI app and agent factory, you can ground agents on enterprise data already in OneLake instead of creating new copies and compounding sprawl.

This eBook explores why a unified, secure platform helps you activate your entire data estate for AI innovation and provides a roadmap to implementation.

From trusted data to trusted action

To move AI from experimentation to impact, you need a data foundation that teams can trust and reuse. For most organisations, that's difficult because data environments evolved over time into specialised systems and silos.

Three strategic shifts help you scale analytics and AI without ballooning cost or risk:

Move from isolated systems to shared data.

AI can't fill in missing context. Reducing fragmentation helps agents reason across operational and analytical realities.


Raise the standard to execution-grade data.

Humans can interpret data that's good enough, but AI and agents require data that's consistent, validated and governed so actions are trustworthy.

Add business meaning beyond schemas.

Semantics matter in AI, which needs shared definitions, relationships and logic so it can interpret data the way your organisation operates.

These shifts turn data from a record of what happened into a foundation for decisions and automated action.



The AI-ready roadmap: unify, govern and activate

A unified data foundation is how you scale AI safely and deliver reliable outcomes. The fastest path follows a straightforward pattern: Unify, govern and activate.

Unify your estate into a single logical lake

OneLake provides the organisation-wide foundation for each Fabric tenant, unifying all of your multicloud and on-premises sources, while also serving as a knowledge source for the agents your developers build in Foundry.

Apply consistent governance

With OneLake security, access policies follow the data, including when it's accessed outside Fabric through authorised external apps and agents. For teams, the OneLake catalogue makes governed data easier to find and reuse by centralising discovery, sensitivity, ownership and lineage.

Unify business meaning across data and AI

Beyond unifying data, you must unify meaning. Fabric IQ adds a semantic layer to OneLake, allowing teams to define business concepts once and apply them consistently across analytics, real-time decisions and AI agents.

Activate the data for decisions and agents


After you unify and govern your data, you can use it to ground agents so teams and automated workflows operate from the same trusted view.

Explore more →

"Microsoft Fabric is helping us evolve our data foundation into a more unified, AI-ready platform. Combined with Power BI and capabilities like Fabric IQ, it enables the enterprise to turn data into intelligence and act on it faster."


Shekhar Gowda

Vice President of Global Marketing Technologies at The Coca-Cola Company



Unify: OneLake for your entire organisation

OneLake makes data available to your business users and all the other teams that need it – including AI developers, data engineers, data scientists, business analysts and database administrators. Built on Azure Data Lake Storage (ADLS) Gen2, OneLake supports open table formats, so multiple workloads can access the same data without duplicating it.



OneLake provides a single governed foundation that serves business intelligence, data science, real-time analytics and AI.

Reduce duplication

You can load data once and use that same copy across Fabric workloads – from SQL and Apache Spark to Real-Time Intelligence and Power BI – while maintaining a single point for discovery and exploration.

Virtualise data across domains, clouds and systems

OneLake shortcuts are symbolic links to data that let teams reference data as if local, even when it lives in other OneLake workspaces or external locations such as Amazon S3, ADLS Gen2, Google Cloud Storage, Dataverse, Azure Blob Storage or OneDrive and SharePoint folders.

Bring databases and catalogues into OneLake – with governance intact

Through *mirroring*, you can add proprietary databases and data warehouses to OneLake without ETL.

Depending on the source, mirroring can replicate the full database or publish governed metadata into OneLake as Delta Parquet tables. Data stays in sync in near real time, giving you a low-cost, low-latency solution. Mirroring replica storage is free up to a capacity-based limit. You can also extend this pattern to your own applications, using open mirroring to bring proprietary data into OneLake.

Connect more data sources

You can bring in or transform even more types of data stores using built-in connectors in Fabric Data Factory. Your teams can automate ingestion and preparation without stitching together point solutions.

Enterprise data in the lake

Bring your existing trusted enterprise data into OneLake so your teams can combine it with the rest of your data estate and accelerate analytics and AI outcomes. Here are just a few common sources:

Dataverse links to OneLake via shortcuts, making Dynamics and Power Apps data available in Fabric for governed, read-optimised analytics and AI – without traditional export pipelines.

SAP data can be continuously replicated into OneLake to support near real-time analytics and AI on operational records. Your key performance indicators (KPIs) and agents stay consistent across SAP and non-SAP systems.

Oracle databases can be mirrored in OneLake for near real-time analysis with minimal latency and cost.

