White paper

From Data Management to Intelligence Activation

Architecting Agentic Systems Through Signal-Oriented Design

Phinean Gardner
Teradata

teradata.

Version 2.0

Table of Contents

Abstract

This paper outlines Teradata's vision for agentic systems, combining current capabilities with strategic roadmap elements to illustrate the path from data management to intelligence activation.

Enterprises are at an inflection point: storing and reporting on data is no longer enough to win in the AI era. Competitive advantage now depends on the ability to detect signals in real time, interpret them in context, and activate them autonomously at previously impossible scale. This white paper outlines Teradata's evolution from a high-performance data platform to an enterprise **system of intelligence**—culminating in agentic architectures where intelligent software agents continuously sense, decide, and act across business workflows.

At the core of this evolution are three interdependent building blocks: **data products** as reusable, multi-dimensional assets; **signals** as meaningful patterns derived from complex relationships; and **agents** as autonomous entities operating across data and application layers. These systems harness both **multi-structured** data (structured, semi-structured, and unstructured) and **multidimensional** data (interconnected attributes analyzed across multiple perspectives), enabling a shift from retrospective analysis to signal-oriented design where intelligence is contextual, composable, and reusable across the enterprise.

We introduce the signal framework as a blueprint for operationalizing intelligence—illustrated through customer intelligence as a comprehensive business architecture that transforms how organizations understand, engage, and act on customer insights in real time. We explore **Vantage Customer Experience** (VCX) as the activation engine connecting signal processing to enterprise applications through scalable publish/subscribe architecture. Finally, we define Teradata's unified **agentic strategy**, integrating natural language understanding (NLU) with Model Context Protocol (MCP) and Teradata's execution engine, creating a unified command layer where MCP serves as the intelligent broker connecting conversational intent to machine-executable workflows. The result is an operational foundation for enterprises to transform data into continuous, autonomous value creation.

Introduction: From data management to intelligence activation

For decades, enterprises have measured data maturity by how efficiently they could store, govern, and query their information assets. Organizations invested heavily in building robust data platforms—data warehouses, data lakes, and data lakehouses—establishing comprehensive governance frameworks and creating reliable reporting systems—all centered on the foundational principle of "store once, use many." This data-driven approach created immense value through standardization, operational efficiency, and consistent access to historical information.

But in the AI era, competitive advantage no longer comes from managing data alone—it comes from the ability to detect and interpret signals in real time, and to activate them autonomously at previously impossible scale—where intelligent agents can be deployed and scaled programmatically beyond the limitations of human-dependent processes. This shift moves the enterprise from data-driven operations to signal-oriented architectures, where intelligence is continuously embedded into workflows, applications, and user interactions.

Signal-oriented design

Teradata's strategic evolution has moved beyond a traditional data-driven platform to a comprehensive system of intelligence, designed to detect signals, execute models, and activate decisions in real time. Rather than viewing this as a departure from Teradata's core strengths in high-performance, parallel processing, this represents the fulfillment of its original vision—enabling complex data relationships to be processed and acted upon at unprecedented scale and speed.

The transformation centers on a shift from retrospective analytics to signal-oriented design—architectures that continuously isolate and reuse meaningful patterns across business functions to enable real-time, context-aware responses. This requires a new data foundation: purpose-built data products that transform raw, unified data into structured, multidimensional assets optimized for intelligence generation. Unlike traditional datasets, a data product is an intellectual capital good—a structured asset that functions like a manufacturing tool, capable of being applied repeatedly to produce different insights across various business outcomes without requiring reconstruction or reengineering.

A signal, by definition, is a unique, recognizable pattern within data—a distinct expression of context, behavior, or intent that can drive action. For example, a customer's transaction velocity, geographic patterns, and device usage anomalies may constitute a fraud risk signal that, when properly isolated by a model from well-structured data products, can trigger automated account reviews or enhanced authentication protocols in real time. These signals, derived from multidimensional data products, serve as the foundation for intelligent agents operating autonomously within both data and application layers—not static rule-based systems, but adaptive response mechanisms that learn and improve over time through outcome feedback. Whether orchestrating data workflows, scoring risk in real time, or personalizing customer engagement at interaction moments, these agents form the core of the agentic system model.

Signals at scale

The power of this approach becomes clear when signals are manufactured and reused at scale. A financial institution's credit deterioration signal, for example, can simultaneously activate lending alerts, collections triggers, marketing exclusions, and risk calculations—all from one intelligent output. In an agentic future, this same signal orchestrates thousands of workflows across millions of customers, executing decisions at speeds and volumes impossible with human-dependent processes.

In this model, intelligence becomes contextual, composable, and reusable—delivered not merely through dashboards and reports, but through active services embedded directly in applications and user interfaces. NLU extends this capability by transforming conversational intent into machine-executable workflows, enabling users to interact with complex systems through natural conversation. For example, a user inquiry like "Show me high-value customers with declining engagement this quarter" would automatically trigger data product selection, apply the appropriate churn models, evaluate signal strength, and route the results to the relevant business applications. The result is adaptive enterprise architecture that systematically reduces operational friction, accelerates decision velocity, and scales intelligence autonomously across business domains.

The stakes of this transformation are significant. Organizations that successfully make this transition will achieve operating leverage previously thought impossible—where intelligence compounds over time, decisions accelerate through automation, and competitive advantage becomes self-reinforcing through previously impossible scale rather than human constraints. Those that remain anchored to traditional data management approaches risk being systematically outpaced by competitors who have learned to activate their intelligence in real time, at scale, across all business operations.

Data Driven		Signal Oriented
Organization		**Organization**
Data product as an output	**Disposition**	Derived intelligence made actionable
Data governance	**Proficiencies**	Model adaptability
Deterministic	**Capabilities**	Emergent
Latent	**Decision-making**	Real time
Largely manual	**Workflows**	Somewhat autonomous
Current operations	**Innovation for**	Future operations
Expense reduction	**Value defined by**	Revenue creation
Return on data asset (e.g., RODA)	**Value measured by**	Yield on signal (e.g., operating leverage)

Figure 1. From Data-Driven to Signal-Oriented Design

Teradata in the era of AI

Building on decades of data processing expertise, Teradata now faces a new imperative: extending its proven scale capabilities from data storage to intelligent application activation.

Teradata has historically defined its value through scale—primarily at the data layer. In the era of AI, however, that scale extends to the application layer. This is not a departure from our data foundation, but an expansion that reflects the deeper interdependence between data infrastructure and intelligent applications.

The application layer encompasses the software systems, interfaces, and services where end users and business processes directly interact with intelligence. In financial services, this ranges from mobile banking apps that provide real-time spending insights to loan origination systems that instantly assess creditworthiness to trading platforms that execute automated risk management. In healthcare, applications extend from electronic health records that surface predictive diagnostics to patient monitoring systems that alert clinicians to deteriorating conditions to treatment recommendation engines that personalize care protocols. In retail, these applications span from dynamic pricing engines to inventory optimization systems to omnichannel customer engagement platforms.

For intelligence to be meaningful at the point of decision, it must be nonlinear—adaptive, contextual, and continuously learning.

Evolving applications and data

Applications have become the primary interface where intelligence is consumed, decisions are made, and actions are taken. Yet this amplification at the application layer is only possible because of the richness, structure, and velocity of the underlying data foundation. For intelligence to be meaningful at the point of decision, it must be nonlinear—adaptive, contextual, and continuously learning. This requires establishing dynamic feedback loops between the application and data layers, where multidimensional representations of customers, risks, and assets are constructed and refined in real time.

