TREDENCE
Beyond Possible

databricks
DATA+AI SUMMIT

Databricks Data + AI Summit 2026

Key Announcements & Strategic Insights

As a leading partner of Databricks, Tredence is at the forefront of redefining enterprise decision-making. Our partnership is built on a shared vision: to enable businesses to move from fragmented insights to autonomous, AI-powered decisions at scale. Together, we are building the future of Decision Intelligence—combining cloud-native innovation, real-time AI, and responsible adoption frameworks.

At Data & AI Summit 2026, Tredence was honored with the **Business Transformation Award,** which showcased Tredence's capabilities in transforming businesses and enabling last-mile delivery to solve business problems with deep expertise in business, data, and AI.

Our Four-Pronged Agentic AI Strategy

- **AI-Native Data Foundations**
  Unified, scalable platforms built on Databricks that accelerate the data-to-AI journey by integrating structured and unstructured data for analytics and AI readiness.
- **Autonomous Agentic AI**
  Autonomous decision-making powered by real-time agents built using Databricks Agents and Models of Choice including the foundation models—enabling continuous optimization of operations and workflows.
- **GenAI-Powered Decision Intelligence**
  From dashboards to copilots—Tredence integrates GenAI to drive contextual, proactive, and action-oriented intelligence across the enterprise. As a launch partner for Agent Bricks and Lakebase, Tredence is actively leveraging various out-of-the-box Databricks tooling to deliver insights, recommendations, and real-time actions on business data.
- **Responsible & Scalable AI**
  Security-first frameworks ensure ethical, compliant, and observable AI adoption—building trust as AI scales across critical business functions.

Databricks Data + AI Summit (DAIS) 2026 brought together thousands of industry leaders, technologists, and innovators to explore the future of AI and data. The summit centered on four core product areas — **Agentic Data, Unified Governance, Agentic Work, and Agentic Apps** — with announcements that collectively signal a fundamental shift in enterprise data and AI architecture with a focus on context, control, choice, and cost. This document summarizes the key platform updates unveiled at the event and offers **strategic observations** for enterprise leaders accelerating their AI adoption journeys.

databricks | Data + AI PLATFORM | Context. Control. Choice.

