White paper

Teradata Autonomous Knowledge Platform

From Insight to Autonomous Outcomes at Enterprise Scale

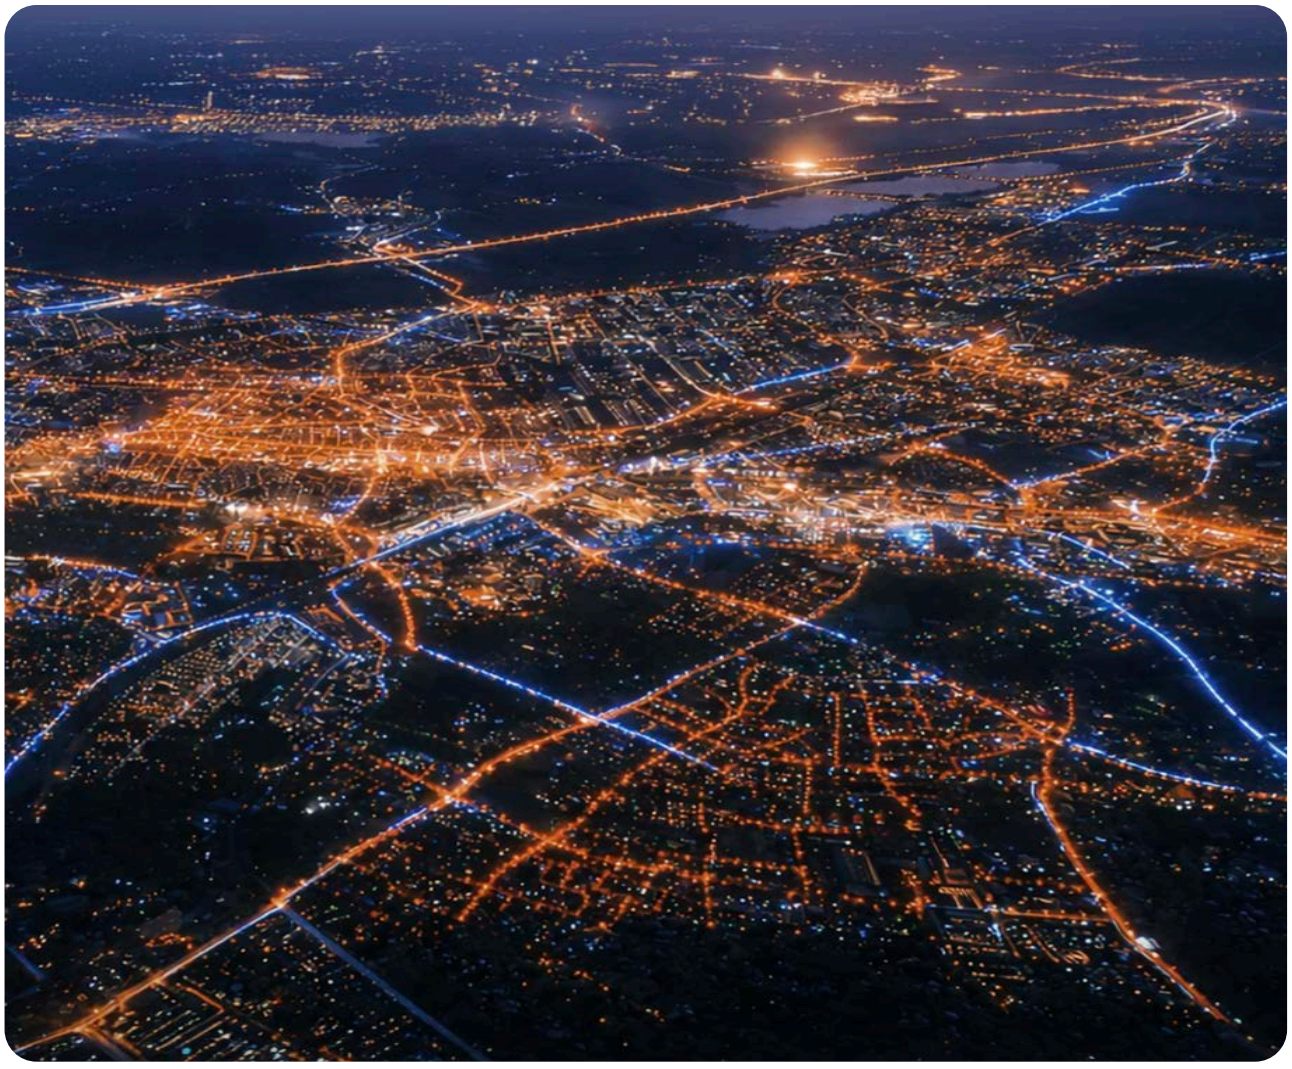

teradata.

Table of Contents

Executive summary

Enterprises are entering a new phase of the intelligence era: AI agents are moving from isolated pilots into continuous, business-critical operations. Unlike human-driven analytics, agentic systems operate 24/7, generate sustained high-concurrency workloads, and require reliable access to trusted enterprise context. This shift is driving a new requirement for platforms that can convert enterprise data into governed, reusable knowledge and then execute continuously—at scale—without constant human tuning.

The Teradata Autonomous Knowledge Platform is designed for this reality. It unifies data, analytics, AI development, agentic execution, governance, and autonomous optimization into a single integrated platform—available as Teradata Cloud, a managed cloud deployment; as Teradata Factory, a fully integrated, on-premises deployment; or as a true hybrid model.

This white paper introduces the autonomous knowledge category, explains why agentic workloads reshape platform requirements, and details the platform components—including AI Studio as the unified interface; active and elastic compute (including elastic compute operating on open table formats); a connected data foundation; Enterprise Vector Store and multi-modal knowledge; AgentStack; and AI Services.

Before exploring the architecture, deployment models, and capabilities that make this possible, it's important to clarify what Teradata means by autonomous knowledge—and how this concept differs from traditional analytics platforms and first-generation AI systems. This foundation provides the context for understanding why agentic workloads require a fundamentally different platform approach.

A practical lens for buyers

The way organizations evaluate and buy technology has changed. Rather than assessing tools in isolation, leading teams now focus on the measurable impact those tools deliver—and they want to see proof quickly, not at the end of a long pilot.