Applications and data must therefore evolve in tandem: applications define the demand for more responsive, context-aware intelligence, while the data layer adapts to deliver the precision and agility those applications require. It's this symbiotic cycle—where applications activate intelligence and the data layer sustains it—that ultimately defines modern AI value creation.

Such an architecture requires a common data foundation shared across all applications to guarantee accuracy, consistency, and transparency. Without it, applications risk operating on disparate versions of truth, undermining the trust and reliability essential for autonomous decision-making.

In this paradigm, Teradata serves as the intelligent foundation that powers this cycle, enabling organizations to transform raw data into actionable intelligence at enterprise scale.

Teradata as a system of intelligence

This application-layer expansion requires a fundamental shift in how data itself is architected—moving from storage-focused platforms to intelligence-ready infrastructure capable of supporting autonomous operations.

As enterprises deploy thousands of autonomous agents, the underlying data architecture becomes the critical bottleneck. While other platforms struggle to scale beyond single-node agent deployments, Teradata's distributed parallel architecture provides the foundation that agentic AI demands—not by accident, but by design.

A system of intelligence differs fundamentally from traditional data platforms in three critical ways: it processes multi-structured and multidimensional data simultaneously rather than in separate pipelines; it executes intelligence operations—model inference, signal detection, and decision logic—directly within the data layer rather than requiring external processing; and it operates continuously and autonomously rather than responding only to explicit queries. This represents a shift from storing data for human analysis to manufacturing intelligence for autonomous systems.

Context-aware decision-making

Modern applications now demand real-time, adaptive intelligence capable of processing complex relationships across large-scale, multi-dimensional datasets—transforming raw information into context-aware decisions instantly. Teradata can finally be utilized as it was originally envisioned—not simply to store or query data, but to power distributed AI applications that process complex, multi-structured, and multidimensional data continuously and at scale.

Delivering this capability requires handling the full spectrum of enterprise data complexity. Multi-structured data encompasses the full spectrum of enterprise information: traditional structured data from transactional systems and databases, semi-structured data such as JSON logs, XML configurations, and API responses, and unstructured data including documents, images, audio files, and sensor feeds. Rather than requiring separate processing pipelines for each data type, Teradata's architecture can simultaneously handle a customer's transaction history (structured), their support chat transcripts (unstructured), and their mobile app interaction logs (semi-structured) within unified analytical workflows.

Multidimensional data processing

Equally important is processing multidimensional data—the complex, interconnected attributes that define business entities across multiple analytical perspectives. A single customer, for example, exists across behavioral dimensions (purchase patterns, channel preferences, engagement frequency), temporal dimensions (seasonal trends, lifecycle stage, recency patterns), contextual dimensions (geographic location, device usage, interaction timing), and relational dimensions (household connections, social network influences, business relationships). These dimensions can be analyzed at various levels of granularity—from high-level aggregates that reveal broad patterns to detailed event sequences that capture nuanced behaviors.

Teradata's parallel execution engine enables precisely this capability. Operating within a connected ecosystem of enterprise systems, it processes, scores, and refines signals directly within the database—executing complex joins, transformations, and model inference across massive datasets in real time. This architecture supports neural, symbolic, and hybrid AI applications that learn from the full complexity of enterprise data landscapes, rather than simplified, flattened representations that lose critical contextual information. The result is the ability to isolate latent signals, respond to context dynamically, and support nonlinear intelligence across multiple decision points simultaneously.

This is not a traditional data platform in the conventional sense. This is the foundation of modern AI infrastructure—a system that operationalizes intelligence at scale, enabling enterprises to act autonomously on complex data relationships in real time.

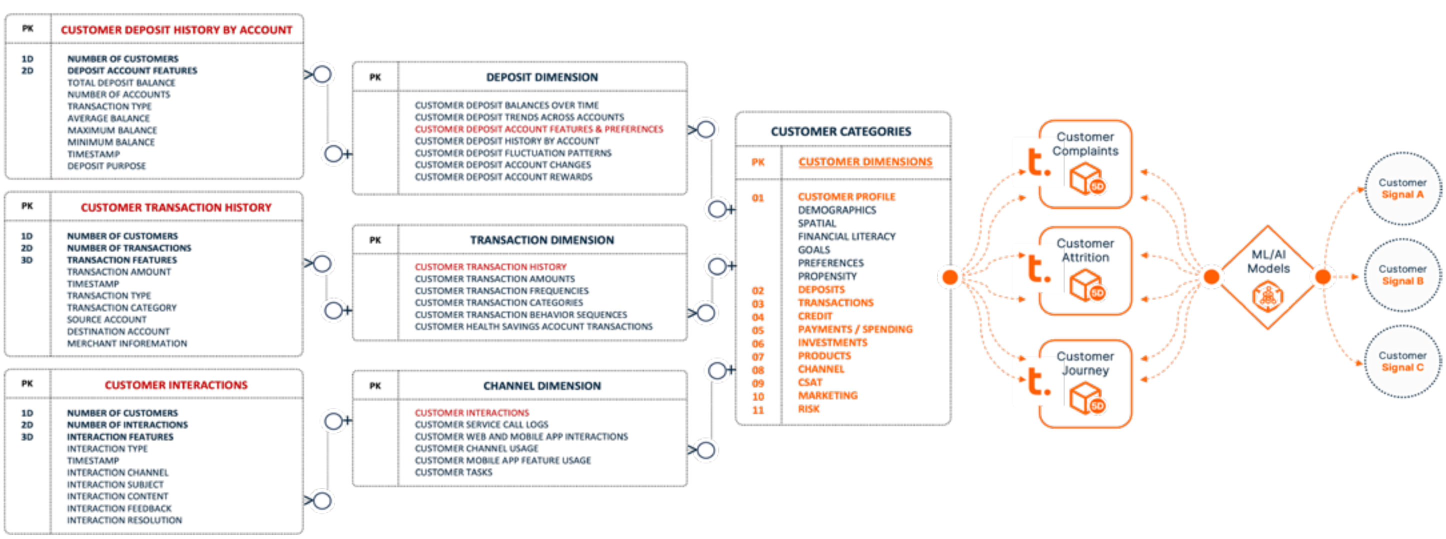

Figure 2. Transforming Signals Into Action With Multidimensional Data

Signals: The unit of intelligence

With the architectural foundation established, the focus shifts to what flows through that infrastructure: transforming static data assets into dynamic intelligence that drives real-time business action.

Traditional data platforms have been built on the principle of "store once, use many." Being data-driven means ensuring data is high quality, well governed, accessible, reusable, and trusted across the enterprise. This approach has led organizations to centralize and curate data assets effectively, but decision-making in this model often remains retrospective and reactive.

While this data platform foundation remains critical, it represents only the first step in the evolution toward intelligent systems. The next phase is becoming signal oriented—an evolution beyond managing data toward activating the intelligence embedded within it.

A signal, by definition, is a unique, recognizable pattern within data—a distinct expression of context, behavior, or intent that can drive action. A model isolates that signal, distinguishing meaningful patterns from noise.
The process of evaluating and interpreting the signal initiates the transformation of raw data into insight and insight into actionable intelligence.

In a signal-oriented approach, organizations go beyond managing and analyzing data. They begin manufacturing intelligence—high-value, real-time outputs derived from signals detected within multidimensional data that can be applied across workflows. These intelligence outputs are scored, prioritized, and reused across teams, systems, and functions. The emphasis shifts from manual, retrospective analysis to semiautonomous, adaptive action.