Azure Databricks Unity Catalogue can be mirrored into Fabric, enabling you to read data managed by Unity Catalogue from Fabric workloads.

And many more sources continue to be added.

A closer look at Azure Databases

Azure provides a comprehensive portfolio of databases built on a common foundation and with native mirroring to OneLake, helping you to scale AI transformation. Analytics, AI and agents can work from governed, near-real-time replicas without impacting transactional workloads. This separation of record and insight helps your teams move faster. Operational systems stay stable, while OneLake becomes the shared foundation for reporting, real-time decisions and AI grounding.

With fully managed services for SQL, NoSQL, PostgreSQL and MySQL, Azure Databases provide AI-ready reliability at scale for critical workloads. You can build new apps, AI agents and data-driven solutions faster through built-in support for GitHub Copilot-assisted development and leading developer, analytics and AI tools.

Design principles

Co-locate data and AI context. Put governed data in OneLake so Power BI can deliver fast, up-to-date insights without extra copies, and agents can ground on the same current enterprise knowledge.

Publish once, manage like a product. Define stable paths, owners and access expectations for shared datasets, so cross-domain reuse stays predictable.

Lead with identity. Use consistent identity-based access end-to-end so governance carries from analytics to agents.

Separate systems of record from systems of insight. Use mirroring to preserve workload isolation as AI adoption grows.

Measure what matters: speed, cost, trust, impact. Tie platform progress to outcomes like fewer pipelines, lower duplication, faster onboarding and governed AI releases.




Govern: Consistent controls for analytics and agents

The OneLake catalogue provides a central view of trusted, endorsed data across the organisation. You can find the right data faster, understand its governance posture and manage access and security roles in one place.

Under the catalogue, OneLake security provides a consistent model for sharing data without exposing sensitive information. You can define roles and permissions at the data item, folder, table or even row or column level, enabling teams to reuse shared data while restricting access to protected fields. These permissions are automatically enforced across all analytics experiences. Everyone sees only what they're authorised to see.

That same governance extends seamlessly into AI. When OneLake files are used as a knowledge source in Foundry, access controls and permissions stay consistent with the policies already in place for analytics. Foundry uses Azure AI Search to securely index OneLake content – including data surfaced through shortcuts or kept current through mirroring – so agents work with up-to-date information that is already governed and approved.

OneLake integrates with **Microsoft Purview** to extend security and governance across the data lifecycle. Using Microsoft Information Protection (MIP) labels and Purview data-loss prevention policies, your organisation can enforce company-wide controls, restricting access to sensitive data, preventing oversharing and helping automate protection. You can also use Microsoft Purview Data Security Posture Management for AI (DSPM for AI) to capture and review prompts and responses, helping you identify potential data exposure risks for Copilot in Fabric and for apps and agents built in Microsoft Foundry.



OneLake catalogue is a centralised place for finding, exploring and using the Fabric items you need, and governing the data you own. Your teams can work with the catalogue from the tools they use every day, including Microsoft Teams, Microsoft Excel and Microsoft Copilot Studio.

Prepare: Turn unified access into execution-grade data

Data preparation converts governed access in OneLake into curated datasets that your teams can reuse across analytics and AI. Without recreating copies or logic each time, Fabric helps you:

Accelerate your data projects. Use role-specific Fabric tools for data integration and engineering, data warehousing, data science, real-time analytics, data modelling and visualisation, in addition to partner-created workloads.

Curate gold datasets. Standardise and validate data, then publish reusable Delta tables in OneLake so business intelligence, machine learning and agents operate on the same definitions.

Integrate and orchestrate for reuse. Use Fabric Data Factory connectors and replication options (including mirroring) to automate ingestion, then shape data with Spark or SQL transformations into governed outputs that multiple teams can trust and extend.

Design for real time. Use Fabric Real-Time Intelligence to ingest streams, events and logs, and land the results in OneLake. Then you can join operational signals with curated OneLake tables so KPIs and agent context use the same current definitions.

Make semantics reusable. Use Fabric IQ to define business concepts once and apply them consistently across OneLake data, so analytics, real-time workflows and agents all operate on the same shared meaning.

Read the full story →



“By choosing Microsoft Fabric, we benefit from consolidated reporting, data engineering, data science and security under one umbrella and architecture.”