These decisions also rarely sit with IT alone. When a platform touches how work gets done across the organization, stakeholders across functions need to be part of the conversation from the start.

That's why it matters to begin with the outcome you're trying to achieve—whether that's accelerating a key process, improving decision quality, or reducing operational costs—and then evaluate whether the platform can reliably deliver it. The questions worth asking:

- How does this connect to what we're already running?
- How does it perform at our scale?
- What does it actually cost to get there?

The enterprise intelligence inflection point

From dashboards to continuously operating systems

The last two decades of enterprise analytics were primarily human-driven: business users and analysts queried data, reviewed dashboards, and acted through separate operational systems. The next decade will be increasingly machine-driven: agents will continuously monitor signals, generate hypotheses, run high volumes of queries, and take actions across systems.

Agents don't sleep: Why platform requirements change

Agentic workloads differ from human workloads in predictable ways: they're always on, API first, and context reasoning by default, and they produce sustained concurrency across mixed workload types. As agents evolve from basic Q&A to outcome-based execution, the workload impact grows sharply—driving higher query volumes, broader data access patterns, and stronger requirements for cost control and governance.

Why most agents fail in production

Many early agent deployments fail or stall when scaled because they lack one or more of the following: trusted context, a governed knowledge foundation, lifecycle controls, and a platform that can sustain always-on economics and reliability. A platform designed for autonomous knowledge must address these failures systematically—by grounding agents in enterprise context, enforcing trust, and optimizing execution autonomously.

