

As AI agents become embedded across enterprise workflows, the need for quality, consistency, and context-rich data will determine whether organizations can scale their AI adoption and ensure trusted, actionable actions and decisions.

The Role of Universal Semantic Layers in the Agentic Data Era

July 2026

Written by: Megha Kumar, Research VP

I. Introduction

Organizations are increasingly moving from AI pilots to production. However, achieving accelerated, scalable adoption depends on modernizing the underlying data foundation with a strong focus on data governance and access and, more importantly, ensuring the data accessed by AI tools and systems is context-rich, enabling more reliable and accurate AI outcomes. In IDC's 2026 *Future Enterprise Resiliency and Spending Survey (FERS Survey)*, 47% of respondents were focused on data modernization, 44% on data intelligence, and 38% on data governance to improve AI data readiness.

However, the journey to AI-ready data reveals significant organizational challenges. Organizations are increasingly dealing with fragmented data environments, requiring unified governance and interoperability. IDC's 2026 *Data Management QuickPoll* reveals the specific barriers organizations face on that journey, with 35% of respondents citing data security and privacy concerns, 22% data quality issues, and 21% data interoperability challenges. While organizations recognize the need for strong data foundations, many struggle to implement proper governance structures and adopt the unified approach required to support AI at scale.

As AI becomes embedded in business processes, organizations are increasingly deploying AI agents across their operations, from enterprise applications to data and analytics processes. IDC's March 2026 *FERS Survey* found that one in three organizations is leaning toward in-app agents, followed by custom and bespoke agents. These deployments aim to deliver improved business value through operational efficiencies, increased employee productivity, and revenue growth.

AI agents are fundamentally shifting how organizations engage in business intelligence and analytics. Rather than delivering insights through dashboards or periodic reports, AI agents now embed insights directly into operational workflows, enabling faster and more collaborative decision-making. IDC's October 2025 *FERS Survey* highlighted the top generative AI (GenAI) and agentic use cases being adopted, including data quality monitoring, integrating insights into

AT A GLANCE

KEY STATS

For 93% of respondents, GenAI/AI agents have increased the focus on semantic layers in business intelligence and analytics (IDC's *FERS Survey Wave 8*, October 2025).

WHAT'S IMPORTANT

Altogether, 47% of organizations say a consistent semantic layer or ontology is critical to trustworthy agentic AI (IDC's *Data Management QuickPoll*, February 2026).

Organizations are adopting semantic technologies within data and analytics environments, with 52% implementing business glossaries or data catalogs, 51% deploying semantic graphs, and 41% adopting semantic layers in data platforms. (IDC's *Data Management QuickPoll*, February 2026).

business workflows, and automated report generation. The effectiveness of these agents depends entirely on the quality, accuracy, context, and governance of the underlying data.

Semantics, which were once peripheral to analytics, have become central to AI success. For years, business intelligence and analytics (BIA) tools operated within bounded environments, where business users accessed dashboards built on known data sources, with definitions managed informally or through tribal knowledge. This approach relied on users' ability to interpret context and handle ambiguity. What was once considered "nice to have" is now a prerequisite for trustworthy agentic AI. IDC's 2026 *Data Management QuickPoll* showed that 47% of organizations view a consistent semantic layer or ontology as critical to trustworthy agentic AI. Even more significantly, 93% of respondents in IDC's October 2025 *FERS Survey* indicated that the focus on semantic layers has grown within BIA due to GenAI/AI agents, with 40% noting that semantic layers help translate terms to business language, enable self-service analytics, and standardize metrics and KPIs; nearly one-third cited their role in reducing hallucinations in AI.

