Microsoft

Secure and Govern your Data with **Microsoft Fabric and Microsoft Purview**

Contents

Rethinking data security and governance for the era of AI

Today, data is the foundation for innovation, decision-making and transformation at scale. Yet for many organisations, managing that foundation has never been more complex. Multi-cloud environments, sprawling data sources, limited visibility and fragmented data security and data governance tools have created disconnected systems that are difficult to secure and govern. At the same time, oversharing and unauthorised access create exploitable vulnerabilities and make it harder to maintain compliance.
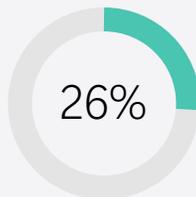
This fragmentation not only poses security risks, but it also has real business consequences: It slows innovation, delays AI adoption and erodes the ability to compete in fast-moving markets. According to

Gartner, only **14%** of security and risk management leaders say they can effectively secure organisational data while also enabling its use to achieve business objectives.[1] Complex, evolving regulations raise the stakes even further, while low data quality and unclear lineage undermine trust in analytics and AI outputs.
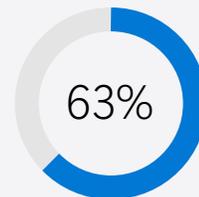
For AI specifically, the problem is even more acute, with **26%** of leaders saying they can't implement AI solutions at all because their data foundation lacks sufficient governance.[1] And without clear oversight, the risks also multiply: **63%** of breached organisations either have no AI governance policy or are still developing one, according to IBM's 2025 Ponemon Institute report.[2]
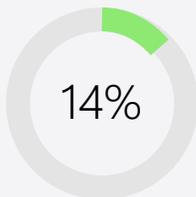
14%	**26%**	**63%**
Only **14%** of security and risk management (SRM) leaders can effectively secure organisational data assets while also enabling the use of data to achieve business objectives.[1]	**26%** of leaders say they can't implement AI solutions due to their data foundation lacking sufficient governance.[1]	In a 2025 report conducted by Ponemon Institute, **63%** of breached organisations either don't have an AI governance policy or are still developing a policy.[2]

[1] Gartner Press Release, Gartner Survey Reveals Only 14% of Security Leaders Successfully Balance Data Security and Business Objectives, February 11, 2025. GARTNER is a registered trademark and service mark of Gartner, Inc. and/or its affiliates in the U.S. and internationally and is used herein with permission. All right reserved.

[2] IBM Report: Cost of a Data Breach 2025

Overcoming these challenges means rethinking how organisations secure and govern their data – moving beyond fragmented tools and ad hoc processes toward a comprehensive approach. This starts with a single, unified data platform that brings all data together securely, consistently and with full visibility. Security should be built in at every layer, governance should let teams move quickly without losing control, and AI should rely on trusted, high-quality data.

With Microsoft Fabric and Purview, organisations can put this strategy into practice – keeping data secure and governed so they can adopt AI confidently, make smarter real-time decisions and stay adaptable in a fast-changing world.

Introducing Microsoft Fabric and Microsoft Purview: Unified platforms built to instil trust in data

Microsoft Fabric and Microsoft Purview are two powerful platforms that, when used together, create a trusted, end-to-end data foundation that lets organisations innovate with AI and analytics while maintaining control, visibility and regulatory alignment.

Let's explore these solutions in greater detail to see how they each contribute to creating a trusted data foundation for successful AI transformation.

Spotlight on Microsoft Fabric

Microsoft Fabric is a fully managed Software-as-a-Service (SaaS) platform that unifies analytics workloads – from engineering and integration to real-time analytics and business intelligence – into a single, pre-integrated experience. Its open, lake-centric foundation (OneLake) makes it easy to ingest and manage both structured and unstructured data from any source, all within a secure and governed environment.

To support a wide range of roles and use cases, Fabric offers a variety of integrated workloads, each tailored to specific tasks and designed to work together within the same platform:

Fabric workloads

Data Factory		Begin by integrating various data sources seamlessly
Data Engineering		Process and prepare data for advanced analysis
Data Warehouse		Store vast volumes of structured data ready for queries
Data Science		Dive deeper into data with machine learning and AI-driven insights
Databases		Streamline AI app development with autonomous SaaS databases
Real-Time Intelligence		Receive immediate insights from live data streams
Power BI		Visualise data and uncover hidden trends and patterns
Partner workloads		Add custom workloads from leading software developers built in Fabric

Having everyone working from a single unified data foundation keeps data consistent, avoids duplication and makes team collaboration faster and smoother. Just as importantly, it makes it easier to enforce robust security and governance at every step.