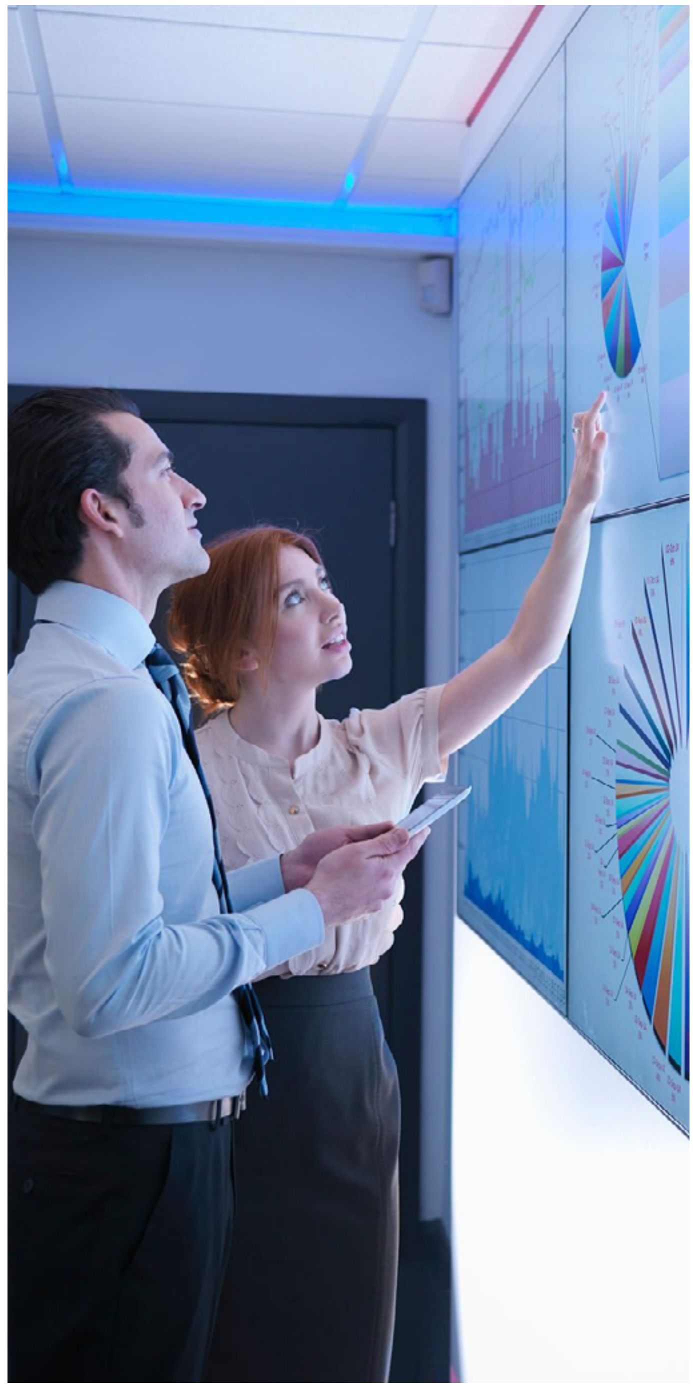

 Signals: The unit of intelligence

From analysis to action

Consider a financial institution that manufactures a "credit deterioration signal" by continuously processing payment history, account balances, transaction patterns, and external economic indicators. Rather than generating monthly credit reports for individual review, this signal is scored in real time (high, medium, low risk), prioritized by exposure level, and automatically distributed across multiple business functions. The lending team receives early warning alerts for portfolio review, the collections team gets proactive intervention triggers, the marketing team excludes high-risk customers from credit offers, and the risk management team updates regulatory capital calculations—all from the same manufactured signal, activated simultaneously across workflows without manual coordination. In an agentic future, this same signal could simultaneously orchestrate thousands of such workflows across millions of customers—executing thousands of simultaneous decisions that would be impossible with human-dependent workflows.

This reuse creates significant operating leverage: the investment in developing and maintaining one sophisticated signal generates value across four distinct business functions, each previously requiring separate analytical efforts. As signal quality improves through continuous learning, all dependent workflows benefit simultaneously, multiplying the return on the original intelligence investment.

The fundamental shift is one of purpose and value creation. Data-driven models primarily seek cost savings through standardization and operational efficiency. Signal-oriented models focus on value generation—creating measurable return on signal intelligence that compounds as multiple applications leverage the same intelligent outputs. This enables organizations to achieve operating leverage, accelerate time to insight, and embed intelligence into the core of business operations—not merely as analysis, but as direct, automated action.

These intelligence outputs are scored, prioritized, and reused across teams, systems, and functions.

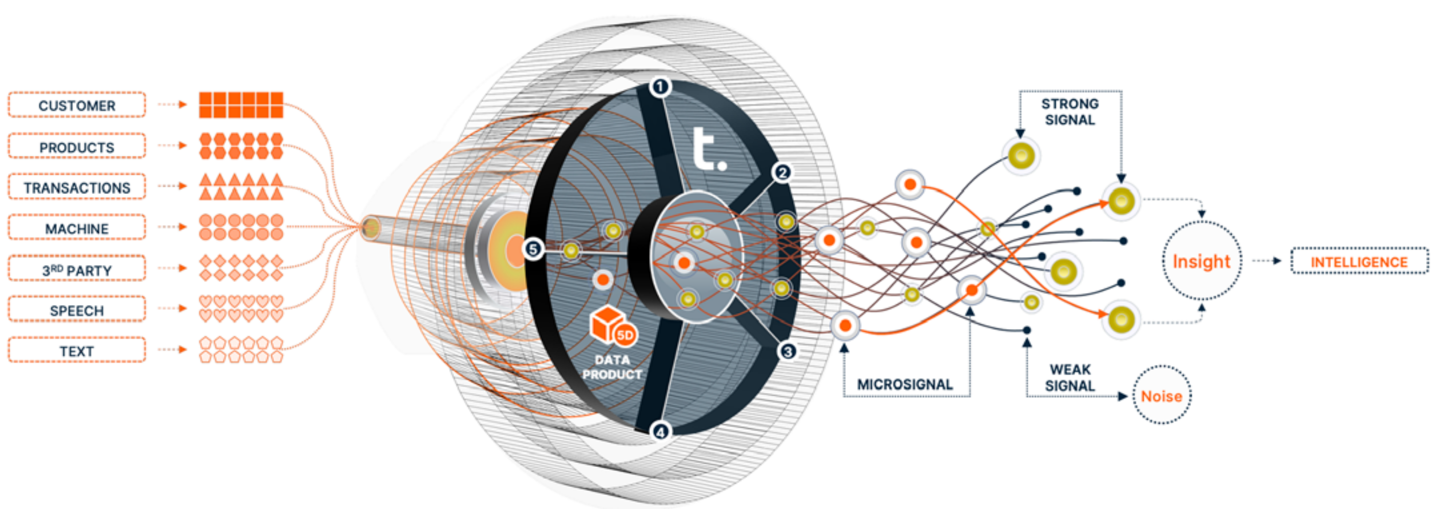

Figure 3. Modeling Intelligence Through Multidimensional Signal Analysis

The signal-oriented framework

Moving from signal theory to practical implementation, organizations need a systematic methodology for manufacturing and deploying intelligence across diverse business contexts.

The signal framework is a business-aligned design pattern for operationalizing intelligence at scale. At its core, it enables organizations to isolate meaningful signals from data and convert them into reusable microservices that can be embedded directly into decision-making workflows. This is how data becomes actionable—and how intelligence drives outcomes.

