

STRATOLA ANALYST REPORT

March 2026

VOLT ACTIVE DATA

Why AI Fails without Decision Products

Why AI Fails Without Real-Time Data Products and How Volt Active Data Fills the Critical Gap

Dinesh Chandrasekhar

Chief Analyst

Stratola

Executive Summary

At the Gartner Data and Analytics Summit in March 2026, Actian's CTO and CMO presented a session in which they advanced the thesis that AI fails without data products and data contracts. The statement is not rhetorical. Research from Actian's 2026 survey of enterprise data leaders provides hard evidence that organizations with mature data product practices are scaling AI at rates that far outpace those of their peers. This report examines the evidence, situates it within the broader enterprise AI landscape, and makes the case that a new category of data product, which we at Stratola call the Decision Product, is the missing data layer that separates AI projects from AI at scale.

AI agents and agentic applications should not be working with old batch data but require the most recent, real-time data for critical decision-making scenarios. The concept of "Data Immediacy" is something that Stratola has been advocating for the last couple of years and it cannot be more pertinent than now when agents are being deployed for true enterprise scenarios.

Volt Active Data, with its sub-10 millisecond ACID-compliant in-memory architecture, is positioned to become the defining platform for Decision Products in the enterprise. This report draws on the Actian 2026 survey data, Deloitte's State of AI in the Enterprise 2026, IDC projections, and Stratola's own analysis of the real-time data market to build that case.



Source for statistics in this summary: 2026 Actian Survey of Enterprise Data Leaders.

The Data Products Landscape Has Crossed a Tipping Point

The concept of the data product has circulated in enterprise data circles since the data mesh principles were introduced in 2019. For years, the conversation remained largely theoretical. Teams debated definitions, architects explored frameworks, and vendors built catalogs. Production deployments were rare. That era is now over.

According to the 2026 Actian Survey, 69 percent of organizations now deploy data products in active operational use, up 21 percentage points from 48 percent in 2024. That rate of acceleration is notable because it occurred across a 13-month window. This indicates an inflection point.

Among organizations that self-identify as data leaders, the adoption rate is even sharper: 93 percent use data products. The question for every CDO today is how fast to move and where to focus within their data strategy.

The motivation for that adoption is instructive. When survey respondents were asked why their organizations implement data products, two answers dominated. Fifty-nine percent cited the need for trustworthy inputs for business decisions. Sixty percent cited the need for trustworthy inputs for AI. Trust is the organizing principle, not speed, not volume, and not cost reduction. Enterprises want to know that the data powering decisions and AI models is accurate, governed, and consistently available.

What Constitutes a Data Product

A data product is a self-contained, productized unit of data that packages data, metadata, business logic, and access interfaces into a governed, consumable asset. Stratola's framework characterizes data products across four maturity tiers. Source-aligned products capture raw event streams and CDC feeds. Curated products deliver entity-resolved and enriched datasets. Consumer-aligned products serve purpose-built analytics use cases. Decision Products, the fourth tier, deliver

real-time datasets with sub-10 milliseconds decisions, enabling data immediacy and supporting use cases where latency is measured in milliseconds and errors carry real financial or operational consequences.

The first three tiers have attracted significant tooling investment. Platforms from Databricks, Snowflake, Microsoft Fabric, and Atlan address discovery, curation, and consumption at analytical speeds. The fourth tier, Decision Products, has remained largely unaddressed by the mainstream data platform market. That is the opportunity.

Data Product Tiers	The Fourth Tier
<p>Source-Aligned (Bronze) CDC streams, raw events, real-time ingestion</p> <p>Curated (Silver) Entity resolution, master data, enrichment</p> <p>Consumer-Aligned (Gold) Purpose-built analytics, BI-ready datasets</p>	<p>Decision Products (Real-Time) Sub-10ms ACID-compliant actions for fraud scoring, dynamic pricing, real-time offer decisions, IoT responses, and agentic AI workflows</p>

The Real-Time Decision Gap and Why It Matters

Enterprise streaming infrastructure has matured considerably. Kafka, Confluent, AWS Kinesis, and Google Pub/Sub have made high-throughput data movement a solved problem. Organizations can move billions of events per day across their data estates. The problem is what happens at the end of that pipeline. Enterprises have solved for data in motion. Most have not solved for decisions in motion.

The dominant path for real-time data today runs through a batch or micro-batch processing layer. Even with Apache Spark Structured Streaming, latency from ingest to decision typically exceeds 100 milliseconds and often runs into seconds. For a fraud detection system, a next-best-offer engine, or a 5G network policy enforcer, that latency is the difference between a decision that matters and a decision that arrives too late.

Enterprises have solved the problem of moving data at high speed. For most, the decisions are still arriving late. The gap between streaming infrastructure and actionable intelligence is the last mile problem of the data era.