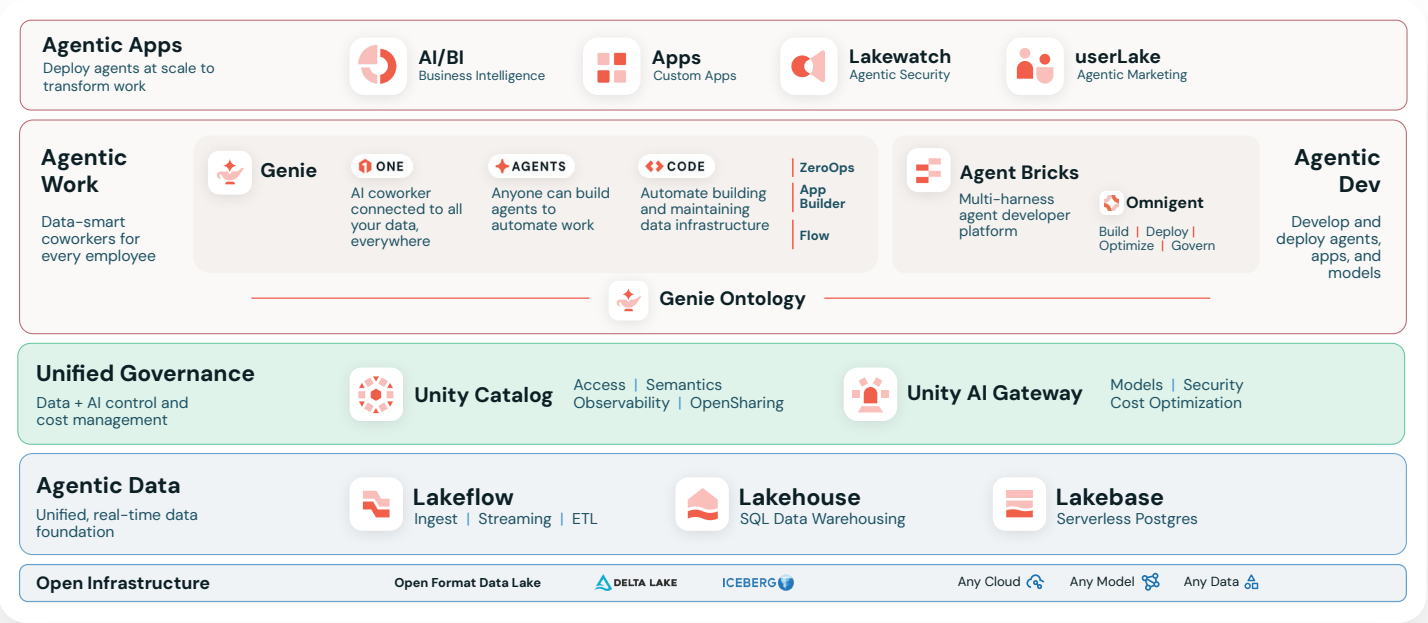

Section 1 – Agentic Data

1. Lake Transactional/Analytical Processing (LTAP)

Summary

Databricks introduced LTAP — a new data processing architecture that unifies transactions, analytics, streaming, and operational data on a **single copy** of lake storage, eliminating ETL pipelines and replicas by design. Databricks claims to be the world's first LTAP platform, combining **Lakebase** — its serverless Postgres engine built on open object storage — with the **Lakehouse** under a single **governance** model and **storage** layer. LTAP resolves the problem at the storage layer itself — enabling **transactional** and **analytical** workloads to scale independently while remaining built on **open standards** like Iceberg, Delta, and Postgres.

Observation

LTAP represents a genuine architectural inflection point for enterprise data infrastructure. For four decades, transactional and analytical workloads have lived in separate systems bridged by brittle CDC pipelines — a trade-off that becomes untenable as AI-assisted development produces roughly 50x more applications, many requiring near real-time data access. **By collapsing this divide at the storage layer rather than the engine layer, Databricks repositions the Lakehouse not merely as an analytics platform but as the operational backbone for AI agents.** The open format foundation — Delta, Iceberg, and Postgres — also meaningfully **reduces lock-in** concerns that have historically slowed enterprise adoption of proprietary platforms.

2. Open Format

Summary

Open data formats were a central theme across multiple DAIS 2026 announcements. The LTAP architecture stores all data natively in Delta and Iceberg without transformation or degradation. Managed Iceberg support in the Databricks Runtime (DBR) is currently in development. During the summit, Databricks announced **Apache Iceberg v3** entering Public Preview, highlighting that Iceberg v3 introduces **Row Lineage, Deletion Vectors, and VARIANT** — ending the performance-vs.-interoperability trade-off, since these gains are part of the open specification and do **not** sacrifice cross-engine compatibility. According to Databricks, the underlying format should be of **least concern** to the business and both Delta and Iceberg should be **interoperable** in the Databricks platform.

Observation

These announcements collectively signal that the open data format landscape has reached a decisive inflection point. The decade-long **conflict** over storage format supremacy has effectively been **resolved** — broad adoption of Iceberg and Delta, bridged by **Databricks' UniForm standard**, has commoditized the storage layer and shifted the strategic battleground toward governance. The addition of multimodal and geospatial support extends open formats far beyond structured tabular data, a move that could dramatically **accelerate** cross-organizational AI collaboration.

3. Agentic Data & Real-Time Foundations

Products: Lakehouse//RT , Lakebase - Real-Time Analytics on Operational Data, Lakeflow Connect

Summary

DAIS 2026 marked a significant leap in real-time data capabilities across multiple products. Lakehouse//RT, powered by the **Reyden engine**, introduces a new warehouse option delivering **millisecond query performance** at high concurrency directly on Delta/Iceberg in Unity Catalog. Lakebase Search was announced as the first agent-native search database, embedding full-text, semantic, and hybrid search natively into Postgres. Real-time analytics on operational data now **enables** complex OLAP queries to run on live OLTP data without degrading application performance. To complete the ecosystem, **Lakeflow Connect** expanded to over **100 native**, fully-managed

Observation

These capabilities directly address one of the most critical hurdles in enterprise AI: agents acting on stale data. By eliminating the need for fragmented real-time serving layers — such as ClickHouse or legacy data warehouses — and removing brittle CDC pipelines, enterprises can **substantially lower total cost of ownership.** For organizations building responsive AI agents, these innovations ensure that agents operate on the **freshest** available business state rather than lagged pipeline data, providing a strong foundation for infrastructure consolidation.