Understanding the process

The framework follows a systematic process that begins by defining a specific business problem and forming a testable hypothesis about how it can be addressed analytically. From there, relevant data is identified, captured, and organized into a structured data product. This data product becomes the foundation for feature engineering, feeding models designed to encode complex relationships between those features and business outcomes.

The model, in turn, isolates the signal—the unique, meaningful pattern in the data that reveals what matters most for the specific business context. This signal is then exposed to business users, who apply domain expertise to interpret its meaning and relevance. Through this interpretation process, the signal evolves into an insight, which is further refined and contextualized into actionable intelligence.

The final step is activation: embedding this derived intelligence into applications as reusable services. These services can be published for consumption or subscribed to by any system that requires them, enabling intelligence to flow seamlessly across the enterprise without duplication or inconsistency.

Benefits of the framework

The strength of the framework lies in its adaptability across industries. In retail, for example, an inventory demand signal integrates point-of-sale data, seasonal trends, supplier lead times, digital engagement patterns, and external factors such as weather. This signal simultaneously drives supply chain restocking, dynamic pricing, marketing promotions, and procurement adjustments—all orchestrated from a single, reusable signal. In healthcare, a patient deterioration signal combines real-time vitals, lab results, medication adherence, and historical patterns to detect early risks. It triggers coordinated actions, including clinical alerts, pharmacy dosing adjustments, family notifications, and intensive care unit (ICU) resource allocation—improving patient safety while ensuring regulatory compliance. In manufacturing, an equipment failure prediction signal analyzes internet of things (IOT) sensor data, maintenance history, production schedules, and environmental conditions to anticipate breakdowns. It activates preventive maintenance, adjusts production schedules, initiates parts ordering, and alerts quality teams—while also optimizing for sustainability and energy efficiency.

Teradata designed the signal framework to eliminate friction throughout this flow—streamlining how signals are generated, transformed into intelligence, and scaled across workflows. It's not just a model for managing data but for activating it operationally, closing the loop between insight and action while maximizing business value in the AI era. Perhaps its most powerful application lies in customer intelligence: by treating customer data as a continuous source of multidimensional signals, enterprises can move beyond static views to enable real-time detection of intent, prediction of needs, and activation of personalized engagement at scale.

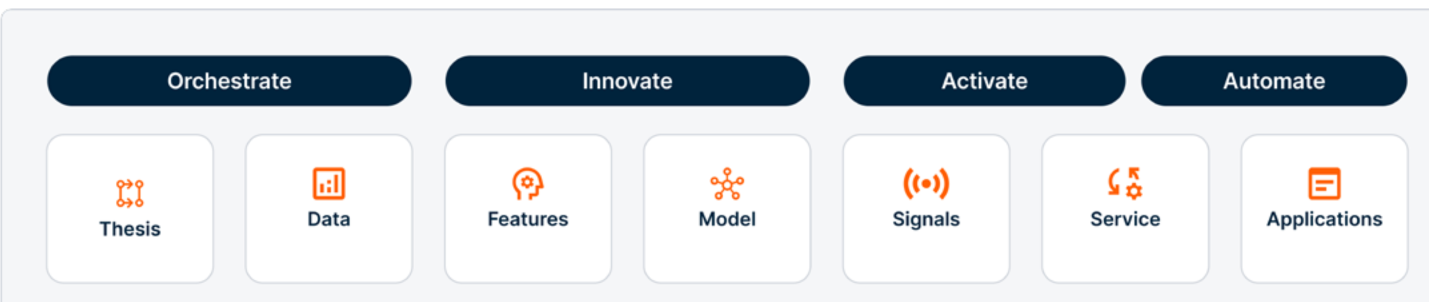

Figure 4. Orchestrating Intelligence: From Signal Detection to Workflow Activation

Customer intelligence as business architecture

The framework's versatility becomes most apparent when applied to the enterprise's most valuable asset: customer relationships. Here, multidimensional signals create unprecedented opportunities for competitive advantage—moving customer intelligence far beyond customer experience (CX) into a true business-wide architecture.

To illustrate how the signal framework operates in practice, customer intelligence represents the continuous detection, interpretation, and activation of multidimensional customer signals. These signals span behavioral, transactional, contextual, and operational dimensions—each providing a unique perspective on the customer journey. When properly integrated and activated, they enable organizations to uncover intent, anticipate needs, and deliver personalized, real-time experiences with unprecedented precision.

It's crucial to understand that Customer 360 is not customer intelligence. A Customer 360 view represents unified customer data—typically consolidated from various systems of record to support reporting and basic decision-making. While foundational and necessary, this represents raw material—not intelligence. Intelligence begins when models are systematically applied to this unified data, isolating signal from noise and detecting patterns of intent, opportunity, or risk that would otherwise remain hidden. This application of modeling initiates the transformation from static data to dynamic insight and from insight to actionable intelligence.

Customer intelligence is therefore more than a single capability—it represents a comprehensive business architecture designed to address one of the most pressing strategic challenges organizations face: developing a deep, dynamic, and actionable understanding of the customer that drives competitive advantage.

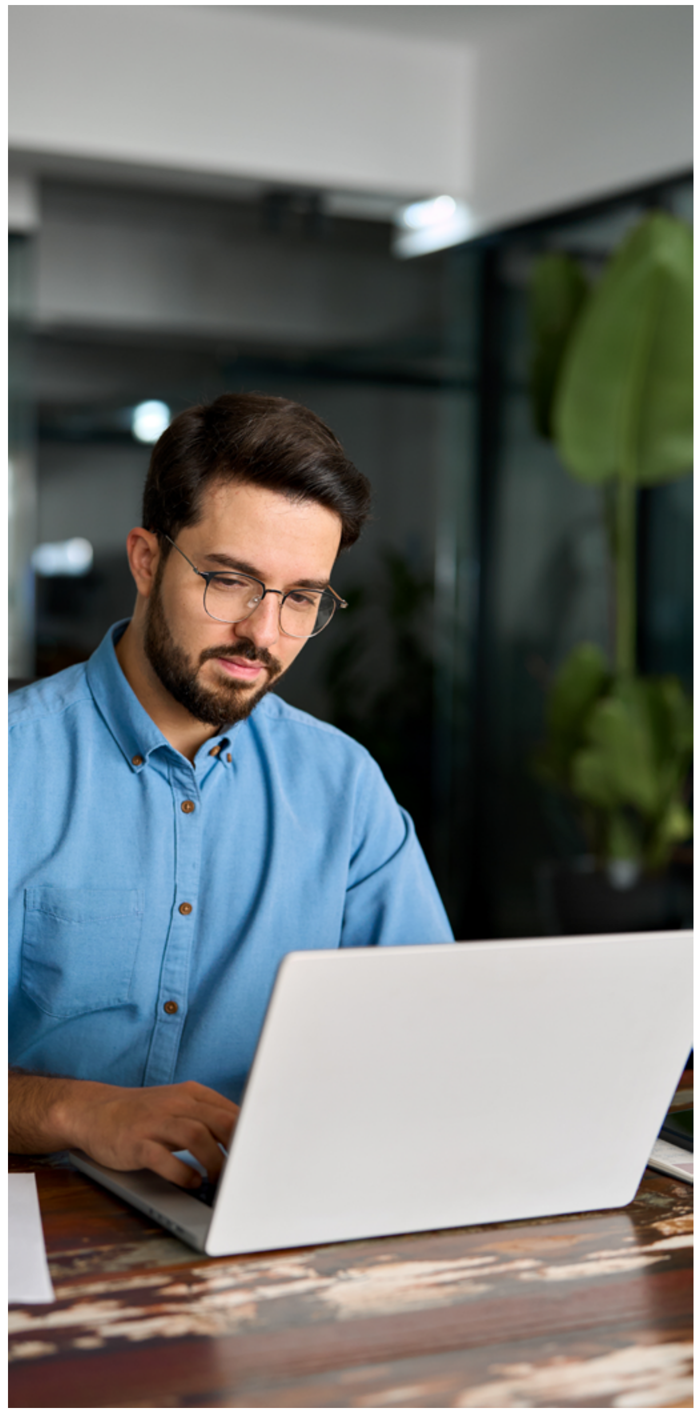

The importance of customer intent

In an environment in which customers interact across multiple channels, devices, and time horizons, traditional approaches that rely on demographic segmentation rather than understanding customer intent prove insufficient. To maintain competitive advantage, businesses must shift from static personas to signal-based intelligence architectures that sense and respond continuously to changing customer contexts. This requires aligning data, models, and application logic around a unified customer signal layer—transforming raw data into derived intelligence that can be reused across workflows, touchpoints, and business functions.

