



# DAIS 2025: KEY ANNOUNCEMENTS AND INSIGHTS



Tredence and Databricks are collectively rewriting the enterprise Data and AI playbook. As a trusted partner of Databricks, we are proud to stand at the forefront of enterprise AI transformation. We are enabling a new era of **zero-friction development** across Data Foundations, Analytics, and Generative AI—focused on **interactive learning loops** that bridge the gap between business users and technologists. Be it Agent Bricks, Lakebridge, Lakeflow or Lakeflow designer – each of these innovations stay true to this vision.

We're thrilled to share a major milestone in this journey: Tredence has been named the **2025 Databricks Retail & Consumer Goods Partner of the Year** for the fourth consecutive year and recognized as the **Growth Partner of the Year – Americas!** This recognition not only affirms our leadership in data and AI transformation but also strengthens our shared mission with Databricks—to co-create the future of enterprise Data & AI.

**Here's what mattered in DAIS 2025: the standout tech updates and the strategic takeaways for Data and AI Leaders.**

## 01. AGENT BRICKS

### SUMMARY:

Agent Bricks is a no-code framework that enables enterprises to build, evaluate, and deploy high-quality AI agents on proprietary data, tailored for common AI use cases. Users can select a use case from a library of **prebuilt "bricks"**—including information extraction, knowledge assistance, custom LLMs for text transformation, and multi-agent supervisor systems—and define tasks in natural language. Agent Bricks then automatically generates task-specific evaluation benchmarks and applies advanced techniques such as fine-tuning, reward modeling, and **Test-Adaptive Optimization (TAO)** to continuously improve agent performance.

A standout feature of Agent Bricks is **Agent Learning from Human Feedback (ALHF)**, which empowers domain experts to steer agent behavior using natural language guidance. This feedback is intelligently converted into technical system-level optimizations, enabling agents to learn and adapt over time.

### KEY CAPABILITIES:

- ◆ Support for connecting to Genie Spaces
- ◆ Automatic index creation for documents in vector databases
- ◆ Seamless MLflow3.0 integration, enabling automatic tracing and metrics logging
- ◆ Agent learning through user provided natural language Q&A formats or with synthetic Q&A generation to enhance quality.
- ◆ The multi-agent supervisor system—an AI system that brings Genie spaces, agent endpoints, and tools using MCP servers—is a true game changer for orchestrating complex, collaborative AI agents.

### OBSERVATION & ENTERPRISE BENEFITS:

Databricks has introduced the boilerplate framework with Agent Bricks, empowering enterprises to build production-grade agentic systems for high-impact use cases like **knowledge assistance, information extraction, custom LLMs for text processing, and fully orchestrated multi-agent supervisors**. In just a fraction of traditional development time, organizations can transform ideas into intelligent, enterprise-ready Gen AI solutions.

With innovations like ALHF, TAO, and MixAttention incorporated, Agent Bricks positions Databricks as a production-grade AI agent development platform for both developers and business users.

## 02. DATABRICKS AI/BI – ENHANCEMENTS

### SUMMARY:

Databricks has introduced powerful enhancements to its AI/BI stack, delivering a more intelligent and interactive analytics experience. Key updates include cross-filtering, drilldowns, calculated dimensions, Sankey diagrams, and multi-page reporting. Users can now schedule and subscribe to dashboards, apply theming, and embed visualizations with ease. Advanced capabilities like Integrated Forecasting, Top Driver Analysis, and Deep Research (Preview) allow business users to answer predictive, diagnostic, and strategic questions directly from the platform.

Additionally, **Genie** introduces real-time knowledge extraction by learning semantics from user interactions. When faced with new questions, Genie clarifies intent, extracts insights, and promotes validated semantics to **Unity Catalog** as governed metrics—ensuring consistency across tools and applications.

### OBSERVATION & ENTERPRISE BENEFITS:

Databricks is redefining business intelligence by enabling BI systems to efficiently answer futuristic and strategic questions like “What will sales look like next month?”, “Why did sales drop?”, and “How can we hit our Q3 numbers?” The key innovation lies in the integration of **line forecasting, top driver analysis, and Genie Deep Research**, bringing predictive and prescriptive power directly into the BI layer.