The architectural reason for this gap is structural. Today's analytical real-time stack requires at minimum five system integrations: a streaming broker, a stream processing engine, a state store, a

serving layer, and an application gateway. Each hop adds latency. Each system boundary introduces failure modes. The aggregate result is architectural complexity that produces seconds of latency where milliseconds are required.

The Need for Data Immediacy

Data immediacy, a term that Stratola introduced a couple of years ago, has slowly gained acceptance and momentum in the industry. It represents the ability for a data layer in your data stack to be able to provide the most current data that is necessary for enterprise decision making at the very moment it is needed. Most of the data stacks today are able to capture real-time data but are unable to make it readily available when needed for decision making. This negates the very purpose of real-time data ingestion in the first place.

The Infrastructure Obstacle for Agentic AI

Agentic AI sharpens this problem considerably. According to Deloitte's 2025 Emerging Technology Trends study, while 30 percent of organizations are exploring agentic options and 38 percent are piloting solutions, only 14 percent have deployment-ready systems and a mere 11 percent are actively running agentic AI in production.

The chief infrastructure obstacle cited by practitioners is the inability of traditional data pipelines to support the real-time, stateful, transactionally consistent data access that autonomous agents require.

Agentic AI systems do not operate on batch data. They observe, reason, and act within continuous data flows. A customer service agent resolving a billing dispute needs a unified customer state that is current to the millisecond. A supply chain agent reprioritizing warehouse fulfillment needs inventory data that reflects the most recent transaction, not the most recent warehouse sync. A fraud detection agent approving or declining a transaction needs a risk score computed within the window of a single API call. None of these requirements can be met by architectures built around analytical latency.

46%

CAGR projected for the AI agent market through 2030 (IDC, 2026)

11%

of enterprises have agentic AI actively running in production today (Deloitte 2025)

38%

are still in pilot phase, blocked by data infrastructure gaps (Deloitte 2025)

The gap between the volume of agentic AI experimentation and the volume of agentic AI in production is a data infrastructure problem. Organizations can build agents. They cannot yet

provide those agents with the real-time, governed, transactionally consistent data foundation those agents need to operate reliably at scale.

Data Products as the Prerequisite for Scalable AI

The 2026 Actian survey data makes a striking empirical case for the relationship between data product maturity and AI outcomes. Among organizations with company-wide data product maturity, 85 percent have three or more AI projects running in production. Among organizations without data product maturity, only 25 percent achieve the same threshold. That is a 3.4x performance gap driven entirely by the presence or absence of governed, productized data.

The pattern repeats for agentic AI. Among organizations with company-wide data products, 77 percent have agentic or autonomous AI running in production. Among organizations that are still experimenting with or not yet using data products, only 23 percent have achieved the same. The 3.3x gap in agentic AI outcomes mirrors the general AI production gap almost exactly. This is not a coincidence. It is causation.

Organizations that have invested in data products are scaling AI at more than three times the rate of those that have not. Data products are more of a pre-requisite for AI and not just a supporting capability.

Understanding why this relationship holds requires examining what data products actually provide to AI systems. A data product is not simply a cleaned dataset. It carries defined ownership and accountability, built-in quality guarantees, standardized schemas and semantics, self-service discoverability, and service-level objectives for latency and availability. When an AI model or an autonomous agent consumes a data product, it does so with a contractual assurance about the reliability and freshness of that data.

Deloitte's 2026 State of AI in the Enterprise survey reinforces this from the organizational side. The research finds that AI leaders are breaking down data silos with domain-owned data products and embedding standards for quality, interoperability, and lineage across their data estates. The firms delivering transformative AI outcomes, rather than incremental productivity gains, are the ones with mature data product strategies in place before deploying their AI models.

Due to the quality assurance that data products deliver for AI, it also lowers the hallucination and improves the accuracy of AI agents. In a way, it also enables AI explainability that we can track back to understanding why a specific decision was made due to which data in the decision product.

Volt Active Data and the Decision Product Opportunity

The market for data products is expanding rapidly. The enterprise data management market is projected to reach 221 billion dollars by 2030, growing at a 12.4 percent CAGR from approximately 110 billion dollars in 2024. The streaming analytics segment, which represents Volt's core operating territory, carries an even more aggressive growth trajectory. The data-as-a-service segment, where productized data assets are bought and consumed, is projected to reach 62 billion dollars by 2030 at a 20 percent CAGR.

Within this landscape, the fourth tier of data products, Decision Products, represents a segment that is almost entirely under-served by current platform vendors. The reasons are architectural. Building a platform that combines streaming ingestion, in-memory transactional state, ACID-compliant processing, embedded decision logic, and five-nines availability in a single unified system is exceptionally difficult. Volt Active Data offers such a platform that does this at sub-10 millisecond latency with full ACID compliance.