Section 2 – Unified Governance

1. Unity AI Gateway

Summary

Databricks announced Unity AI Gateway as a unified governance layer across all AI assets — both on Databricks and externally hosted — **covering models, agents, MCP services, skills, and enterprise tools.** Built on Unity Catalog, it **extends** governance beyond data and static assets to the runtime interactions between AI components, delivering centralized security controls, cost management, and agent monitoring across all AI providers and custom systems. Six capability areas were announced:

- **Agent Registry :** Enables **registering agents** created within and outside the Databricks platform to serve as a repository of agents developed throughout the enterprise
- **Access Control : Governs** agents with **access control** tied to the enterprise identity and management solution
- **Contextual Policy :** Policies that enable governing agents and serve as **agent guardrails** at the enterprise level
- **Budgets :** Enables budget allocation at the agent level for **cost observability** and governance
- **Smart Routing :** Based on the complexity of each request, smart routing **intelligently routes** queries to smaller or larger models, enabling significant token and **cost savings** for the enterprise
- **Agent Tracing :** End to end **observability** and **traceability** of agent interactions and actions that enable better governance and control

Observation

One of the biggest challenges enterprises face today is building end-to-end governance for agents being built at scale. Unity AI Gateway may be the most strategically consequential announcement of the entire summit. It enables enterprises to **address the most complex challenge in agent governance.** Databricks is making a calculated bet that the biggest blocker to enterprise AI at scale is not model quality or cost — it is governance fragmentation. What sets it apart from prior governance tools is its scope: it is not a policy layer bolted onto a single model provider, but a **runtime control plane** that spans every model, agent, MCP service, and skill an enterprise runs, regardless of where they are hosted.

2. Other Governance Capabilities offered by Databricks for data and AI

Products: Global UC, Governance Hub, UC Glossaries, OpenSharing

Summary

Databricks significantly expanded Unity Catalog to address the governance complexities introduced by AI agents and global enterprise footprints.

Global UC extends governance across metastore, region, and cloud under a single namespace without requiring data replication.

The **Governance Hub** provides a centralized **command center** for monitoring security posture, risk signals, and cost.

UC Glossaries delivers governed, machine-readable business context — a **semantic layer** for AI and BI tools to draw upon. Together with Unity AI Gateway, these capabilities form a comprehensive governance stack for the agentic era.

OpenSharing enables sharing AI assets, including models, agents, and skills across clouds without vendor lock-in.

Observation

Next-generation governance capabilities allow organizations to build a **robust governance** solution across data and AI, spanning multiple regions, and establish a context layer that gives BI tools the enterprise context needed to deliver accurate, relevant responses to user queries.

Section 3 – Agentic Work

1. Omnigent

Summary

Databricks announced Omnigent as an **open-source meta-harness** for building and running agents across models, frameworks, and coding tools — including Claude Code, OpenAI Agents SDK, and Codex. Its defining design principle is the enforcement of governance policies at the orchestration layer, outside the model entirely: each agent runtime is wrapped in an isolated sandbox with a uniform API, and access controls, context sharing across developers, network boundaries, cost ceilings, and human-in-the-loop approval requirements are enforced before any action reaches the underlying LLM. Databricks also announced a **managed version of Omnigent** within Agent Bricks, enabling teams to orchestrate multiple agent harnesses with horizontal autoscaling to Databricks Apps and full Unity AI Gateway governance.

Observation

Omnigent addresses one of the most pressing unsolved problems in enterprise agentic AI: governance that actually holds. Prompt-based controls break because when rules live inside the prompt, the agent can reason around them. With Omnigent, cost ceilings and human-in-the-loop approvals are enforced before any action reaches the underlying LLM — unlike prompt-based guardrails, which model can circumvent. **By moving controls to the orchestration layer, Omnigent represents a structurally sounder approach to agent safety** than anything widely deployed to date. The Apache 2.0 open-source release follows the same playbook that made Apache Spark, Delta Lake, and MLflow category-defining projects — establishing Omnigent as the community standard that enables **ease of adoption** across enterprises.

2. Agentic Work, Dev & Platforms

Products: AI Runtime, Sandbox for Agent Platform, Memory for Agent Platform

Summary

To operationalize AI development and applications, Databricks announced a suite of platform capabilities alongside Omnigent. The **Sandbox** for Agent Platform provides a serverless, isolated Python compute environment for agents to **safely** execute code with downscoped permissions. **Memory** for Agent Platform offers persistent, per-user memory governed by Unity Catalog. **AI Runtime** provides serverless, on-demand access to A10 and H100 GPUs for deep learning and agent fine-tuning directly on Lakehouse data — removing the operational burden of managing custom GPU infrastructure.