Fabric embeds security at every layer. Role-based access controls, built-in data protections and centralised monitoring help prevent unauthorised access, data leaks and compliance gaps. At the same time, it supports decentralised data ownership: teams can work independently with their own data while governance insights and administrative controls maintain consistent oversight across the organisation.

Designed for the era of AI

Being ready for AI today means more than just using it – it also means being prepared to build AI solutions that drive real business value. Fabric is designed for both: it equips your teams to leverage AI immediately while providing the foundation to develop and scale AI-driven innovations.

Fabric powers agentic AI and Copilot experiences using curated, high-quality data, enabling smarter, more responsible insights and automation. Teams can leverage AI tools directly within Fabric for any data project – and extend the platform with industry-specific features as needed.

> Fabric isn't just a SaaS analytics platform, it's a secure, AI-ready foundation for your data-driven future. With protections built in from the start, your data is unified, accessible and safeguarded at every step. Backed by robust governance and compliance, organisations can confidently innovate with AI knowing their data, and the insights it fuels, remain protected, private and trusted.

Built for the era of AI

→ **Copilot in Fabric**
Automate routine tasks and get smart suggestions to work faster across the entire data lifecycle, and use natural language to generate code, build pipelines, create reports and uncover insights.

→ **Open and AI-ready data lake**
Fabric lets you tap into your entire multi-cloud data estate through one single data lake, work with consistent data across different analytics tools and get your data ready to fuel AI innovation.

→ **AI-enabled business users**
AI-powered Q&A and rich visuals in Microsoft 365 apps let everyone – not just data experts – explore data and uncover insights quickly. Ask Copilot to get answers and understand key findings in a Power BI report instantly.

→ **Mission-critical foundation**
Confidently deploy and manage Fabric with category-leading performance, instant scalability, shared resilience and built-in security, governance and compliance.

What is OneLake?

OneLake is the unified data lake at the heart of Fabric, designed to simplify data management by providing a single, secure and scalable storage layer for all your analytics workloads. But OneLake is more than just centralised storage. It's built on an open, governed lakehouse architecture that combines the flexibility of a data lake with the structure and performance of a data warehouse. This modern foundation enables seamless data access, governance and interoperability across diverse platforms and formats.

Additionally, integrating Fabric Data Factory with OneLake turns your unified data lake into a complete, end-to-end platform. Teams can ingest, transform and orchestrate data from hundreds of sources directly within the lake, all while maintaining consistent governance, security and visibility. This seamless integration accelerates time to insights, simplifies architecture and ensures your data is immediately ready to power analytics and AI initiatives.

Key features of OneLake

→ **Delta Parquet open file format** supports ACID transactions, schema evolution and time travel, allowing Fabric engines like Power BI, Data Factory and Spark to read and write data natively without costly format conversions.

→ **OneLake catalogue** enables seamless metadata discovery across both structured and unstructured data, making data management easier regardless of format.

→ **OneLake security** defines consistent roles and access permissions across the entire data estate, ensuring strong protection without locking you into proprietary systems.

→ **Unity Catalogue integration** supports governance for Delta tables stored in Azure Data Lake by mirroring Unity Catalogue-managed datasets directly into OneLake. This allows real-time analytics and semantic modelling on Databricks data without duplication, streamlining cross-platform collaboration.

→ **Open data lake** messaging uses open APIs and protocols to enable smooth communication between Fabric and external platforms like Azure Data Lake and Databricks. Features like shortcuts and mirroring synchronise metadata and credentials, maintain cross-cloud interoperability, reduce complexity and allow flexible access without unnecessary data movement.

Organisations need flexible ingestion methods to securely connect and synchronise data from a wide range of sources, including on-premises systems, cloud platforms, real-time streams and external applications.

Fabric offers several secure ingestion methods designed to handle diverse data sources and use cases efficiently:

→ **Shortcuts**
Access external data stores like ADLS Gen2, Amazon S3 and Dataverse directly without moving or duplicating data, reducing latency and storage costs.

