News Release

Teradata Brings Enterprise AI Directly to Microsoft OneLake

2026-09-02

Customers can now run Teradata enterprise AI on OneLake data where it lives, without replicating or moving it

SAN DIEGO, Sept. 2, 2026 /PRNewswire/ -- **Teradata** (NYSE: TDC) today announced the integration between the Teradata Autonomous Knowledge Platform and Microsoft OneLake, enabling enterprise customers to run Teradata's high-performance enterprise AI directly on data stored in OneLake -- without extract, transform, and load (ETL) pipelines, data duplication, or migration.

Built on open Apache Iceberg standards, the integration allows Teradata users to query OneLake tables in place using standard Iceberg APIs, with cross-platform authentication and access controls handled natively. Enterprises with data in both environments can now join and analyze it without first consolidating it into a single system or rebuilding governance from scratch.

Why it matters

Moving data between analytics platforms has long been one of the more expensive and time-consuming parts of running enterprise AI. Teams routinely maintain separate ETL pipelines to synchronize copies of the same data across systems, introducing latency, governance gaps, and storage costs that compound as AI workloads scale.

This integration eliminates that step for customers running both Teradata and Microsoft Fabric environments. Data stays in OneLake where Teradata reads it directly.

This means that data engineers spend less time maintaining pipelines, AI teams can begin feature engineering without waiting for data to be replicated, and governance — including security policies and lineage — travels with the data rather than being rebuilt for each copy.

What it does

The integration supports read access to OneLake tables via Iceberg APIs, with Microsoft Entra ID handling cross-platform authentication. Customers can run Teradata AI analytical workloads, including complex joins, AI feature pipelines, and SLA-bound tactical queries, against data stored in OneLake without modifying that data or requiring it to exist anywhere else.

For organizations already invested in Microsoft Fabric, Teradata is now accessible as an enterprise AI platform within that ecosystem, bringing the performance, workload management, and enterprise SLAs that mission-critical AI demands. For existing Teradata customers, it expands access to the full Microsoft Fabric and Azure environment without requiring data migration.

Enabling AI Where Data Already Lives

The OneLake integration extends the Teradata Autonomous Knowledge Platform vision of enabling enterprise AI wherever data already lives. The Autonomous Knowledge Platform is designed to provide the business context, governance, and performance backbone that AI agents require to operate reliably at scale, including across hybrid and multi-cloud environments. Bringing Teradata AI and analytics to OneLake data directly supports that architecture by removing data movement as a prerequisite for building governed AI workflows.

Executive Commentary

"Moving data to analyze it is a constraint that adds cost and complexity without adding value. Bringing Teradata enterprise AI directly to Microsoft OneLake means customers can run the workloads that matter, on the data where it already lives, inside the governance model they've already built. This is a concrete step in making the Autonomous Knowledge Platform work wherever enterprise data exists."

— Sumeet Arora, Chief Product Officer at Teradata

"Microsoft OneLake is built to be the unified data foundation for the AI era, connecting organizations with the platforms, tools, and data ecosystems they already trust. Teradata's integration through Apache Iceberg exemplifies the openness and interoperability at the heart of Microsoft Fabric, enabling customers to bring their data together and accelerate their AI transformation without the complexity or rearchitecting around a single vendor."

— Dipti Borkar, Vice President, Microsoft OneLake and ISV Ecosystem at Microsoft

Availability

Read access to Microsoft OneLake tables via Teradata is now available.

The integration will be demonstrated at the **European Microsoft Fabric + SQL Community Conference,** Sep 28 – Oct 1, 2026, in Barcelona, Spain

About Teradata

Teradata empowers enterprises to turn intelligence into autonomous action, grounding AI agents in deep business context and trusted data. As AI agents multiply, Teradata is the context foundation, governance layer, and performance backbone that companies need now. The Teradata Autonomous Knowledge Platform puts AI into production across cloud, on-premises, and hybrid environments.

MEDIA CONTACT

January Machold

january.machold@teradata.com

View original content:**https://www.prnewswire.com/news-releases/teradata-brings-enterprise-ai-directly-to-microsoft-onelake-302866131.html**

SOURCE Teradata Corporation