Observation

These releases target the **last mile of AI deployment.** Most enterprise AI agents today remain trapped in isolated, ungoverned development environments — functional in demos but not suitable for production. The Dev platform and agentic capabilities enable teams to **transform** agents from non-governed software into governed, enterprise-grade assets that can execute autonomously **without** compromising production data. AI Runtime abstracts the significant operational overhead of GPU cluster management, allowing teams to focus on building applications and fine-tuning models rather than maintaining infrastructure.

3. Genie Ontology

Summary

At the core of the Genie product family is Genie Ontology — **a continuously self-improving context layer** that builds a comprehensive knowledge web of an organization from its own data: documents, tags, content, applications, and people. It is designed to solve one of enterprise AI's hardest problems: fully understanding a business from its own signals. The Genie Ontology has about **4.5 million ontology snippets** and enables answering enterprise questions through the Ontorank algorithm similar to Google's PageRank algorithm for the web, built by the Databricks R&D team.

Genie Ontology automatically extracts and **continuously updates** business knowledge from Databricks and connected workplace tools, including files, tickets, chats, and meetings.

Observation

Genie Ontology represents a meaningful shift in how enterprise AI systems acquire and retain institutional context. By **learning** through organizational data and incorporating human-annotated business semantics, the ontology enables agents to answer fundamental and **contextualized** business questions such as when the fiscal year begins, who leads a function, what a 'churned customer' means in a specific context. The Genie ontology is transparent to the enterprise, **learns on its own,** and enables contextualized responses to user queries.

This directly addresses a chronic failure mode in enterprise AI: generic agents that repeatedly underperform because they lack the institutional knowledge required for grounded, accurate responses.

4. Genie & AI Coworkers

Products: Genie One, Genie Agents, Genie App Builder, Genie ZeroOps

Summary

Databricks evolved the Genie suite into a comprehensive AI workforce platform.

Genie One delivers contextual intelligence directly into business **workflows** via Slack, Teams, and mobile applications through the Genie Ontology.

Genie Agents allow users to convert successful chat interactions into scoped, reusable AI **coworkers** without writing code.

Genie App Builder introduces native vibe-coding to **deploy** secure, scale-to-zero internal applications in plain English.

Genie ZeroOps places data and AI operations on autopilot — continuously monitoring pipelines and **automatically** diagnosing failures to free data engineers from routine maintenance.

Observation

This suite fundamentally **democratizes AI** for both business and data teams. Genie Ontology and Genie Agents allow subject matter experts to encode and share domain knowledge in a **governed, reusable** form — turning institutional expertise into a scalable enterprise asset. Genie App Builder lowers the barrier for non-technical users to **deploy** production-grade applications in plain language. For data engineers, Genie ZeroOps is a genuine productivity shift: it transitions time away from reactive pipeline maintenance toward **higher-value** architectural work.

Section 4 – Agentic Apps

1. Lakewatch

Summary

Databricks debuted Lakewatch as a new open, agentic **SIEM** (Security Information and Event Management) designed to defend against increasingly sophisticated agent-based threats. Lakewatch **unifies** security, IT, and business data into a single governed environment for AI-driven threat detection and response. At DAIS 2026, its integration with the broader platform was deepened: unified agent tracing in Unity AI Gateway captures model and MCP activity in one governed telemetry layer, with **traces directly analyzed in Lakewatch.** It also serves as the investigation layer for AI security incidents — enabling teams to detect suspicious activity, investigate policy violations, and accelerate AI security investigations from a single surface.

Observation

Lakewatch signals Databricks' entry into the enterprise security market. Traditional SIEMs were built for human-generated log data, not the high-velocity, multimodal telemetry that agentic AI systems produce at scale. By building **Lakewatch natively on the Lakehouse** using open formats, Databricks positions it as a fundamentally different class of SIEM — one that can scale to **petabyte-level** threat detection **without proprietary data lock-in.** The claimed **TCO reduction of up to 80%** compared to legacy SIEM solutions is a compelling proposition for enterprises already running workloads on Databricks, and the deep integration with Unity AI Gateway means Lakewatch functions as an **active governance layer** rather than a purely reactive forensics tool.