What changes when agents become the primary data consumers

- Always-on operation (24/7) replaces episodic usage (8/5)
- Concurrency becomes a baseline requirement—not an edge case
- Mixed workloads (SQL + vector + ML + multi-modal) run side by side
- Context is required for every decision (semantics, lineage, policy, KPIs)
- Governance must be embedded, because agents can act—not just answer

The Teradata Autonomous Knowledge Platform

What the platform is

The platform is a single, integrated system designed to activate autonomous enterprise intelligence. It unifies data, analytics, AI workflows, governance, and execution so organizations can move beyond isolated insights to continuously operating outcomes.

Core architectural components

At a high level, the platform is organized around a unified access layer, AI Studio as the primary interface, a knowledge and context layer, a compute layer with always-on and on-demand options, a connected data foundation, and an execution layer where autonomous agents optimize operations. The sections that follow detail each component and how they work together.

Defining autonomous knowledge

What autonomous knowledge means

Autonomous knowledge is the ability of an enterprise platform to transform raw structured, unstructured, and multi-modal data into trusted, governed understanding. It captures meaning (semantics), relationships, lineage, and constraints so that knowledge can be reused safely across people, applications, and agents.

Autonomy at two levels: System and knowledge

In the Autonomous Knowledge Platform, autonomy operates across two complementary layers. System-level autonomy manages execution—predicting workload behavior, placing and sizing compute, and optimizing performance, cost, and data locality. Knowledge-level autonomy enables continuous creation and evolution of enterprise knowledge—leveraging agents to curate and refine data products and context over time.

Autonomous by design

The Teradata Autonomous Knowledge Platform unifies system-level autonomy with enterprise-grade knowledge so AI agents can operate continuously, make trusted decisions, and turn insight into governed action at scale—with minimal manual intervention.

AI Studio:
The unified platform interface

The interface for every platform and every persona

AI Studio is the place where creators build, activate, and manage AI outcomes at scale. It provides a single, consistent interface across cloud, on-premises, and hybrid deployments—so teams do not need separate toolchains for different environments.

What AI Studio provides

AI Studio brings together the capabilities required to create AI outcomes, including agent development, model operations, AI marketplaces, analytics and machine learning (ML), Python compute, and accelerated compute. It's designed to unify exploration, development, deployment, and governance so teams can move from prototypes to production with repeatability.

These capabilities are deliberately sequenced to build on one another, forming a complete enterprise AI platform that delivers tangible customer value at every layer. The foundation is enhanced Python compute, providing an industry-standard Jupyter Notebook experience with native Teradata machine learning connectivity. This allows data engineers and data scientists to work directly with governed Teradata data using Python—without exporting data to external platforms that may not scale, perform, or meet enterprise security and privacy requirements. Building on this foundation, conversational AI democratizes access to data and insights through a natural language interface, enabling self-service analytics for nontechnical users while ensuring the accuracy, repeatability, and governance required for enterprise grade business intelligence (BI).

Through Teradata AgentStack, agent capabilities introduce autonomy through procode Python orchestration, leveraging frameworks such as LangChain and LlamaIndex alongside ModelOps to industrialize the full AI lifecycle, including deployment, monitoring, and retraining. Partner integrations such as WisdomAI and Karini further enrich AI Studio—powering agentic BI experiences and accelerating intelligent automation while remaining fully grounded in Teradata's governed data foundation.

Finally, the vector store connects unstructured knowledge with structured enterprise data, enabling retrieval augmented generation and agentic workflows at scale. Taken together, these capabilities position AI Studio as a unified, governed AI platform—rather than a collection of disconnected tools.

Example workflows enabled by AI Studio

A data scientist prototypes a retrieval-augmented agent, connects it to governed enterprise knowledge, tests it against a repeatable evaluation set, then deploys and monitors it—using one workspace.

An analyst uses natural language to explore governed data products and publish an outcome to an operational workflow, while platform teams preserve policy enforcement and lineage.