However, as organizations aim to scale GenAI/AI agents beyond BIA solutions, semantic inconsistencies can impede adoption. Business terms rarely have a single universal meaning. "Revenue," "customer," or "active user" can mean different things to finance, sales, and product teams. When AI agents query data without a governed semantic layer, they risk operating on inconsistent definitions, producing outputs that are technically correct but contextually wrong. At scale, this results in a loss of trust. Organizations must consider additional compounding challenges when it comes to semantics:

- » **Fragmented definitions across systems undermine organizational alignment.** IDC's April 2026 *CIO QuickPoll* found that 36% of respondents identified inconsistent definitions across systems as a common data quality issue impacting AI initiatives. This creates a dependency on technical teams when AI agents are operating across multiple business functions, departments, and platforms.
- » **Semantic drift compounds over time.** Business logic evolves constantly as organizations restructure, launch products, or enter new markets. Without a centralized semantic layer, definitions become fragmented across dashboards, pipelines, and models. AI agents trained on stale or conflicting business contexts will degrade in accuracy over time, and identifying the source of that drift becomes increasingly difficult as agent deployments multiply.
- » **Multisystem environments multiply the problem.** Most enterprises operate across multiple cloud platforms, data warehouses, and SaaS applications. Semantic definitions maintained in one environment rarely propagate consistently to others, meaning agents operating across heterogeneous ecosystems encounter conflicting or absent context, depending on where they are querying.
- » **Governance, context, and access control cannot be separated from meaning.** Semantic context does not exist independently of who is allowed to use it. An AI agent that understands what a metric means but does not enforce who can access the underlying data creates significant compliance and security risks. Managing semantics in isolation from governance introduces dangerous gaps when agents execute actions autonomously.
- » **Human-in-the-loop validation cannot scale as a permanent fix.** Organizations today compensate for semantic inconsistency by keeping humans in the loop to validate agent outputs before taking action. This is reasonable as a transitional posture, but it caps the value of agentic AI and reintroduces the bottlenecks agents are designed to eliminate.

- » **Semantic inconsistency carries hidden costs.** As AI and agentic workloads scale, organizations must also manage the cost implications of expanded data access, model usage, and compute-intensive analytics. Without a consistent semantic layer, duplicated business logic across pipelines, dashboards, and models creates unnecessary redundancy and inaccurate contextual information. At the same time, poorly defined data forces AI agents to run broader, more expensive queries to resolve ambiguity.

These compounding challenges make it clear that organizations need a comprehensive enterprisewide approach to semantic consistency. According to IDC's February 2026 *IT Data Management Poll*, organizations are adopting semantic technologies within data and analytics environments, with 52% implementing business glossaries or data catalogs, 51% deploying semantic graphs, and 41% adopting semantic layers in data platforms. However, scaling AI agents and delivering insights across enterprise operations demands more. A universal semantic layer across the entire data estate is no longer optional but critical, enabling AI agents to reason on consistent, governed, and context-rich data.

A universal semantic layer provides a consistent business definition layer across distributed data platforms and tools, enabling consistency, powering self-service analytics, delivering context, improving data governance, and allowing AI to scale across different systems. By reducing duplicated logic, optimizing query execution, and enabling governed workload routing, a universal semantic layer shifts investment toward a more efficient architecture as AI agent deployments grow.

II. Definitions

IDC defines a semantic layer as a business representation of data, providing a common set of terms, definitions, and relationships that business users can understand and use to access and analyze data. It abstracts the technical complexity of underlying data systems, enabling the consistent and accurate interpretation of data across tools, platforms, and users.

IDC defines a universal semantic layer as an enterprisewide abstraction and governance framework that sits between raw data infrastructure and end consumers, translating physical data assets into standardized business-meaningful concepts, such as metrics, KPIs, and domain entities. It operates across multiple data platforms, tools, and ecosystems, ensuring that definitions, calculations, and business logic are authored once and consistently applied across heterogeneous and distributed data environments, whether consumed by human-facing analytics tools or AI-driven workloads.

