News Release

Teradata Brings Elastic Python Compute to AI Studio, Unifying Data, Governance, and AI at Enterprise Scale

2026-10-14

New distributed Python capability accelerates large-scale AI operations while keeping data and compute within the Teradata platform

SAN DIEGO, Oct. 14, 2026 /PRNewswire/ -- **Teradata** (NYSE: **TDC**) today announced the general availability of Elastic Python Compute within AI Studio. The capability gives enterprise data science teams a managed, distributed Python environment for scaling AI compute and machine learning workloads through parallel Python execution for data engineering, machine learning, and model development.

Elastic Python Compute is built on Ray, a distributed computing framework, and integrates natively with the Teradata Autonomous Knowledge Platform and AI Studio. Data scientists can parallelize model training, execute custom Python applications at scale, and serve high-scale ML and AI workloads without the overhead of manual cluster setup or managing Ray clusters. Elastic Python Compute supports both single-server and cluster deployments, giving organizations the flexibility to scale from individual workloads to enterprise-scale operations, with faster development cycles, stronger governance controls, and lower infrastructure costs.

Keeping Data and AI Compute Within the Platform

Elastic Python Compute supports local model hosting and governance, with no external API dependencies and no token-based inference costs. AI workloads run entirely within the Teradata platform, including AI Studio, with no data movement required. Elastic Python Compute reads and writes data directly to the Teradata Database, S3 object storage, and Open Table Format catalogs, such as Apache Iceberg and Delta Lake, enabling data science teams to incorporate files and images for model training and scoring. Flexible cluster management, including the ability to run separate clusters for different workloads, allows teams to optimize cost and performance by selecting CPU or GPU compute based on the task at hand. Massive parallel processing supports large-scale AI and data workloads across cloud, on-premises, and hybrid deployments.

The platform also supports Arrow Flight SQL, enabling parallel, columnar data transfer between applications and the Teradata platform **at speeds up to 84x faster than traditional row-based transfer methods,** based on Teradata benchmark testing. *

Customer Impact

A retail bank had been exporting tens of millions of customer records to local Python environments for credit feature engineering, a practice that created data governance gaps and put the organization at risk of breaching data residency regulations. With AI Studio's in-database analytics, credit scoring pipelines covering payment history, utilization ratios, and delinquency flags now run three times faster while remaining within the governance perimeter. A complete audit trail from raw data through feature engineering to final model output allows regulators to replay any step in the process on demand.

For a major telecommunications provider, the challenge was scale. Its data science team had been exporting billions of call detail records and network telemetry events to external Python clusters, a slow, costly process that created data sovereignty risk. Feature engineering across call drop rates, signal quality indices, and ARPU (average revenue per user) trends now runs directly inside the Teradata platform. The churn prediction model trains entirely within the provider's own infrastructure, satisfying both GDPR obligations and national data residency requirements. Feature engineering runtime fell from four hours to twenty minutes, and model accuracy improved by twelve percent as a result of training on the full dataset rather than a constrained sample.

In the public sector, a government maritime agency needed to track vessels switching off their transponders to obscure their movements. Using Elastic Python Compute, the agency's data science team combined satellite imagery, radar signals, AIS data, and commercial imaging to reconstruct vessel movements that would otherwise go undetected. AI-driven analytics fused these diverse data sources into actionable maritime intelligence, improving visibility and enabling faster, more effective responses to security threats at sea.

Executive Perspective

"Enterprise data science teams have spent years working around the fundamental constraint of moving data outside their platforms to build and train AI models, accepting the governance risk and cost that comes with it. Elastic Python Compute removes that constraint. Teams can now run distributed AI workloads at scale where their data already lives, and get to production faster without compromising on governance or cost control."
— Sumeet Arora, Chief Product Officer, Teradata

Availability

Elastic Python Compute is generally available within AI Studio on AWS and Microsoft Azure. Availability on Google Cloud is expected in Q4 2026, with on-premises available now with Teradata Factory.

* Performance results may vary. Benchmarks were conducted in a specific test environment and may not reflect results in your infrastructure.

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.

The Teradata logo is a trademark, and Teradata is a registered trademark of Teradata Corporation and/or its affiliates in the U.S. and worldwide.

MEDIA CONTACT

January Machold

January.Machold@Teradata.com

View original content:**https://www.prnewswire.com/news-releases/teradata-brings-elastic-python-compute-to-ai-studio-unifying-data-governance-and-ai-at-enterprise-scale-302908570.html**

SOURCE Teradata Corporation