2. CustomerLake (userLake)

Summary

Databricks announced CustomerLake — an agentic CDP that equips marketing and data teams with AI agents that **continuously analyze** customer behavior, make decisions, and act across channels. The product launched in Private Preview with early customers including HP, Circle K, AB InBev, and Getnet by Santander. Central to the announcement is the concept of "infinity campaigns" — **agent-driven engagement loops** that continuously analyze customer signals and act across channels, replacing the static, one-off campaign model.

Observation

CustomerLake represents Databricks' most direct challenge yet to the established martech industry, taking aim at standalone CDP solutions by arguing the CDP should not exist as a separate system. Because CustomerLake runs inside an organization's governed storage boundary, **customer data, AI models, and governance remain co-located — eliminating** the data movement and duplication inherent in maintaining a separate CDP stack. The broader strategic implication is significant: as consumers increasingly deploy agents to research, evaluate, and transact on their behalf, traditional martech is not designed for agent-to-agent interaction. The new capability enables enterprises to build their **martech stack natively on Databricks,** where data lives and agents are governed through Databricks' own solutions.

Conclusion

The announcements at Databricks DAIS 2026 represent more than incremental product updates — they reflect a **coherent and ambitious architectural vision for the enterprise AI era.** Across every major announcement, three themes are consistent: **open standards** that reduce lock-in and accelerate adoption; **governance** embedded at the infrastructure layer rather than bolted on as an afterthought; and a deliberate **blurring** of the boundary between data platform and operational AI backbone. For enterprise leaders, the strategic implication is clear: the organizations best positioned to scale agentic AI are those that have built governed, real-time, open data foundations. Databricks is betting heavily that the Lakehouse becomes that foundation — and DAIS 2026 made that bet more credible than ever.

Yet the distance between a compelling platform vision and realized business value is where most enterprises struggle — and where Tredence plays a decisive role. As a leading Databricks partner, **Tredence bridges that last mile:** accelerating data modernization by leveraging Databricks capabilities. On the agentic development side, Tredence can help engineering teams operationalize Omnigent and Agent Bricks with **enterprise-grade guardrails** enabled through Databricks Unity AI Gateway, translating Databricks' orchestration architecture into governed, auditable agent workflows. For business functions, Tredence brings **proven accelerators** across industries to stand up CustomerLake-powered marketing and Lakewatch-driven security operations — moving organizations from **pilot to production at speed.** Tredence's value proposition lies in building accelerators that **augment** Databricks' native capabilities to enable faster time-to-insight, lower total cost of ownership, and AI systems that earn the trust of the enterprise.