A unified workspace

AI Studio brings together:

- Humans
- Apps and agents
- AI marketplace
- ModelOps
- AgentStack
- Generative AI and model services
- Analytics and ML
- Python compute (scale-up and scale-out)
- CPU/GPU accelerated compute

Active and elastic compute (including open table formats)

Active compute for always-on workloads

Always-on compute is designed for mission-critical, latency-sensitive workloads and continuously operating agents. It's optimized for high concurrency and predictable execution characteristics for workloads that run around the clock.

Elastic compute for innovation and burst workloads

Elastic compute provides on-demand scalability for bursty, exploratory, and compute-intensive tasks, such as feature engineering, model training, and large-scale background processing. It allows customers to spin up dedicated engines for specific tasks and shut them down when complete to improve cost efficiency.

Elastic compute with open table formats (OTFs)

Elastic compute is particularly powerful when paired with OTFs (like Iceberg or Delta) stored in object storage. This enables teams to operate directly on open-format data for large-scale processing and innovation workloads without duplicating data into separate systems.

Example: An organization runs continuous fraud scoring on always-on compute while using elastic compute over Iceberg tables to run periodic large-scale feature engineering jobs during peak cycles.

Workload placement guidance

Use always-on compute for:

- Production analytics
- SLAs
- Continuously operating agents
- Latency-sensitive workloads

Use elastic compute for:

- Exploration
- Feature engineering
- Large-scale transformation
- Background training
- Ad-hoc workloads

It's one platform with two compute modes —no duplicated data or governance gaps.

Connected data foundation

What it is
The connected data foundation is the platform's unified data layer that connects data, vectors, semantics, and metadata across environments. It enables consistent access to data stored on block storage and object storage, including support for OTFs, and in Enterprise Vector Store.

Why it matters
A connected foundation matters because agents require complete enterprise context. If structured, unstructured, and open-format data live in separate systems, organizations pay a tax in data movement, duplicated governance, and operational inconsistency. The connected data foundation is designed to reduce that fragmentation while preserving customer ownership of data and metadata.

To ensure access to both structured and unstructured data in a single system, the connected foundation includes Enterprise Vector Store with vector capabilities designed to operate within the same enterprise governance environment as core data. This supports retrieval augmented generation (RAG), hybrid retrieval approaches, and agent grounding in enterprise truth. Plus, through Teradata and Pinecone's partnership, users can enable a variety of low-latency use cases for real-time results.

Example: One query surface across data types
The connected data foundation, along with the integration of Unstructured, is designed to ingest and reason over documents, images, audio, and other content types so agents can operate on the full enterprise information estate. With this knowledge available, a team can run analytics over mission-critical structured data while also accessing OTF data in object storage and unstructured content represented through vectors—without shifting between multiple governance models.

Open by design. Trusted in production.
Organizations increasingly adopt open table formats to avoid lock-in. The connected data foundation pairs openness with enterprise-grade governance (identity, access control, lineage, and policy enforcement) so open data remains usable for production AI and analytics.