## 03. NATIVE GEMINI

### SUMMARY:

Databricks announced a strategic partnership with Google Cloud to integrate Gemini models natively into the Databricks Data Intelligence Platform. This enables organizations to build and scale domain-specific AI agents using Gemini’s advanced reasoning and language capabilities—directly on their enterprise data—without requiring data movement or external integrations. Gemini models, including Gemini 2.5 Pro and

Flash, will be accessible via SQL and model endpoints within Databricks and will be billed through existing Databricks contracts.

#### OBSERVATION & ENTERPRISE BENEFITS:

Databricks customers gain access to advanced features like *Deep Think mode* for multi-step reasoning, *enhanced coding capabilities*, and *native multimodality*—all within a Databricks platform, leveraging Google Models. The integration supports up to a *1 million+ token context window*, making it ideal for handling complex enterprise tasks and large documents with superior context retention.

## 04. DATABRICKS LAKEBASE

#### SUMMARY:

Lakebase is a next-generation operational database architecture from Databricks, built on fully managed, open-source PostgreSQL. It combines the power of traditional relational databases with the scalability and flexibility of Lakehouse.

#### KEY FEATURES:

- ◆ **Fully Managed PostgreSQL:** Lakebase integrates deeply with the Databricks Lakehouse and supports full PostgreSQL compatibility—including SQL, PostGIS, pgvector, and Postgres-native tools.
- ◆ **Separation of Compute and Storage:** Designed with decoupled compute and storage, Lakebase delivers sub-10ms latency and supports 10K+ QPS without performance degradation.
- ◆ **Instant Serverless Auto-Scaling:** Lakebase's serverless engine launches in under one second and scales to zero when idle, removing the need for manual provisioning.
- ◆ **Pay-As-You-Go Model:** Costs are based on active compute usage. When idle, only storage costs apply—ensuring cost efficiency.
- ◆ **Branching for Safe Experimentation:** Zero-copy database branching enables safe experimentation and point-in-time testing with minimal storage impact—ideal for both developers and AI agents.
- ◆ **Available on Databricks Workspaces:** Lakebase is now available in Public Preview (General Availability coming soon) and can be enabled directly within Databricks workspaces.



## OBSERVATION & ENTERPRISE BENEFITS:

Traditionally, Databricks has been viewed as a platform optimized for advanced analytics, data engineering, and machine learning. For low-latency, sub-second response times, data was often offloaded to external relational databases, requiring complex ETL pipelines, change data capture (CDC) logic, and ongoing synchronization.

With Lakehouse architecture powered by Delta Lake, Databricks now supports low-latency query performance natively. This eliminates the need for external systems and enables seamless integration with web applications and real-time reporting for end-user consumption.

## 05. DATABRICKS LAKEFLOW

### SUMMARY:

Databricks introduces Lakeflow and Zerobus for Unified Data Engineering. Databricks has launched Lakeflow, a unified data engineering framework, and Zerobus, a high-throughput direct-write API—both natively integrated with Unity Catalog. These capabilities streamline data ingestion, transformation, orchestration, and real-time streaming, significantly simplifying the development and management of modern Lakehouse architectures.

### KEY FEATURES:

- ◆ **Ingest via Lakeflow Connect:** Seamlessly integrates with a wide range of sources including Salesforce, Workday, Oracle, Google Analytics, S3, SharePoint, SQL Server, Snowflake, Teradata, and more.
- ◆ **Transform with Declarative Pipelines:** Production-grade, fully Spark-compatible pipelines built with an IDE tailored for data engineering workflows.
- ◆ **Orchestrate with Lakeflow Jobs:** Enables end-to-end orchestration across tools like Power BI, dbt Labs, and Snowflake, streamlining complex workflows.
- ◆ **AutoCDC:** Real-time CDC with a single line of code.
- ◆ **Zerobus – Direct Write API:**
  - High-throughput, low-latency streaming from cloud, edge, and IoT sources.
  - Eliminates the need for intermediaries like Kafka or message buses. Enables real-time ingestion with lower operational overhead and total cost of ownership.