A properly implemented customer intelligence architecture does not merely support customer experience functions such as marketing and service—it becomes a strategic enterprise asset that drives value across the entire organization. In product innovation, customer signals inform needs identification and roadmap prioritization. Risk and fraud teams leverage behavioral patterns for anomaly detection and compliance monitoring. Operations teams use customer demand signals for forecasting and supply chain optimization. Finance teams apply customer intelligence for profitability analysis and investment prioritization.

Customer intelligence, therefore, isn't simply about knowing your customer—it's about designing your entire business to act intelligently on that knowledge, in real time, at every point of customer interaction. At the heart of making this vision operational is the data product. It's the foundation that transforms unified customer data into a reusable, intelligence-ready asset— ensuring that insights aren't isolated within individual projects but consistently applied across models, workflows, and functions—making intelligence scalable, repeatable, and enterprise-wide.

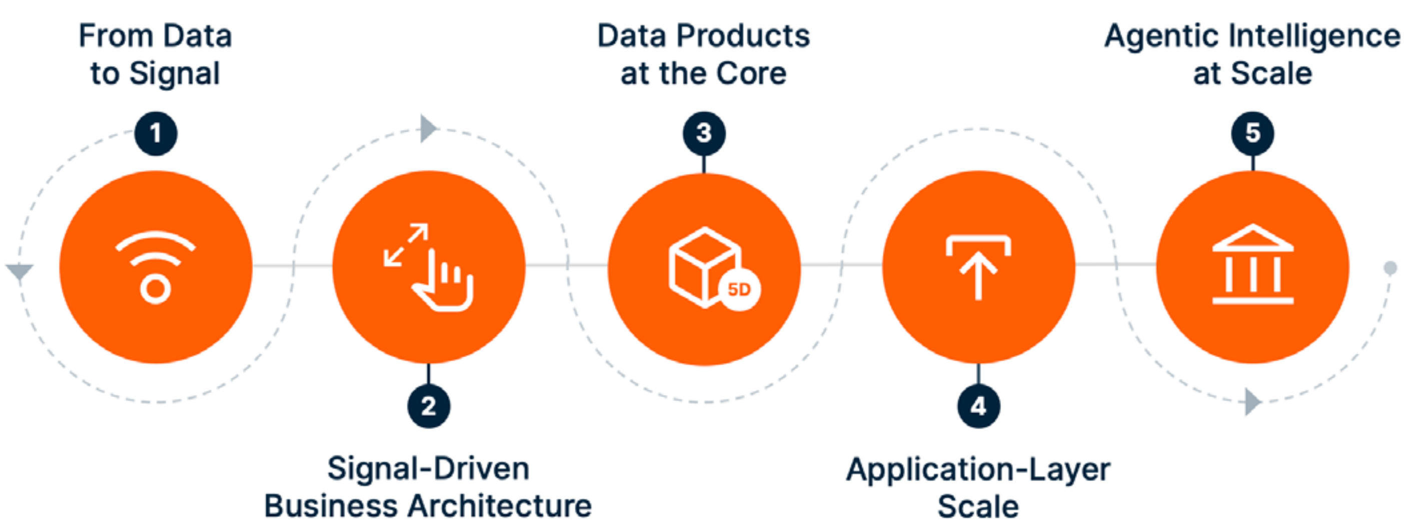

Figure 5. The Pillars of Customer Intelligence Activation

The role of data products in customer intelligence

Operationalizing customer intelligence requires more than unified customer data—it demands purpose-built assets engineered specifically for continuous signal detection and activation.

At the foundation of customer intelligence lies the data product—a critical building block that brings structure, governance, and reusability to complex, multidimensional customer data. While Customer 360 serves as the raw, unified data layer, the data product transforms that raw material into something purpose-built for intelligence: organized, dimensional, and optimized for signal detection at scale.

A data product isn't a report, dashboard, or one-off dataset. It's a reusable, logically organized asset—typically composed of multiple source systems and structured around a consistent business entity such as the customer, household, or account. It encapsulates data, metadata, lineage, governance policies, and embedded business logic, making it consistently usable across teams, models, and applications. Think of the data product as an intellectual capital good—a structured asset that functions like a manufacturing tool, capable of being applied repeatedly to produce different insights across various business outcomes without requiring reconstruction or reengineering.