→ **Mirroring**
Keep data in sync with operational databases like Azure SQL, Cosmos DB and Snowflake in near real-time, eliminating traditional ETL overhead.

→ **Fabric Data Factory**
Use 170+ built-in connectors to efficiently ingest data from a wide range of sources, including databases, SaaS apps and cloud storage.

→ **Fabric Real-Time Intelligence**
Stream data into OneLake with minimal latency, powering real-time analytics for scenarios like anomaly detection and live dashboards.

Spotlight on Microsoft Purview

Microsoft Purview offers a modern, unified approach to help organisations secure and govern data across their heterogeneous data estate. With integrated data security, governance, compliance and privacy solutions for the era of AI, Microsoft Purview helps reduce risk and complexities, increasing team productivity and improving overall data security, governance and regulatory compliance. By bringing discovery, classification and protection tools into a single platform, Purview ensures that sensitive information – whether in databases, file systems, SaaS applications or AI platforms – is visible, secure and well-governed.

Microsoft Purview is purpose-built for Data Security teams and admins, Governance offices (including data stewards and data owners) and Risk and Compliance teams (including Compliance and Privacy offices). It enables organisations to apply consistent governance policies across all data, enforce protection with sensitivity labels, access controls and data loss prevention and meet regulatory requirements with built-in compliance and audit capabilities.

Just as importantly, Purview also helps reduce risk by monitoring how data is accessed and used, ensuring that innovation never comes at the expense of privacy, security or trust.

Purview covers three main areas to help organisations secure and govern their data effectively.

Data Security

Information Protection, Data Loss Prevention, Insider Risk Management, Data Security Investigations and Data Security Posture Management

Data Governance

Unified Catalogue (Curation, Discovery, Governed Access) and Data Management (Data Quality, MDM and Health Controls)

Data Compliance

Compliance Manager, eDiscovery and Audit, Communication Compliance, Data Lifecycle Management, Records Management

By staying aligned with evolving regulatory requirements, Purview enables responsible AI adoption while safeguarding what matters most: your data.

Combining capabilities for end-to-end security and governance

Together, Fabric and Purview form a powerful duo that addresses the core challenges of modern data management. Fabric provides a unified, AI-ready data platform that simplifies and accelerates your data workflows with security and governance designed for Fabric users built in. At the same time, Purview adds an estate-wide security and governance layer to ensure all of your data stays secure, compliant and well-managed.
With these solutions working together, you can securely manage your entire data estate, activate your data with confidence knowing it's trustworthy and governed and uphold federated governance to meet evolving regulatory and AI compliance demands.

Let's take a closer look at the specific features and capabilities that make all this possible and how they can transform your organisation's data strategy.

Embedded security at every layer

Fabric and Purview work together to provide robust security tools for your data. Fabric offers network, data, item and workspace security in the form of granular RLS, CLS, OLS and encryption and Identity security with Entra ID integration. Purview adds data security to discover risks and protect data consistently across your estate. Together, they help organisations detect and prevent data loss and discover hidden risks. At the same time, they balance flexibility with control, making it easier to innovate securely across different departments and locations.

Here's how this integration allows organisations to secure their data estate.

→ Microsoft Fabric comes with end-to-end **inbound and outbound network security** so you can connect securely to any data source. At the tenant level, you can use Microsoft Entra ID to authenticate every inbound request. For specific scenarios, you can use Azure Private Links to send traffic over Microsoft's own private network. You also enable specific workspaces like Azure Private Links, Outbound Access Protection, Customer Managed Keys, Workspace IP Filtering, Trusted Workspace Access and more – giving you precise control over the isolation level of each Fabric workspace based on the sensitivity of its data.

→ Once your data is in Fabric, you can use **OneLake security** to define security on specific folders, rows and columns on your data items, so you can share data items without exposing sensitive data. This security then ladders up to roles which can be defined at every layer of Fabric, including domains, workspaces and items, and enforces it uniformly across all Fabric engines, so users only access the data they need.