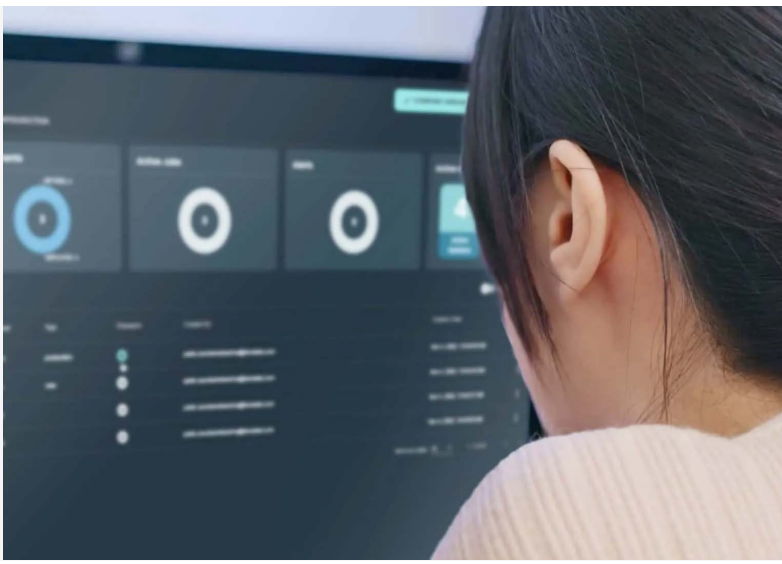

Autonomous execution and platform intelligence

Autonomous execution agents

At platform scale, manual tuning becomes a bottleneck. The platform is designed to apply autonomy to execution: observing workload behavior, optimizing placement and sizing, and continuously balancing performance, cost, governance, and data locality.

Why autonomy matters

Autonomy enables continuous operation. It reduces operational overhead and allows teams to focus on outcomes rather than infrastructure decisions—especially as agent workloads grow in volume and complexity.

Deployment models: Cloud, factory, and hybrid

Teradata Cloud

The managed cloud deployment combines always-on active compute with on-demand elastic compute in a single governed system. It's designed to support continuous, production-grade analytics and agentic workloads while enabling experimentation and innovation on demand.

Teradata Factory

The on-premises deployment of the platform is delivered as an integrated system. It's designed for predictable economics; data sovereignty; private AI requirements; and high-density, always-on workloads. It includes integrated infrastructure with graphics processing unit (GPU) acceleration.

Although the software stacks are the same for the cloud and factory deployments, elastic compute is not yet supported with the factory deployment, so compute variations differ.

True hybrid operation

The cloud and factory deployments can be used together as a hybrid model. Organizations can keep regulated, mission-critical workloads on premises while using the cloud for global analytics, experimentation, or burst processing—without changing platform architecture or governance.

Deployment decision guide

- Choose Teradata Cloud when you need managed operations, elastic innovation, and global scale.
- Choose Teradata Factory when you need predictable economics, sovereignty, private AI, or on-premises modernization.
- Choose hybrid when you need both—one platform operating model across environments.

AI Services: Accelerating business outcomes

Why services matter in the agentic era

Many organizations struggle to move from pilots to production. AI Services help customers design, build, deploy, and scale AI and agentic outcomes using platform-native capabilities and proven delivery patterns.

How AI Services complement the platform

AI Services are designed to accelerate time to value and reduce operational and governance risk. They don't replace platform autonomy, but help customers operationalize it through strong solution design, implementation rigor, and adoption enablement.

Possible outcomes

AI Services can improve data quality and large-scale text intelligence for public-sector transformation, accelerate insight generation for retail promotions, enable computer vision for operational effectiveness, and deploy agent-driven decision support and onboarding optimization.

Partner ecosystem: Extending without fragmentation

Partner-first ecosystem

The platform is designed to work within connected ecosystems. Partnerships help organizations bring flexible choice in model providers, ingestion and processing technologies, and infrastructure acceleration while preserving enterprise governance.

Partner roles and examples

Partners can support a range of capabilities—from processing unstructured documents and integrating agentic frameworks to accelerating GPU-based training and inference. The platform is built to flex across hybrid and multi-cloud environments, fitting naturally within the major cloud ecosystems.

Partner positioning

Partners work seamlessly within the platform so organizations never have to manage disconnected tools or vendors. This provides one consistent experience with consistent trust and governance and a partner ecosystem delivering expanded capabilities and faster results.
Trust, governance, and enterprise readiness

Trust, governance, and enterprise readiness

Trust built in—not bolted on

Autonomous intelligence is only as valuable as its trustworthiness. The platform is designed to apply security, access control, lineage, and policy enforcement consistently across data, AI workflows, and agentic execution.

Operating autonomous intelligence safely

As agents scale, auditability and compliance become central requirements. A governed foundation ensures decisions and actions remain explainable, traceable, and aligned with policy constraints.