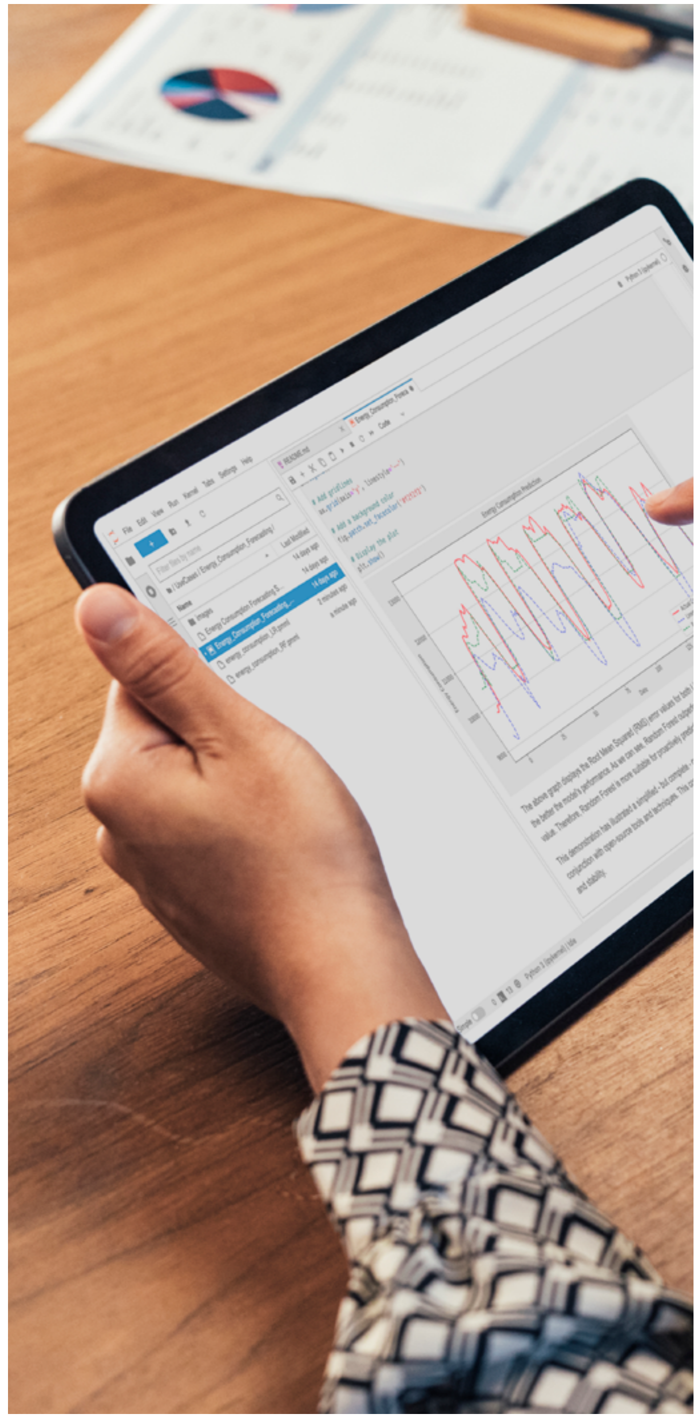

Customers across multiple dimensions

What makes the data product central to customer intelligence is its ability to represent customers across multiple dimensions: behavioral, transactional, demographic, interactional, and contextual. These dimensions aren't flat representations—they can be modeled in both low-dimensional spaces (aggregates, summaries, patterns) and high-dimensional spaces (detailed, event-level, sequence-based data).

This dimensional flexibility is what enables sophisticated signal detection at enterprise scale. Low-dimensional representations compress complex customer behaviors into interpretable patterns—risk scores, propensity segments, or lifecycle stages—that enable broad signal detection and strategic decision-making across large populations. High-dimensional representations preserve the full granular context—detailed transaction sequences, interaction histories, and relationship networks—that enable precision decisioning and personalized responses at the individual level.

The analytical power lies in seamlessly moving between these dimensional spaces: using low-dimensional views to identify signals of interest, then accessing high-dimensional detail for contextual action. This dual capability means the same data product can simultaneously support portfolio-level risk management and individual customer personalization; population-level trend analysis and real-time behavioral prediction.

A three-dimensional perspective

Consider a data product as analogous to a three-dimensional Rubik's Cube—a structured, multidimensional object that contains many possible perspectives. The cube itself represents the reusable foundation: engineered for intelligence and complete with all relevant customer dimensions properly integrated. When you rotate the cube—changing the axis of evaluation—you expose a new perspective and reveal different patterns. Each rotation represents a different use case or analytical approach.

In this metaphor, the data product is the cube—integrated, reusable, and engineered specifically for signal activation. The use case represents the perspective—the specific slice or rotation of the cube applied to address a particular business challenge. This distinction is fundamental: a use case is an application—a contextual expression of value for a specific purpose. A data product is a durable, logically organized asset designed to support multiple use cases over time without requiring reconstruction.

Consider how a customer transaction data product demonstrates this flexibility in practice. When rotated toward fraud detection, the same data product surfaces anomaly patterns—unusual spending locations, transaction timing, or merchant categories that deviate from established behavior. Rotated toward credit scoring, it reveals payment consistency, account balance trends, and financial stability indicators. For personalization, the data product exposes preference signals—brand affinities, seasonal purchasing patterns, and channel preferences that drive recommendation engines. When applied to churn prediction, it highlights engagement trends—declining transaction frequency, reduced spending levels, or shifts in product usage that indicate customer disengagement risk.

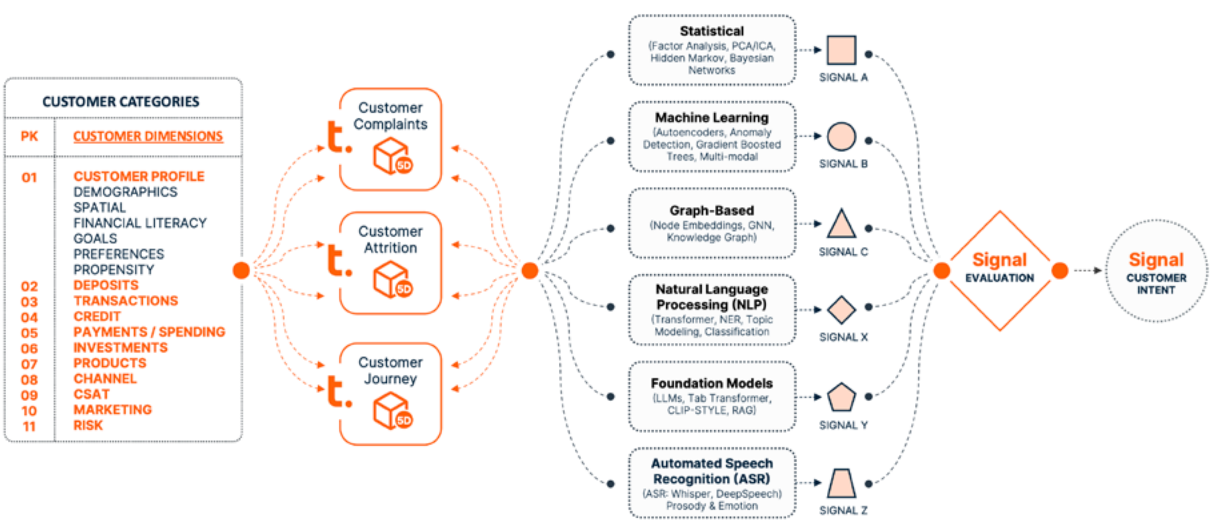

Figure 6. Unlocking Latent Signals With Multidimensional Model Architectures

Architectural flexibility

This single data product supports four distinct business outcomes without requiring separate data engineering efforts. Each “rotation” activates different dimensional perspectives of the same underlying customer data, enabling specialized teams to extract relevant signals while maintaining data consistency and governance across all applications.

This architectural flexibility is where Teradata provides unique differentiation. With its legacy of high-performance, multidimensional data processing, Teradata enables organizations to build and operationalize data products that scale not just in volume, but in complexity—spanning billions of interconnected records, diverse data structures, and concurrent analytical perspectives. These products can be scored continuously in place without data movement, rotated across multiple analytical axes to isolate originating signals, and evaluated for signal strength, durability, proximity, and influence in real time. This is scale in every dimension—data volume, analytical depth, operational concurrency, and real-time adaptability.

This capability—supporting both broad signal detection and fine-grained decision execution simultaneously—empowers Teradata to operationalize customer intelligence in ways other platforms can't match. The platform provides not just data storage but engineered data products that activate intelligence at speed, in context, and at enterprise scale. But intelligence only delivers value when it's activated—when it flows from these data products into the applications and workflows where real-time decisions are made. Bridging that gap between intelligence creation and intelligence consumption is the role of Teradata's Vantage Customer Experience (VCX) platform.

The data product transforms that raw material into something purpose-built for intelligence: organized, dimensional, and optimized for signal detection at scale.

VCX: The activation layer

With intelligence-ready data products established, the critical challenge becomes delivering signals seamlessly across enterprise applications where real-time decisions create business value.

VCX is the enterprise-scale activation layer that operationalizes customer intelligence by delivering signals from data products directly into application workflows. It transforms intelligence from a stored asset into a real-time operational advantage, serving as the critical bridge between Teradata's signal processing capabilities and enterprise applications.

In the modern enterprise, applications have evolved beyond passive endpoints to become dynamic execution environments for intelligence. They represent the critical interface where intelligence is consumed, decisions are made, and actions are taken in real time. Yet traditional point-to-point integrations create fragmented intelligence delivery, where each application requires custom connections and separate analytical efforts.