### **OBSERVATION & ENTERPRISE BENEFITS:**

The introduction of Lakeflow and Zerobus marks a major shift in how organizations ingest and process data. Traditionally, data ingestion required third-party tools and complex integrations to connect external systems with Databricks, resulting in higher development and maintenance overhead.

With native capabilities for large-scale ingestion, declarative SQL-based transformation, orchestration, and seamless deployment—all integrated into the Databricks Lakehouse—ETL pipeline development becomes significantly more streamlined. This reduces operational complexity and lowers support and maintenance costs.

Additionally, Zerobus enables direct, high throughput streaming into Databricks, eliminating the need for intermediate layers such as Kafka or message buses. This drastically reduces infrastructure footprint and supports costs for real-time data processing at scale.

## **06. LAKEFLOW DESIGNER**

### **SUMMARY:**

Databricks now offers a no-code ETL solution that empowers non-technical users to build production-grade data pipelines using natural language through Databricks Assistant and GenAI. These pipelines are powered by Lakeflow, ensuring scalability, reliability, and performance.

### **KEY FEATURES:**

- ◆ A drag-and-drop interface for defining data transformations.
- ◆ Side-by-side views of logic and intermediate outputs for better visibility.
- ◆ Hybrid editing capabilities—switch seamlessly between visual canvas and direct SQL input.

### **OBSERVATION & ENTERPRISE BENEFITS:**

This marks a significant step toward democratizing data engineering on Databricks, especially for data analysts and business users without deep technical expertise. By enabling pipeline development through natural language and a visual interface, users can now create complex, production-ready ETL processes without writing Python or SQL.

The result is a faster time-to-value, lower development and maintenance overhead, and greater alignment between technical and non-technical teams, unlocking broader adoption of the Databricks Lakehouse across the enterprise.

## 07. UNITY CATALOG UPDATES

### a) Iceberg Managed Tables

Databricks has introduced Iceberg Managed Tables in Unity Catalog, expanding its native support for open table formats. This capability allows organizations to govern, access, and manage Apache Iceberg tables directly within Unity Catalog, enabling seamless interoperability across a broad ecosystem of engines — including Databricks, Snowflake, Amazon EMR, Trino, and more.

With this enhancement, Unity Catalog evolves from a Delta-native governance solution to a truly multi-format catalog, empowering enterprises to adopt open standards while maintaining centralized data governance, lineage, and access controls across platforms.

### b) Unity Catalog Metrics: Centralized Metric Governance for Lakehouse

A unified, centralized framework within Databricks to define, govern, and operationalize business metrics across data engineering, BI, and AI workloads.

#### KEY FEATURES:

- ◆ **Centralized Metric Definitions:** Establish a single source of truth for business metrics—ensuring consistency across teams, tools, and use cases.
- ◆ **End-to-End Governance:** Enable certification, versioning, and governance of metrics to ensure trust, accuracy, and compliance.
- ◆ **UI Based Metric Management:** Intuitively manage metrics through the UI—tag, certify, document logic, manage dimensions/measures, and assign ownership.

### c) Unity Catalog Discoverer: Curated Marketplace for Certified Data Products

Unity Catalog Discoverer is a curated, internal data marketplace that enables data consumers to easily discover, access, and trust certified data assets—including tables, dashboards, and insights—across the organization.

#### KEY FEATURES:

- ◆ **Searchable, Governed Catalog:** A centralized, searchable interface for discovering certified data products with rich metadata, ownership details, and usage context.
- ◆ **Bridging Producers and Consumers:** Streamlines collaboration by connecting data producers with consumers, promoting transparency and reusability of trusted assets.
- ◆ **Promotes Data Democratization:** Encourages self-service analytics, reduces data silos, and fosters a data-driven culture across departments.

## 08. DATABRICKS LAKEBRIDGE

### SUMMARY:

Databricks has launched “Lakebridge”, a free, open, and AI-powered migration platform designed to modernize enterprise data warehouses by accelerating and simplifying migration from legacy systems into the Databricks Lakehouse. This launch is further enhanced by the integration of BladeBridge, a specialized code conversion company recently acquired by Databricks.