The Volt Architecture Advantage

Volt collapses the five-layer real-time data stack discussed earlier. Streaming ingestion, in-memory state management, ACID transactions, stored procedure logic, and output serving are unified in a single platform. The result is ingest-to-action latency measured in single-digit milliseconds. The platform carries five-nines availability through patented cross-datacenter replication. It scales linearly by adding nodes. It supports embedded ML inference through Java stored procedures. It integrates with existing Kafka deployments without requiring upstream changes.

Volt's positioning is not as a database. It is as the Decision Layer that sits between your streaming infrastructure and your real-world outcomes. That reframes matters enormously for how enterprise buyers evaluate and procure the platform.

Decision Products in Practice

The Decision Product category spans several high-value use cases where Volt's architecture is most distinctively suited. In financial services, a fraud score API must return a risk score, a decision, and contributing factors within 10 milliseconds and maintain 99.99 percent uptime. The inputs are transaction events and device signals processed continuously. The output is a governed, versioned product consumed by downstream payment authorization systems.

In retail and e-commerce, a dynamic pricing Decision Product must evaluate inventory levels, competitor price feeds, and demand signals to return an optimal price per SKU within 8 milliseconds at 50,000 requests per second. In telecommunications, a 5G BSS policy enforcer must

evaluate subscriber state and usage patterns to make real-time policy decisions across millions of concurrent sessions.

Each of these use cases shares a common profile: sub-10 millisecond latency requirements, ACID consistency requirements, high availability requirements, defined schema and SLA contracts, and direct coupling to revenue or operational outcomes. They are, in every meaningful sense, data products. They are Decision Products. And they are the use cases for which Volt's architecture was built.



Stratola's Guidance for CDOs and CAIOs

Data product maturity is the single strongest predictor of AI success available to enterprise organizations today. The correlation between data product adoption and production AI outcomes holds across general AI deployments and is even sharper for agentic AI. The enterprises winning the AI race are the ones with the most trusted, most governed, most accessible data.

That insight has direct implications for the strategic positioning of every vendor operating in the data and AI infrastructure space. It also has implications for enterprise buyers who are still building their data product strategies and wondering where to prioritize.

For Enterprise Data and AI Leaders

The imperative for CDOs and Chief Data and AI Officers is to treat data product maturity as a prerequisite investment for AI scaling, not as a parallel track. Organizations that deploy AI models on top of ungoverned, unstandardized data will continue to struggle with model drift, hallucinated outputs, and AI systems that work in development but fail in production.

The prioritization question within data product strategy should include explicit attention to the Decision Product category. Analytical data products deliver value through improved reporting and model training. Decision Products deliver value through direct operational outcomes: fraud prevented, revenue captured, customers retained, risks mitigated. The return on investment is more immediate, more measurable, and more defensible to boards and executive teams.

Organizations that have already invested in Kafka or Confluent streaming infrastructure are particularly well positioned to move into Decision Products. The streaming layer is in place. What is missing is the unified in-memory state and decision logic layer that converts high-speed data movement into high-speed, trustworthy decisions. That is precisely the gap that Volt Active Data fills.

For the Vendor Ecosystem

The data product market is stratifying rapidly. Platform vendors such as Databricks, Snowflake, and Microsoft Fabric are consolidating their positions in the analytical data product tiers. Catalog and governance vendors such as Atlan and Collibra are building the metadata infrastructure that makes data products discoverable and trustworthy. The Decision Product tier, however, remains open. As consolidations happen across the ecosystem, data vendors must pay attention to this need and either start building this muscle on their own or partner with the likes of Volt Active Data.

Looking Ahead

The trajectory of enterprise AI adoption in 2026 and beyond will be smooth for those organizations that can build their AI infrastructure on reliable decision products. As Deloitte observed in its 2026 State of AI in the Enterprise report, the leaders who are seeing transformative AI outcomes are the ones with unified, trusted data strategies that connect operational, experiential, and external data flows into governed platforms that AI systems can rely on.

The agentic AI wave will make real-time data product infrastructure even more critical. Agents operating autonomously across enterprise workflows will require data that is not just fast, but contractually reliable, transactionally consistent, and available at the millisecond granularity those agents need. Once again, data immediacy! The organizations that build Decision Product infrastructure today will be the ones whose agentic AI systems scale in 2026 and 2027.

Data products are the infrastructure requirement for enterprise AI. Decision Products are the real-time tier of that requirement.

About This Report

This report was authored by Dinesh Chandrasekhar, CEO and Chief Analyst at Stratola. Statistical data is drawn from the 2026 Actian Survey of Enterprise Data Leaders, Deloitte's State of AI in the Enterprise 2026, IDC research, and Stratola's proprietary analysis of the real-time data market. Stratola advises enterprise technology vendors and practitioners on go-to-market strategy, pricing and packaging, and technology research.

Stratola | stratola.com | March 2026