VCX fundamentally changes this paradigm by creating a unified activation layer—a scalable publish/subscribe architecture that enables seamless orchestration between intelligence generation and action execution. Originally designed as a real-time decisioning component for customer interactions, VCX delivers AI-powered, context-aware responses by operationalizing signals derived from multidimensional data processing. It supports real-time decisioning across both inbound and outbound customer touchpoints, ensuring the right insight is delivered to the appropriate application at precisely the right moment.

What makes VCX strategically significant is its ability to bridge customer intelligence with broader industrial applications. While its origins lie in marketing and service personalization, VCX has the architectural capability to serve as an enterprise-wide activation gateway, delivering derived intelligence across multiple business domains simultaneously.

Enterprise intelligence fabric

Consider how customer intent and behavior signals can be simultaneously activated across diverse enterprise applications: Pega receives these signals for marketing personalization and next-best-action campaigns; Salesforce leverages the same signals for service optimization, enabling agents to anticipate customer needs and personalize interactions; product management applications like FIS use these signals to configure dynamic loan products and real-time credit decisions; fraud detection platforms such as FICO Falcon consume behavioral signals to identify anomalous patterns and trigger investigation workflows; and credit risk applications like Moody's RiskCalc apply these signals for dynamic scoring and adaptive lending decisions. This signal orchestration across industrial applications transforms VCX from a marketing tool into an enterprise intelligence fabric that activates the same intelligence across all business domains simultaneously.

This activation layer does more than simply deliver responses—it interprets intent, evaluates signal strength relative to context, and aligns intelligence with specific application requirements. VCX ensures that decisioning is not only fast but situationally aware, driving relevance, precision, and speed across enterprise operations.

Signal orchestration

By evolving VCX into a general-purpose signal delivery utility, Teradata enables a new architecture for enterprise activation—one in which signals are orchestrated rather than siloed and intelligence moves dynamically with the customer rather than remaining static behind them. This transformation elevates customer intelligence from a passive capability to an active operational advantage, making every application intelligent by design.

The next frontier is moving from intelligent activation to autonomous orchestration. As VCX ensures that the right intelligence reaches the right application at the right moment, the opportunity now is to empower the applications—and the enterprise as a whole—to act on that intelligence without manual intervention. This is the domain of agentic intelligence: systems that not only deliver insights but also autonomously initiate, coordinate, and adapt actions through intelligent software agents.

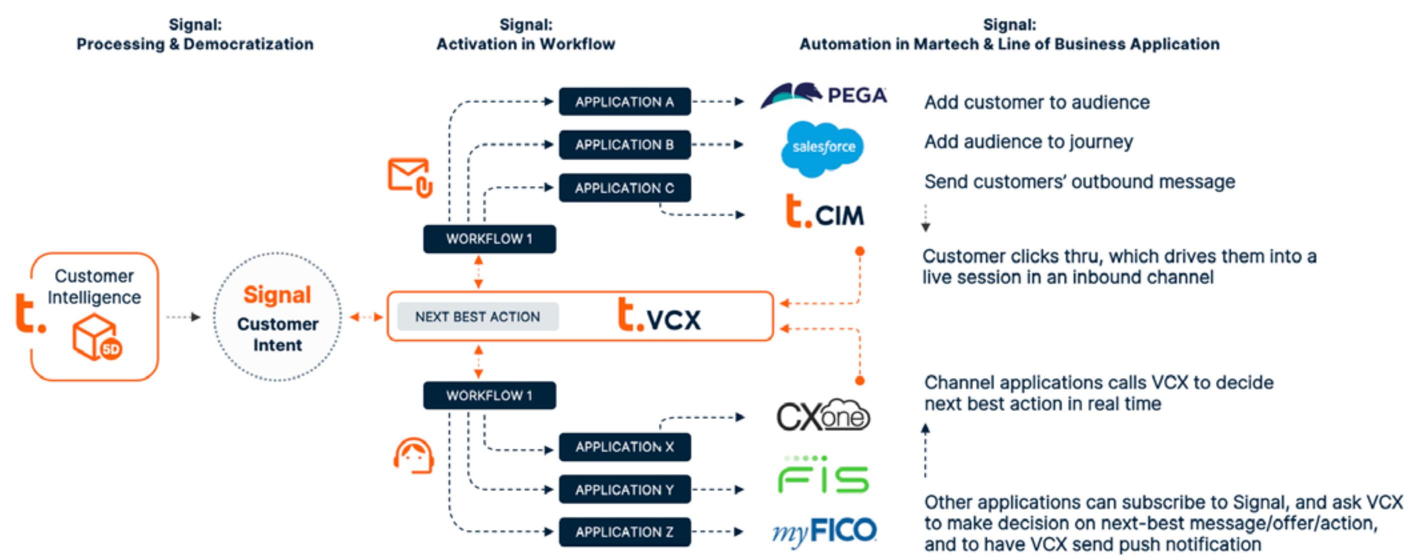

Figure 7. Activating Customer Intelligence With VCX

Agentic intelligence: A three-layer architecture

The convergence of data products, signals, and activation capabilities enables the ultimate evolution: autonomous systems that continuously sense, decide, and act across enterprise operations without human intervention.

As enterprises evolve from data-driven operations focused on storing, curating, and querying data toward signal-oriented decisioning centered on detecting, evaluating, and activating real-time intelligence, the next strategic frontier is agentic systems. Unlike rule-based systems, which follow static logic and fixed triggers ("if this, then that"), agentic systems are adaptive: they learn from new signals, adjust decisions based on feedback, and evolve their behavior over time. Think of a rule-based system as a global positioning system (GPS) with preset directions—it will follow the planned route no matter what. An agentic system is more like a self-driving car—constantly sensing the environment, rerouting when traffic changes, and improving its navigation from every trip.

Introducing the layers

Teradata operationalizes agentic intelligence through a three-layer architecture that forms the foundation of autonomous enterprise systems:

1. **Agents within the data layer:** These agents focus on orchestrating, transforming, and preparing multidimensional data continuously. They automate the construction of signal-ready data products, dynamically optimize processing workflows, and maintain governance policies without human intervention.
2. **Agents within the application layer:** These agents consume derived intelligence and autonomously execute contextual actions at the moment of business decision. They interpret signals, trigger compound workflows, and adapt responses across marketing, fraud detection, risk management, and product domains.
3. **NLU as the command layer:** This layer unifies the agentic system by enabling conversational interaction with both data and application layer agents. It creates a natural language interface in which business users can transform intent into machine-executable actions through integrated MCP.

Together, these three layers create a comprehensive topology for autonomous intelligence in which Teradata is uniquely positioned to lead, grounded in its deep expertise in high-performance, multidimensional data processing. This foundational capability—spanning structured, semi-structured, and unstructured data—enables agents to fluidly organize, abstract, and interpret complex signals with the scale, precision, and speed required for autonomous operation.

A continuous feedback loop

This three-layer architecture reinforces the deep interdependence of the data and application layers, reflecting their necessary evolution in tandem: as the data layer adapts to support real-time, multidimensional signal processing, the application layer responds by executing intelligence with greater precision, speed, and autonomy. The result is a continuous feedback loop in which application demands drive data-layer innovation, while enhanced data capabilities enable more sophisticated application behaviors.

In this model, Teradata transcends traditional insight generation to operationalize intelligence, transforming enterprise architecture into an adaptive system of agents capable of autonomous decision-making and action execution in real time.

Let's examine how each layer operationalizes autonomous intelligence, beginning with the foundation that makes it all possible.