### KEY FEATURES:

#### ◆ Analysis

- A deep scanner performs comprehensive scope assessment.
- Assists in evaluating migration complexity and estimating total cost of ownership

#### ◆ Code Converter

- Utilizes an LLM-powered converter to transform legacy SQL into Databricks SQL-compliant ANSI code.
- Includes a business logic verifier and A/B testing framework to ensure functional accuracy.

#### ◆ Data Migration

- Lakeflow Connect facilitates seamless ingestion of source data into the Databricks Lakehouse.

#### ◆ Data Validation

- A robust reconciler ensures correctness of both transformed logic and migrated data, maintaining data integrity.

#### ◆ Supported sources – Teradata, Oracle Exadata, SQL Server, IBM DB2, Netezza, Vertica, Cloudera, Azure Synapse, Snowflake, Amazon Redshift, Google BigQuery.

### OBSERVATION & ENTERPRISE BENEFITS:

Traditionally, migrating to Databricks from legacy data platforms has been complex, time-consuming, and costly, requiring multiple disconnected tools for analysis, data migration, transformation, and validation.

With Lakebridge, organizations gain a unified, out-of-the-box migration solution that streamlines the entire process—from scope assessment to data validation. This significantly reduces effort, cost, and risk, enabling faster, scalable, and more reliable migrations to the Databricks Lakehouse.



## OTHER ANNOUNCEMENTS:

- ◆ **MLflow 3.0:** Enables end-to-end agent observability across any deployment environment—cloud or on-prem—with added support for prompt versioning through a built-in prompt registry.
- ◆ **Multi-Modal AI Functions in SQL:** `ai_parse_document` makes it effortless to extract structured information from complex documents.
- ◆ **Storage-optimized endpoints for Mosaic AI Vector Search:** Enable petabyte-scale retrieval with multi-billion vector capacity, lower cost, faster indexing, and SQL-style filtering—powered by decoupled storage and Spark-native parallelism.
- ◆ **Serverless GPU Compute (Beta):** Making high-performance AI workloads more accessible with fully managed, on-demand A10g, H100 GPUs. Integrated with Unity Catalog, it supports training, inference, and data transformation directly from notebooks or jobs, all with built-in governance.
- ◆ **Apache Spark 4.0:** Apache Spark 4.0 introduces a major evolution of the Spark engine, focused on real-time processing, simplified data engineering, and enhanced developer experience. It includes innovations in SQL, streaming, and data pipeline orchestration, and is now available as an open-source release with tighter integration into the Databricks Lakehouse ecosystem.
- ◆ **Model Context Protocol (MCP) on Databricks:** Provides managed MCP servers for Vector Search, Genie Spaces, and Unity Catalog functions. Additionally, it supports hosting custom or third-party MCP servers as Databricks apps, offering flexibility to extend agent capabilities across any system.
- ◆ **Databricks One:** A new product experience designed for business users, offering a simple, intuitive interface to interact with data and AI without needing technical expertise. It enables users to explore AI/BI dashboards, ask questions in natural language, use custom apps, and access domain-specific content, all in a clean, streamlined layout.
- ◆ **Next Gen SQL Warehouse:** Databricks unveiled its Next-Gen Data Warehouse, delivering 25% performance improvement at no extra cost. This upgrade builds on the core lakehouse architecture and introduces deeper AI integration, better interoperability, and further cost-efficiency—positioning Databricks SQL Warehouse as a leading choice for enterprise-grade analytics.

## CONCLUSION:

DAIS 2025 marks a transformative leap for Databricks as it boldly redefines the future of enterprise AI and data with a clear focus on **low-code/no-code innovation** and **human-in-the-loop intelligence**. By empowering domain experts and business users—not just data scientists, data engineers, and analysts—to directly build, manage, and optimize data, BI, and AI assets, Databricks is breaking down the traditional barriers to technology adoption.

Reach out to us for a complimentary workshop to explore how these latest announcements can accelerate your Data & AI strategy.

**Schedule a 1:1 with our experts to dive deeper.**

Follow us at [in](#) [X](#) [YouTube](#) [f](#)