→ Once your data is in Fabric, you can use **OneLake security** to define access to specific folder, rows and columns, so you can share data items without exposing anything sensitive. Data owners can view policy violations, receive alerts and investigate them. DLP ensures that data is secure and not forgotten, irresponsibly scattered throughout the organisation. Users view DLP value in Fabric through policy tip triggering (lakehouses, structured data, mirrored items in OneLake) and restriction of access (lakehouse items)

→ Fabric users can now leverage **Microsoft Purview Information Protection** to apply sensitivity labels to Microsoft Fabric workspaces. This new functionality allows users to manually label Fabric items based on sensitivity levels, with access controls automatically enforced according to pre-defined protection policies by administrators. This empowers Fabric users to protect their Fabric items consistently with the same sensitivity labelling framework used within their M365 unstructured data estate.

Discover and mitigate hidden risks to data and AI usage

In addition to securing data and preventing oversharing, Purview provides an additional layer of unique security value by enabling Fabric users to detect and mitigate insider threats. Here is how Purview allows this:

→ **Insider Risk Management (IRM)** in Purview detects potential insider risks and risky user behaviour, such as data leaks or policy violations within Fabric. User risks are assessed with a score, which allows security professionals to act on malicious user activity to safeguard an organisation's data.

→ **Microsoft Purview Data Security Posture Management for AI (DSPM for AI**) supports Copilot in Power BI, enabling Fabric users to identify data risks, including sensitive information appearing in Copilot prompts and responses. Actionable recommendations are provided through DSPM for AI reports to help address these risks. Additionally, users can govern Copilot interactions using tools like Audit, eDiscovery and retention policies, while also detecting non-compliant usage to promote responsible AI practices.

Confident data activation

The goal is to activate data for insights and AI. To help teams accelerate data innovation, Fabric and Purview allow users to confidently work with protected data within governed experiences that meet the business where it is. This integration enables organisations to centrally configure organisation-wide policies while delegating granular management to those who need it through a flexible, federated data mesh. Fabric and Purview also provide complementary catalogues that ensure data teams and business users have the governance and access they need for their specific roles.

Comprehensive visibility

→ **The OneLake catalogue is an operational catalogue** that helps Fabric discover and manage trusted data, providing governance for data owners with valuable insights, recommended actions and tooling. It's the default source of data for over 30 million monthly active Power BI and Fabric users, most of whom are business users. The OneLake catalogue is also embedded in the apps people use every day, including Microsoft Teams, Microsoft Excel, Microsoft Copilot Studio and hundreds of others.

→ **Unified Catalogue is an enterprise catalogue** that provides a source of truth for enterprise data discovery, allowing business users to discover metadata from Fabric and other multicloud, multiplatform environments easily and across more than 50 data sources across your data estate. Curate data assets in the form of Data Products (business use cases) that allow business users to act upon the data they are searching for regulatory and AI compliance demands.

Grow your data confidence

Building strong confidence in your data starts with knowing your data's quality, where it lives and how it flows across your organisation. These elements work together to ensure your data is trustworthy, well-governed and ready to power reliable AI and analytics.

→ **OneLake catalogue data quality** provides insights and recommended actions to help data owners govern their data. Unified Catalogue provides a comprehensive data quality solution for your federated data governance practice.

→ **Unified Catalogue Data Quality** empowers data owners to oversee the quality of their data ecosystem and facilitate targeted actions for improvement.

→ **Data Quality rules in Purview** enable checks across domains, data products and data assets – as rules flow through your environment. They enable asset-specific queries to check for de-duplication, repetition and empty entries to help improve your data quality.

→ **Deep data quality scans in Purview** for Fabric allow data quality managers and the Chief Data Officer to scan assets in Fabric, such as Delta, Iceberg, Parquet, Avro and ORC, to quantify whether assets are of good quality or not for responsible AI innovation and business usage.

Know where your data goes

→ Fabric **data lineage** is purpose-built for data teams to see how data flows through Fabric projects and allows users to perform impact analysis to assess the impact of data changes.

→ Purview Unified Catalogue **data lineage and data observability** enable you to complement the lineage built into Fabric OneLake and expand on the enterprise oversight that an enterprise catalogue provides for true end-to-end lineage of 50+ data sources.

Each of these is key to successfully activating one's data estate for data usage and AI innovation. Without proper governance, organisations are left to navigate a highly regulatory environment without the necessary precautions.