An agentic system is more like a self-driving car—constantly sensing the environment, rerouting when traffic changes, and improving its navigation from every trip.

Agents within the data layer

The most profound breakthroughs often come not from creating something entirely new but from building on the strengths already in place. Teradata's parallel processing architecture provides exactly such a foundation: hundreds of Access Module Processors (AMPs) operating as persistent, coordinated nodes create the ideal infrastructure for supporting autonomous agents at unprecedented scale. While AMPs themselves aren't agents, their proven architecture of distributed, coordinated processing creates the foundation needed to deploy and scale intelligent agent workloads across the data layer—enabling thousands of agents to operate simultaneously without the memory and coordination constraints that limit other platforms.

This architectural serendipity transforms proven infrastructure into autonomous processors that intelligently orchestrate, abstract, and prepare complex, multidimensional data. Rather than building from scratch, the existing parallel architecture enables intelligent agent workloads to be deployed across distributed nodes that operate continuously—scanning for changes, triggering transformations, optimizing performance, and maintaining governance while keeping data products perpetually signal ready.

A distributed computing foundation

This evolution manifests across four critical operational domains:

- **Construction**: Building data products from multisource inputs with automated quality validation and consistency checking
- **Rotation and re-indexing**: Dynamically adjusting dimensional perspectives to isolate and strengthen signal detection
- **Evaluation**: Continuously assessing signal strength, durability, proximity, and influence across multiple business contexts
- **Publication**: Curating and distributing intelligence to downstream applications based on relevance and priority

These agent capabilities leverage Teradata's massively parallel architecture to operate at scale. Where traditional single-node systems might support dozens of concurrent agents, Teradata's distributed architecture can support thousands of specialized agents working simultaneously across different data products and business domains.

This represents a fundamental advantage: while other platforms struggle to scale agent workloads beyond single nodes, Teradata's proven parallel architecture provides the distributed computing foundation that agentic AI demands for enterprise-scale deployment.

Through these coordinated capabilities, the data layer transforms from a passive storage and processing engine into an active intelligence environment—one capable of autonomously supporting continuous signal production, refinement, and preparation without constant human oversight.

Agents within the application layer

While data layer agents focus on intelligence creation, application layer agents specialize in intelligence activation. These systems consume the refined signals and insights produced upstream and initiate contextually appropriate actions at the critical moment of engagement—where decisioning occurs and business value is realized. They interpret contextual signals, trigger compound workflows, and adapt responses in real time based on changing conditions and outcomes.

Agents within the application layer illustrate how signals are operationalized into action. In fraud detection, agents consume continuously updated risk scores and automatically launch investigation workflows—assigning cases, gathering evidence, and notifying stakeholders in real time. In product operations, agents adapt configurations based on shifting customer behavior, dynamically optimizing the user experience without manual intervention. In the supply chain, agents monitor weather, traffic, and logistics signals to reroute shipments when disruptions occur, proactively adjusting schedules and notifying customers. In financial markets, trading agents respond to volatility signals by rebalancing portfolios, adjusting risk exposure, and executing hedging strategies at speeds impossible for human traders. In customer service, agents analyze requests using sentiment and complexity scoring to intelligently route cases, escalating high-value or frustrated customers to senior representatives while directing routine inquiries to the right teams.

Sophisticated yet intuitive

Together, these examples demonstrate how application-layer agents transform signals into real-time, context-aware actions—embedding intelligence directly into business workflows. These aren't static rule-based systems—they represent adaptive response mechanisms rooted in live, orchestrated signal intelligence that learns and improves over time through outcome feedback.

But how do humans direct this autonomous intelligence ecosystem? While data and application layer agents operate autonomously, enterprises still need a way for business users to initiate, monitor, and guide these intelligent systems. The challenge is creating an interface sophisticated enough to orchestrate complex agent workflows, yet intuitive enough for natural human interaction. This is where conversational intelligence becomes the command layer—unifying autonomous agents under human direction.

NLU interfaces as the command layer

Rather than requiring users to understand complex system architectures or agent protocols, Teradata integrates NLU interfaces with its MCP to create a unified command layer. These conversational interfaces enable business users, analysts, and operators to direct the entire agentic ecosystem through natural language—transforming intent into coordinated action across both data and application layer agents.

An intelligent toolbox

It's important to distinguish between MCP's role as a broker and the broader command layer capabilities this integration creates. MCP itself functions as an intelligent toolbox—a standardized protocol that enables agents to discover and connect to available tools and data products. The "command layer" emerges from the integration of three components: NLU interprets conversational intent, MCP brokers connections to appropriate tools and data products, and Teradata's parallel processing engine executes complex workflows at scale. This integrated system transforms natural language into coordinated autonomous action. Through this integration, users can:

- **Trigger data orchestration**: "Update the customer churn data product with this week's interaction data"
- **Initiate signal evaluation**: "Evaluate the strength of our retention signals for high-value customers"
- **Activate models**: "Run the cross-sell model for customers who opened accounts in the last 30 days"
- **Route intelligence**: "Send personalization signals to the mobile app for users logging in now"

Which high-value customers are showing early signs of churn?

This query would trigger an agent within MCP to identify and access the appropriate data product, apply the churn prediction model, evaluate signal strength and confidence levels, and publish the resulting intelligence to VCX for immediate activation across relevant customer touchpoints.

Activated workflows

In this workflow, MCP functions as the broker—discovering and connecting to the appropriate data products and models—while Teradata provides the high-performance execution environment and VCX handles activation. The "command layer" is the seamless orchestration of all these capabilities through natural language.

By integrating NLU with MCP and Teradata's processing capabilities, conversational intent becomes machine executable—transforming natural language into activated workflows. MCP serves as the intelligent broker, connecting agents to the right tools. The command capabilities emerge from the full NLU, MCP, and Teradata integration, dynamically connecting the appropriate model, tool, or pipeline based on user intent and context. Teradata functions as the high-performance execution engine—processing multidimensional data, scoring signals, and publishing outputs at scale.

VCX completes this intelligence loop as the delivery fabric, streaming real-time insights into enterprise workflows and applications. Together, this integrated architecture elevates the chat interface from a passive question-and-answer tool into an active agentic command layer—bridging human intent and machine execution through intelligence that is contextual, interpretable, and operationally embedded throughout the enterprise.

This represents the future of enterprise intelligence: a seamless integration of human insight, machine processing, and autonomous action that transforms how organizations create value from their data assets.

Conversational intent becomes machine executable—transforming natural language into activated workflows.

Conclusion

The enterprise no longer wins by storing data— it wins by scaling intelligence at previously impossible scale. In the agentic future, organizations succeed by deploying autonomous systems that can execute millions of decisions simultaneously, breaking free from the human bottlenecks that limit traditional business processes. Without systems for continuous activation from multidimensional data, organizations aren't building adaptive, nonlinear intelligence—they're simply extracting static outputs in response to fixed questions. That's not artificial intelligence. That's business intelligence with a time stamp pretending to be strategic.

Teradata's agentic strategy redefines this competitive edge, shifting from traditional data platforms to dynamic signal activation. By embedding signal detection, model execution, and autonomous agents across data and application layers, organizations can continuously sense, decide, and act at scale. The integration of natural language interfaces with MCP completes this loop, democratizing intelligence activation by enabling business users, analysts, and operators to turn conversational intent into machine-executable outcomes. This isn't just a new architecture—it's the realization of enterprise intelligence as a unified system, transforming how organizations create value through autonomous, continuous intelligence activation.

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

teradata.