Why Teradata

Teradata's Autonomous Knowledge Platform activates enterprise intelligence by design. It unifies the workspace layer, AI layer, agent layer, knowledge layer, and foundation layer into a single platform—built for always-on execution, strong economics, and hybrid flexibility. This isn't simply about running AI models closer to data. It's about manufacturing trusted, decision-ready knowledge that autonomous agents can use, learn from, and continuously improve.

Most enterprises already possess vast amounts of data and analytics capability yet struggle to operationalize AI at scale. The gap isn't access to data—it's the absence of a repeatable, governed way to transform data into knowledge that agents can autonomously discover, understand, and act upon. Teradata addresses this directly by treating knowledge as a designed and manufactured asset and not an emergent byproduct of disconnected tools.

In an AI-native enterprise, knowledge doesn't arise spontaneously from queries or dashboards. It's produced through an intentional data model design that encodes meaning, relationships, memory, and learning directly into the platform. Teradata enables this through AI-native data products—a standardized, modular approach for turning enterprise data into knowledge that is usable by agents without human mediation.

AI-native data products bring together six tightly coordinated capabilities—each one necessary to produce knowledge at scale:

- Domain models core business entities as authoritative sources of truth, with temporal consistency and reuse
- Semantic makes meaning and relationships machine discoverable, enabling autonomous navigation and correct joins
- Memory captures architectural decisions, prior agent actions, and learned strategies so intelligence persists over time
- Observability records outcomes and quality signals, closing the loop between action and learning
- Search enables semantic similarity and retrieval across enterprise data
- Prediction delivers engineered features and model outputs with point-in-time correctness for trustworthy decisions

Together, these capabilities transform static data into a living knowledge system—one that compounds in value as agents interact with it, learn from outcomes, and share that learning across use cases and domains.

Why fragmented stacks fall short

While many architectures attempt to assemble these capabilities using lake houses, vector databases, feature stores, catalogs, and external memory systems, fragmentation introduces structural limits. Semantic context drifts across systems, features lose temporal integrity, and learning loops break when outcomes cannot reliably feed back into agent memory. Knowledge becomes duplicated, brittle, or ephemeral.

Autonomous decision-making requires cross capability joins at decision time—for example, combining core business entities with semantic context, vector similarity results, engineered features, prior decisions, and outcome signals in a single execution path. This is not feasible when these elements live in separate systems with different governance models, performance characteristics, and latency profiles.

Teradata colocates all AI-native capabilities on a single, enterprise-grade platform. All six AI-native data product modules operate directly on Teradata, enabling efficient joins across domain data, semantics, vectors, features, memory, and observability—without data movement or duplication.

This colocation preserves semantic consistency, governance, and point-in-time correctness while supporting the complex access patterns required by autonomous agents. Teradata's shared-nothing architecture, mature workload management, and analytic performance make it possible to execute these AI-driven joins efficiently—even as data volumes, agent activity, and concurrent workloads scale.

Teradata doesn't merely integrate AI components—it structurally enables the manufacturing of enterprise knowledge at scale.

Critically, this approach builds on what businesses already have. Your existing data, on and accessible to Teradata, is the asset. Organizations can progressively uplift existing data models by introducing semantic discovery and memory alongside current workloads, immediately making data usable by agents and preserving institutional knowledge without disruption.

Over time, as observability, search, and prediction capabilities are added, enterprises evolve toward a true knowledge platform: one where every new AI initiative benefits from prior learning, shared context, and governed reuse.

Conclusion

The future of enterprise intelligence is autonomous. As agents become continuous enterprise workers, platforms must provide trusted knowledge, governed execution, and predictable economics at scale. The Teradata Autonomous Knowledge Platform is designed to enable that future, bringing together the systems and controls required to turn insight into autonomous outcomes.

17095 Via Del Campo, San Diego, CA 92127 Teradata.com

teradata.