Federated governance and AI compliance

With Fabric and Purview, organisations can innovate quickly without compromising control. Fabric's decentralised data ownership model comes with built-in governance insights and admin controls, giving teams the freedom to move fast while staying compliant with global standards like FedRAMP, SOX, GDPR, EUDB, HIPAA and ISO certifications. Over 54 global data centres ensure you can meet residency requirements, while Purview's federated governance model meets data professionals where they work, maintaining centralised oversight.

Every activity in Fabric is automatically logged in Purview Audit and accessible via APIs, making it seamless to support security reviews, forensic analysis and investigations. Together, these solutions strike the perfect balance between agility and control – enabling secure, compliant innovation across every business unit and geography.

Compliance for the era of AI

AI's potential in the enterprise depends on more than just powerful models – it depends on having the right guardrails in place. The following tools and solutions each bring unique capabilities to AI innovation, and when integrated with Purview, they ensure those capabilities are delivered securely, compliantly and with full transparency.

→ **Fabric Data Agents** enable users to query structured data – such as lakehouses, warehouses or Power BI models – using natural language. When connected to Purview, they automatically respect sensitivity labels and access controls, so governed data stays governed, even in conversational AI scenarios.

→ **Mosaic AI** provides an enterprise-grade platform for building, testing and deploying generative AI applications with governance and observability built in. Through integrations with Unity Catalogue and Purview, it can trace both data and model lineage, enforce access policies and capture detailed telemetry for compliance reporting.

→ **Azure AI Foundry and Copilot Studio** make it easy to build custom AI agents that work across your teams, Microsoft 365 or even public websites. By using Fabric Data Agents as knowledge sources, Azure AI Foundry and Copilot Studio can help you create data-rich agents while enforcing Purview policies on the data in every interaction.

→ **Microsoft Purview Data Security Posture Management for AI support for Copilot in Power BI** allows Fabric users to establish AI boundaries by applying sensitivity labels that prevent Copilot and other AI tools from summarising or generating content from sensitive data. Users can discover data risks such as sensitive information in Copilot for Power BI's prompts and responses, with actionable recommendations surfaced in Microsoft Purview Data Security Posture Management for AI (DSPM for AI) reports. Users can also govern Copilot interactions using Audit, eDiscovery, retention policies and identifying non-compliant usage to support responsible AI usage.

Enhancing interoperability and governance:
Strategic partner integrations

Strategic partner integrations are vital for organisations aiming to unify data from diverse platforms while maintaining strong governance and seamless interoperability. Although Snowflake, Oracle and MongoDB aren't custom-built Fabric workloads, Fabric integrates closely with these external systems to bring their data into a governed, interoperable environment.

Snowflake

Enable bidirectional data access between Snowflake and Fabric and analyse your data in any engine within either platform.

MongoDB

Mirroring and connectors allow MongoDB data to sync in near real-time with OneLake, bringing NoSQL data under Fabric's governance framework and making it accessible for analytics and AI workloads.

Oracle

Oracle data is ingested and transformed using Fabric's dataflows and pipelines, enabling relational data to flow into Fabric with governance controls intact for reporting and modelling.

Architect success
with Microsoft Fabric
and Microsoft Purview

Fabric and Purview combine to provide a secure, unified data foundation that connects, governs and protects data across your organisation. This trusted environment makes data discoverable and auditable, which is key for driving confident AI and analytics innovation.

With built-in security controls and consistent governance policies, they reduce risk by preventing oversharing, detecting risky behaviour and enabling secure collaboration. Plus, their federated governance model empowers teams to innovate quickly while maintaining control and meeting compliance requirements.

Together, Fabric and Purview ensure AI relies on high-quality, compliant data, accelerating responsible AI adoption and building trust with stakeholders. This integrated approach helps turn your data into a strategic advantage, ready to fuel innovation and growth.
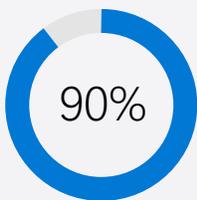
90% of data security, governance, compliance and privacy leaders say their organisation will adopt a unified solution to mitigate data-related risks.[3]

Next steps

→ Explore **Microsoft Fabric** and **Microsoft Purview** hands-on for free.

→ Assess your organisation's **readiness for secure AI**.

→ **Talk to an expert** about federated governance at scale.

[3] Microsoft Customer Requirements Research 2024