Ben Nott, Head of Global Data and Insight, Flora Food Group

Activate: AI agents and real-time execution with Foundry

Activation is where unified, governed data becomes action. AI developers can discover trusted data in the OneLake catalogue, then use Foundry to take that context the last mile – building, evaluating and governing AI apps and agents with clear enterprise guardrails.

Fabric and Foundry together provide an execution layer for agentic workloads. Live signals, shared business meaning and trusted organisational knowledge are combined to ground agents on current, governed sources. Foundry's built-in evaluation, monitoring and security help your teams move from pilot to production faster – while reducing risk by ensuring agents use governed data, follow access rules and remain traceable and auditable. Agents inherit the same access controls and definitions used across Fabric, ensuring AI applications reason over current, enterprise-approved data without creating new copies or bypassing governance.

The result: A repeatable path for turning data into reliable, real-time AI outcomes at scale.





Trusted knowledge for grounded answers

You can ground agents in enterprise knowledge using Foundry IQ knowledge bases, connecting them to sources such as OneLake, SharePoint and Azure Blob Storage with permission-aware retrieval and citations. This approach reduces speculative responses and gives agents traceable, auditable context – so decisions are based on trusted information, not guesswork.

Validation and readiness for production

Foundry adds evaluation and tracing to validate agent behaviour before automation is scaled on live signals. These controls help teams test and diagnose workflows early, reducing risk as AI moves into production. A centralised model catalogue streamlines model selection and deployment for real-time AI scenarios.

Safe execution in the agent loop

With Microsoft Agent Framework, you can orchestrate tools and multi-step processes to run agents reliably in real-world scenarios. Foundry Agent Service then provides the production layer – adding governance, orchestration and scale – so agent workflows operate safely and predictably as usage grows.

Read the full story [→](#)

“We’re enabling university leaders to dive deep into multiple sources of data and understand student engagement, experience and outcomes. The fundamental elements for this are Microsoft Fabric and Microsoft Foundry.”

Robin Gibson, Director of External Affairs, Kortext

Roadmap to execution

With Fabric and Foundry, you can lay the right foundation, connect more of your data estate with zero-copy and real-time capabilities and launch your first AI-powered experience.

Connect your data estate

Create an organisational and technical baseline so every new dataset is discoverable, owned and governed from day one.



Deliverables

- ✓ Create domain-aligned OneLake workspaces with clear ownership, define initial shared data assets with access and usage expectations and set service-level objectives.
- ✓ Add shortcuts to key sources and cross-domain datasets such as ADLS and S3.
- ✓ Configure mirroring for critical Azure Databases tables.
- ✓ Deliver operational OneLake catalogue, including sensitivity labelling, owners and basic governance posture visible to stakeholders.

Prepare and transform your data

Make more of your data estate available across teams without creating new copies or pipelines.



Deliverables

- ✓ Transform and prepare your data using Spark and Fabric Data Warehouse.
- ✓ Add one or two live Real-Time Intelligence streams.
- ✓ Create a contextual layer using semantic models and ontologies in Fabric IQ.

Build data-driven agents

Move from insights to trusted AI experiences that can be governed, evaluated and improved.

Deliverables

- ✓ Use Foundry to build a proof of concept, such as a RAG copilot or lightweight agent, grounded on OneLake gold tables.
- ✓ Evaluate and gate the release for safety and quality before deploying, then monitor and measure outcomes against your KPIs.

Upskill your teams

Develop knowledge and skills faster with the latest resources and expert insights from Microsoft Learn.

 [Get started with Microsoft Fabric](#)

 [Develop agents on Azure](#)

Next steps

AI is changing the standard for enterprise data. At the leading edge, trailblazers are using their enterprise knowledge for end-to-end AI transformation. A unified data platform based on Fabric, OneLake and Foundry helps you make the most of your data wherever it lives – across clouds, accounts, databases, domains and engines. Your teams can explore, govern and secure curated datasets in one place, then activate decisions with Real-Time Intelligence and build AI apps and agents in Foundry grounded on governed sources.

Together, Fabric, OneLake and Foundry create a repeatable path from opportunity to impact.

- ➔ [Get started with a complete data platform using Microsoft Fabric](#)
- ➔ [Learn how to unify your data with OneLake](#)
- ➔ [Discover, govern and secure with OneLake catalogue](#)
- ➔ [Act on streaming data with Real-Time Intelligence](#)
- ➔ [Build and evaluate AI agents and apps with Foundry](#)

¹[The State of AI in the Enterprise](#), Deloitte, January 2026.