III. Benefits

Leveraging universal semantic layer environments has several benefits, including the following:

- » **Consistent business definitions, KPIs, and governed context across tools and platforms:** Business terms, KPIs, metrics, and governed context are defined once and applied consistently across all tools, platforms, user groups, and AI tools, eliminating discrepancies that lead to conflicting reports and unreliable decisions.
- » **A reduction in hallucinations:** By grounding AI and large language model (LLM) interactions in well-defined, governed business concepts, organizations can reduce the risk of AI generating inaccurate or misleading outputs, which is a particularly important consideration as agentic AI takes on more autonomous roles.
- » **The enablement of self-service analytics capabilities:** With a shared semantic layer, business users can independently query, explore, and visualize data without relying on data engineering or IT teams, thus reducing bottlenecks and improving time to insight.

- » **Reduced vendor lock-in through interoperability:** A universal semantic layer decouples business logic from individual platforms, enabling organizations to swap or add data tools without rebuilding their analytical models or losing definitional consistency.
- » **The faster development and deployment of data products:** Semantic models can be reused across multiple data products, reducing duplicated effort and accelerating the development and deployment of new analytics capabilities.
- » **Enhanced support for AI and agentic use cases:** AI agents depend on contextual understanding to reason accurately and take appropriate actions. A universal semantic layer provides the governed, context-rich data foundation that enables agents to function reliably across diverse enterprise workflows.
- » **Increased productivity for data and analytics teams:** By providing a single source of truth for definitions and metrics, the semantic layer eliminates the need for data teams to reconcile conflicting definitions or rebuild logic from scratch repeatedly, freeing them for higher-value work.
- » **Improved scalability across hybrid and multicloud environments:** Organizations operating across hybrid and multicloud environments benefit from a semantic layer that applies consistent definitions regardless of where data resides, enabling scalable and governed analytics at enterprise scale.
- » **Better alignment between business and technical teams:** When business users and technical teams share the same vocabulary and definitions, collaboration improves, misinterpretations are reduced, and the organization can move faster to data-driven decisions.
- » **The enablement of real-time and embedded analytics within workflows:** By embedding consistent semantic context within operational systems and applications, organizations can surface insights directly in the tools and workflows where decisions are made, rather than requiring users to navigate separate BI environments.
- » **Cost efficiencies:** Eliminating duplicated semantic logic, redundant data preparation, and repeated metric reconciliation across tools and teams directly reduces operational overhead and frees up resources for strategic initiatives.
- » **Compute and model-cost optimization:** Enabling governed workload routing, query abstraction, caching, and reuse of trusted semantic definitions across analytics and AI use cases allows organizations to optimize resource consumption and reduce unnecessary compute spend as agentic workloads scale.

IV. Trends

Several trends are shaping the growing importance of universal semantic layers:

- » **A shift from dashboard-centric BI to embedded workflow-driven analytics:** Organizations are moving away from dashboard-centric BI toward insights that are surfaced directly within operational workflows and business applications. This shift demands a semantic layer that can reliably deliver consistent context and definitions wherever insights are consumed, whether in a financial application, a supply chain tool, or a conversational AI interface.

- » **The increasing importance of trust, governance, and explainability:** As AI-generated insights and agent-driven actions grow more pervasive, trust, transparency, and explainability become fundamental requirements. Organizations must be able to trace how an insight was generated and verify that it is grounded in governed, well-defined data. A universal semantic layer meets these requirements by providing a clear, auditable link between business definitions and the data that powers AI outputs.
- » **Rising demand for interoperability and open standards:** The proliferation of data tools, platforms, and ecosystems is driving demand for open, interoperable approaches to semantic modeling. Organizations are seeking semantic layer solutions that can integrate across vendor landscapes, thus avoiding proprietary lock-in and enabling consistent definitions across the entire data estate. Initiatives such as the Open Semantic Interchange (OSI) are a direct response to this need. OSI is a collaborative open source effort dedicated to standardizing how semantic metadata is exchanged across analytics, AI, and BI platforms, providing a vendor-neutral single source of truth for semantic data that any tool or agent in the stack can work from, regardless of the underlying platform or vendor.
- » **The expansion of data products and domain-specific data consumption models:** Organizations are increasingly adopting data product architectures, where domain teams take ownership of their data and expose it as governed, reusable assets. A universal semantic layer is a foundational component of this approach, ensuring that data products carry consistent business context and can be consumed reliably across the organization.
- » **The increasing need for context-aware systems that leverage ontologies and knowledge graphs:** As AI agents and analytical systems grow more sophisticated, demand increases for data environments that go beyond structured definitions and capture richer semantic relationships. Ontologies and knowledge graphs have emerged as tools for encoding business context in ways that support more intelligent reasoning, inference, and discovery by both AI systems and human users.
- » **The adoption of Model Context Protocol (MCP) for governed AI context:** MCP servers enable semantic interoperability across distributed data environments by providing a standardized protocol for agents to discover, access, and leverage semantic definitions while decoupling metadata from individual platforms. This approach enables faster insights, reduces vendor lock-in, and allows AI agents to operate reliably across heterogeneous data landscapes with a consistent governed context.

V. Vendor Profile

Through its Mosaic offering, Strategy provides a universal semantic layer designed to address organizations' complex challenges in unifying fragmented data ecosystems. Strategy via Mosaic aims to deliver the structured business context and trusted meaning that AI agents need to reason and act reliably.

Mosaic solves the universal semantic layer challenge through five foundational capabilities:

- » It delivers a semantic layer that goes beyond basic data models or catalogs, enabling organizations to define consistent business definitions, metrics, security rules, and AI-ready context to abstract technical complexity. Unlike siloed BI tool solutions or platform-specific semantic layers, Mosaic is vendor-neutral and portable, decoupling business semantics and security from underlying data storage technologies. This means organizations can leverage their preferred BI tools while maintaining a single source of truth across all platforms.

- » It is focused on AI-driven efficiency to accelerate semantic layer development. Through Mosaic Studio, an advanced AI-orchestrated modeling environment, organizations reduce data modeling time from weeks to minutes. The platform conducts a comprehensive analysis of source metadata and sample datasets, automatically identifying entities, determining semantics, mapping relationships, and defining business metrics. This AI-powered approach includes features such as automatic data relationship detection, semantic enrichment, and data preparation that requires no extract, transform, and load processing, accelerating time to value.
- » It implements a universal semantic layer architecture that serves as an abstraction layer of interconnected objects with unique definitions, providing structured knowledge representation that significantly reduces hallucination risks common in direct LLM-to-database queries. The platform protects sensitive data from exposure to LLMs while enabling AI-powered features such as auto-generated attributes, modeling suggestions, and natural language querying. Enterprise-grade data governance includes automated metadata enrichment, relationship validation, and lineage tracking for AI-generated insights, ensuring the full auditability of how AI systems derive conclusions.
- » It delivers portability and cost optimization across multicloud environments. By decoupling business semantics and security from data storage, organizations can migrate between cloud providers or database vendors without costly re-architecture. The platform connects to over 200 data sources and operates across hyperscalers and hybrid on-premises environments. Its high-performance in-memory calculation engine and intelligent caching capabilities enable sub-second response times while reducing egress charges and optimizing query costs. Organizations benefit from "cost arbitrage" by switching to cheaper database vendors while maintaining consistent business definitions across their entire data estate.
- » Its native MCP server enables governed agentic AI by providing AI agents with direct access to certified business-meaningful metrics rather than raw database schema. Governance is enforced at query-compile time, ensuring agents operate within defined boundaries and unauthorized or invalid queries return explicit, auditable errors rather than plausible but incorrect outputs. This is reinforced by Mosaic Sentinel, which monitors every agent query in real time, enforces access policies, and maintains tamper-proof audit trails across models and access points. This will allow organizations to scale agentic deployments across heterogeneous AI environments without the need to rebuild governance controls for each model or workflow.

Mosaic is delivered as a fully managed cloud service through Strategy Cloud, eliminating infrastructure overhead and providing instant accessibility once provisioned. The solution is designed with comprehensive governance controls, including row-level and object-level access control, data masking, audit logging, and flexible authentication methods, ensuring consistent policy enforcement across all access points.

Challenges

Despite the significant benefits of universal semantic layers, several challenges can constrain their adoption and effectiveness:

- » **Lack of standardized approaches to semantic modeling:** No universally accepted standard currently exists for defining, structuring, or exposing semantic layers across different data platforms and vendor ecosystems. This fragmentation makes it difficult for organizations to implement a truly universal semantic layer without significant custom integration, creating risks around long-term portability and interoperability. OSI is a positive step in this direction.

- » **Semantic sprawl across tools and workflows:** Business logic often already lives across BI tools, data platforms, spreadsheets, applications, and AI workflows, making consolidation and governance difficult. Organizations must identify and understand where existing semantic definitions reside before consolidating them into a unified layer.
- » **Difficulty quantifying cost savings:** Organizations struggle to establish clear baseline metrics for duplicated data preparation, reconciliation work, query execution, model usage, and infrastructure spend. Without visible cost baselines, it becomes challenging to demonstrate ROI and justify investments in universal semantic layer solutions, particularly when competing with other data initiatives.

VI. Conclusion

Organizations stand at a critical juncture where semantic consistency has evolved from being nice to have to an architectural imperative for trustworthy, scalable agentic AI. AI agents are increasingly embedded into business processes, shifting analytics from dashboard-centric to workflow-driven models. In this environment, the ability to provide context-rich, governed data becomes essential to reliable automation and decision-making. Universal semantic layers address the compounding challenges of fragmented definitions, semantic drift, and multisystem complexity by ensuring that AI agents operate in consistent, well-defined business contexts across heterogeneous data environments. It essentially becomes the organization's context layer. Organizations embracing a universal semantic layer position themselves to scale AI safely and efficiently while maintaining contextual integrity.

About the Analyst



Megha Kumar, Research VP, Analytics and AI

Megha Kumar is research vice president in IDC's Worldwide AI, Automation, Data, and Analytics organization and global research lead for business analytics, enterprise intelligence, and decisioning solutions. Megha's research focuses on providing insights into business analytics, enterprise performance management, and decision intelligence technology trends, adoption, and vendor strategies.

MESSAGE FROM THE SPONSOR

As AI agents scale, semantic context becomes the foundation.

Strategy Mosaic is a universal semantic layer that serves as the governed context layer between enterprise data and the AI agents, analytics tools, and applications that consume them. Built on more than three decades of enterprise semantic technology, Mosaic enables organizations to define business metrics, security policies, and data relationships once and apply them consistently across any cloud, database, or AI framework.

Mosaic includes a native Model Context Protocol server, allowing AI agents to query governed metrics directly rather than the raw database schema. Mosaic Sentinel provides real-time monitoring of every data access event, maintaining audit trails and enforcing access policies across all models and access points.

Organizations can explore how Mosaic supports trusted, scalable AI grounded in a governed business context at <https://software.strategy.com/strategymosaic>

 IDC Custom Solutions

The content in this paper was adapted from existing IDC research published on www.idc.com.

IDC Research, Inc.

One Beacon Street
Suite 33100
Boston, MA 02108, USA
T 508.872.8200
F 508.935.4015
blogs.idc.com
www.idc.com

IDC Custom Solutions produced this publication. The opinion, analysis, and research results presented herein are drawn from more detailed research and analysis that IDC independently conducted and published, unless specific vendor sponsorship is noted. IDC Custom Solutions makes IDC content available in a wide range of formats for distribution by various companies. This IDC material is licensed for external use, and in no way does the use or publication of IDC research indicate IDC's endorsement of the sponsor's or licensee's products or strategies.

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for information technology, telecommunications, and consumer technology markets. With more than 1,300 analysts worldwide, IDC offers global, regional, and local expertise on technology and industry opportunities and trends across more than 110 countries. IDC's analysis and insight help IT professionals, business executives, and the investment community make fact-based technology decisions and achieve their key business objectives.

©2026 IDC. Reproduction is forbidden unless authorized. All rights reserved. [CCPA](#)