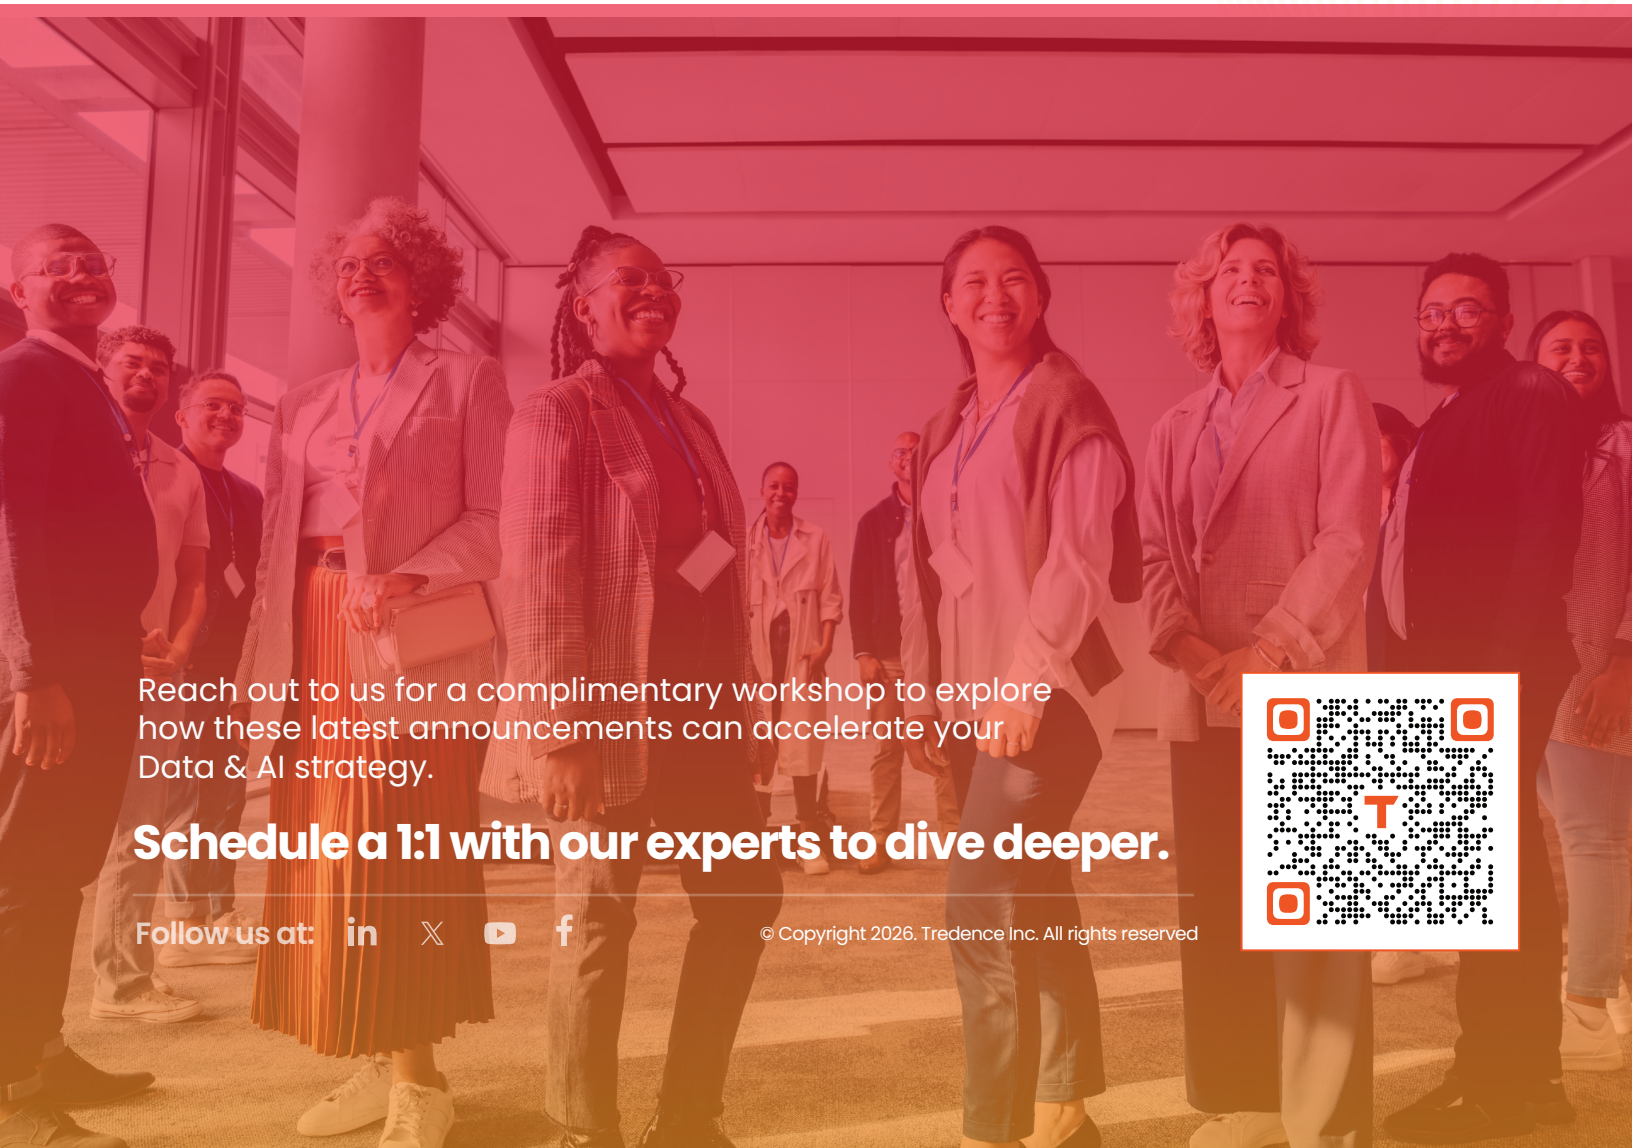

Reach out to us for a complimentary workshop to explore how these latest announcements can accelerate your Data & AI strategy.

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

Follow us at: